

Ultra-High Energy Particle Detection with the Lunar Cherenkov Technique

Clancy William James
BEc, BSc(Hons)



School of Chemistry and Physics,
University of Adelaide

March 27, 2009

Contents

Contents	i
List of Figures	iii
List of Tables	vi
Abstract	ix
Disclosure	x
Acknowledgements	xi
1 Introduction	1
1.1 Outline of the Thesis	1
1.2 Cosmic Rays, Neutrinos, and the Lunar Cherenkov Technique	3
2 Simulations of the Parkes Lunar Radio Cherenkov Experiment	13
2.1 Introduction	16
2.2 Parkes Experiment	18
2.3 Simulations	19
2.4 Results	22
2.5 Conclusions	24
3 Isotropic Apertures of Past and Future Experiments	27
3.1 Introduction	29
3.2 The Lunar Cherenkov Technique	31
3.3 Description of Modelling	32
3.4 Preliminary Investigation on Secondary Muon and Tau Interactions	41
3.5 Observational Phenomenology	44
3.6 Effective Aperture of Past and Future Experiments	48
3.7 Neutrino Limits from Past Experiments	49
3.8 Flux Predictions and Future Experiments	50
3.9 UHE Cosmic Ray Detection with the SKA	56
3.10 Discussion	58

4	Directional Dependence	61
4.1	Introduction	63
4.2	Lunar Cherenkov Observations	65
4.3	Directional Aperture	68
4.4	Potential Exposure of Future Experiments	74
4.5	Conclusions	76
5	Implementation of a Thinning Algorithm in the ZHS Code	77
5.1	Introduction	79
5.2	Preliminaries	80
5.3	Thinning Electromagnetic Showers	82
5.4	Testing Thinning Methods	92
5.5	Model and Parameterisation for Coherent Cherenkov Radiation	100
5.6	Secondary Peaks	106
5.7	Conclusion	108
6	Data Analysis for LUNASKA Lunar Observations with the ATCA	111
6.1	Introduction	111
6.2	Description of the Experiment	112
6.3	Dead-time and Efficiency	120
6.4	Dispersive and Sampling Effects	121
6.5	Relative Timing Calibration with Astronomical Point Sources	128
6.6	Sensitivity Calibration	158
6.7	Results	169
6.8	Summary and Conclusion	187
A	Comparison of Simulation Results using the Same Methods	191
B	Further Computation of Secondary Muon and Tau Interactions	195
C	Data Tables from Chapter 6	199
	Bibliography	207

List of Figures

1.1	Huygens' Wavelet illustration of the Cherenkov condition.	5
1.2	Cherenkov radiation from a single particle track.	6
1.3	Coherency condition for Cherenkov radiation.	7
1.4	Comparison of radiation from hadronic and electromagnetic cascades.	10
2.1	Effective aperture of Parkes compared to other experiments.	23
2.2	Limits on a UHE neutrino flux from Parkes.	24
3.1	Small- and large-scale lunar surface roughness.	34
3.2	Cosmic ray and local surface interaction.	34
3.3	Contributions to effective aperture of secondary interactions.	42
3.4	Effects of increasing bandwidth on the effective aperture.	45
3.5	Effects of antenna diameter and pointing position on effective aperture.	47
3.6	Effective apertures to UHE neutrinos of past experiments.	47
3.7	Effective apertures to UHE neutrinos of future experiments.	48
3.8	Existing limits on a UHE neutrino flux.	50
3.9	Models of and expected limits on a UHE neutrino flux.	52
3.10	Effective aperture to UHE CR of the SKA AA.	57
4.1	Interaction geometries of neutrinos.	66
4.2	Neutrino mean free path compared to the lunar radius.	66
4.3	Effective aperture per arcmin ² of lunar disk of the Parkes experiment.	68
4.4	Coordinate definitions in the antenna-Moon system.	69
4.5	Normalised effective area of the Parkes experiment to 10 ²² eV neutrinos.	70
4.6	Combined exposure vs. UHE ν arrival direction of Parkes and Goldstone.	73
4.7	Potential exposure of future experiments.	75
5.1	Subdivision of particle tracks in ZHS.	80
5.2	A simple 'doubling' shower under different thinning methods.	85
5.3	Simulated spectrum under approximation A.	86
5.4	Quality of total excess tracklength under the standard algorithm.	94
5.5	Comparison of tracklength quality under two thinning algorithms.	95
5.6	Variation of shower parameters with thinning level.	97
5.7	Tracklength quality in extreme media.	98
5.8	Box model of shower development.	100

5.9	Comparison of spectra for 100 EeV showers.	105
5.10	Secondary peaks in 100 EeV showers.	107
6.1	Diagram of the signal path at each antenna.	115
6.2	Observed noise-cal pulse and origin-dependent predictions.	118
6.3	Trigger rates and efficiency over May 18 th .	122
6.4	Two spectra used to estimate sampling and dispersive effects.	125
6.5	A simulated coherent pulse sampled at three rates.	126
6.6	Effects of sampling rate on the detected pulse height.	126
6.7	Fraction of signal peak detected as a function of ionospheric dispersion.	127
6.8	Ionospheric influence over our observation periods.	130
6.9	Expected offsets of buffer centres for the sensitivity calibration.	133
6.10	Effects of data manipulation in the correlation program.	137
6.11	Consistency check of correlation times for the 2008 data.	140
6.12	Delay vs UT for two-fold coincident triggers during Feb. 2008.	148
6.13	Delay vs UT for two-fold coincident triggers during May 2008.	150
6.14	Delay vs UT for Three-fold coincident triggers.	151
6.15	Histograms of the change in offsets $\delta\Delta t_{ij}$ for the Feb. '08 run.	155
6.16	Histograms of the arrival directions of 'far-field' events.	156
6.17	Plan of the central area of ATCA.	158
6.18	Predicted positional reconstruction of sample on-site RFI sources.	159
6.19	Reconstructed positions of some near-field events.	160
6.20	Beam power patterns and received lunar emission.	164
6.21	Raw vs cleaned CA05 'nothingness' spectra for May 18 th .	165
6.22	Comparison of $k(\nu)$ between antenna and polarisation channel for May 18 th .	165
6.23	Comparison of calibration functions $k(\nu)$ obtained for the CA01 A channel.	166
6.24	Measured power spectra on all data channels for May 2008.	167
6.25	Measured and fitted calibration functions $k_{1A}(\nu)$ for 2008.	167
6.26	One of the four remaining candidate events from May 2007.	171
6.27	Antenna delay of Moon-consistent candidates vs all candidates.	172
6.28	A narrow-band-RFI-dominated three-fold trigger.	173
6.29	A narrow time structure event from February 27 th .	174
6.30	Best by-eye alignment of a candidate event.	175
6.31	Ratio (lunar/terrestrial origin) of candidate pulse heights.	176
6.32	Simulated signal height as a function of STEC.	177
6.33	Effective apperture of LUNASKA ATCA observations.	179
6.34	Limit on an isotropic flux of neutrinos from LUNASKA ATCA observations.	179
6.35	Fractional apertures for each polarisation orientation.	180
6.36	The effect of polarisation orientation for different neutrino spectra.	181
6.37	Effective area to UHE ν as a function of arrival direction.	182
6.38	Effective area of 2008 ATCA observations along a 'slice' profile.	183

6.39	Exposure of the 2008 LUNASKA ATCA observations.	185
6.40	Combined exposure of UHE ν detection experiments.	186
A.1	Comparison with published apertures for GLUE and Kalyazin.	192
A.2	Comparison between published and reproduced apertures for LOFAR.	193
B.1	Effect of varying y_{\min} on the total aperture and contributions thereto.	196
B.2	Fractional contributions to the effective aperture of the SKA dishes.	197

List of Tables

3.1	Lunar density profiles.	36
3.2	Shower and radio Cherenkov parameters.	39
3.3	Parameters of radio instruments used in the modelling.	40
3.4	Expected yearly neutrino event rates.	54
4.1	Directional properties of past lunar Cherenkov experiments.	71
5.1	Definitions of variables describing excess tracklength.	92
5.2	Variation in shower parameters under different thinning levels.	99
5.3	Fitted shower parameters.	103
5.4	Parameters fitting \bar{k}_L .	103
5.5	Parameters fitting σ_{k_L} .	104
6.1	ATCA antenna configurations.	113
6.2	Maximum trigger rates vs buffer length for 2008.	120
6.3	Efficiency of the 2008 observations.	122
6.4	Calculation of mean experimental efficiency for May 2007.	122
6.5	Interpolated VTEC values during the observations.	123
6.6	Estimated fractions of the peak signal detected during ATCA observations.	130
6.7	Description of symbols.	131
6.8	Cable delays and time offsets for the Feb. '08 observations.	133
6.9	Start times and duration for Feb. 2008 observations.	142
6.10	Start times and duration for May 2008 observations.	143
6.11	Number of two-fold and three-fold coincidences for February 2008.	145
6.12	Number of two-fold and three-fold coincidences for May 2008.	146
6.13	Calibration information gleaned from Figs. 6.12, 6.13, 6.14 only.	152
6.14	Alternative trial offsets for a pulse search.	160
6.15	Parameters for a fit to the beam shape.	163
6.16	Effective RMS electric field strengths per unit bandwidth.	169
6.17	Experimental exposures to UHE neutrinos from Cen A and Sgr A.	186
C.1	Calibration times for the 2008 observations.	200
C.2	The raw hexadecimal input into the trigger levels.	201
C.3	Trigger levels (s.u.) for all three observation runs.	202
C.4	Buffer RMS values over the entire bandwidth.	203

C.5	Buffer RMS values over the nominal bandwidth.	203
C.6	Calibration data for the observation periods in February and May 2008.	204
C.7	Fitted values of $k(\nu)$.	205
C.8	Effective thresholds over all data channels for all observation periods.	206

Abstract

The lunar Cherenkov technique is a promising method to resolve the mystery of the origin of the highest energy particles in nature, the ultra-high energy (UHE) cosmic rays. By pointing Earth-based radio-telescopes at the Moon to look for the characteristic nanosecond pulses of radio-waves produced when a UHE particle interacts in the Moon's outer layers, either the cosmic rays (CR) themselves, or their elusive counterparts, the UHE neutrinos, may be detected. The LUNASKA collaboration aims to develop both the theory and practice of the lunar Cherenkov technique in order to utilise the full sensitivity of the next generation of giant radio telescope arrays in searching for these extreme particles. My PhD project, undertaken as part of the collaboration, explores three key aspects of the technique.

In the first three chapters, I describe a Monte Carlo simulation I wrote to model the full range of lunar Cherenkov experiments. Using the code, I proceed to calculate the aperture to, and resulting limits on, a UHE neutrino flux from the Parkes lunar Cherenkov experiment, and to highlight a pre-existing discrepancy between existing simulation programs. An expanded version of the simulation is then used to determine the sensitivity of past and future lunar Cherenkov experiments to UHE neutrinos, and also the expected event rates for a range of models of UHE CR production. Limits on the aperture of the Square Kilometre Array (SKA) to UHE CR are also calculated. The directional dependence of both the instantaneous sensitivity and time-integrated exposure of the aforementioned experiments is also calculated. Combined, these results point the way towards an optimal way utilisation of a giant radio-array such as the SKA in detecting UHE particles.

The next section describes my work towards developing accurate parameterisations of the coherent Cherenkov radiation produced by UHE showers as expected in the lunar regolith. I describe a 'thinning' algorithm which was implemented into a pre-existing electromagnetic shower code, and the extensive measures taken to check its veracity. Using the code, a new parameterisation for radiation from electromagnetic showers is developed, accurate for the first time up to UHE energies. The existence of secondary peaks in the radiation spectrum is predicted, and their significance for detection experiments discussed.

Finally, I present the data analysis from three runs of LUNASKA's on-going observation program at the Australia Telescope Compact Array (ATCA). The unusual nature of the experiment required both new methods and hardware to be developed, and I focus on the timing and sensitivity calibrations. The loss of sensitivity from finite-sampling of the electric field is modelled for the first time. Timing and dispersive constraints are used to determine that no pulses of lunar origin were detected, and I use my simulation software to calculate limits on an UHE neutrino flux from the experiment.

Disclosure

This work contains no material which has been accepted for the award of any other degree or diploma in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text.

I give consent to this copy of my thesis, when deposited in the University Library, being made available for loan and photocopying, subject to the provisions of the Copyright Act 1968.

The author acknowledges that copyright of published works contained within this thesis (as listed below) resides with the copyright holders of those works.

C.W. James, R.M. Crocker, R.D. Ekers, T.H. Hankins, J.D. O'Sullivan, R.J. Protheroe, *Limit on ultrahigh energy neutrino flux from the Parkes Lunar Radio Cherenkov experiment*, Monthly Notices of the Royal Astronomical Society, **379** (2007) 1037-1041. ©2007 The Authors. Journal compilation ©2007 RAS.

C.W. James, R.J. Protheroe, *The sensitivity of the next generation of lunar Cherenkov observations to UHE neutrinos and cosmic rays*, Astroparticle Physics **30** (2009) 318-332. ©2008 Elsevier B.V.

NOTE:

Statements of authorship appear in the print copy of the thesis held in the University of Adelaide Library.

Acknowledgements

First thanks goes to my supervisor, Ray, for sticking with me over both an Honours and a PhD project – and thank you most of all for your confidence and trust.

To my friend Jaime – it’s been fantastic working with you, even if the only neutrinos we saw came from Bimbadgen Estate’s ’06 merlot.

Thanks to Tanja and Roland, who provided both good advice and good friendship; also to my office-mate Victor, who kept me sane by having nothing to do with this work what-so-ever.

To my fellow lunatics, Rebecca, Chris, Paul, and Justin: I enjoyed working with you all, even when things went wrong at 03:00 am.

Throughout my PhD, I have been the fortunate recipient of a Ferry Scholarship, and I would like to thank both the Trustees for granting the award, and especially the late Cedric Ferry for bequeathing it.

Thanks to my parents, Chris and Bob, for your love and support.

To the irresistible force of Ron: your expertise, supervision, and general life-coaching have been invaluable. Thanks for corrupting me into a radio astronomer – I hope you’re happy!

Finally, to everyone I’ve met and worked with: the crew at USC, the radio types at the ATNF, the (mostly now former) Swinburne people, and everyone who at various times has graced level 4 of the Oliphant Wing, Physics Building, University of Adelaide: thank you all. I’ve enjoyed every minute.

Introduction

1.1 Outline of the Thesis

My PhD Thesis – "Ultra-High Energy Particle Detection with the Lunar Cherenkov Technique" – deals with three aspects of the ultra-high energy (UHE; $\gtrsim 10^{18}$ eV) particle (cosmic ray and neutrino) detection technique known as the lunar Cherenkov technique, in which Earth-based radio-telescopes are pointed at the Moon to search for the characteristic short-duration pulses of radio-wave radiation generated by UHE particle interactions in the Moon's outer layers. Loosely, these aspects can be described as simulations of the entire detection process, simulations of UHE particle cascades in particular, and experimental work at the Australia Telescope Compact Array (ATCA). The work was conducted as part on the LUNASKA collaboration (Lunar UHE Neutrino Astrophysics with the Square Kilometre Array), and deals with virtually all aspects of the collaboration's activities to a greater or lesser degree.

The Thesis is organised as follows. This chapter is intended to give sufficient introduction that subsequent chapters – which present my original work – can be understood. Firstly, I outline the motivation for a project such as LUNASKA, which boils down to the unknown origin of the highest-energy cosmic rays. I describe the link between cosmic rays and neutrinos at the highest energies, and the physical mechanism – the 'Askaryan effect' – by which they may be detected. The method and history of lunar Cherenkov technique is described, and its current status and near-future prospects are outlined.

The next three chapters – Chapters 2, 3, and 4 – are each articles (the first two published, the second submitted) dealing with my simulations of the lunar Cherenkov technique. These works are all based on extensions of my Honours project, in which the fundamentals of the simulation were developed. Chapter 2 presents my calculations of the sensitivity of the Parkes lunar Cherenkov experiment, which was the first to utilise the technique. It is a short paper, and the limits from the experiment were neither expected nor found to be significant at the time of publication. Beyond completeness however, it does contain several useful results. Firstly, that the sensitivity obtained when pointing the antenna at the lunar limb is very high, and secondly, the paper discusses for the first time an apparent discrepancy

between pre-existing simulation results. Chapter 3 is based upon essentially the same simulation as in Chapter 2, though with some significant modifications. The chapter describes the simulation techniques used in detail, and calculates the sensitivity of a wide range of experiments – including a likely model of the SKA, the Square Kilometre Array – to an isotropic flux of UHE neutrinos. It also, for the first time, analyses the effects of observing over a very broad bandwidth, and also estimates the effects on the sensitivity of the interaction between cosmic rays and large-scale lunar surface features. Importantly, the sensitivity estimates contained therein indicate that the SKA could be an extremely powerful instrument for UHE cosmic-ray and neutrino detection. Chapter 4 takes the results on UHE particle detection from Chapter 3 and asks the question: ‘How do we turn UHE particle detection into UHE particle astronomy?’. It thus examines the directional-dependence of the lunar Cherenkov technique, both in terms of the instantaneous sensitivity, and the integrated exposure over an observation period. The key results indicate to which areas of the sky are current limits strongest in the $> 10^{20}$ eV energy range, which regions will be accessible in the near future, and how to choose observations times and the antenna pointing position in order to maximise the sensitivity to particular regions of the sky.

In Chapter 5, I describe a joint investigation between myself and J. Álvarez-Muñiz into the use of ‘thinning’ techniques to extend the energies at which full simulations of electromagnetic showers are possible into the EeV range, as required to model experiments such as those described in Chapters 2–4. This chapter is also based heavily on an article – Ref. (18) – in the process of being submitted. However, I am not the primary author, and thus it only includes sections where I made a significant contribution. Using thinning techniques, we modify an existing code to allow us to develop updated parameterisations of the radiation from UHE electromagnetic showers, which for the first time includes a method to simulate the variation in shower length introduced by LPM fluctuations.

Finally, Chapter 6 describes the data analysis for the LUNASKA lunar observations at the Australia Telescope Compact Array (ATCA). Both because this chapter is the only one not based on an article, and because it must serve as a record for the rest of the collaboration, it is by far the longest chapter of the thesis, despite much of the technical data being placed in an Appendix. The chapter described our unique experimental set-up, and the problems of dedispersion, sampling, triggering, sensitivity, and time-alignment are dealt with in detail. For the first time, the effects of dispersion over a broad bandwidth, and the effects of a non-infinite sampling rate, are dealt with quantitatively. The two key problems, that of determining the experimental sensitivity, and aligning the times at each of the antennas, are solved with substantial detail, using unique methods developed specifically for this experiment. A pulse search is described, which eliminates all candidate pulses. My simulation (described mostly in Chapter 3) is used to determine the instantaneous aperture and integrated limit from the experiment, both to an isotropic flux, and to neutrinos originating from the nearby radio-galaxy Centaurus A and the galactic centre. I find that while the sensitivity to an isotropic flux is not the most limiting over the whole range of neutrino-hunting experiments, both the instantaneous aperture and limits are the most sensitive of all lunar

Cherenkov experiments in the $\lesssim 3 \times 10^{22}$ eV range.

Finally, the three Appendices contain respectively a comparison of my simulation results to those from other experiments, further calculations applicable to Chapter 3, and data tables from the experimental analysis of Chapter 6.

Together, this work represents significant progress towards the stated aims of the LUNASKA collaboration, that of developing both the theory and practice of the lunar Cherenkov technique. Where the results are not conclusive, they highlight key areas for further treatment; where they are conclusive, they provide insight into how a giant radio array such as the SKA might be utilised in the search for the highest energy neutrinos and cosmic rays. Most importantly, the results are promising, and suggest that investigations such as the LUNASKA project should continue.

1.2 Cosmic Rays, Neutrinos, and the Lunar Cherenkov Technique

The UHECR mystery

In their 2003 report ‘Connecting quarks with the cosmos’, the Turner committee named "How do cosmic accelerators work and what are they accelerating?" as one of the Eleven Science Questions for the New Century (114). A key motivator behind this was the mystery of the ultra-high energy (UHE; $E \gtrsim 10^{18}$ eV) cosmic rays (CR): microscopic particles – thought to be primarily protons and atomic nuclei – with macroscopic energies of ten joules or more. Impacting the Earth at a rate one per square kilometre per year, the UHE CR are the highest-energy particles in nature. Their source remains a mystery.

The origin of the UHE CR remains obfuscated because cosmic rays are primarily charged particles, so that, unlike photons, their paths are bent in the magnetic fields pervading the universe. While more relativistic particles are deflected less, only recently has an experiment – the Pierre Auger Observatory (PAO) – shown that at the very highest energies (5.6×10^{19} eV and above) do UHE CR retain some directional information (1). However, at these very highest energies, the flux of UHE CR is even lower, at approximately 1 per km^2 per century. Thus the 3000 km^2 PAO had less than 30 events from which to extract information on the arrival direction, and only a correlation with the broad-scale distribution of galaxies in the nearby universe was observed.

One logical way to proceed is thus to build bigger UHE CR detectors. The proposed Northern site for the PAO observatory is planned to cover $20,000 \text{ km}^2$ (34) and target UHE CR above 10^{19} eV, while the EUSO experiment – planned to be installed onboard the International Space Station – will have an instantaneous coverage of order $100,000 \text{ km}^2$. It remains to be seen whether improved statistics will be able to pinpoint the source(s) within this distribution, or whether the deflection is such that this will remain impossible even with the vastly improved statistics offered by these planned experiments.

UHE neutrinos

Another method to determine the origin of the UHE CR are observations of the as-yet unseen UHE neutrinos. During propagation, cosmic rays will interact with background radiation fields, for example via pion photoproduction on the CMBR: $p_{\text{CR}} + \gamma_{\text{CMBR}} \rightarrow \pi^+ + n$, as predicted by Greisen (52) and Zatsepin and Kuzmin (124). Inevitably, the decay of interaction products will produce neutrinos, for instance $\pi^+ \rightarrow \bar{\mu} + \nu_{\mu}$ and subsequently $\bar{\mu} \rightarrow e^+ + \nu_e + \bar{\nu}_{\mu}$ (27). Some models of UHE CR production also predict a flux of UHE ν from UHE CR production sites – therefore, some flux of UHE ν is almost guaranteed to accompany UHE CR.

Observations of UHE ν have been proposed (45; 91) as holding the key to the origin of UHE CR because they are both uncharged and weakly interacting: they will travel in straight lines, and the source flux will be observed directly. The observed spectrum is expected to be sensitive to the UHE CR acceleration mechanism (90), and their arrival directions could point back to their source. More generally, observations of UHE ν would open up a new window on the highest-energy astrophysical processes, and the potential discovery of entirely new phenomena.

The Askaryan Effect

While UHE ν observations are very attractive, they are also very difficult. Though the flux may be more than an order of magnitude above that for UHE CR, an extremely low interaction rate (penetrating of order 200 km of water at 10^{20} eV) means that a massive detector volume is required. Directly implementing such a volume with (say) scintillating material – as used to detect cosmic ray μ secondaries – would be prohibitively expensive, and so remote detection methods have been proposed. Experiments use either the optical (e.g. Ice-Cube (133), Nemo (132), and ANTARES (131)), acoustic (e.g. the AMADEUS sub-array in ANTARES), and/or radio (e.g. RICE (71; 72), ANITA (134)) signatures expected to be produced by the cascades of secondary particles resulting from a primary UHE neutrino interactions. Of the three, the radio technique is the most attractive at the very highest energies.

The production of radio emission from UHE particle showers is known as the Askaryan Effect, after G. A. Askaryan who predicted the mechanism in a series of papers in the 1960s (22). A high-energy particle interaction in a non-tenuous medium (i.e. not in space) will rapidly evolve into a cascade of numerous secondary particles, so that instantaneously the shower appears as a ‘pancake’ of particles moving at very close to c , the speed of light in a vacuum – and almost certainly faster than the speed of light in the interaction medium c/n , where n is the refractive index. This will result in the emission of Cherenkov radiation – the electromagnetic analogue of a sonic boom – at the characteristic Cherenkov angle $\theta_C = \cos^{-1}(1/n)$ for highly-relativistic $v \approx c$ particles. Askaryan’s effect describes how some of the secondary particles will collide with electrons in the interaction medium, knocking them into the shower, while positrons in the cascade may annihilate with electrons in the medium.




Figure 1.1: Huygens' Wavelet illustration of the Cherenkov condition.

The gain of electrons and loss of positrons in the cascade causes a charge imbalance, so that at wavelengths comparable to the instantaneous size of the pancake, the cascade appears as one large concentration of negative charge. At these wavelengths, the Cherenkov radiation from each individual charge will add coherently, with radiated power scaling as the square of the number of excess electrons. In a dense ($\rho \sim 1 \text{ g/cm}^3$) medium, the instantaneous size of the shower will be a few tens of cm, so that the coherent part of the radiated spectrum extends into the GHz (microwave-radio) regime. If the medium is comparatively transparent to radio waves, this radiation – which manifests itself as a short-duration pulse – can escape and be detected remotely. Since the electron excess in a high-energy cascade is proportional to the number of shower particles, which in turn is proportional to the primary particle energy, the power in the coherent Cherenkov radiation scales with the square of the primary energy, so that very energetic showers produce extremely energetic pulses which can be detected at large distances. Viewing a dense, radio-transparent medium with a radio antenna thus allows huge volumes of material to be utilised as UHE particle detectors.

1.2.1 Cherenkov radiation

Radiation from a single particle track

Cherenkov radiation can best be understood using a Huygens' Wavelet diagram such as that in Fig. 1.1. A charged particle travelling to the left at velocity v travels a distance $d = vt$ in the same time as radiation from its original position travels ct/n , where n is the refractive index of the material. Provided $v > c/n$, the wavelets add up along a wavefront propagating at an angle $\theta = \cos^{-1}(c/vn)$. When $v \approx c$, the angle is the characteristic Cherenkov angle of the medium, $\theta_C = \cos^{-1}(1/n)$. Note that the refractive index n will in general be frequency-dependent – unless otherwise noted I use that applicable at radio-frequencies, where it is approximately constant over the frequency ranges of all experiments considered here.

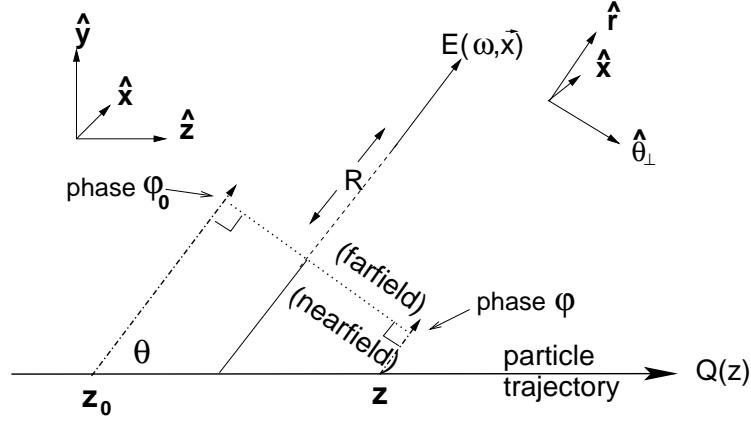


Figure 1.2: Cherenkov radiation of electric field strength $E(\omega, \vec{x})$ for a single particle track, as seen by an observer in the farfield at distance R and angle θ to the z (particle) axis. The charge varies in general as $Q(z)$ from an initial position z_0 .

A derivation based on Maxwell's equations (see Ref. (58)) leads to the formula for the radiation electric field strength from a single linear track from a particle of charge q given in Eq. (1.1), where the geometry is defined in Fig. 1.2. The definitive paper on the radiation produced from a charged-particle shower is that of Zas, Halzen, and Stanev (1992) (123); Eq. (1.1) can be made equivalent to Eq. (12) from that work by setting $\vec{v} = c\hat{z}$ and $t_1 = 0$.

$$E(x, \omega) = \frac{i\omega\mu}{2\pi} \frac{e^{ikR}}{R} \hat{\theta}_\perp \int \sin\theta dz q e^{ik(1/n - \cos\theta)z} \quad (1.1)$$

Though apparently for a simplistic case, Eq. (1.1) contains most of the essential properties of Cherenkov radiation from a high-energy cascade. The phase factor $ik(1/n - \cos\theta)z$ comes from the phase ϕ from the particle at z being advanced by $k/n(z - z_0)$ relative to ϕ_0 due to the radiation being emitted later, but retarded $k(z - z_0)\cos\theta$ due to having less distance to travel. The well-known Cherenkov condition for a particle travelling at $v = c$ can be derived by setting this phase factor $(1/n - \cos\theta)$ equal to zero. The linear dependence of the emission on the frequency ν (here I use the angular frequency $\omega = 2\pi\nu$) in the coherent regime is also present. The $\sin\theta dz$ factor is most usefully interpreted as stating that the peak field strength at the Cherenkov angle is proportional to the total tracklength as seen by the observer – the ‘total projected tracklength’. The polarisation is in the $\hat{\theta}_\perp$ direction, i.e. the radiation is linearly polarised in the particle-track/observer plane. Finally, the field strength is proportional to the charge q , so the total power radiated goes as q^2 .

At higher frequencies and longer tracklengths, the phase factor becomes non-negligible for smaller deviations away from the Cherenkov angle, and for tracklengths of only a few wavelengths, the radiation is emitted in a ‘cone’ about an angle θ_C to the particle track. The cone width decreases with frequency: at very low frequencies where the wavelength comparable to the tracklength, the radiation is almost isotropic (bar the $\sin\theta$ factor), while at very




Figure 1.3: Coherent and incoherent conditions for Cherenkov radiation, and the scaling of the power at the Cherenkov angle.

high frequencies (i.e. visible light) where the tracklength might be a near-infinite number of wavelengths, the emission forms a δ -function at the Cherenkov angle. The reduction in field strength away from θ_C due to the changing phase-factor represents destructive interference along the length of the particle track, and is termed ‘longitudinal decoherence’.

Other results pertaining to Cherenkov radiation are not relevant to a calculation of the coherent emission from a high-energy cascade. For a derivation of the radiated power per frequency/solid angle, and other expressions similar to the above formula, refer to Ref. (76). Quantum effects cause Eq. (1.1) to break down when the wavelength becomes comparable to the particle localisation, so that as $\omega \rightarrow \text{inf}$ the field strength remains finite. For radiation from a single particle track, the peak power is emitted in the UV part of the spectrum, though the rising power with frequency causes Cherenkov light (e.g. from radioactive particles in a nuclear reactor’s coolant water) to appear blue to the human eye. Coherent Cherenkov radiation however has completely different effects governing the peak emission, which is the subject of the next section.

Coherent Cherenkov radiation from high energy cascades

Coherent Cherenkov radiation occurs when the emission from multiple particle tracks adds coherently. Intuitively, this will occur when the distance between the tracks is small compared to the wavelength in question. In the case of a high-energy particle cascade, the relevant distance is the instantaneous size of the shower, which will look like a thin pancake travelling with velocity $v \approx c$. As shown in Fig. 1.3, for small wavelengths, the emission is incoherent, and the power p is proportional to the sum of the square of the charges. For long wavelengths, the emission is coherent from all the charges, and the power is proportional to the sum of the charge squared.

The largest instantaneous dimension of a shower is its width, the characteristic size of which is the Moliere radius R_M , which is of order 10 cm in media of comparable density to water. Therefore, in such media, emission can be expected to remain coherent into the GHz

regime. The charge excess predicted by Askaryan has since been determined by experiment and simulation to be of order 10% of the particles in an electromagnetic shower (or 20% more electrons than positrons), while the number of particles in a 1 EeV shower might be of order 10^{11} . The coherent part of the radiation therefore has power proportional to $(10\% \times 10^{11})^2$ – which is a very big number indeed. It is this quadratic scaling of radiated power with energy in the coherent regime that suggests that at some appropriately high energy (which turns out to be around 10^{16} - 10^{17} eV), the low (radio-microwave) frequency emission from a particle cascade surpasses the high (optical-UV) frequency emission in strength.

The coherent Cherenkov emission from a particle cascade has a similar form to that from a single particle track: the total length of the shower replaces that for a single particle track, and the excess charge $Q(z)$ replaces the unit charge q . Two new factors are introduced into the formula, the first being the large increase in peak field strength, which for a shower of energy E_s will appear as a constant factor $C E_s$ multiplying Eq. (1.1). The constant C will mostly represent the number of excess electrons per unit energy of E_s , altered slightly to reflect effects such as some particles travelling with $v < c$. The most important term however is a ‘lateral decoherence factor’, to reflect interference over the width of the shower. This factor – which is usually expressed as a multiplier to Eq. (1.1) – must by definition be 1 for $\lambda \gg R_M$, and ~ 0 for $\lambda \ll R_M$, and it provides the high-frequency cut-off to the coherent part of the spectrum.

1.2.2 Interaction phenomenology and the radiated spectrum

Any high-energy particle interacting in a dense medium will initiate a cascade of secondaries leading to coherent Cherenkov radiation. There are two types of cascade processes relevant to high-energy showers in dense media – electromagnetic cascades, and hadronic cascades. Electromagnetic cascades consist of γ and e^\pm , progressing via bremsstrahlung ($e^\pm \rightarrow e^\pm + \gamma$) and pair production ($\gamma \rightarrow e^+ + e^-$). Hadronic cascades propagate by shower hadrons interacting with atomic nuclei, and consist primarily of both charged and neutral pions, and also more exotic hadrons. It is the electromagnetic component which, at MeV energies, gives rise to the charge excess and results in coherent Cherenkov radiation.

Cosmic rays will always initiate a hadronic cascade. Neutrinos will also, interacting either via a ‘neutral current’ (transfer of momentum to a nuclear quark via a Z^0) or a ‘charged current’ (transfer of momentum to a nuclear quark via a W^\pm). In each case the quark will recoil with high momentum and initiate a hadronic cascade. A charged-current interaction will also result in a leptonic product corresponding to the neutrino type, which can lose energy via (for instance) pair-production and bremsstrahlung. Only in the case of an e^\pm from a charged-current $\nu_e/\bar{\nu}_e$ interaction will this energy loss will be rapid enough to consider the lepton as initiating a single electromagnetic cascade, rather than a series of widely-spaced sub-showers of mixed hadronic and electromagnetic type. Therefore only a fraction of UHE neutrino interactions, and no UHE CR interactions, will begin purely cascades with a 100%

electromagnetic component (primary γ will too, though I do not consider these here), and even these will coincide with a shower of initial hadronic component.

Depending on the density of the medium and the energy of the particles, different mechanisms exist to transfer energy from a hadronic cascade into an electromagnetic component. The most relevant for a dense medium is the decay of neutral pions: $\pi^0 \rightarrow \gamma + \gamma$. Although the more numerous charged pions do not decay to produce electromagnetic products, due to time dilation and a longer lifetime they decay only at much lower energy scales than their neutral counterparts, so that in each generation more energy in the hadronic component is converted into an electromagnetic component. For ultra-high energy primaries, most of the energy ends up in the electromagnetic component, so that the amount of charge excess from all cosmic-ray and neutrino initiated cascades is expected to be similar for a given cascade energy. See Ref. (9) for a further description of this mechanism. Note that at the highest energies, both photons and electrons can interact with medium nucleons to introduce a hadronic component into an initially purely electromagnetic shower – however, the cross-section is uncertain, and I ignore this effect here.

While all UHE particle interactions are expected to lead to coherent Cherenkov radiation, the distribution of this radiation depends on three key shower properties. Loosely, each shower can be described as developing inside a ‘box’, with characteristic length L , width W , and total excess tracklength T , each of which governs key properties of the radiation. The width of a shower of particles – the Moliere radius – is largely governed by low-energy processes, since only then is the longitudinal momentum small enough to allow any significant angular deviation from the direction of the primary particle. Therefore, for a given medium, the width (and hence high-frequency cut-off) is almost independent of the nature and energy of the original particle. A similar argument applies to the excess tracklength, though here this is proportional to the primary particle energy – both are discussed in more detail in Chapter 5.

It is therefore the shower length L which differentiates the two types of cascades. The longer the shower, the more significant the longitudinal decoherence term, and the thinner the width of the Cherenkov cone. Although all showers increase in length with primary energy, the growth for hadronic showers is approximately logarithmic, since each generation of particles will develop over a similar distance (9). For high-energy electromagnetic showers however, the Landau-Pomeranchuk-Migdal (LPM) effect leads to a much more rapid growth in length with energy (73; 80), approximately as \sqrt{E} . The LPM effect occurs when the formation length for bremsstrahlung photons (there is a similar effect for a pair-producing γ) becomes sufficiently long that there is a significant probability of a second interaction occurring over this length, suppressing the emission. The LPM effect operates only at very high energies, though the denser the medium, the lower the energy (the LPM energy, E_{LPM} , is typically of order 1 PeV in dense media) at which it becomes significant. Above this energy, the cross-sections for the interactions propagating an electromagnetic cascade reduces with energy, leading to an extremely long shower. As a result, the width of the Cherenkov cone for such showers is extremely thin, making them less detectable

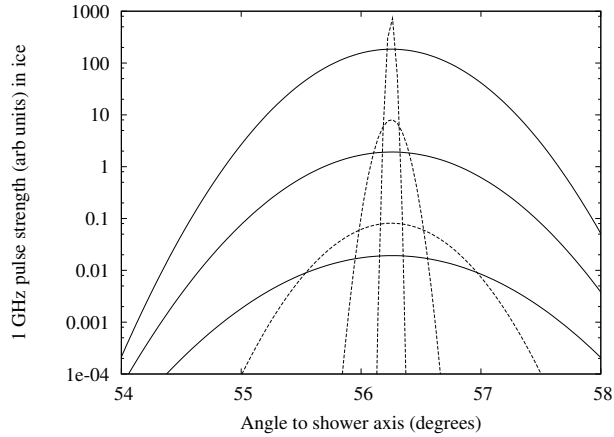


Figure 1.4: Radiated spectrum (arbitrary units) from typical hadronic (solid) and electromagnetic (dashed) cascades in ice, for primary neutrino energies of (top to bottom) $E_\nu = 10^{22}$, 10^{20} , and 10^{18} eV, as a function of angle from the shower axis. Shower energy E_S is assumed to be $0.2E_\nu$ for hadronic showers, and $0.8E_\nu$ for electromagnetic showers.

than their hadronic counterparts (11; 123). Fig. 1.4 gives approximate radiation patterns as a function of angle from the Cherenkov angle for electromagnetic (dashed) and hadronic (solid) cascades, using the parameterisations of *Álvarez-Muñiz et al.* (11; 9), with hadronic shower width extended beyond 1 EeV as per Williams (121).

1.2.3 The lunar Cherenkov technique

Utilising the radio technique requires a large volume of radio-transparent material. The primary condition for radio-transparency of a medium is dryness, and suitably large volumes of such material exist naturally on Earth. Both salt domes (101) and the permafrost of the Arctic tundra (84) have been proposed as suitable media. The most useful volume on Earth however is the Antarctic ice sheet, as utilised by both the RICE (72) and ANITA (50) experiments.

One of the original media proposed by Askaryan as suitable for the detection method was the lunar regolith, being composed of the fragmented surface layer of debris from meteorite impacts. It is also comparatively radio-transparent, and has a massive volume of approximately, $400,000 \text{ km}^3$ (conservatively) assuming a 10 m depth. Askaryan's original idea was to use radio-detectors on the Moon itself, and this idea has not entirely died (63). However, the sole technique of current experimental interest was proposed by Dagkesamanskii and Zheleznykh (37), who suggested using Earth-based radio-telescopes to observe the entire visible volume of the regolith at once. Known as the lunar Cherenkov technique, its advantages are that the detectors – radio-telescopes – already exist, and that the entire visible volume of the regolith can be observed at once. The disadvantage is that the Moon

is over one light-second away, resulting in a big hit to sensitivity from the $1/R^2$ law. Consequently, the technique is best suited to the detection of the very highest energy particles only, although this is somewhat offset by the high sensitivity of modern radio telescopes.

Another advantage of the technique is its (theoretical) ability to detect UHE CR also. UHE showers initiated by UHE CR will generally appear similar to those initiated by UHE neutrinos, though for Earth-based experiments the atmosphere provides a shield against the CR. There is no such shield on the Moon, which also provides a challenge – that of distinguishing between UHE CR and ν . However, since observations of either UHE CR and/or UHE ν promise to aid in solving the UHE CR mystery, it should be seen as a positive that the technique is sensitive to particles of both types.

Experimental history

The first attempt to detect coherent Cherenkov radiation from UHE particle interactions in the lunar regolith was made at Parkes in New South Wales in 1995 by Hankins et al. (53). Using the 64 m Parkes radio-antenna, the Moon was observed for approximately 10 hours. Five 100 MHz (centred at ~ 1.5 GHz) data-channels at each of two circular polarisations were returned to the control room, for a combined area-bandwidth-polarisation product of $3200 \text{ m}^2 \text{ GHz}$. RFI discrimination was based off the expected ionospheric time-delay between the signals in two 100 MHz sub-bands. However, only two of the ten data channels could be used in real-time to form a trigger, and then only in coincidence; also, the majority of the observation time was spent pointing at the lunar centre, since it was not realised at that time the majority of events would appear to come from the limb of the Moon. Of the order 300 candidates, all were eliminated as potential events. Nonetheless, the experiment did show that the expected delay in the ionosphere could be used to discriminate against terrestrial RFI.

The most well-known lunar Cherenkov experiment was the Goldstone Lunar UHE neutrino Experiment, or GLUE (47). Run from 2000 to 2003, GLUE utilised two antennas – of 70 m and a 32 m diameter – at NASA’s Goldstone Deep Space Tracking Station in California. Each antenna had an approximate bandwidth of 150 MHz at 2.2 GHz at a single polarisation (RCP), each of which was generally split into two 75 MHz bands. Also, the 70 m antenna returned a 40 MHz LCP band, and a global trigger required a five-fold coincidence over all bands. Three observing modes were used, which are fairly self-explanatory: limb, centre, and half-limb. For the latter two modes, the 70 m antenna was defocused so that its beam matched that of the 34 m antenna. Over 100 hours of raw observation time was accumulated - much of it on the limb – with the separation between the two antennas used to rule out terrestrial RFI. Again, no events were positively identified. Importantly, simulations were used to determine the effective aperture of the experiment to UHE ν , and the resulting limit was used to place constraints on some of the more optimistic models of UHE CR and neutrino production.

Simultaneously with GLUE, a Russian group used the 64 m antenna at Kalyazin in an

ongoing observing program, approximately 30 hours of which were reported by Beresnyak *et al.* (26) in 2006. Like Parkes, observation bands at multiple frequencies were used to eliminate RFI, with the main detector bandwidth being 128 MHz at 2.2 GHz. All the observation time was spent pointing at the limb, and again, no candidates passed all acceptance criteria.

This was the status of experiments at the beginning of my PhD in 2005. Since then, the LUNASKA collaboration has made observations at the ATCA (Australia Telescope Compact Array) (61), and a group have begun an experiment ('NuMoon') using low-frequency observations at the WSRT (Westerbork Synthesis Radio Telescope) in the Netherlands (103). Neither have detected plausible candidate pulses, although the experimental sensitivities are ahead of the aforementioned experiments. Proposals for further observations with both are in the pipeline, while there has also been interest in using three antennas at the NRAO Green Bank observatory (74); the (currently) decommissioned 25 m antenna at Stockert (65); the Low Frequency Array, or LOFAR (104); and with Australia's SKA pathfinder, ASKAP (64). The ultimate goal for all these collaborations will be to utilise the Square Kilometre Array, even though this is the stated goal of our (LUNASKA) collaboration only.

The Square Kilometre Array and LUNASKA

The Square Kilometre Array (SKA) is a giant radio-array to be built in either Southern Africa or Western Australia beginning around 2015. The name derives from the total planned collecting area of 1 km², giving it a sensitivity (collecting area over system temperature) of ten times that of the ~300 m dish at Arecibo. The planned frequency range is from 70 MHz up to at least 5 GHz (and possibly up to 40 GHz), with a large instantaneous bandwidth. For more information, see the project website www.skatelescope.org. Suffice to say that in the foreseeable future, the SKA will be the ultimate instrument for a wide range of radio-astronomy, including the lunar Cherenkov technique.

The LUNASKA project (Lunar UHE Neutrino Astrophysics with the SKA) is, in the words of R. J. Protheroe, "a theoretical and experimental project for UHE neutrino astrophysics using a giant radio observatory". The goal is to develop the theory and practice of the lunar Cherenkov technique in order to enable an instrument such as the SKA to be used at its full potential for UHE particle searches. The LUNASKA methodology was to develop simulations to determine the ideal configuration of current/future experiments and the sensitivity of past ones, while using the Australia Telescope Compact Array (ATCA) as a test-bed for a much larger array such as the SKA. My original role in the project was only on simulation development, however eventually I became very involved in the experimental side also. It is my hope therefore that the work contained in this thesis has gone some way towards realising ultra-high energy particle astronomy with the SKA.

Simulations of the Parkes Lunar Radio Cherenkov Experiment

Limit on UHE Neutrino Flux from the Parkes Lunar Radio Cherenkov Experiment

C. W. James¹, R. M. Crocker¹, R. D. Ekers², T. H. Hankins³, J. D. O'Sullivan², R. J. Protheroe¹

¹Department of Physics, School of Chemistry & Physics, University of Adelaide, Adelaide
SA 5005, Australia

²Australia Telescope National Facility, PO Box 76, Epping NSW 1710, Australia

³Physics Department, New Mexico Tech, Socorro, NM 87801, USA

Monthly Notices of the Royal Astronomical Society 379, pp 1037-1041 (2007)

STATEMENT OF AUTHORSHIP

LIMIT ON UHE NEUTRINO FLUX FROM THE PARKES LUNAR RADIO
CHERENKOV EXPERIMENT

Monthly Notices of the Royal Astronomical Society 379, pp 1037-1041 (2007)

James, C.W. (Candidate)

Wrote computer code; generated and interpreted data; wrote manuscript and produced figures; acted as corresponding author.

I hereby certify that the statement of contribution is accurate

Signed: _____ *Date:* _____

Crocker, R.M.

Manuscript editing and advice.

I hereby certify that the statement of contribution is accurate

Signed: _____ *Date:* _____

Ekers, R.D.

Conception of the work; performed original experiment.

I hereby certify that the statement of contribution is accurate

Signed: _____ *Date:* _____

Hankins, T.H.

Performed original experiment.

I hereby certify that the statement of contribution is accurate

Signed: _____ *Date:* _____

O'Sullivan, J.D.

Performed original experiment.

I hereby certify that the statement of contribution is accurate

Signed:

Date:

Protheroe, R.J.

Conception of the work; manuscript editing and advice.

I hereby certify that the statement of contribution is accurate

Signed:

Date:

Limit on UHE Neutrino Flux from the Parkes Lunar Radio Cherenkov Experiment

C. W. James¹, R. M. Crocker¹, R. D. Ekers², T. H. Hankins³, J. D. O'Sullivan², R. J. Protheroe¹

¹Department of Physics, School of Chemistry & Physics, University of Adelaide, Adelaide SA 5005, Australia

²Australia Telescope National Facility, PO Box 76, Epping NSW 1710, Australia

³Physics Department, New Mexico Tech, Socorro, NM 87801, USA

Abstract

The first search for ultra-high energy (UHE) neutrinos using a radio telescope was conducted by Hankins, Ekers & O'Sullivan (53). This was a search for nanosecond duration radio Cherenkov pulses from ultra-high energy (UHE) neutrino interactions in the lunar regolith, and was made using a broad-bandwidth receiver fitted to the Parkes radio telescope, Australia. At the time, no simulations were available to calculate the experimental sensitivity and hence convert the null result into a neutrino flux limit.

Proposed future experiments include the use of broad-bandwidth receivers, making the sensitivity achieved by the Parkes experiment highly relevant to the future prospects of this field. We have therefore calculated the effective aperture for the Parkes experiment and found that when pointing at the lunar limb, the effective aperture at all neutrino energies was superior to single-antenna, narrow-bandwidth experiments, and that the detection threshold was comparable to that of the double-antenna experiment at Goldstone. However, because only a small fraction of the observing time was spent pointing the limb, the Parkes experiment places only comparatively weak limits on the UHE neutrino flux. Future efforts should use multiple telescopes and broad-bandwidth receivers.

Keywords: neutrinos – instrumentation: detectors – telescopes.

James, C.W., Crocker, R.M., Ekers, R.D., Hankins, T.H., O'Sullivan, J.D. & Protheroe, R.J. (2007) Limit on Ultrahigh Energy Neutrino Flux from the Parkes Lunar Radio Cherenkov Experiment
Monthly Notices of the Royal Astronomical Society, v. 379 (3), pp. 1037-1041

NOTE:

This publication is included on pages 16-25 in the print copy of the thesis held in the University of Adelaide Library.

It is also available online to authorised users at:

<http://dx.doi.org/10.1111/j.1365-2966.2007.11947.x>

Isotropic Apertures of Past and Future Experiments

**THE SENSITIVITY OF THE NEXT GENERATION OF LUNAR CHERENKOV
OBSERVATIONS TO UHE NEUTRINOS AND COSMIC RAYS**

C.W. James¹, R.J. Protheroe¹

¹ Department of Physics, The University of Adelaide, Adelaide, SA 5005, Australia

Astroparticle Physics 30, 6, pp 318-332 (2009)

STATEMENT OF AUTHORSHIP**THE SENSITIVITY OF THE NEXT GENERATION OF LUNAR CHERENKOV
OBSERVATIONS TO UHE NEUTRINOS AND COSMIC RAYS**

Astroparticle Physics 30, 6, pp 318-332 (2009)

James, C.W. (Candidate)

Wrote computer code; generated and interpreted data; wrote manuscript and produced figures (except where noted below); acted as corresponding author.

I hereby certify that the statement of contribution is accurate.

Signed:

Date:

Protheroe, R.J.

Manuscript editing; produced Figs. 3.8, 3.9(a) and 3.9(b); wrote Sec. 3.8.1.

I hereby certify that the statement of contribution is accurate.

Signed:

Date:

The sensitivity of the next generation of lunar Cherenkov observations to UHE neutrinos and cosmic rays

C.W. James¹, R.J. Protheroe²

Department of Physics, The University of Adelaide, Adelaide, SA 5005, Australia

Abstract

We present simulation results for the detection of ultra-high energy (UHE) cosmic ray (CR) and neutrino interactions in the Moon by radio-telescopes. We simulate the expected radio signal at Earth from such interactions, expanding on previous work to include interactions in the sub-regolith layer for single dish and multiple telescope systems. For previous experiments at Parkes, Goldstone (GLUE), and Kalyazin we recalculate the sensitivity to an isotropic flux of UHE neutrinos. We find the published sensitivity for the GLUE experiment to be too high (too optimistic) by an order of magnitude, and consequently the GLUE limit to be too low by an order of magnitude. Our predicted sensitivity for future experiments using the Australia Telescope Compact Array (ATCA) and the Australian SKA Pathfinder (ASKAP) indicate these instruments will be able to detect the more optimistic UHE neutrino flux predictions, while the Square Kilometre Array (SKA) will also be sensitive to all but one prediction of a diffuse ‘cosmogenic’, or ‘GZK’, neutrino flux.

Outstanding theoretical uncertainties at both high-frequency and low-frequency limits currently prevent a reliable estimate of the sensitivity of the lunar Cherenkov technique for UHE cosmic ray (CR) astronomy. Here, we place limits on the effects of large-scale surface roughness on UHE CR detection, and find that when near-surface ‘formation-zone’ effects are ignored, the proposed SKA low-frequency aperture array could detect CR events above 56 EeV at a rate between 15 and 40 times that of the current Pierre Auger Observatory. Should further work indicate that formation-zone effects have little impact on UHE CR sensitivity, observations of the Moon with the SKA would allow directional analysis of UHE cosmic rays, and investigation of correlations with putative cosmic ray source populations, to be conducted with very high statistics.

Keywords: UHE neutrino detection, UHE cosmic ray detection, coherent radio emission, lunar Cherenkov technique, UHE neutrino flux limits

James, C.W. & Protheroe, R.J. (2009) The sensitivity of the next generation of lunar Cherenkov observations to UHE neutrinos and cosmic rays
Astroparticle Physics, v. 30 (6), pp. 318-332

NOTE:

This publication is included on pages 29-60 in the print copy of the thesis held in the University of Adelaide Library.

It is also available online to authorised users at:

<http://dx.doi.org/10.1016/j.astropartphys.2008.10.007>

Directional Dependence

4

**THE DIRECTIONAL DEPENDENCE OF APERTURES, LIMITS AND
SENSITIVITY OF THE LUNAR CHERENKOV TECHNIQUE TO A UHE
NEUTRINO FLUX**

C.W. James, R.J. Protheroe

¹ Department of Physics, The University of Adelaide, Adelaide, SA 5005, Australia

Astroparticle Physics (submitted March 2008)

STATEMENT OF AUTHORSHIP**THE DIRECTIONAL DEPENDENCE OF APERTURES, LIMITS AND SENSITIVITY OF THE LUNAR CHERENKOV TECHNIQUE TO A UHE NEUTRINO FLUX.**

Astroparticle Physics (submitted March 2008)

James, C.W. (Candidate)

Wrote computer code; generated and interpreted data; wrote manuscript and produced figures (except where noted below); acted as corresponding author.

I hereby certify that the statement of contribution is accurate.

Signed:

Date:

Protheroe, R.J.

Manuscript editing; produced Figs. 4.1 and 4.2.

I hereby certify that the statement of contribution is accurate.

Signed:

Date:

The Directional Dependence of Apertures, Limits and Sensitivity of the Lunar Cherenkov Technique to a UHE Neutrino Flux

C.W. James¹, R.J. Protheroe²

Department of Physics, The University of Adelaide, Adelaide, SA 5005, Australia

Abstract

We use computer simulations to obtain the directional-dependence of the lunar Cherenkov technique for ultra-high energy (UHE) neutrino detection. We calculate the instantaneous effective area of past lunar Cherenkov experiments at Parkes, Goldstone (GLUE), and Kalyazin, as a function of neutrino arrival direction, finding that the potential sensitivity to a point source of UHE neutrinos for these experiments was as much as thirty times that to an isotropic flux, depending on the beam pointing-position and incident neutrino energy. Convoluting our results with the known lunar positions during the Parkes and Goldstone experiments allows us to calculate an exposure map, and hence the directional-dependence of the combined limit imposed by these experiments. In the 10^{21} – 10^{23} eV range, we find parts of the sky where the GLUE limit likely still dominates, and areas where none of the limits from either Parkes, GLUE, or experiments such as ANITA or FORTE are likely to be significant. Hence a large anisotropic flux of UHE neutrinos from these regions is not yet excluded.

We also determine the directional dependence of the aperture of future planned experiments with ATCA, ASKAP and the SKA to a UHE neutrino flux, and calculate the potential annual exposure to astronomical objects as a function of angular distance from the lunar trajectory through celestial coordinates. We find that the potential exposure of all experiments at 10^{20} eV and below, integrated over a calendar year, is flat out to $\sim 25^\circ$ from the lunar trajectory and then drops off rapidly. The region of greater sensitivity includes much of the Supergalactic Plane, including M87 and Cen A, as well as the Galactic Centre. At higher energies this high-sensitivity region becomes broader, and we find that the potential exposure of the SKA at 10^{21} eV and above is almost uniform over celestial coordinates.

Keywords: UHE neutrino detection, coherent radio emission, lunar Cherenkov technique, UHE neutrino flux limits

James, C.W. & Protheroe, R.J. (2009) The Directional Dependence of Apertures, Limits and Sensitivity of the Lunar Cherenkov Technique to a UHE Neutrino Flux
Astroparticle Physics, v. 31 (5), pp. 392-398

NOTE:

This publication is included on pages 63-76 in the print copy of the thesis held in the University of Adelaide Library.

It is also available online to authorised users at:

<http://dx.doi.org/10.1016/j.astropartphys.2009.04.002>

Implementation of a Thinning Algorithm in the ZHS Code

This chapter contains segments from the following submitted paper:

THINNED SIMULATIONS OF EXTREMELY ENERGETIC SHOWERS IN DENSE MEDIA FOR RADIO APPLICATIONS

J. Álvarez-Muñiz¹, C. W. James², R. J. Protheroe², E. Zas¹

¹ Univeristy de Santiago de Compostela, Santiago de Compostella, Spain

² School of Chemsitry and Physics, University of Adelaide, Adelaide, SA 5005, Australia

Astroparticle Physics (submitted February 2009)

STATEMENT OF AUTHORSHIP**THINNED SIMULATIONS OF EXTREMELY ENERGETIC SHOWERS IN
DENSE MEDIA FOR RADIO APPLICATIONS***Astroparticle Physics (submitted February 2009)***James, C.W.** (Candidate)

Primary author of Secs. 5.4.1, 5.5.1, and 5.5.2. Produced all included figures and tables. Co-author of Secs. 5.4, 5.4.3, 5.5, and 5.7. Developed methodology and ran simulations for ‘quality of thinning technique’ (Sec. 5.4.1).

I hereby certify that the statement of contribution is accurate

Signed:

Date:

J. Álvarez-Muñiz

Original conception of the work. Co-author of Secs. 5.4, 5.4.3, 5.5, and 5.7. Contributing author of Secs. 5.4.1 and 5.5.1. Developed methodology and ran simulations for the ‘simultaneous thinning’ technique (Sec. 5.4). Ran simulations and produced data for fitted model parameters (Sec. 5.5.1, and Tables 5.3, 5.4, and 5.5).

I hereby certify that the statement of contribution is accurate

Signed:

Date:

R. J. Protheroe

Manuscript editing.

I hereby certify that the statement of contribution is accurate

Signed:

Date:

E. Zas

Manuscript editing.

I hereby certify that the statement of contribution is accurate

Signed:

Date:

5.1 Introduction

Assessing the viability of UHE neutrino detection experiments relies upon accurate simulations of the hadronic and electromagnetic showers produced by various mechanisms from the initial interaction, and the coherent Cherenkov radiation produced at cm-m scales. Because the number of particles grows linearly with the energy, full simulations (where every particle is tracked in three spatial and one time dimension) of showers become computationally prohibitive at energies above 10^{15} eV, whereas fluxes of UHE particles have been predicted up to 10^{23} eV. Future simulations, which may take into account transition radiation arising from partially developed showers, and/or the near-field roughness effects of transmission through the lunar–vacuum interface, will only increase the time required for any given shower energy. This motivates a search for a method to speed up such codes.

This chapter deals with a joint investigation by myself and J. Álvarez-Muñiz of the University of Santiago de Compostela (USC), Spain, into the use of thinning algorithms to enable a faster computation of purely electromagnetic showers. Using the thinning technique, only a fraction of particles are followed by a simulation, with the remainder having an increased statistical weighting in the calculations to compensate. By reducing the number of particles which must be tracked, the computing time required for calculations of Cherenkov radiation – or indeed any other calculation – is decreased. As a result of this investigation, a dual-threshold thinning algorithm – optimised for the calculation of both the frequency- and angular-distribution of Cherenkov radiation – was developed and implemented in the purely electromagnetic ZHS (Zas, Halzen, Stanev – see Ref. (123)) code. Using this algorithm, we analysed the structure of radio-emission from UHE electromagnetic showers, and developed new parameterisations for the radiation in different media, accurate to 10^{20} eV – six orders of magnitude higher than those previously attainable.

The results of this work – titled "Thinned simulations of extremely energetic showers in dense media for radio applications" by J. Alvarez-Muñiz, C. W. James, R. J. Protheroe, and E. Zas – have been submitted as a paper to *Astroparticle Physics*. Since my contribution is not a majority of the paper, it would be inappropriate to include the whole as a stand-alone chapter in this thesis. Rather, I have included segments only where I feel my contribution was sufficient to warrant this - this is noted in the relevant sections, figures, and tables. For further information, please see the declaration by myself and my co-authors at the start of this chapter.

The following section – Sec. 5.2 – introduces the ZHS code and outlines UHE shower phenomenology. Section 5.3 describes the thinning method, specifies its implementation in the ZHS code, and presents a formalism to quantify the effects of thinning. Section 5.4 is taken from the aforementioned paper Ref. (18), and describes my method of optimising the thinning parameters. Appended to the section are further details which did not make the paper, but are nonetheless useful in aiding the understanding of thinning behaviour. Section 5.5 is also taken from the aforementioned paper, and presents our model parameterisation, fitting results, and discussion thereof, although the fitted values themselves (which were




Figure 5.1: Diagram illustrating the difference between the subdivision of tracks under approximations A, B, and C in the ZHS code.

generated by my co-author) are left for the submitted work. Before the obligatory conclusion, Sec. 5.6 discusses qualitatively the effects of significant secondary peaks in the radiated spectrum, which for the first time can be simulated by our code.

5.2 Preliminaries

5.2.1 The ZHS code

ZHS is an electromagnetic code developed by Zas, Halzen, and Stanev (123) specifically for calculating the far-field coherent Cherenkov radiation from a high-energy electromagnetic shower. The code takes an initial electron, positron, or γ photon, and follows it and any subsequent secondaries down to energies at which they can no longer produce (or in the case of γ , no longer give sufficient energy to electrons in the medium to produce) Cherenkov radiation. The code runs in a uniform medium – originally only ice – and calculates the far-field radiation pattern as a function of a user-defined grid in angle-frequency space.

The goal of the code is to generate linear charged-particle tracks to pass to the routine `mpsum`, which then calculates that track's contribution. Exactly how the code subdivides the tracks therefore influences the final calculation of the emitted radiation, with finer subdivisions being valid to higher frequencies. Three methods – approximations A, B and C – have been implemented in the code, and are illustrated in Fig. 5.1. In approximation A, a particle track is defined between the points when the particle enters the shower – is created, knocked from the surrounding medium, or is injected as the primary particle – to when it leaves – either falling below the Cherenkov threshold or annihilating. Approximation B divides the tracks as being between every ‘significant’ interaction – that is, between interactions which are modelled as being discrete by the code, and which typically add new particles to the shower. Under approximation C, the `mpsum` routine is called at every step in the multiple-scattering routine, so that low-energy particles may have thousands of tracks. Because the `mpsum` routine typically dominates the time taken to run the simula-

tion, the times taken to run the simulation are in the ratio 1:6:30 for approximations A:B:C. For a comparison of the accuracies of the three methods, see Alvarez-Muñiz *et al.* (2000) (13); to summarise, in ice, a spectrum generated under approximation A is identical to one generated under B or C until 500 MHz, and one from B is identical to one from C up to 5 GHz, above which only approximation C gives accurate results.

5.2.2 Electromagnetic shower phenomenology

The minimum electron kinetic energy required to produce Cherenkov radiation in ‘normal’ media ($n \sim 2$, e.g. ice, the regolith, and salt) is of the order of 100 keV, so that the code must trace particles down to an energy E of $E = mc^2 + 100$ keV: ~ 600 keV for e^\pm , and ~ 100 keV for γ . Interesting primary particle energies range into the hundreds of EeV; unsurprisingly the interaction phenomenology changes significantly over this range, and it is worthwhile briefly discussing how the phenomenology changes with energy.

Firstly, the ‘normal’ regime, in which Heitler’s toy model of shower development ((55)) can successfully explain many properties. Here, the dominant radiation process for electrons is bremsstrahlung, while for γ it is pair-production; the energy-loss distance – the ‘radiation length’, χ_0 – of the former approximately equals the interaction length for the latter, and these distances are approximately constant with energy, so that the total shower length increases logarithmically with primary energy. Losses to the medium in this regime are small, so hence energy conservation dictates that the number of particles of energy E grows as E^{-1} . The shower profile (number of particles vs. shower depth) can be described well using Greisen’s parameterisations (51). Also, the charge is balanced.

The lower limit to this regime is given by the ‘critical energy’ E_{crit} , often defined as the energy below which the ionisation energy-loss rate for electrons is greater than that for bremsstrahlung (see Ref. (19) for a discussion of the various definitions of and approximations to E_{crit}). Loosely, below E_{crit} – typically O~50 MeV for media considered here – effects such as Compton, Moeller, and Bhabha scattering (where, respectively, γ , e^- , and e^+ knock a medium-electron into the shower), and electron-positron annihilation, become important. It is in this energy regime that the charge-excess which produces coherent Cherenkov radiation develops.

At energies of a few MeV, interactions producing secondaries with energy above the Cherenkov threshold become scarce, so for simulation purposes the shower ceases to grow. Particle tracks in this regime tend to be strongly affected by multiple-scattering from atoms in the medium, sometimes causing particles to travel backwards. This regime is still important though, since the charge-excess has already developed, and the number of particles is large, so that while individual tracks are typically short, the total contribution to excess tracklength is significant.

The high-energy limit to Heitler’s model is determined by the Landau-Pomeranchuk-Migdal (LPM) Effect (73; 80). For an e^\pm of energy E near the characteristic energy E_{LPM} (typically a few hundred TeV to a couple of PeV), the formation length for low-energy

bremsstrahlung photons becomes large enough that there is a high probability that the emitting particle will undergo a significant perturbation over the formation-length of the radiation, thereby suppressing the emission. Here, a ‘significant perturbation’ is typically one (or many) multiple scatterings which deviate the particle by an angle of order $1/\gamma_p = m_p/E_p$ (67). A similar effect occurs for pair-production by high-energy γ . At energies much higher than E_{LPM} , effectively all emission / pair-production is suppressed, with the cross-section decreasing as $E^{-0.5}$. The result is both that the shower length increases rapidly with primary energy, and that the profile becomes ‘lumpy’, being composed of relatively isolated sub-showers. It is this LPM-elongation which causes the very narrow width of the Cherenkov cone from the highest-energy electromagnetic showers.

The very highest energy regime cannot be modelled by the ZHS code. With a rapidly-dropping pair-production cross-section with energy due to the LPM-effect, eventually the photon-nucleon photohadronic cross-section (and also the less-well-known electron-nucleon photohadronic cross-section) will become important, and initially purely-electromagnetic showers will have a significant hadronic component. For most predictions of the cross-sections – Ref. (98) – this is expected to occur for photon energies around 10^{19} eV, or approximately e^\pm energies of about 10^{20} eV, so that above this energy, ZHS-based simulations of UHE electromagnetic cascades will become unrealistic.

5.3 Thinning Electromagnetic Showers

5.3.1 Thinning basics

The basic idea of thinning (also known as ‘thin sampling’) is simple - track fewer particles, while assigning those kept a statistical weighting to compensate for the lost particles. Each time a particle interacts to produce multiple secondaries, only a representative sample of the secondaries are followed to save computing time, but with an increased weight to compensate for the missing particles. Thus the thinning algorithm is called at each vertex in a simulation code (where `is called under approximation B`).

A thinning algorithm is defined by how it chooses which particles to keep and which to ignore, and what weights to give the remaining particles. The choice should be effectively random to ensure an unbiased selection, so that for a general interaction $A \rightarrow B+C+D+\dots$, particle B (and similarly for C, D etc.) is followed with probability p_B and weight $W_B = w_B W_A$, where w_B is called the ‘weight increase’. For an unthinned shower therefore, particle P has $W_P = w_P = p_P = 1$ always.

The goal is to minimise both computing time and artificial variation – that is, to get quick and accurate results. Except perhaps in very unusual circumstances, any acceptable thinning algorithm should obey conservation of particle number, i.e. $w_P = 1/p_P$. This ensures that basic properties like energy, momentum etc. are at least statistically conserved,

e.g. as demonstrated below for particle A with energy E_A :

$$\langle E_{\text{final}} \rangle = \langle W_B E_B \rangle + \langle W_C E_C \rangle + \dots \quad (5.1)$$

$$= p_B W_B E_B + p_C W_C E_C + \dots \quad (5.2)$$

$$= p_B w_B W_A E_B + p_C w_C W_A E_C + \dots \quad (5.3)$$

$$= W_A E_B + W_A E_C + \dots \quad (5.4)$$

$$= W_A E_A \quad (5.5)$$

$$= E_{\text{initial}} \quad (5.6)$$

where $\langle f(x) \rangle$ denotes the expectation value of a random quantity f , and we assume for now that $E_B + E_C + \dots = E_A$. From here on therefore, we only consider thinning algorithms where $w_P = 1/p_P$ always.

While statistical conservation of energy can be assured by setting $w_P = 1/p_P$, it is possible to do even better than this by the following choice of probabilities/weights:

$$p_B = E_B/E_A \quad (5.7)$$

and choosing to follow one of $B, C, D \dots$ only. Eq. (5.7) is so useful because energy is conserved explicitly at each vertex, not just statistically, as shown below, assuming arbitrarily that B is chosen:

$$E_{\text{final}} = W_B E_B \quad (5.8)$$

$$= p_B^{-1} W_A E_B \quad (5.9)$$

$$= W_A E_A. \quad (5.10)$$

Some procedure must also be devised to determine when and when not to thin, since applying the above (or any) thinning algorithm over the entire life of a shower will produce very variable results, since only one particle will ever be tracked. Hillas (56) used thinning thresholds based on particle energy: if an interaction took place below some critical energy E_{thin} , it was thinned; otherwise, it was not. This is known as single-level thinning. More commonly, it is usual to proceed as Kobal (69), and define two thinning thresholds, E_{min} and E_{max} , so that the thinning algorithm only gets called if the interaction energy lies between E_{min} and E_{max} . Setting $E_{\text{min}} = E_{\text{thresh}}$ recovers the original Hillas algorithm. The main advantage of a two-level algorithm is that it allows the weights to be limited, so that the greatest weight that any given particle can obtain is given by the thinning level $f_L = E_{\text{max}}/E_{\text{min}}$, thus preventing undue statistical fluctuations.

5.3.2 Thinning and ZHS

While thinning techniques have been successfully applied to high-energy air-shower simulations (56; 69; 44), we expect the calculation of coherent Cherenkov radiation to require a

new optimisation. Note that all electromagnetic interactions in ZHS can be characterised as $A \rightarrow B + C$ (only two interaction products), which greatly simplifies any thinning algorithm. Since the charge-excess is produced at low energies, and we wish to accurately model the effects of LPM elongation above ~ 100 TeV, using a two-threshold algorithm makes intuitive sense. The goal will be to optimise the parameters E_{\max} and E_{\min} as (in general) a function of shower energy and the interaction medium. But first, some subtleties.

Thinning and approximation A

Approximation A defines a ‘track’ as being from the point when a particle enters the shower to when it leaves. If the particle is an electron or positron, it will generally undergo multiple significant interactions during this time, so that if the particle is kept, the weight will change over this tracklength. When the particle leaves the shower and its track is passed to the Cherenkov algorithm `mpsum`, the final weight will be (incorrectly) applied to the entire track. Fig. 5.2 demonstrates this in the case of a simple ‘doubling’ shower, where each particle interacts after a uniform distance to produce two identical secondaries. The effect of thinning under approximation A is an obvious artificial increase in the effective number of tracks. Note that this also cannot be corrected for by a simple re-weighting by a factor of (in this case) $32/41$, since both the centre and width of the shower have changed (moved left and increased, respectively). Assuming the tracks at each generation are at positions 0, 1, 2, 3 and 4, the centre under Approx. A has shifted from 3.16 to 2.86, while the standard deviation about the centre has increased from approximately 1.1 to 1.2, i.e. the shower is thinner by 10%.

To test this model, a simple single-threshold thinning algorithm was implemented in the code, with the simulated spectra at the Cherenkov angle plotted in Fig. 5.3. Evidently, as thinning level increases (and E_{\min} decreases for a given thinning level), the simulated strength increases significantly. Therefore we conclude that approximations B and C only are appropriate for a thinning implementation. As the time taken for approximation C is so much longer with a gain in accuracy only at the highest frequencies (where no experiments plan to observe anyway), we only considered thinning under approximation B.

Threshold-crossing interactions

It is not possible to use both a well-defined thinning threshold and explicitly conserve energy at each vertex as per Eq. (5.7). Consider the case $A \rightarrow B + C$ where $E_A > E_B > E_{\max} > E_C$. Applying Eq. (5.7) would result in a particle with energy $E_B > E_{\max}$ being either discarded or having its weight increased; also, the weight of E_C , if it is followed, could exceed the weight limit. However, not applying Eq. (5.7) would result in an inefficient algorithm with many more low-energy particles C being tracked, possibly giving a bias to those resulting from high-energy interactions. Here, we present and analyse two solutions to this problem.

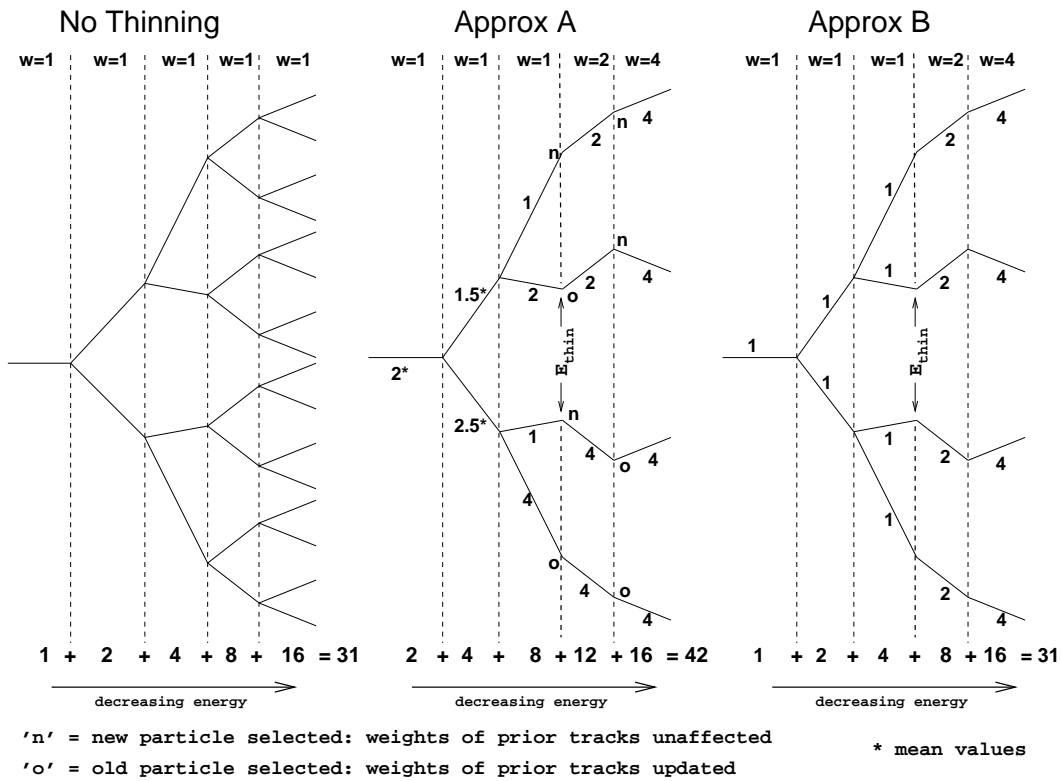


Figure 5.2: A simple 'doubling' shower under different thinning methods and thinning threshold E_{thin} (energy is decreasing from left to right). The unthinned shower (left) gives the correct number of tracks as 31. Thinning under approximation B (right) reproduces this number correctly, since the weights compensate for the reduced particle numbers at low energies. Thinning under approximation A (centre), however, artificially increases the weighting of previous tracks whenever the old ('o') particle is followed, since the new weighting is applied to previous tracks. Here this is assumed to occur with 50% probability, so that the total number of tracks is overestimated at, on average, 42.

One possible thinning algorithm, referred to in our paper as the 'alternative' routine though also dubbed 'exclusive thinning' since interactions remain unthinned unless all their particles lie inside the thinning thresholds, is defined as follows:

1. IF $E_{max} > E_A, E_B, E_C > E_{min}$, choose *one* of B or C with $p_{B[C]} = E_{B[C]}/E_A$.
2. ELSE do not thin.

The standard procedure however is to maintain the nice behaviour of the weights with particle energy at the expense of explicit energy conservation. From here-on called the 'standard algorithm', it treats both thinning thresholds in the same manner as Hillas's treatment of the single threshold in his well-known algorithm (56) – that is, by only thinning off the energy 'within' the thinning regime:

1. IF $E_{max} > E_A, E_B, E_C > E_{min}$, choose *one* of B or C with $p_B = E_B/E_A$,




Figure 5.3: Radiated spectrum at the Cherenkov angle calculated for 100 GeV showers thinned under Approximation A for three combinations of thinning parameters (five showers simulated for each combination).

2. ELSE IF $E_A, E_B, E_C > E_{\max}$ OR $E_A, E_B, E_C < E_{\max}$ do not thin,
3. ELSE choose either, none or both B and C *independently*, with probability of retaining particle P given by:

$$p_P = \begin{cases} 1 & \text{if } E_p > E_{\max}, \\ E_p/E_{\max} & \text{if } E_A > E_{\max} > E_p > E_{\min}, \\ E_{\min}/E_A & \text{if } E_{\max} > E_A > E_{\min} > E_p, \\ E_{\min}/E_{\max} & \text{if } E_A > E_{\max}, E_{\min} > E_p. \end{cases}$$

It is only in ‘threshold-crossing’ interactions (case 3) where the two algorithms differ, but it is an important one. Energy cannot be conserved explicitly in the standard routine since for such threshold-crossing interactions the tests to retain B and C are necessarily independent. However, a quick inspection will show that the weight of a particle will be completely determined by its energy as per Eq. (5.11):

$$W_P = \max(1, \min(f_L, E_{\max}/E_P)). \quad (5.11)$$

In the case of the alternative routine this imposes only an upper limit on the weight. We initially considered both routines.

The base thinning energy

The above discussion of energy conservation avoids a basic question: which energy do we want to conserve? The ZHS simulation works off the total (kinetic plus rest mass) energy, and at first this seems like the obvious answer. But there are two important reasons why *not* to use the total energy of massive particles in determining probabilities and weights. The first is that it is generally the kinetic energy of the particles which is the energy available to produce observable results – this is especially true of a relativistic process such as Cherenkov radiation. The second is in fact to ensure energy conservation in Moeller, Bhabha and Compton scatterings, in which electrons from the medium are knocked into the shower. Including rest-mass energy means that the final energy $E_B + E_C$ is greater than the initial energy E_A by the 511 keV rest-mass energy of the knock-on electron, while the kinetic component is conserved. Admittedly, this logic does not follow for a shower-positron annihilating with a medium-electron, where the rest mass is converted into ‘useful’ energy, since the γ -ray products can Compton-scatter medium-electron into the shower with energy above the Cherenkov threshold. However, e^+/e^- annihilation contributes only of order 10% to the charge excess (123), and does so only for positrons at low energies – for which a thinning base energy excluding the rest mass is preferable anyhow.

5.3.3 Definitions and approximations

In the following I take time to define some basic relations which aid in understanding the behaviour of the ZHS code with a thinning algorithm implemented. Throughout, a superscript of ‘0’ indicates the properties of an unthinned shower and an implicit dependence on E_0 , while any quantity with a ‘*thin*’ superscript indicates it applies to a thinned shower with an implicit dependence on E_0 , f_L , and E_{\min} . Note also that in general I use the set $\{f_L, E_{\min}\}$ to describe a choice of thinning parameters rather than the equivalent $\{E_{\max} = f_L E_{\min}, E_{\min}\}$, for reasons which will become obvious later.

Number of particles in the shower

The behaviour of thinning – time saved, amount of artificial variation etc – can be mostly explained with the simple model of electromagnetic showers development used by Hillas (56) and illustrated in Fig. 5.2. Note that in this model, the ratio of particles at any given instant $n_{e^-} : n_{e^+} : n_\gamma$ is maintained at 1:1:1 throughout, though at low energies we expect the charge excess $n_{e^-} > n_{e^+}$. This process stops once the particles reach some critical energy E_{cts} where continuous energy-loss mechanisms (ionisation losses) dominate, so that the total number of particles N_p in an unthinned shower of primary energy E_0 is given by:

$$N_p^0 = 1 + 2 + 4 + \dots + E_0/E_{\text{cts}} \quad (5.12)$$

$$= 2E_0/E_{\text{cts}} - 1. \quad (5.13)$$

In reality E_{cts} is not be a well-defined physical boundary below which no discrete interactions take place, but rather a characteristic energy below which very few new particles are added. However, counting the total number of particles in a simulated shower is easy, so that E_{cts} is most usefully defined as per Eq. (5.14) by approximately inverting Eq. (5.13):

$$E_{\text{cts}} = \frac{E_0}{2N_p^0}. \quad (5.14)$$

Note that a post-bremsstrahlung e^\pm is counted as a new particle. Also, E_{cts} will generally be less than E_{crit} , since many electrons are knocked into the shower below this energy, while E_{cts} must be at lowest equal to E_{thresh} , since particles below E_{thresh} are either not counted or treated as continuous losses. Running simulations indicated that for ice, $E_{\text{cts}} \approx 915$ keV. The low value is largely due to the number of very low-energy knock-on electrons, since these outnumber all other particles; using $N_p^0 = 3N_{e^+}$ (as in Hillas's toy model) gives $E_{\text{cts}} \approx 2.25$ MeV.

Number of particles in a thinned shower

First let me make a distinction in the notation: a lower-case 'n' denotes the number of particles/particle tracks (in Approximation B the two are equivalent) at a *particular* energy, while an upper-case 'N' denotes the total number of particles in some energy range. With this defined, I proceed:

Thinning at energies above E_{cts} will reduce the number of particles in the shower, with the total number of particles N_p^{thin} given by separate contributions from each of three energy regimes:

$$N_p^{\text{thin}} = N_{E_0 > E > E_{\text{max}}} + N_{E_{\text{max}} > E > E_{\text{min}}} + N_{E_{\text{min}} > E > E_{\text{cts}}} \quad (5.15)$$

$$\equiv N_1 + N_2 + N_3. \quad (5.16)$$

Note that N_2 and N_3 will be dependent on the thinning method – the following discussion applies to the standard algorithm only. For the alternative method, it may be possible to define $E'_{\text{max}} < E_{\text{max}}$ and $E'_{\text{min}} > E_{\text{min}}$ where the following relations approximately hold – however, I do not investigate this explicitly.

The contribution N_1 from the pre-thinning (high-energy) part of the shower can be estimated trivially as:

$$N_1 = 2E_0/E_{\text{max}} - 1 \quad (5.17)$$

since the high-energy component develops exactly like an unthinned shower with E_{cts} replaced by E_{max} . The contribution N_3 from the post-thinning (low-energy) part of the shower is similarly trivial: $n_3 = E_0/E_{\text{max}}$ particles are 'injected' into this regime with individual energies $E = E_{\text{min}}$, each of which subsequently develops a sub-shower over the energy range $\{E_{\text{min}}, E_{\text{cts}}\}$. The only complexity is that the particles appearing at E_{min} must themselves be

excluded, since these are counted in N_2 , or N_1 for the unthinned case. Therefore N_3 is given by Eq. (5.18):

$$N_3 = \frac{E_0}{E_{\max}} \left(2 \frac{E_{\min}}{E_{\text{cts}}} - 2 \right). \quad (5.18)$$

The contribution N_2 from the thinning regime is more complex to calculate. If $N_2 = g_2 n_2$, where n_2 ($\approx E_0/E_{\max}$) is the (constant) number of particles in the thinning regime, and g_2 is the mean number of interactions after which a followed particle of initial energy E_{\max} has energy E_{\min} , then in the simplest case of $A \rightarrow B + C$ where $E_B = E_C = 0.5E_A$ we find $g_2 = \log_2(f_L)$ and hence get:

$$N_2 = \log_2(f_L) \frac{E_0}{E_{\max}}. \quad (5.19)$$

In general however, $E_B \neq E_C$; the thinning routine will be likely to keep the more energetic product, and the number of interactions in the thinning regime will be increased; thus we expect $g_2 > \log_2(f_L)$. If the mean energy fraction of a followed particle is \bar{y} , we find it useful to define $q = 1/\bar{y}$, so that $g_2 = \log_q(f_L)$, and hence:

$$N_2 = \log_q(f_L) \frac{E_0}{E_{\max}}. \quad (5.20)$$

Observe that $1 < q < 2$, where the minimum of $q = 1$ arises from one secondary always containing all the primary energy. In this (obviously unrealistic) case, the energetic secondary will always be followed, the shower will not be reduced in energy, and hence it will have an infinite number of tracks. The maximum of $q = 2$ occurs for the equipartition ($E_B = E_C$) assumption above, while for a flat energy distribution of the secondaries, it can be shown that $q = \sqrt{e}$. The actual value of q thus depends on the distribution of product energy-fractions and so will be energy-dependent.

Combined, Eqs. 5.17, 5.18, and 5.20 give the total number of particles in a thinned shower N_p^{thin} as:

$$N_p^{\text{thin}} = \frac{E_0}{E_{\max}} \left(2 \frac{E_{\min}}{E_{\text{cts}}} + \log_q(f_L) - 1 \right). \quad (5.21)$$

Provided $E_{\min} \gg E_{\text{cts}}$ and/or $\log_q(f_L)$ is small, N_p^{thin} can be approximated as:

$$N_p^{\text{thin}} \approx 2 \frac{E_0}{E_{\max}} \frac{E_{\min}}{E_{\text{cts}}} \quad (5.22)$$

$$\approx \frac{N_p^0}{f_L}. \quad (5.23)$$

That is, the total number of particles is approximately reduced by a factor f_L .

Time-saving and efficiency

The goal of thinning is to save calculation time, which is approximately proportional to the number of particle tracks. We define the ratio of times t^0/t^{thin} as f_t , where t^0 (t^{thin}) is the actual time taken to run an unthinned (thinned) simulation at some energy. Since the time taken will be platform-dependent, a more useful (and easily-measured) quantity is the ratio N_p^0/N_p^{thin} , defined here as f_n . Almost every aspect of the simulation code takes time proportional to N_p , and we always use the proxy $f_n \sim f_t$, which should be accurate to a good degree of accuracy. From Eq. (5.23), we also expect $f_L \approx f_n$, and hence the time saved should be proportional to the thinning level f_L .

The actual time reduction factor f_t will likely be greater than the expectation value f_L for several reasons. Most obviously, if the dropped logarithmic term in Eq. (5.21) becomes significant, more tracks will be generated than N^0/f_L . Also, for the alternative thinning algorithm, interactions spanning the thinning boundaries will not be thinned, so the effective thinning level f'_L will be less than f_L . Finally, any continuous energy-losses in the thinning regime will reduce the effective thinning level, since thinning can only act on significant interactions.

To measure the effects of all of deviations from the simplified model above, the efficiency ξ is defined as the ratio between the actual and expected time reduction factors as per Eq. (5.24):

$$\xi = \frac{f_n - 1}{f_L - 1}. \quad (5.24)$$

This definition, rather than $\xi = f_n/f_L$, was chosen so that if $f_n = 1$ (no time reduction) and $f_L \neq 1$, we find $\xi = 0$; also, this measure is sensitive at small values of f_L .

Measures of variability

It is important to understand and clearly define the different sources of variability associated with thinned HE showers, and their measurement. In general, two types of variability are discussed: the intrinsic random variation between fully simulated ('real') showers, and the artificial variation introduced by thinning. The first is defined as the mean variation between an ensemble of initially identical showers, while the second is the expected difference between a thinned shower and its otherwise-identical unthinned (real) counterpart. Mathematically, I use the largely-intuitive definitions described below.

For N_{shwr} independent, unthinned showers, each simultaneously thinned (see Sec. 5.4) N_{thin} times, $X_{i,j}$ is defined as the estimate of a parameter X from the j^{th} simultaneous thinning of the i^{th} shower, and the estimate of X from the i^{th} unthinned shower is denoted X_i^0 . Note that the only difference between the simultaneously thinned showers j for a given i is randomisation within the thinning algorithm – outside this algorithm, the showers are identical. Also note that in the following I do not work at the individual particle level, although necessarily each $X_{i,j}$ is calculated as a function of the individual particles $p_{i,k}$ (k^{th} particle,

i^{th} shower) and weights $W_{i,j,k}$ (weight applied to the k^{th} particle due to the j^{th} thinning result in the i^{th} shower).

Using this formalism, we can now define the ‘intrinsic variation’ V_X^0 in X *between* showers by first calculating the mean of X over all showers, \bar{X}^0 :

$$\bar{X}^0 = \frac{1}{N_{\text{shwr}}} \sum_{i=1}^{N_{\text{shwr}}} X_i^0 \quad (5.25)$$

$$V_X^0 = \sqrt{\frac{\sum_{i=1}^{N_{\text{shwr}}} (X_i^0 - \bar{X}^0)^2}{N_{\text{shwr}} - 1}}. \quad (5.26)$$

$$(5.27)$$

To calculate the accuracy of the particular combination of thinning parameters in question, we define the variation V_{Xi} in X for the i^{th} shower according to Eq. (5.28), and consequently evaluate the average variation \bar{V}_X caused by thinning as per Eq. (5.29):

$$V_{Xi} = \sqrt{\frac{\sum_{j=1}^{N_{\text{thin}}} (X_{i,j} - X_i^0)^2}{N_{\text{thin}}}} \quad (5.28)$$

$$\bar{V}_X = \frac{1}{N_{\text{shwr}}} \sum_{i=1}^{N_{\text{shwr}}} V_{Xi}. \quad (5.29)$$

\bar{V}_X therefore measures the mean artificial variation induced by thinning, which we have decided to tolerate up to the 10% level. At high energies where fully-simulated (unthinned) showers are not possible, Eq. (5.28) must be approximated by replacing V_{Xi} by V'_{Xi} , where V'_{Xi} is defined as the value of V_{Xi} calculated by estimating the value X^0 from the thinning estimates themselves:

$$\bar{X}_i = \frac{1}{N_{\text{thin}}} \sum_{j=1}^{N_{\text{thin}}} X_{i,j} \quad (5.30)$$

$$V'_{Xi} = \sqrt{\frac{\sum_{j=1}^{N_{\text{thin}}} (X_{i,j} - \bar{X}_i)^2}{N_{\text{thin}} - 1}}. \quad (5.31)$$

A value of \bar{V}_X calculated by replacing V_{Xi} by V'_{Xi} in Eq. (5.29) is correspondingly denoted \bar{V}'_X .

Note that only by using the simultaneous-thinning algorithm can we determine if there is a bias introduced by the thinning algorithm. We found no evidence for such bias in any case investigated.

Quantity	T	$\langle r \rangle$	M_r	\bar{z}	$\langle z \rangle$	M_z
Def: \sum_i	T_i	$\frac{T_i}{T} \sqrt{x_i^2 + y_i^2}$	$\frac{T_i}{T} (x_i^2 + y_i^2)$	$\frac{T_i}{T} z_i$	$\frac{T_i}{T} z_i - \bar{z} $	$\frac{T_i}{T} (z_i - \bar{z})^2$

Table 5.1: (Reproduced from Table 1 of Ref. (18)) Definitions of variables used to assess the accuracy of our thinning routine. The sums in the formulas run over all particles in the shower. T_i is the tracklength of each particle –ve for positrons, +ve for electrons, so the excess is a positive quantity located at mean position (x_i, y_i, z_i) , where $x = y = 0$ corresponds to a particle travelling along the shower axis.

5.4 Testing Thinning Methods

I implemented both the standard thinning algorithm (energy not conserved explicitly) and the alternative thinning algorithm (energy forced to be conserved) into the ZHS code. The thinning routine is called after every discrete interaction, returning a flag to keep or discard each particle, and an updated weighting for that particle. Preliminary tests indicated that thinned simulations gave reasonable results – the next stage was to define a way of quantifying the best thinning routine.

The remainder of this section – Sec. 5.4 – comes from Sec. 3 of Ref. (18).

Our goal with the introduction of thinning in the simulations of Cherenkov radio emission from high energy showers is to reproduce the spectrum and angular distribution of Cherenkov radiation predicted in full simulations as accurately as possible while at the same time minimising the computing time. Since the spectrum of Cherenkov radiation is approximately the Fourier transform of the distribution of excess negative charge (13; 33), we assess the goodness of the thinning algorithms based on their ability to reproduce this distribution, reflected by the parameter set defined in Table 5.1. Of these, total excess tracklength T is the most useful, since it has very low intrinsic variability, and scales with the primary particle energy. Also defined is the shower “centre” at $(\bar{x} = 0, \bar{y} = 0, \bar{z})$, and the means (linearly-weighted variation) and “moments” (quadratically-weighted variation) of both shower length ($\langle z \rangle, M_z$) about the shower centre \bar{z} , and shower width ($\langle r \rangle, M_r$) about the shower axis ($x = 0, y = 0$) in the direction perpendicular to the shower axis, all weighted by excess tracklength.

We measure the artificial variation in any given quantity introduced by the thinning algorithm using two methods. One is to generate a large number of thinned and unthinned showers, and compare the variation in (for example) excess tracklength T within the thinned showers to that of unthinned showers. Following (69), this allows the definition of the quality parameter Q_X for an arbitrary quantity X and a particular thinning level as $Q_X \propto (t^{0.5} V_X)^{-1}$, where t is the CPU time taken, and V_X is the relative variation (standard deviation) in the quantity X . Both t , V_X and hence Q_X depend on the thinning level f_L . Here, we use the number of particle tracks n as proxy for CPU time, and normalise quality

by that for an unthinned shower. Our definition then becomes:

$$Q_X = \frac{n_0^{0.5} V_X^0}{n^{0.5} V_X} \quad (5.32)$$

where V_X^0 stands for the intrinsic variability of the quantity X in the set of unthinned showers. For a purely random quantity X , $Q_X = 1$, since the variation V_X will go as $n^{-0.5}$ and t goes as n . If the chosen thinning parameters are such that $Q_X > 1$ then \bar{X} (the mean value of X) is better estimated in a given time from the thinned showers, and on the contrary $Q_X < 1$ means that the fully simulated showers give the best estimation. The selection process for a suitable thinning algorithm using this method involves choosing a set of thinning parameters (thinning level f_L and E_{\min}) corresponding to computer-time constraints, and selecting those which give the highest quality parameter estimates for the set of relevant observables in Table 5.1.

The above method is relevant for estimating the mean properties of electromagnetic showers at a given energy. It cannot however determine the absolute accuracy in estimating a quantity X for an individual shower. To do this we have simultaneously run multiple copies of the same shower each of them with different thinning parameters (f_L, E_{\min}) and/or algorithms. In this way we minimise the effect of the intrinsic shower fluctuations and we are confident that the differences in the values of the observables are solely due to the different thinning levels. For this purpose the ZHS code was modified to allow exactly the same shower tracks to be assigned different weights according to different thinning parameters and/or algorithms. Observables such as those in Table 5.1 estimated from the thinned showers can then be compared to their “true” (i.e. fully simulated) values. Optimal thinning parameters and/or algorithms can then be selected based on their attaining a desired level of accuracy and maximising time reduction. Here, we require a minimum level of accuracy of 10% between the thinned and the fully simulated showers, which is at the level of agreement between different cascade simulations below PeV energies (16).

5.4.1 Quality of thinning

This section – Sec. 5.4.1 – comes from Sec. 3.1 of Ref. (18).

We have calculated the quality factor corresponding to different combinations of thinning parameters (f_L, E_{\min}) and to the observables in Table 5.1 for primary shower energies of $E_0 = 1$ TeV to $E_0 = 10$ PeV, using one hundred showers for each energy and thinning parameter set.

The quality of length-related parameters $Q_{(z)}, Q_{M_z}$ was found to be the highest among the observables in Table 5.1 for all thinning parameters and algorithms, since length is determined by the high-energy component which remains unaffected by thinning. Conversely, both the quality of total excess tracklength Q_T , and measures of width $Q_{(r)}$ and Q_{M_r} , were found to be the lowest, since these are governed by the low-energy (and thus most highly thinned) component on the shower. Excepting a very few cases, we found $Q_T < \{Q_{(r)}, Q_{M_r}\}$,

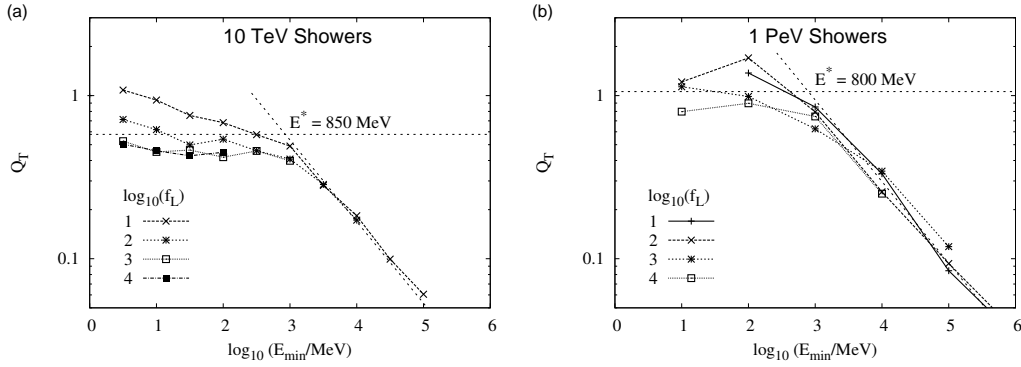


Figure 5.4: (Reproduced from Fig 1 of Ref. (18)) Quality of total excess track-length Q_T in (a) 10 TeV and (b) 1 PeV showers for various f_L as a function of E_{\min} using the standard algorithm. The (dotted) trend lines are found by fitting a constant k only to the range $E_{\min} \leq 300$ MeV data for $Q_T = k$ (flat line), and $E_{\min} \geq 3$ GeV for $Q_T = k E_{\min}^{-0.5}$ (sloped line) — not all fitted points are shown for purposes of clarity. The $Q_T = k E_{\min}^{-0.5}$ trend is clearly present for $E_{\min} \gtrsim 1$ GeV.

and therefore in this section we choose the worst possible case to select the optimal thinning parameters and use Q_T exclusively, keeping in mind that the quality factor of the other observables is almost always better.

In Fig. 5.4, we have plotted our results for Q_T against E_{\min} for 10 TeV and 1 PeV showers in ice. Two regimes for E_{\min} are clearly visible. Q_T decreases as $E_{\min}^{-0.5}$ for $E_{\min} \gtrsim 500$ MeV, and is stable with E_{\min} for $E_{\min} \lesssim 500$ MeV. These two regimes are apparent in the whole energy range we have explored, $1 \text{ TeV} < E_0 < 10 \text{ PeV}$.

This behaviour arises because increasing E_{\min} reduces the thinning level – and hence the artificial variation due to thin sampling – proportionately, while the artificial variation associated with not enforcing energy conservation is essentially unaffected. Since increasing E_{\min} also proportionately increases the number of particles that are explicitly tracked (and hence the time taken), Q_T stays stable if the dominant component of artificial variation is due to thin sampling. However, for sufficiently large E_{\min} , variability due to not enforcing energy conservation will dominate ¹, so that Q_T will decrease as $E_{\min}^{-0.5}$. The turnover energy between $E_{\min} = 500$ MeV and $E_{\min} = 1$ GeV observed in Fig. 5.4 gives the point where the two components of artificial variation are equal. Altering either E_0 or E_{\max} changes the contribution of both sources of variability equally, so the critical energy at which they equate is independent of both parameters.

Though we have a condition on E_{\min} so that the “standard thinning algorithm” produces maximum values of quality which are stable with E_{\min} , the question remains as to whether or not it is in fact a superior thinning algorithm. Since the “alternative thinning algorithm” forces conservation of energy by not thinning over the thresholds, the time reduction will clearly be less than in the case of the standard thinning algorithm for a given set of thinning

¹Energy is only conserved explicitly when E_A, E_B and $E_C \in (E_{\min}, E_{\max})$




Figure 5.5: (Reproduced from Fig. 2 of Ref. (18)) Comparison of quality in total excess tracklength Q_T in 100 TeV showers between the two thinning algorithms, for the full range of thinning parameters (f_L, E_{\min}), plotted as a function of f_n (see text). For each E_{\min} , f_L is increasing to the right in factors of 10, beginning at $f_L = 10$.

parameters. From our simulations we found that the alternative thinning routine forcing energy conservation explicitly takes between one and two orders of magnitude longer than the standard routine.

Given the difference between f_L and f_n (the actual reduction in computing time) for the two thinning algorithms, the quality Q_T of excess tracklength must be compared in the context of f_n , since this is an actual measure of the time saved. We therefore plot Q_T against f_n in thinned 100 TeV showers using both algorithms, shown in Fig. 5.5. For the same reduction in computing time, the standard thinning algorithm provides a higher quality in T than the alternative algorithm, provided the condition $E_{\min} \lesssim 500$ MeV is met, and even when E_{\min} is slightly greater, e.g. 1 GeV. These results held throughout the 1 TeV–10 PeV energy range studied. Therefore, we reject the alternative thinning algorithm, and proceed to analyse the standard thinning algorithm only for $E_{\min} \lesssim 1$ GeV by simultaneously running the same shower with different thinning parameters.

At this stage we briefly summarize our findings:

- When $E_{\min} < 1$ GeV, the quality factor $Q_T \sim O(1)$ regardless of thinning level f_L .
- When $E_{\min} > 1$ GeV for the standard algorithm, the quality factor $Q_T < 1$ and declines rapidly with increasing E_{\min} .
- In the standard thinning algorithm (energy conserved statistically) the reduction in time is larger than in the alternative thinning algorithm (energy conserved explicitly at every vertex) for any given E_0, E_{\min}, f_L .

- For the same reduction in time, quality is better in the standard thinning routine (provided $E_{\min} \leq 1$ GeV).
- The above conclusions hold for a wide energy range (1 TeV to 10 PeV) and different media, and the behaviour of the standard thinning routine is stable over this range.

5.4.2 Addendum to the testing of thinning methods

Comparison of quality for different parameters

In Sec. 5.4.1 I claimed that the quality of total excess tracklength Q_t is the least of the five observables, thus justifying it as a proxy for the accuracy of all. This claim was made from results similar to that presented in Fig. 5.6, which plots the artificial variation in the five observables. The variation in the moments M_r and M_Z have been halved to compensate for the fact that these are squared quantities. Except for $E_{\min} = 1$ GeV, the total excess tracklength T is the quantity with the poorest estimate. Note that increasing both the primary energy and thinning level by a factor of 10 while keeping the E_{\min} constant produces essentially the same level of artificial variation, indicating that we can increase primary energy and thinning level proportionately while keeping both accuracy and time constant.

The critical energy E_{\min}^* for different media

In Fig. 5.7, I show plots of Q_T for 10 TeV showers in four different media. While the densities are the same as in ice, the variation in the composition (atomic-number density) will strongly affect the relative cross-sections of different interaction processes. Therefore, I use both pure lead and pure beryllium as extreme examples of high- and low-Z materials, with zirconium being a mid-range example, and ice included for comparative purposes. The thinning level chosen was only $10^{0.5}$, in order to be more sensitive to a particular energy range.

From Fig. 5.7, quality is generally decreasing downwards with increased E_{\min} even at low energies. That no $Q_T \approx 1$ behaviour is seen is typical of low values of f_L – the same results are evident in Fig. 5.4 (a) and (b). However, turnover energies for E_{\min} evidently exist between 500 MeV and 1 GeV depending on the medium, and in all cases for $E_{\min} \gtrsim 1$ GeV, the $Q_t \propto E_{\min}^{-0.5}$ trend is obvious. Most importantly, the quality is good for $E_{\min} < 500$ MeV, so that points 1 and 2 of the summary in Sec. 5.4.1 hold for these extreme media. Therefore we concluded that our findings could be generalised to all media.

Accuracy of parameter estimates

Using the simultaneous-thinning method developed by my colleague, I was able to calculate various measures of the variation as given in Tbl. 5.2 (taken from our paper, Ref. (18)). The table gives the three measures of variability as defined in Sec. 5.3.3 for 100 TeV and 1 EeV showers in ice. The moment of radius M_r shows the greatest amount of artificial variation,

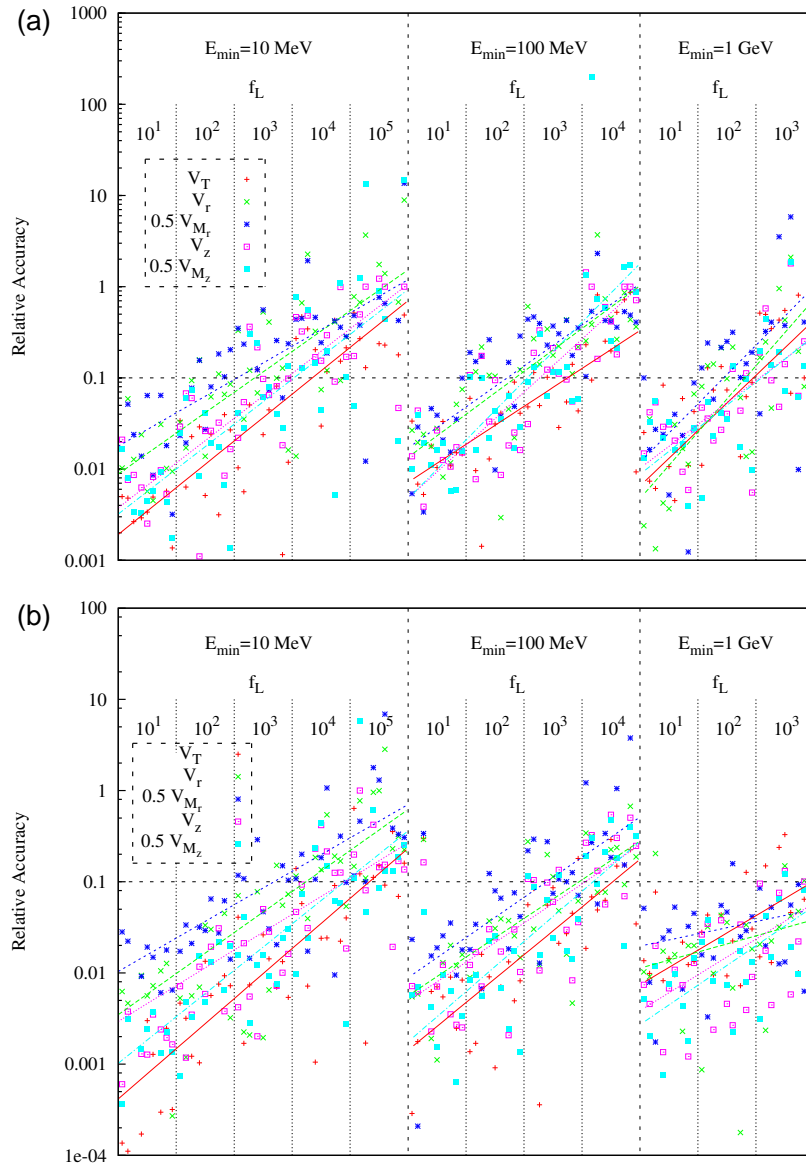


Figure 5.6: Relative variation of thinned simulations from their simultaneously unthinned counterparts, for (a): 10 TeV, and (b): 100 TeV showers in ice. At each energy, three independent shower simulations were run, each thinned three times at each of many combinations of thinning parameters. Within each bin, the horizontal axis is meaningless, and is only included for purposes of clarity. The simple least-squared fits in log-log space are plotted through the central x -value of each bin.

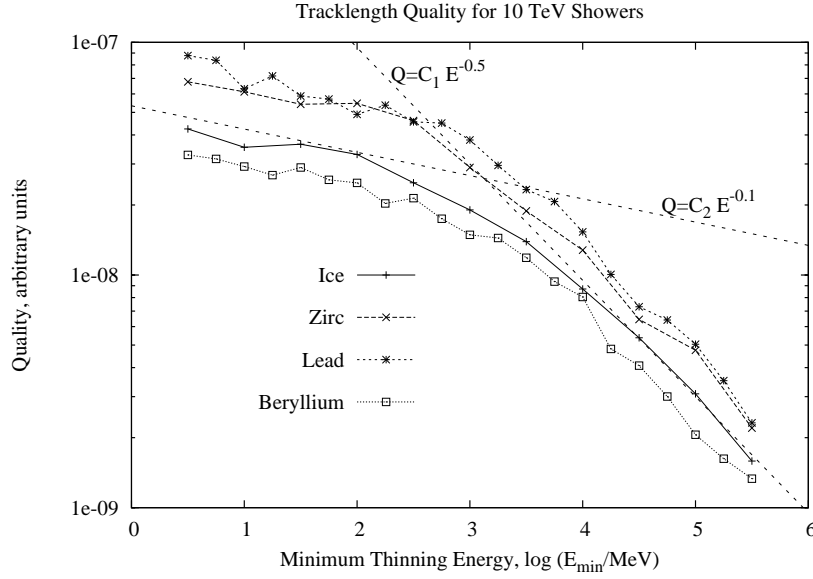


Figure 5.7: Q_T (arbitrary units) against E_{\min} for $f_L = 10^{0.5}$ in various media. Also plotted are by-eye fits to the high- E_{\min} and low- E_{\min} parts of the graph. Each point represents the average of 100 simulations.

which pushes against our limit of a 10% accuracy requirement when $f_L = 10^{-6}(E_0/1 \text{ MeV})$ and $E_{\min} = 10 \text{ MeV}$, and breaks it when $E_{\min} = 100 \text{ MeV}$. Generally we therefore run at $f_L = 10^{-7}E_0/1 \text{ MeV}$.

A note on further findings

Further tests were carried out by J. Álvarez-Muñiz using the simultaneous-thinning technique, which is described broadly in Sec. 5.4. This was almost exclusively his work, and I refer readers to Sec. 3.2 of our paper – Ref. (18) – for the details.

Another important technique used by J. Álvarez-Muñiz to assess the validity of thinned simulations of UHE showers at energies where full simulations become infeasible is to use ‘hybrid’ simulations. In such a simulation, showers below some energy (which must be below the LPM energy) have their tracklength approximated by a one-dimensional Greisen parameterisation (51; 66). While a hybrid simulation can only estimate the total excess tracklength via an approximation such as $T_{e^-e^+} = kT_{e^-+e^+}$ (excess tracklength proportional to the total tracklength) for some constant k , the ZHS-thinned code can exactly calculate the 1D distribution of total tracklength, allowing a comparison. Again, this showed for energies at 1 EeV and over that our thinning method was accurate.

5.4.3 Summary of thinning methodology

This section – Sec. 5.4.3 – comes from Sec. 3.2 of Ref. (18)

f_L	E_{\min}	$\frac{\Delta T}{T}$ (%)	$\frac{\Delta r}{r}$ (%)	$\frac{\Delta M_r}{M_r}$ (%)	$\frac{\Delta z}{z}$ (%)	$\frac{\Delta M_z}{M_z}$ (%)
$E_0 = 10^{14}$ eV		\bar{V}_X^0 : intrinsic variation between fully simulated showers.				
1	0	0.21	0.19	0.44	4.1	7.8
$E_0 = 10^{14}$ eV		\bar{V}_X : variation of thinned from fully simulated showers.				
10	10	0.19	0.7	1.5	0.15	0.3
10	100	0.5	1.0	3.4	0.25	0.5
100	10	0.7	2.2	7.8	0.5	0.95
100	100	1.9	3.3	10.5	0.9	1.7
$E_0 = 10^{18}$ eV		\bar{V}'_X : variation between thinned showers only.				
10^5	10	0.23	0.65	2.4	0.16	0.3
10^5	100	0.55	0.95	3.3	0.28	0.55
10^6	10	0.65	2.2	7.5	0.5	0.9
10^6	100	1.55	2.9	9.5	0.65	1.4

Table 5.2: (Table 3 from Ref. (18)) Variation in shower parameters (see Table 5.1) for sample optimal thinning parameters. \bar{V}_X^0 is the variation in parameters between unthinned showers (the ‘intrinsic variation’; see Eq. (5.26)); \bar{V}_X the variation of the parameters from simultaneous thinnings of the same shower about the unthinned (fully simulated) values (‘artificial variation’; see Eq. (5.29)); \bar{V}'_X – used only at energies where full simulations become impossible – is the variation between simultaneous thinnings of the same shower (i.e. it estimates the artificial variation; see Eq. (5.31)).

Summarising, we have found that the thinned showers generated by our thinning algorithm with optimal combinations of parameters reproduce the Cherenkov radiation spectrum from full simulations with an excellent degree of accuracy up to energies of ~ 1 PeV. Above this energy full simulations become infeasible, but we have shown that the frequency spectra are internally consistent to within a few percent for individual showers. Moreover, a totally independent hybrid simulation gives results for the longitudinal development in very good agreement with the optimally thinned simulations. The time-reduction thinning allows at the highest energies is huge: whereas the radiated spectrum from a 100 EeV shower would have taken a few years to calculate with the standard ZHS routine, using our ZHS-thinned code we can now achieve acceptable accuracy in less than an hour. This enables us to run numerous showers of energies above 1 EeV, and we therefore develop parameterisations for the Cherenkov radiation produced based on these thinned showers for practical applications.




Figure 5.8: (Fig. 6 from Ref. (18)) The “box model” of shower development over the lifetime of the shower. A shower has an effective length L and width R , giving characteristic time delays of δt_L and δt_R . For $\delta t_L = 0$, we obtain the condition $L \cos \theta = L/n$ for the Cherenkov angle $\theta = \theta_C$.

This is the subject of the next section.

5.5 Model and Parameterisation for Coherent Cherenkov Radiation

This section – Sec. 5.5 – comes from Sec. 4 of Ref. (18)

The physical basis for our parameterisation of the radiated spectrum is the “box model” of shower development (Fig. 5.8), which has been used to explain the radiated spectrum from fully simulated showers at sub-LPM energies in a variety of media (17; 123; 33). In this model, the distribution of particle tracks making up the shower over its lifetime has a characteristic length L and width R , from which the resultant time delays δt_L and δt_R from radiation emitted over the length and breadth of the shower – with respect to a particle travelling along the shower axis at the speed of light – can be calculated with simple geometry (see Fig. 5.8) for any viewing angle. The condition $\delta t_L = 0$ (no delays associated with the longitudinal spread of the shower) defines the Cherenkov angle θ_C , so that interference over the width of the shower only (lateral decoherence) determines the spectral shape when the shower is viewed near θ_C , while far from θ_C , interference over the shower length (longitudinal decoherence) dominates (17). The radiated spectrum $\vec{E}(\nu)$ can then be described as a fully coherent amplitude A — increasing linearly with frequency — multiplied by lateral and longitudinal decoherence factors d_R and d_L , defined such that

$d_{R,L} \leq 1$ with $d_{R,L} \rightarrow 1$ as $\nu \rightarrow 0$. Also, $d_L \equiv 1$ at the Cherenkov angle $\theta = \theta_C$. In general, A , d_L and d_R will be dependent upon the shower energy E_0 , frequency ν , and viewing angle θ with respect to the shower axis, leading us to the following functional form:

$$r|\vec{E}(E_0, \theta, \nu)| = A(E_0, \theta, \nu) \times d_L(E_0, \theta, \nu) \times d_R(E_0, \theta, \nu) \quad (5.33)$$

where r is the (far-field) observation distance. In this paper, we always include both decoherence terms, although only for a small range of angles will both be relevant at the same time. We use the forms of Álvarez-Muñiz et al.(17) — given in Eq. (5.34) — for the decoherence factors, which have been shown to provide a good fit to the spectrum for showers with primary energies below E_{LPM} :

$$d_{R[L]} = \frac{1}{1 + (\nu/\nu_{R[L]})^{\bar{\alpha}[\bar{\beta}]}}. \quad (5.34)$$

Here, ν_R and ν_L are characteristic frequencies at which lateral and longitudinal decoherence become important, while $\bar{\alpha}$ and $\bar{\beta}$ give the strength of the decoherence. The frequencies ν_L and ν_R will be inversely proportional to the time delays δt_L and δt_R , which are in turn functions of L and R . The amplitude A of the fully coherent (low-frequency) component will be proportional to the total excess tracklength T . The proportionalities of L , R , and T with properties of the interaction medium are well-established theoretical results (123; 33; 11; 96; 97). We use the formalism of Álvarez-Muñiz et al. (17):

$$\nu_L(\theta) \approx \frac{c}{L|1 - n \cos \theta|} = \frac{\rho}{\bar{k}_L X_0 |1 - n \cos \theta|} \quad (5.35)$$

$$\nu_R(\theta = \theta_C) \approx \frac{c/n}{R} = \frac{\rho}{\bar{k}_R R_M \sqrt{n^2 - 1}} \quad (5.36)$$

$$A(E_0, \theta, \nu) = \bar{k}_E \nu T \sin \theta \approx \bar{k}_E \frac{E_0 X_0}{E_C \rho} \nu \sin \theta \quad (5.37)$$

where the medium properties X_0 , n , ρ , R_M and E_C are respectively the radiation length (g cm^{-2}), refractive index, density (g cm^{-3}), Moliere radius (g cm^{-2}), and critical energy (MeV). \bar{k}_L , \bar{k}_R , and \bar{k}_E are proportionality constants. Throughout, we use the approximation $\nu_R(\theta) = \nu_R(\theta_C)$, since lateral decoherence is only important near the Cherenkov angle. Simulations by Álvarez-Muñiz et al. (17) in ice, salt, and the regolith have shown that variation in \bar{k}_L and \bar{k}_R between these media for showers at sub-LPM energies is of the order of 10%, while for \bar{k}_E it is of order 30%. Also, and as we have previously mentioned, L is dependent upon the shower energy E_0 , so that the “constant” \bar{k}_L is in fact energy-dependent. For sufficiently low-energy showers, the dependence of L on E_0 is small, while at energies where LPM-elongation is significant (typically for $E_0 \gtrsim E_{LPM}$), L increases rapidly with

E_0 (109; 11; 93; 70). Thus we expect \bar{k}_L to be fit well by the following relation:

$$\bar{k}_L = \begin{cases} k_0 (E/E_{\text{LPM}}^*)^{\gamma_0} & E < E_{\text{LPM}}^* \\ k_0 (E/E_{\text{LPM}}^*)^{\gamma_1} & E > E_{\text{LPM}}^* \end{cases} \quad (5.38)$$

where k_0 , γ_0 and γ_1 are constants, and E_{LPM}^* is an effective energy – typically larger than E_{LPM} itself – at which LPM-elongation of the shower longitudinal profile begins to affect the emitted electric field. The quantities \bar{k}_E and \bar{k}_R are energy-independent however, since both excess tracklength and shower width are determined by low-energy interactions. Likewise, we expect $\bar{\alpha}$ to be constant, while in general $\bar{\beta}$ may vary at high energies as the shape of the shower changes. Note that while our toy model of shower development describes the shower as a box, our parameterisation Eqs. (5.33), (5.34) do *not* assume this shape, with the parameters $\bar{\alpha}$ and $\bar{\beta}$ related to the precise distribution of tracklength within the box.

Our fitting procedure is thus as follows. For each shower, we fit Eqs. (5.33) and (5.34) to the simulated spectrum at the Cherenkov angle for the parameters \bar{k}_E , ν_R , and $\bar{\alpha}$ by setting $d_L = 1$ ($\nu_L = \infty$), and we then relate ν_R to \bar{k}_R via Eq. (5.36). Given these constraints, we allow ν_L to vary, and then fit for ν_L and $\bar{\beta}$ at various angles away from the Cherenkov angle. Plotting the fitted ν_L against θ gives \bar{k}_L according to Eq. (5.35). We repeat this many times at each energy, and vary the energy in multiples of $\sqrt{10}$ over the energy range 1 TeV to 100 EeV. Above this energy, the cross-section for photo-nuclear interactions becomes comparable to the pair-production and bremsstrahlung cross-sections (68), which is not included in the (purely electromagnetic) ZHS code. However, the energy range applicable to our calculations includes the full range of predictions for a GZK-flux of UHE neutrinos (88; 41; 7), while the relevance of UHE neutrino interactions above ~ 100 EeV reduces as limits on models predicting a flux of UHE ν at energies above that from GZK interactions become stronger (23; 47).

We repeat the above procedure in ice, salt, and the lunar regolith and megaregolith. The parameters for the first three are those of (17), while the megaregolith is treated as per Ref. (62) as regolith with refractive index 2.68 and density 3.0 g/cm^3 .

5.5.1 Results for different media

This section – Sec. 5.5.1 – comes from Sec. 4.1 of Ref. (18)

The fitted parameters \bar{k}_E , \bar{k}_R , $\bar{\alpha}$, $\bar{\beta}$ for each medium are given in Table 5.3 (Table 4 of Ref. (18)). For \bar{k}_L , we find the form given in Eq. (5.38) to be suitable, and the parameters defined therein are summarised in Table 5.4 (Table 5 of Ref. (18)). We also find an upward trend in $\bar{\beta}$ with energy for all media, consistent with the shower becoming more elongated and less singly-peaked. Since the effect is small, due to the emission falling rapidly for $\nu > \nu_L$, we do not include this trend in our parameterisation. While all the fitted parameters vary on a shower-to-shower basis, variations tended to be small, so that approximating $k_E \approx \bar{k}_E$, $\beta \approx \bar{\beta}$, etc. for individual showers is appropriate. The exception is k_L , where we find (not unexpectedly (11; 93; 70)) that variations from the mean fitted values \bar{k}_L are

Medium	\bar{k}_E	\bar{k}_R	$\bar{\alpha}$	$\bar{\beta}$
Ice	$4.65 \cdot 10^{-16}$	1.54	1.37	2.74
Salt	$3.12 \cdot 10^{-16}$	1.39	1.32	2.70
Regolith	$3.37 \cdot 10^{-16}$	1.50	1.32	2.79
Megaregolith	$3.46 \cdot 10^{-16}$	1.47	1.32	2.68

Table 5.3: (Table 4 from Ref. (18)) Fitted shower parameters, as defined by Eqs. (5.35)-(5.37). Units for \bar{k}_E are V/cm/MHz². \bar{k}_L (not shown in this table) was fitted to Eq. (5.38) with parameters given in Table 5.4.

Medium	$\log_{10} k_0$	$\log_{10} (E_{\text{LPM}}^*/\text{eV})$	γ_0	γ_1
Ice	1.52	16.61	$5.59 \cdot 10^{-2}$	0.39
Salt	1.71	16.44	$8.21 \cdot 10^{-2}$	0.44
Regolith	1.58	16.40	$6.71 \cdot 10^{-2}$	0.43
Megaregolith	1.57	16.38	$6.39 \cdot 10^{-2}$	0.46

Table 5.4: (Table 5 from Ref. (18)) Parameters of the fit for \bar{k}_L , as defined by Eq. (5.38).

large, particularly at high energies. These fluctuations will dominate the radiation pattern in the regime where longitudinal decoherence dominates, which at the highest energies is everywhere except within a small fraction of a degree from θ_C . We assume that at a given energy k_L is log-normally distributed about the mean \bar{k}_L as given by Eq. (5.39), with an energy-dependent variance in log-space of $\sigma_{k_L}^2$:

$$\log_{10}(k_L) \sim N\left(\log_{10}(\bar{k}_L), \sigma_{k_L}^2\right). \quad (5.39)$$

Fitting for σ_{k_L} in each medium, we find that above a characteristic energy E_σ , the variation in shower length increases rapidly, while below E_σ the variation is approximately constant and very low. This leads us to use a parameterisation for σ_{k_L} — Eq. (5.40) — similar to Eq. (5.38) for k_L itself:

$$\sigma_{k_L} = \begin{cases} \sigma_0 + \delta_0 \log_{10}(E/E_\sigma) & E < E_\sigma \\ \sigma_0 + \delta_1 \log_{10}(E/E_\sigma) & E > E_\sigma \end{cases} \quad (5.40)$$

with parameters σ_0 , E_σ , δ_0 , and δ_1 again fitted separately to each medium, with results given in Table 5.5 (Table 6 of Ref. (18)). These parameters complete the parameter set with which we characterise electromagnetic showers.

Medium	σ_0	$\log_{10}(E_\sigma/\text{eV})$	δ_0	δ_1
Ice	$3.39 \cdot 10^{-2}$	14.99	~ 0	$2.25 \cdot 10^{-2}$
Salt	$4.98 \cdot 10^{-2}$	14.99	~ 0	$2.40 \cdot 10^{-2}$
Regolith	$5.04 \cdot 10^{-2}$	14.98	~ 0	$2.44 \cdot 10^{-2}$
Megaregolith	$5.19 \cdot 10^{-2}$	14.99	~ 0	$2.50 \cdot 10^{-2}$

Table 5.5: (Table 6 from Ref. (18)) Parameters relating to σ_{k_L} (the variation of $\log_{10}(k_L)$ about $\log_{10}(\bar{k}_L)$) as defined by Eq. (5.39) and Eq. (5.40).

5.5.2 Comparison with other calculations and discussion

This section – Sec. 5.5.2 – comes from Sec. 4.2 of Ref. (18).

Comparisons of our parameter estimates with those from previous simulations are only possible for ice over the entire energy-range, and for salt and the regolith for $E_0 < 100$ TeV. The parameter estimates obtained by Álvarez-Muñiz et al. (17) for ice and the regolith at low energies agree with our results to within $\sim 3\%$ for \bar{k}_E , and within 10% for \bar{k}_R and $\bar{\alpha}$ in all media. The larger differences in the shape parameters are not unexpected, since these vary more on a shower-to-shower basis, and the fitted values are partially subject to the fitting procedure.

Comparisons of \bar{k}_L and $\bar{\beta}$ with the estimates in (17) must be viewed with caution. Here these parameters have a slightly different meaning to those obtained in (17), since we retain the transverse decoherence term d_R from Eq. (5.33) at all angles. A better comparison is to the one-dimensional results of Álvarez-Muñiz et al. (11), who found a parameter set approximately corresponding to $\{\log_{10}(E_{LPM}^*), \gamma_1, \gamma_2\} = \{16.15, 0.03, 0.3\}$, whereas here we find (for our full 3-D simulation) the set $\{16.69, 0.04, 0.45\}$. The most important update to this parameterisation is the steeper increase in shower length with energy, and hence narrower emission at the highest energies.

In Fig. 5.9, we plot the radiated emission at various angles from a 100 EeV shower in ice (Fig. 5.9(a)) and a 100 EeV shower in the megaregolith (Fig. 5.9 (b)), which is the least “ice-like” medium. For each shower, we show (i): the spectra from a particular simulated shower as calculated by ZHS-thinned, obtained using three simultaneous calculations with $(E_{\min}, f_L) = (100 \text{ MeV}, 10^7)$ – we also plot (ii): the fitted spectrum for that particular shower using the functional forms and fitting procedure described in this paper, (iii): the parameterisations presented here, using three values of k_L ($\bar{k}_L - \sigma_{k_l}$, \bar{k}_L , and $\bar{k}_L + \sigma_{k_l}$), and (iv): the parameterisation of Álvarez-Muñiz et al. in (13) for ice and extrapolated from its known range of validity (100 TeV - 1 PeV at most), used as published in ice in the top panel, and scaled as per James et al. (62) for the megaregolith. Also shown (v) is the spectra calculated by Fourier-transforming the distribution of excess tracklength from each of the three simulated showers (13), normalising by our fitted constants \bar{k}_E , and multiplying by the lateral decoherence factor d_R .

The broad range of expected spectra given by the variation in k_L is immediately ap-

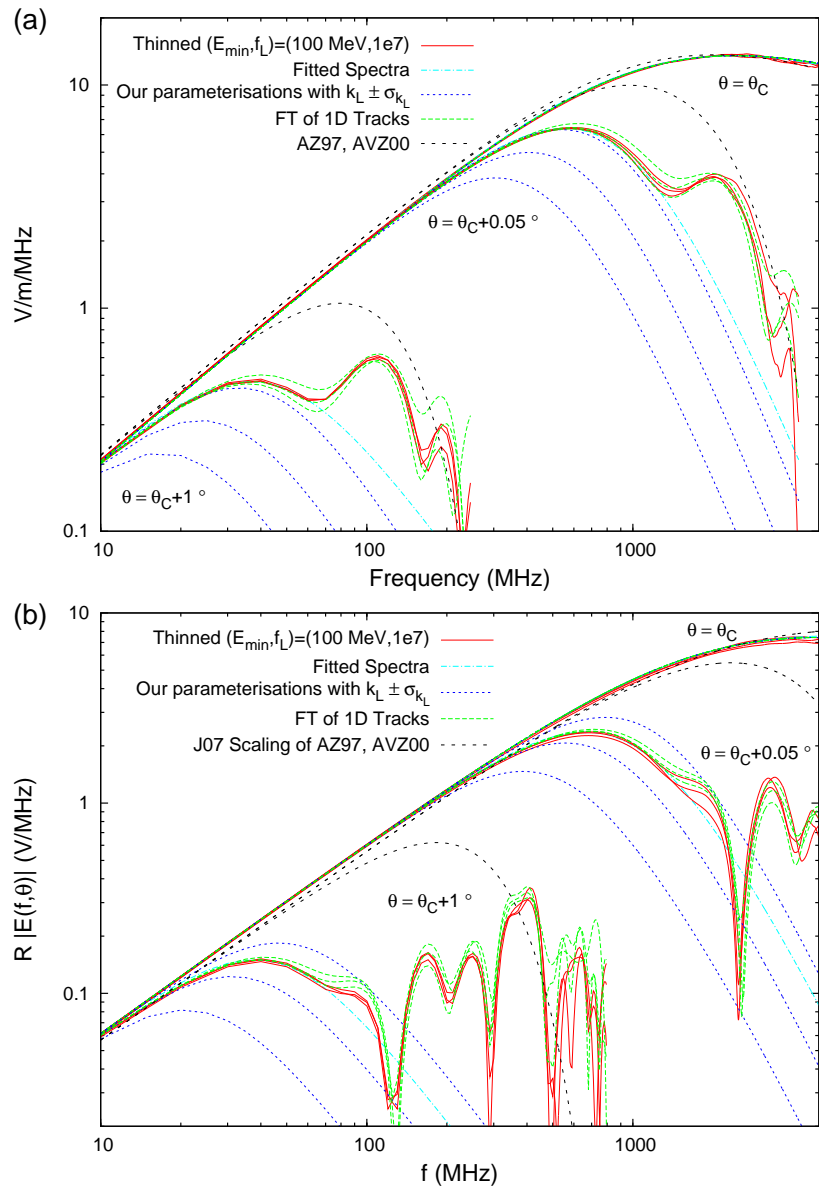


Figure 5.9: 100 EeV electron showers in ice (top panel), and the lunar megaregolith (bottom). Shown are both direct calculations of the spectra and parameterisations (see text for details).

parent: whereas we have chosen a typical shower in the megaregolith, the shower in ice is significantly shorter than average – more than 1σ – as evinced by the high turn-over frequencies. It is also clear that previous parameterisations which were not valid away from the Cherenkov angle for energies above E_{LPM} overestimate the width of the Cherenkov cone at the highest energies. This emphasises the need for direct simulations such as those presented in this work. The overestimation is also greater in the megaregolith than in ice. The effect on the sensitivity of neutrino-hunting experiments is not likely to be large however, since the shorter hadronic (and subsequent electromagnetic) cascades expected from neutral-current ν -interactions are more readily detectable.

The similarity of our parameter estimates for the regolith and megaregolith (identical chemical composition, different densities and refractive indices) reflects the suitability of the scaling inherent in the box-model of shower development. The largest relative parameter difference between these media was 7%, and in particular, the estimated values for k_E (which have shown the greatest variation (17)) were within 3%. This compares to differences between salt-regolith parameter estimates (20%) and ice-anything ($> 30\%$), indicating that the variation is due to subtle compositional effects, rather than easily-characterisable bulk properties. We therefore expect that the box-model can be used with confidence to scale results for standard ice to variants thereon found deeper in the Antarctic ice sheet, or equally from regolith/megaregolith to intermediate/transitional layers.

5.6 Secondary Peaks

The remarkable agreement of thinned simulation results (solid lines in Figs. 5.9) allow us to examine in more detail the radiation pattern away from the primary maxima. Note that calculating the radiated spectra from the 1-D tracklength distribution also provides a reasonable estimate of these secondary peaks, but not as good as with a direct calculation of the electric field from the tracks themselves. However, the time saved is significant, and far from the Cherenkov angle the approximation improves.

Previous parameterisations of coherent Cherenkov radiation — as well as ours presented here — model only the emission in the primary maxima, whereas secondary fringes of various heights are to be expected and are accurately reproduced by our code (to within 10% in the case of the second maxima using standard thinning levels). Indeed, for a long, thin shower, with a uniform distribution of excess tracklength, we expect the spectrum to go approximately as $|E(\theta, \nu)| \sim \cos(\nu/(\nu_0 \Delta\theta))$, where $\Delta\theta = \theta - \theta_C$ and ν_0 is some constant. That is, the height of all secondary peaks should be equal to the height of the primary peak. From Fig. 5.9(b), it is evident that the $|E| \sim \cos(\nu)$ behaviour is reproduced at least qualitatively: the frequencies of the first three minima at $\Delta\theta = +1^\circ$ are approximately equally-spaced at 130, 290, and 490 MHz. Importantly, the power contained in the secondary maxima in both the ice and megaregolith showers at $\Delta\theta = +1^\circ$ is greater than that in the primary maxima.

To investigate the behaviour of the radiated spectrum beyond ν_L (i.e. that part containing

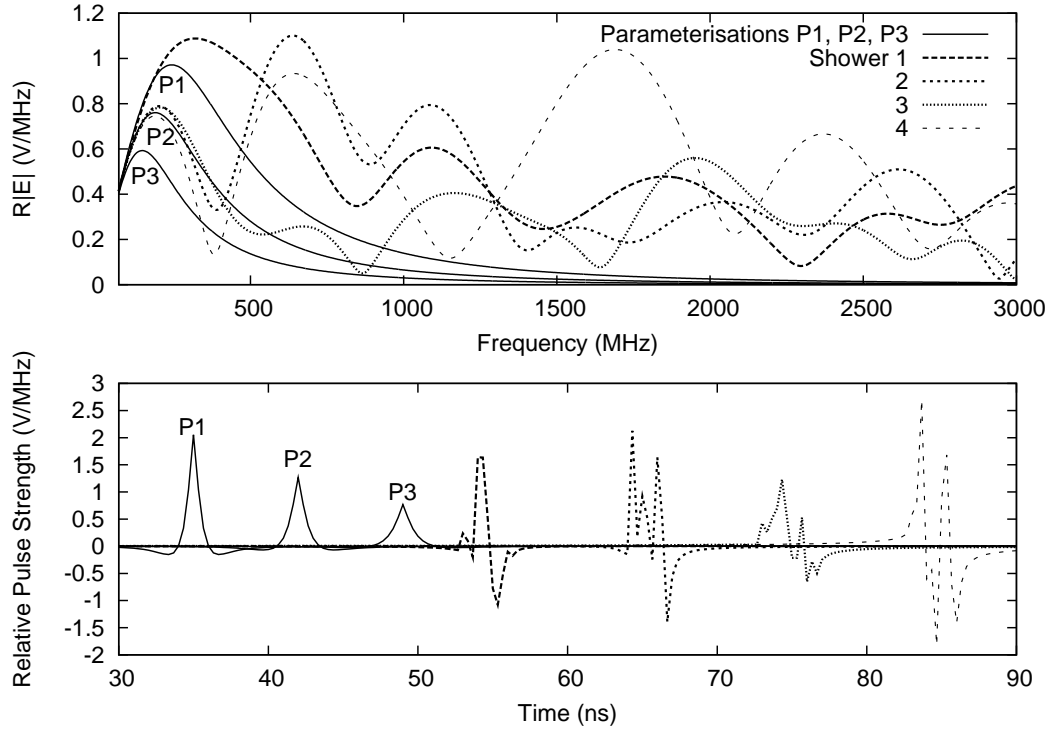


Figure 5.10: Four 100 EeV electron showers in the lunar megaregolith. Top panel: frequency spectra at $\theta = \theta_C + 0.2^\circ$, together with our parameterisations from this work. Bottom panel: time-domain transforms of the above spectra over a 0–3000 MHz bandwidth, with arbitrary vertical scale and horizontal offsets. The parameterisations P1, P2, and P3 respectively use $k_L = \bar{k}_L - \sigma_{k_L}$, \bar{k}_L , and $\bar{k}_L + \sigma_{k_L}$.

the secondary maxima), we ran a small number of very detailed simulations of 100 EeV showers, calculating the spectra over a small range of angles only. Our results for four showers in (arbitrarily) the megaregolith at $\Delta\theta = +0.2^\circ$ only, together with the expected time-domain pulses over a 0–3000 MHz bandwidth, are shown in Fig. 5.10. Note that we have not plotted the spectra on the usual log-log scale, since this tends to mask the significance of high-frequency peaks. Since the variation in the frequency spectrum for $\nu > \nu_L$ between the showers plotted in Fig. 5.10 is very large, we do not here attempt to parameterise the spectrum beyond the primary maxima. Instead, we discuss the results qualitatively, and leave such a parameterisation (or more likely, random generation method) to a future work.

In two of the four showers shown in Fig. 5.10 (upper panel), the power contained in at least one of the secondary maxima at $\Delta\theta = +0.2^\circ$ is greater than that of the primary maxima — in general, we find the fraction to be slightly higher than this. To first order, the spectrum scales with angle as $E(\nu, k\Delta\theta) \approx (1/k)E(k\nu, \Delta\theta)$, so that the shape appears constant with angle to within a scale factor. However, the height of the secondary peaks relative to the primary maximum does increase away from θ_C , so that at $\theta_C + 1^\circ$ virtually every shower

in the megaregolith exhibits a secondary peak of greater power than the primary maximum. The width of each secondary peak tends to be similar to that of the primary maxima, i.e. of width $\sim 2\nu_L$, with a tendency for the width to slowly increase with the order of the maxima (as observed previously in Fig. 5.9(b)). Finally, in the regime $\nu > \nu_R$, lateral decoherence renders all secondary maxima negligible, and the aforementioned scaling breaks down.

Our results suggest that the sensitivity of detection experiments to purely electromagnetic showers in the frequency range ($\nu_R > \nu > \nu_L$) could be significant, whereas the steep spectral fall-off beyond ν_L used in current parameterisations (and hence simulations) would suggest an effectively zero detection probability in this range. Naively, there will approximately be a 50-50 chance of a $\nu_R > \nu > \nu_L$ observation falling on a peak in the spectrum rather than a trough, so the detection probability in this range will be of order half that at low ($\nu < \nu_L$) frequencies.

The lower panel of Fig. 5.10 shows the time-domain transforms of the spectra in the upper panel, covering the full frequency range 0–3000 MHz. For this broad frequency range, the height range of the simulated pulses is covered by our range of parameterisations, though this would not be true for experiments observing only for $\nu > \nu_L$. However, the simulated pulses do not appear as the clean impulse-like signals which parameterisations predict, but rather exhibit oscillations due to interference between the secondary peaks. Such “noisy” signals would be expected for any broad-bandwidth observation, where the definition of a “broad” bandwidth in this context is $\Delta\nu \gtrsim 2\nu_L$, i.e. greater than the approximate full width of the secondary peaks, and thus dependent on the observation angle. Experiments which are anticipated to be sensitive to events at angles at which the observation bandwidth becomes broad should therefore take care not to too readily classify such signals as RFI.

5.7 Conclusion

The following quoted segment comes from Sec. 5 of Ref. (18)

“We have used thinning techniques to develop a robust version of the ZHS code able to calculate the frequency spectrum of coherent Cherenkov radiation from ultra-high energy (UHE) electromagnetic showers in an arbitrary dense medium. Using this code, we have identified a range of thinning parameters able to produce accurate results with a large reduction in computing time, and extended the energy range at which full calculations of the radiated spectrum are possible from 100 TeV to 100 EeV. This has enabled us to analyse the structure of highly-elongated LPM-affected showers in unprecedented detail, and the parameterisations we have developed should prove a useful and accurate tool in simulations of UHE neutrino detection using the radio technique.

In particular, we have shown that the spectral fall-off due to longitudinal decoherence occurs at a lower frequency than scaling of results from lower energies

would imply, and for the first time give a method of including LPM-induced fluctuations into the parameterisation by varying the length parameter k_L ."

Also, we have drawn attention to the existence of secondary maxima in the frequency spectra beyond the characteristic roll-over, which will have implications for choosing the "optimum" frequency range for experiments which are sensitive to UHE electromagnetic showers.

Data Analysis for LUNASKA Lunar Observations with the ATCA

6.1 Introduction

This chapter describes my work as part of the LUNASKA collaboration's use of the Australia Telescope Compact Array (ATCA) to observe the Moon using the lunar Cherenkov technique. It is broken into eight sections, including this introduction. Almost all collaboration members contributed significantly towards the LUNASKA lunar observations, and at times distinguishing between my contribution (that which should be counted towards my PhD) and that of others can be difficult. This applies chiefly to the first non-trivial section (Section 6.2), which gives an outline of our equipment, methods, and our observations – essentially an overview of LUNASKA's experimental activities towards lunar pulse detection. The work described in this section can therefore be assumed to be a joint collaborative effort included in my thesis for purposes of clarity; I have noted those instances where the work is chiefly that of a particular individual, myself or otherwise.

My involvement was chiefly in the data analysis (although I also participated in the observations and – more importantly – planned the observation dates), which is the subject of the remainder of the chapter, and is entirely my own work using (with minor exceptions noted in text) methods developed entirely by myself. The first 'analysis' section – Sec. 6.3 – calculates the efficiency in the use of the observation time, via estimating the dead-time caused by a trigger event. The next (Sec. 6.4) is also in a way an efficiency calculation, being the efficiency in the use of our raw (collecting-area-limited) sensitivity, which is reduced due to our use of a finite sampling rate and only approximate dedispersion. While the calculations in Sec. 6.3 are standard procedure, those in Sec. 6.4 are – to the best of my knowledge – a first, with no previous experiment taking these effects into account.

The longest section is Sec. 6.5, which describes my work towards aligning the times on each antenna, and is vital to determining the arrival direction of any detected events and therefore their status as lunar or otherwise in origin. It was anticipated pre-experiment that this procedure would be simple; however, due to a programming error in the sampler boards, it proved extremely complex, and much of the chapter is devoted to undoing the complexity. The last purely 'analysis' section, Sec. 6.6, describes my work towards deter-

mining the sensitivity of the experiment to lunar Cherenkov pulses, using the Moon itself as our calibrator.

The results of the experiment are presented in Section 6.7. I describe the search for lunar pulses in each of the observation periods, and the methods to eliminate false RFI events. It should be said now that this search detected no lunar pulses, so I calculate apertures and limits of the experiment based on this result. Finally, I summarise the significance of our observations, and what we have learned, in the Conclusion (Sec. 6.8).

6.2 Description of the Experiment

The ATCA (Australia Telescope Compact Array) is an aperture synthesis telescope located near Narrabri, NSW, Australia, consisting of six identical 22 m dishes. While the West-most antenna is fixed in position, the other five can be moved along a 3 km East-West baseline, and also a ~ 150 m North-South baseline. Antennas are numbered ‘CA01’, ‘CA02’, . . . , ‘CA06’ in order from East to West, and I use these labels throughout.

The ATCA was chosen as an ideal SKA test-bed for the Lunar Cherenkov technique for four major reasons. Firstly, the small antenna size was comparable to the expected dish size for the SKA, and also the beam-size matches the lunar disk at 1 GHz. Secondly, the CABB (Compact Array Broad-Band) upgrade provided us with 600 MHz of bandwidth, and programmable hardware capable of performing pulse-detection in real time. Thirdly, it is an array (albeit a mostly one-dimension one), providing both strong timing discrimination against terrestrial RFI, and suggesting that methods and techniques useful for the ATCA would be applicable to the SKA. Fourthly, it was in Australia, so that the scientists and engineers who were most familiar with the telescope could be accessed to enable such a non-standard experiment to go ahead in the first place.

6.2.1 Observations

The observations described here cover three observing periods spanning a year: May 2007, February 2008, and May 2008, all at the ATCA. Other experimental work at the ATCA has been carried out by the LUNASKA collaboration – both to measure ionospheric dispersion, and also to ensure correct lunar tracking by the antennas – but I was not involved in this. In none of the three periods was the set-up the identical, though the evolution has been sufficiently gradual to describe all observations as one experiment. The aim of the observations reported here was to actually detect lunar Cherenkov pulses, or at least improve our methods in the attempt.

Antenna configurations

Table 6.1 specifies which antennas at what station were used in each observation period. The station numbers are given in the format Wxxx, where ‘W’ stands for the East-West baseline, and ‘xxx’ is the number of ‘station increments’ (units of 15.306 m) West of the

Date		Antenna	Stations			Baselines (m)		
May 2007	5 th -7 th	CA03/04/05	W173	W190	W195	260.2	336.8	76.5
Feb 2008	26 th -28 th	CA01/03/05	W98	W113	W148	229.6	765.5	535.8
May 2008	16 th -18 th	CA01/03/05	W147	W172	W195	382.7	734.8	352.1
	19 th	CA01/03/05	W102	W109	W125	107.2	352.1	244.9

Table 6.1: Positions of the antennas used during the LUNASKA lunar observations, and the baselines in order 1-3, 1-5, 3-5 (3-4, 3-5, 4-5) for 2008 (2007).

East-most station. We attempted to get short baselines if possible, since this reduces the window in the time-domain (caused by the large angular diameter of the Moon) from which true pulses could arrive.

In all my calculations, I assume that the baseline is exactly East-West, and does not follow the curvature of the Earth but is straight. Neither are true, but the associated errors (e.g. 5 cm in altitude) are negligible. The antennas' centre is taken from the ATNF website (Ref. (130)) as being $30^{\circ} 18' 52.048''$ S, $149^{\circ} 32' 56.327''$ E, and 209.3 m above sea level.

Observation modes

We had three observing modes during these observations. In all three, the data returned was first processed by our pulse-detection hardware. This hardware analysed the full bandwidth in real time and triggered (returning data kept in a running buffer) when a signal of appropriate strength was detected. The modes are described below:

- **Pulse detection mode.** In pulse detection mode, we tracked the Moon (centre or limb – see Sec. 6.2.1) with the noise-calibration diode off to reduce false events (see Sec. 6.2.3), trigger rates set at 5-10% of the maximum rate to balance efficiency with sensitivity (see Sec. 6.3), and buffer lengths typically short (except in 2007 – see Sec. 6.2.2) to increase the maximum trigger rate. All available time not spent in other modes was dedicated to pulse detection.
- **Sensitivity calibration mode.** In sensitivity calibration mode, buffer lengths were set to maximum to obtain the highest total data rate (trigger rate times buffer length), with trigger levels set to zero to ensure unbiased samples were obtained. Again, the noise diode was turned off. The antennas were pointed successively at the Moon, and two points in the sky containing no strong sources, one with a similar galactic latitude as the Moon, and one far from the galactic plane. This calibration was performed once per observation period.
- **Timing calibration mode.** In timing calibration mode, trigger levels were set sufficiently high that only the noise diode (which was turned on) could be detected, buffer lengths were maximised to give the greatest possible overlap in the signal between antennas, and the antennas were set to track a strong point-like source, being either M87

(only approximately point-like on our longer baselines) or 3C273. The timing calibration was performed about twice per night in 2008, while it was not implemented in 2007.

Observation times and antenna pointings

We had two main considerations in choosing our observation times, the first being to confine the observations to within the approximate hours of 10 pm to 6 am, in order to have a stable ionosphere (see Sec. 6.2.2). The requirement of the Moon being visible meant that this gave us a window of perhaps five days in every 29.5-day synodic lunar month where the Moon would be sufficiently visible during this period to warrant observations. The second requirement was that the Moon be ‘near’ (within $\sim 35^\circ$, beyond which the ATCA sensitivity drops markedly) particular regions of sky deemed to be of interest. This occurs once per 27.3-day lunar orbit, so that combined, we typically had three good and two marginal periods of a few days each year in which to observe any given source.

The 2007 and February 2008 runs were tailored to ‘target’ a broad ($\gtrsim 20^\circ$) region of the sky near the galactic centre, harbouring the closest supermassive black hole to Earth, and a potential accelerator of UHE CR. Preliminary calculations showed that for beam-sizes similar to that of the ATCA, the greatest total effective aperture (and hence sensitivity to an isotropic or very broadly-distributed flux) is achieved in centre-pointing mode. Since any UHE neutrino flux from this region is likely to be broadly-distributed, we pointed the antennas at the lunar centre for these runs. Our May 2008 observing period targeted Centaurus A only, a nearby active galaxy which could potentially account for two of the UHE CR events observed by the Pierre Auger observatory (1). Regardless of their source, this suggested the likelihood of an accompanying excess of UHE neutrinos, and we could not exclude the possibility of seeing the UHE CR themselves. Since this source is point-like compared to the angular size of the ‘moon-beam’ (effectively the beam-shape of the entire antenna-Moon system – see Sec. 6.7.3), we could achieve maximum sensitivity by targeting the portion of the lunar limb closest to Centaurus A.

Software to track either the lunar centre or an arbitrary point on its surface already existed at the ATCA, though we still had to determine which part of the limb to track. I wrote software to determine the direction of Centaurus A from the Moon at a given point in time, and hence the desired antenna pointing position on the limb (being that closest to Cen A). Converting to approximate selenocentric coordinates allowed NASA’s Horizons data system (129) to produce a pointing solution. System temperature measurements at the time were consistent with these pointing positions, indicating that the procedure worked.

6.2.2 Specialised hardware

Fig. 6.1 gives a diagram of the hardware and signal path at each antenna. In order to perform a search for short-duration lunar pulses, we had to build specialised hardware to detect and store candidate events in real time. For this we used the digital, FPGA-based sampler boards

NOTE:
This figure is included on page 115 of the print copy of
the thesis held in the University of Adelaide Library.

Figure 6.1: Diagram of the signal path at each antenna, adapted from Ref. (79) with the permission of R. A. Robinson.

developed by the ATNF (Australia Telescope National Facility) for the CABB (Compact Array BroadBand) upgrade, each of which could digitise and perform simple logic (from a computer's perspective) on two data streams at a rate of 2.048 GHz. As well as this, we required software to be run on the control-room computers to interface with the CABB sampler boards and record candidate events, and also a hardware method to correct for the dispersive effects of the Earth's ionosphere. Mostly this hardware is one-of-a-kind, so although I had little to no involvement in developing the equipment, I include the details below, which are essential for understanding the experiment.

Data channels and signal path

For each observation period, our specialised hardware was implemented on both linear polarisation output streams from the standard L-band signal path on three antennas. This gave a nominal bandwidth of 600 MHz between 1.2 and 1.8 GHz. For 2007 and February 2008, no filtering was applied to this output, while for May 2008 a filter was implemented to remove the $\nu < 1.2$ GHz component, which was the source of much RFI. Aside from the relatively clean cut-off provided by this filter, our bandpass was by no means a clean box-function over the nominal bandwidth.

Each polarised data stream was fed through an analogue dedispersion filter before sampling, and then into a CABB sampler board. This operated at 2.048 GHz (really 8×256 MHz) at 8-bit effective precision. Since the (then) 100 Mb/s connection to the

control room was inadequate to return the full data rate (2×2048 MB/s per antenna), we had to reduce our data rate by triggering independently at each antenna and returning only chunks of data containing candidate events. The signal was thus copied into both a running buffer and passed to a trigger algorithm (Sec. 6.2.2), and on fulfilling the trigger conditions a portion of the buffer was returned to the control room and recorded.

Trigger logic and levels

The trigger algorithm was set up to be a simple threshold trigger at each antenna – if the square of any single sample on either the A or B polarisation data streams was above a certain value, both buffers were returned. The thresholds were set by specifying a four-digit hexadecimal number, giving a range from 0 to 65,534. Since the squares of the digitised data could only take on 129 values ($0, 1, 2^2, \dots, 128^2$), most values were redundant, so that two slightly different trigger levels would in practice be the same. The thresholds were set so that both polarisations on all the antennas were generating a similar (40-50 Hz) trigger rate. No real-time noise-discrimination (such as not triggering if too many samples in a short time interval were above threshold) was implemented. The trigger levels were generally adjusted only when the slowly-varying background temperature (e.g. from the ground) increased, since periods of intense RFI were short-lived and of sufficient strength that even a very high trigger level would still result in many false triggers.

Dedispersion filters

The dedispersion filters – designed by P. Roberts of the ATNF – were variable-width waveguides of approximately one metre in length, with the output being the continuous sum of reflections along the length. Upon reflection, high frequencies experienced a greater delay than low-frequencies, with the design such that this cancelled out the delay due to ionospheric dispersion at low frequencies. Thus an in-phase signal (e.g. coherent Cherenkov radiation from a UHE particle interaction in the Moon) entering the top of the ionosphere should appear in-phase after dedispersion, provided the correct dispersion measure was used.

R. A. Robinson used NASA data (128) to predict typical values of the ionospheric dispersion at Narrabri, with results reported in Ref. (79). The results show that between the hours of 10 pm and 6 am the ionosphere is comparatively stable, with low and predictable VTEC measures around 7 ± 1.3 TECU. Using Eq. (6.9), this corresponds to a vertical delay of 3.65 ns over a 1.2-1.8 GHz bandwidth. Since the actual delay will reflect the slant angle, we chose to build the filters assuming a 5 ns delay over the band, i.e. a TEC along the line of sight (slant TEC, or STEC) of 9.67 TECU; this is also equivalent to the mean VTEC of 7 TEC with a lunar elevation of 47° . Therefore we expected to lose some sensitivity when the Moon is directly overhead, and also very near the horizon. The sensitivity lost due to deviations of the actual VTEC from the mean and variations in lunar elevation is discussed in Sec. 6.4.

Baseband recording

During our May 2007 observations, we did not have timing information available on each event, and had to rely on the computer clock. Our method to recover the timing was to utilise the baseband recording equipment developed for pulsar astronomy at Swinburne University, which was capable of writing up to four streams of 64 MHz data with 2-bit sampling continuously. This data would have accurate timing information, so that by correlating the appropriate spectral components of the returned buffers with the sub-band data, the correct time could be extracted. For this reason, we set the buffer length to the (then) maximum of 1024 samples, to give the greatest possible correlation. The data acquisition hardware was set to record only one polarisation channel from each antenna, since the two polarisation channels from each antenna are aligned with respect to each-other.

Ultimately, this technique proved impractical, for two reasons. Firstly, although the correct time produced a significant 4σ correlation, the search space was of order ten milliseconds. Thus each event typically produced of order 40 potential candidate correlations, with no further way to distinguish between them. The second was a failure on the CA04 sampler board which changed the sampled frequencies, making the procedure of extracting the sub-band from the returned buffers very difficult. Though for our 2008 observations we could prove the computer clocks were generally accurate to within a few tens of microseconds (which would have reduced the search space sufficiently), this was not necessarily the case during 2007, where the set-up of the network was different.

Since these observations therefore could not produce highly accurate times, we had a very large window within which a three-fold coincidence could potentially be of lunar origin, and consequently our rejection power was much less strong. Nonetheless, we found we were able to exclude the vast majority of candidates, as discussed in Sec. 6.7.1.

6.2.3 Miscellaneous phenomena

Under the following sub-headings I discuss miscellaneous but important issues relevant to the analysis.

Obstruction of CA01 by CA02 during May

The configuration used for May 19th 2008 positioned CA02 partly in the path of CA01, thus possibly obscuring CA01's sensitivity to the Moon. In order to avoid this, CA02 was pointed at a target 90° away from the Moon, so that the actual dish would be out of CA02's path. For such a configuration, no pre-existing shadowing data were available, so that success or otherwise of this effort remains partially unknown. However, the sensitivity calibration of May 19th was performed during this period, which will include a contribution of thermal emission from the CA02 antenna, and exclude a contribution from the Moon. While there is no easy method to distinguish between the two, the similarity of the May




Figure 6.2: (a): A typical noise-cal pulse – in this case, from CA01 A on Feb 26th;
 (b): Predictions of an observed impulse of different origins.

19th results to those of May 18th in Fig. 6.23 suggest there was little effect. Here this is ignored.

The noise diode

During testing, we observed numerous strong pulses similar to that in Fig. 6.2, at a rate of approximately 8 Hz – a typical pulse is shown in Fig. 6.2. It was discovered that these pulses were generated by the switching of the noise-calibration diode, located in the receiver to calibrate the power. Thankfully, the switching of the noise diode could be turned on and off from the control room, avoiding a cost to experimental efficiency, albeit at the cost of measurements of T_{sys} for parts of our bandwidth.

This seemingly trivial piece of information proved exceptionally important at various stages of the experiment. Firstly, from Fig. 6.2, there is obviously a pre-pulse to the noise-cal pulse, arriving approximately 18 ns (peak-to-peak; start-to-start is closer to 14 ns) earlier, and with a height approximately fixed as a fraction of the main pulse. This agrees with the interpretation of the ‘wobbly’ spectrum (seen in Sec. 6.6) given in Sec. 6.2.3. Secondly, since this signal is obviously terrestrial in origin, it allows us to test the use of the predicted dispersion in the Earth’s ionosphere to determine the origin of the signal as terrestrial (direct-path), terrestrial (satellite bounce), or from within the solar system. I wrote a quick program to simulate the effects of the dedispersion filters on an impulse in each of these cases – the output is shown in Fig. 6.2(b). Obviously the prediction for a terrestrial pulse is correct. At the time of this analysis however, we did not know the origin, so that the comparison in Fig. 6.2 gave us our first hint of what was causing the pulse. It was only the next day, after this analysis, that we discovered the periodicity and then the noise-cal as the source.

A third – and ultimately essential – use for the noise diode was to create approximately coincident false triggers between the three antennas, since the time difference between the turning on and off of the noise diode between each antenna was small ($< 1 \mu\text{s}$) and the

variability smaller (~ 200 ns). The strength of the pulses also meant that trigger levels could be set above the level of almost all RFI (and well above the RMS noise level) so that only these pulses were detected.

Impedance mismatch in the dedispersion filters

It was discovered during the May 2007 run that each filter suffered from an impedance mis-match at the connection between the filter and the cable to it. This caused a significant reflection which, not passing through the dedispersion filter, arrived earlier than that part of the signal passing through the filter. The clue was the observed pre-pulse from the noise-cal – the approximate 20 ns time delay was as expected given the extra path-length.

A useful calculation is to determine the effect that this reflection should have on a signal in the frequency-domain – that is, the bandpass. The recorded signal $E'(t)$ can be represented in terms of the standard (unreflected, dedispersed) signal as $E(t)$ as follows:

$$E'(t) = E(t) + kE(t - \delta t) \quad (6.1)$$

where δt is the time difference between the fast and standard signal, and k is the ratio of strengths of the two components. Both k and δt can be measured easily in either the time- or frequency-domains: from Fourier-theory, Eq. (6.1) implies that the recorded spectrum $E'(\nu)$ is given by:

$$E'(\nu) = (1 + k \exp(-2\pi\nu\delta t))E(\nu).. \quad (6.2)$$

The mathematics simplifies when finding the magnitude of the recorded spectrum $|E'(\nu)|$ satisfies the following relations:

$$|E'(\nu)|^2 = (1 + k \cos(-2\pi\nu\delta t))^2 + (k \sin(-2\pi\nu\delta t))^2 \quad (6.3)$$

$$|E'(\nu)| = \sqrt{1 + 2k \cos(-2\pi\nu\delta t) + O(k^2)} \quad (6.4)$$

$$\approx 1 + k \cos(-2\pi\nu\delta t). \quad (6.5)$$

Thus an approximate sinusoidal ‘wobble’ is predicted in the spectrum, while in the time-domain a pre-pulse will be expected from an impulse-like RFI event or true signal. The accuracy of measuring k and δt in the time-domain is determined by the size of any given pulse compared to the noise, and thus is limited when the pulse height exceeds the maximum value of the analogue-digital converter (ADC); the accuracy in the frequency-domain depends only on the total length of buffers captured, which could, in theory, be infinite. Using the recorded spectra (see Section 6.6.2 and, for example, Fig. 6.22), the spectral half-amplitude of ~ 0.15 gives k , while the periodicity of ~ 50 MHz gives δt as 20 ns. The pre-pulse in fact appears at slightly more than 15% of the noise-cal pulse since it is not dedispersed and thus more condensed.

Since the interference between the standard and reflected signals alternates between be-

Size	16256	8192	4096	2048	1028	512	256	128	64
Rate	22	42	83	163.5	317.5	581	1040	1690	2450

Table 6.2: Maximum trigger rates (Hz) as a function of buffer length for 2008.

ing constructive and destructive, the mean power remains that of a standard signal, whereas the height of the Cherenkov pulse will be that of the unreflected, dedispersed part only. This is accounted for in the sensitivity calibration of Sec. 6.6 by averaging the ‘wobble’ in the frequency spectrum over the bandwidth in the electric-field (rather than power) domain, since the mean electric field strength is unaffected.

Amplification and quantisation

One important lessons learnt from the first run in May 2007 was that 8-bit sampling is barely sufficient for this type of experiment. The RMS noise level during these first observations was set quite low, of order 3 (the sampler ran from -127 to 128), meaning that the trigger thresholds could only be adjusted in increments of approximately 0.3σ . When triggering at levels of approximately 5σ , this meant the rates could only be altered in factors of 5, whereas obviously a finer adjustment is required. Changing the sampler levels for an RMS of 8 or so to allow for finer adjustment meant that a 5σ threshold was at a level of 40. Since the full range at 8 bits is ± 128 , the dynamic range was severely limited. Ultimately we decided that sacrificing dynamic range was the lesser of two evils, on the basis that having a real event with an unknown pulse height would be a problem we’d like to have. The measured RMS rates are given in Appendix C.

6.3 Dead-time and Efficiency

A certain degree of dead-time loss is suffered for every trigger. This is important since setting the thresholds too low (trigger rates too high) can make the effective observation time negligible, especially so when considering that all three antennas need to be ‘on’ to record an event. This dead-time can be easily measured by setting the thresholds to zero and recording the maximum trigger rate for a given buffer length. Such measurements were performed at each observation period, and the results for the 2008 run are recorded in Table 6.2.

For the May 2007 run, the maximum trigger rate was highly variable. The computers used to receive the buffers are not isolated, and are thus subject to traffic from other users with clogged the network. Thus while the ‘maximum maximum’ trigger rate was 300 Hz per antenna, the actual maximum rate varied between 40 and 300 Hz. This problem was fixed for 2008, so that the maximum trigger rate was both higher and more stable. The

exception was May 17th 2008, where network errors on antennas CA01 and CA03 reduced the maximum trigger rate by a factor of approximately 5.

Using Table 6.2, the dead-time per trigger for a buffer length of 256 samples (the standard in 2008) was approximately 1 ms. The efficiency of the experiment can be defined as the time-fraction when all three antennas are sampling and ready to trigger. For a sampling rate r_i (Hz) on antenna i , max rate R_i , and purely random trigger events, the efficiency ξ is given by:

$$\xi = \prod_i \left(1 - \frac{r_i}{R_i}\right) \quad (6.6)$$

where the i multiplies over all three antennas. A non-random fraction to the trigger rate requires an alteration to Eq. (6.6), since the dead-times on all three will tend to coincide. If the three-fold coincident rate due to non-random triggers is r_3 , with the random rate being $r'_i = r_i - r_3$, then the efficiency equation becomes:

$$\xi = -\frac{r_3}{R_{low}} + \prod_i \left(1 - \frac{r'_i}{R_i - r_3}\right) \quad (6.7)$$

where here R_{low} is the lowest maximum trigger rate of all antennas, since in general the R_i could be different. This equation assumes that the dead-time is much larger than the window to accept triggers as being ‘coincident’, which is the case here. See Sec. 6.5.4 for an in-depth discussion of coincident events.

Using the observed trigger rates on all antennas, Fig. 6.3 plots the trigger rates (Hz) and efficiency ξ for the 2008 period, assuming a constant $R_i = 1040$ from Table 6.2. Triggers are considered ‘three-fold’ if they arrive within an 8 μ s window – again, see Sec. 6.5.4. The calculations in Fig. 6.3 estimate r_3 from the observed three-fold rate, whereas some of the three-fold events will actually be due to random coincidence. Since random coincidences will either be a small fraction of r_3 , or r_3 itself will be small, this approximation is sound.

The effective observation time t_{eff} can be found by integrating the efficiency over the observation time t_{obs} , so that $t_{eff} = \bar{\xi} t_{obs}$. These are given in Table 6.3 for 2008.

A summary of data rates and efficiency for May 2007 is given in Table 6.4. Note that the efficiency estimate is a minimum, since it assumes complete independence of the triggers. However, the maximum rates are only approximate and tended to vary – those in the table are estimated from nineteen minutes of data on the 6th, between 19:02 and 19:21 UT.

6.4 Dispersive and Sampling Effects

6.4.1 Ionospheric dispersion

To our nominal experimental bandwidth of 600 MHz at 1.5 GHz, the effects of dispersion in the Earth’s ionosphere are significant. The dispersion is due to a frequency-dependent refractive index caused by free (ionised) electrons in the ionosphere. Using the standard




Figure 6.3: Trigger rates (lower lines) and efficiency (upper lines) from Eq. (6.7) for the night of May 18th, 2008.

Date	t_{obs}	$\bar{\xi}$	t_{eff}
May 2007 7 th	254	82.5	210
26 th	239	85.5	204
Feb. 2008 27 th	319	86.9	277
18 th	314	87.4	274
17 th	324	69.0	224
May 2008 18 th	376	72.8	274
19 th	440	71.9	316

Table 6.3: Raw observation time t_{obs} (minutes), mean efficiency $\bar{\xi}$ (%), and effective observation time t_{eff} (minutes), for all observation periods.

	CA03	CA04	CA05
Triggers	99556	155654	121494
Rates (Hz)	6.53	10.2	7.97
Max Rate (Hz)	120.6	147.1	126.9
Dead Time (ms)	8.29	6.80	7.88
Efficiency (%)	94.6	93.1	93.7
Global Efficiency:			82.5%

Table 6.4: Calculation of mean experimental efficiency for May 2007. No correlation between trigger times is assumed – such a correlation would imply an increased overlap of dead-times between triggers, and hence increased effective observation time when all buffers were untriggered.

Date\Time	UT	10:00	12:00	14:00	16:00	18:00	20:00
	AEST	20:00	22:00	00:00	02:00	04:00	06:00
May 2007	5 th	4.5	5.0	6.1	7.5	6.1	5.1
	6 th	5.6	5.5	6.6	6.3	6.1	5.1
	7 th	7.0	7.0	6.9	8.8	8.8	6.3
Feb 2008	26 th	6.4	5.4	6.0	5.5	5.2	5.8
	27 th	11.0	10.0	9.5	8.7	7.1	7.4
	28 th	8.9	6.9	7.7	6.8	5.7	6.0
May 2008	17 th	5.6	5.5	6.6	6.2	6.7	5.8
	18 th	5.2	6.0	5.9	7.9	7.7	6.9
	19 th	10.7	10.3	9.8	11.1	8.2	6.4

Table 6.5: Values of VTEC (Vertical Total Electron Content) in TECU ($10^{16} \text{ e}^- \text{ cm}^{-2}$) as interpolated from GPS measurements (95; 128), both in UT, and AEST, the local (non-daylight savings) time.

measure for the number of electrons (total electron content units, or TECU: $10^{16} \text{ e}^-/\text{cm}^2$), the time-delay δt relative to a vacuum for a frequency ν is given by Eq. (6.8):

$$\delta t = 1.34 \cdot 10^{-7} \text{ TECU } \nu^{-2}. \quad (6.8)$$

Of more use is the dispersion over a bandwidth Δt , given by Eq. (6.9):

$$\Delta t = 1.34 \cdot 10^{-7} \text{ TECU } (\nu_{\min}^{-2} - \nu_{\max}^{-2}) \quad (6.9)$$

$$\approx 2.68 \cdot 10^{-7} \text{ TECU } \Delta \nu \bar{\nu}^{-3} \quad (6.10)$$

if $\Delta \nu \ll \bar{\nu}$, where $\bar{\nu}$ is the mean frequency $(\nu_{\min} + \nu_{\max})/2$. Note that in terms of phase delay, the correction goes as $\bar{\nu}^{-2}$.

For typical night-time values of TEC (approximately 7 TECU), the induced delay is 3–4 ns over our bandwidth, which is significant. The effects of this dispersion (or equivalently, incorrect dedispersion, since we partially correct for these effects) will depend not just on deviations of the line-of-sight (LOS) TEC from that assumed in our dedispersion filters, but also on both the sampling rate and the intrinsic pulse spectrum. All of these effects are dealt with here.

Measured VTEC

As discussed in Sec. 6.2.2, our dedispersion filters were only set for a typical TEC (total electron content) at an elevation of 47° , so that the actual differential delay (and hence sensitivity) will be affected by the elevation and real VTEC values. The actual dispersion measurements over the observation dates (supplied by R. Robinson from the NASA Crustal Dynamics Data Information System (128)) are given in Table 6.5.

As well as the TEC itself, The effects of ionospheric dispersion will depend both on the sampling rate, and on the spectral shape. I therefore treat each effect individually, before proceeding to estimate the combined effect of only an approximate dedispersion on our sensitivity.

6.4.2 Quantifying the loss of sensitivity

Sample spectra

The range of possible spectra is very broad. A general characterisation can be made based on the relative strengths of the high- and low-frequency components. At one extreme is a purely coherent pulse with $E(\nu) \propto \nu$, which gives the greatest possible weight to a high-frequency component – this is the case plotted in Fig. 6.5. The other extreme is given by a high-energy shower viewed at a large angle away from the Cherenkov angle, but at zero depth: a frequency turnover due to decoherence is generally sharper than that caused by absorption¹. Choosing increasingly high-energy showers viewed increasingly far from the Cherenkov angle makes the spectrum tend towards a δ -function at the lowest frequency, so that a highest ‘realistic’ energy must be chosen. Taking this as 10^{23} eV, and a spectral shape of the form $E(\nu) \propto \exp(-C\nu^2)$ (which has been shown to approximately fit the spectrum far from the Cherenkov angle for hadronic showers (9)), allows the constant C to be solved for by letting the total power in the bandwidth equate to that of a fully-coherent pulse at 10^{20} eV:

$$A \int_{\nu_1}^{\nu_2} \nu d\nu = B \int_{\nu_1}^{\nu_2} \nu \exp(-C\nu^2) d\nu \quad (6.11)$$

$$A \left[0.5\nu^2 \right]_{\nu_1}^{\nu_2} = B \left[-0.5 \frac{1}{C} \exp(-C\nu^2) \right]_{\nu_1}^{\nu_2} . \quad (6.12)$$

$$C \left(\exp(-C\nu_1^2) - \exp(-C\nu_2^2) \right) = \frac{A}{B} \left(\nu_2^2 - \nu_1^2 \right). \quad (6.13)$$

Since $B = (10^{23}/10^{20})A$, solving Eq. (6.13) numerically for $\nu_1 = 1.2$ GHz, $\nu_2 = 1.8$ GHz gives $C = 3.515$ for ν in GHz. The two frequency-spectra are plotted in Fig. 6.4, and used throughout this section.

The effects of finite sampling only

Our sampling rate of 2.048 GHz, while less than the Nyquist rate of twice the maximum frequency, was sufficient to prevent aliasing over the frequency range 1.024 to 2.048 GHz. Within this range – and in the absence of signals outside this range – the full spectrum can be reconstructed, since there can be no ambiguity between frequencies, and no information is lost. Detection experiments (such as this one at the ATCA) however tend to rely on a simple

¹The extrema should in fact be defined in terms of the most and least broadband pulses. While the least broadband pulse is the maximally-decoherent one described in text, the most broadband pulse will in fact have some small decoherence term to reduce the high-frequency component. I ignore this minor distinction here.




Figure 6.4: The two extreme spectra used in the estimates of sampling and dispersive: a fully coherent spectrum (dotted line), and a spectrum with a decoherence factor of $\exp(-Cv^{-2})$ (solid line).

algorithm based on the peak pulse height as captured by the sampling threshold. For any finite sampling rate (including the Nyquist rate), there will be a random offset in the phase of the sampler between the peak of the pulse and the (approximately instantaneous) sampling times. Fig. 6.5 plots the peak pulse height as a function of the arbitrary phase offset for three sampling rates: the minimal rate of 1.2 GHz required to avoid aliasing for a 600 MHz bandwidth, with the nominal bandpass of 1.2-1.8 GHz downconverted to 0 – 600 MHz; the actual rate of 2.048 GHz as used in our experiment, with the downconverted bandpass at 248-848 MHz; and a 4.096 GHz sampling rate with no downconversion, as might be implemented in the final CABB upgrade. Since in this case the pulse peak is forced to lie at the central sampling point, the sampling here makes only a cosmetic difference.

To first illustrate the effects of sampling before delving into dedispersive effects, in Fig. 6.6(a) and Fig. 6.6(b) are plotted the maximum measured values in the time domain relative to the peak value as a function of the arbitrary phase offset between the pulse peak and the sampling values, for (a) a fully coherent pulse, and (b) an extremely incoherent pulse. Immediately the advantages of a faster-than-minimal sampling rate when using simple threshold trigger logic becomes apparent: whereas on average the peak measured value at the minimal rate of a coherent (incoherent) signal is only 76.5% (81%) of the peak (and it can fall below 60% in both cases), for 4.096 GHz sampling the average is 93% (95.5%) with a minimum of 84% (89%). For an incoherent pulse the effect is less since the power is concentrated over a smaller frequency range.

Including (de)dispersion

For each combination of intrinsic spectrum, base phase offset, and dispersion measure, I calculate the peak pulse strength in the time domain. Averaging this over all base phase




Figure 6.5: The same simulated pulse (in the coherent case) over a 1.2-1.8 GHz bandwidth, sampled at three rates.

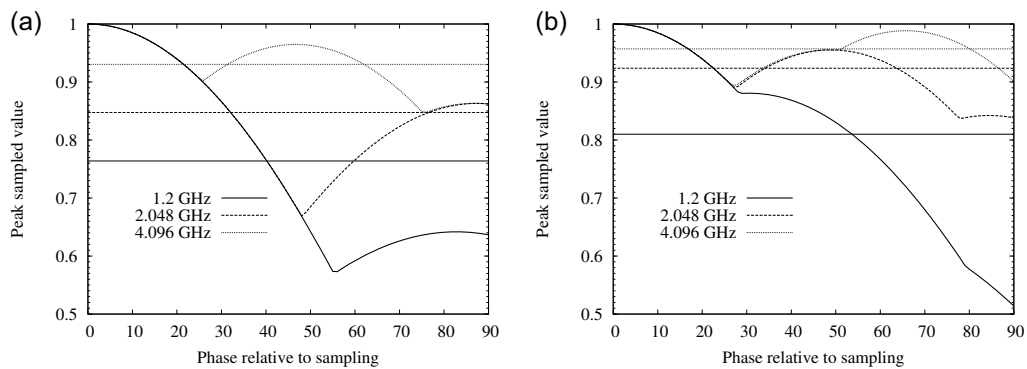


Figure 6.6: Maximum sampled value as a fraction of the true pulse height as a function of the arbitrary phase offset for three different sampling rates and two pulse types (left: coherent, right: incoherent).

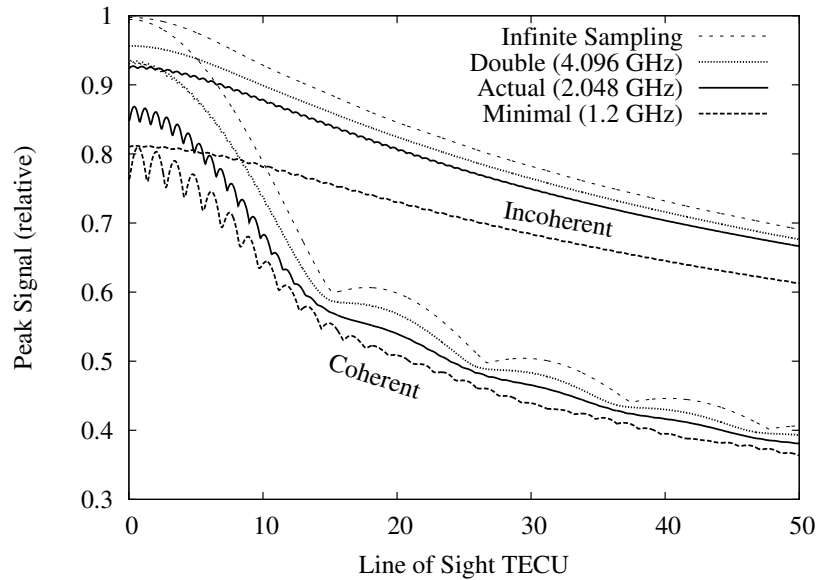


Figure 6.7: Greatest detected amplitude of the time-domain signal relative to the maximum undispersed amplitude, for two signal spectra and four sampling rates, as a function of the line-of-sight total electron content, in TECU.

offsets (which will be random) and dividing by the magnitude of the peak undispersed pulse at zero phase offset produces Fig. 6.7. The upper and lower sets of lines are for incoherent and coherent pulses respectively. The coherent pulse is most adversely affected by ionospheric dispersion, since the signal is spread over the greatest frequency range. The rapid oscillations in average peak amplitude with changing TECU for the coherent pulse are more evident at slow sampling rates. The means from Fig. 6.6 are given by the intercept with the y-axis, allowing a comparison between the effects of finite sampling and dispersion. By comparison with the estimates assuming an infinite sampling rate, the effects of finite sampling and dispersion can be separated. For incoherent signals, it is interesting to note that for a minimal sampling rate, the loss due to slow sampling of 19% at 0 TECU is greater than that from dispersion only (infinite sampling rate) for a TEC less than 25 TECU. For our actual experimental sampling rate, the extra loss from dispersion becomes equal to that from finite sampling at ~ 11 TECU for the incoherent pulse, while for a coherent pulse the two losses equate at approximately 8 TECU.

6.4.3 Loss of experimental sensitivity

The dedispersion filters were set to correct for a dispersion over the bandwidth of 5 ns, equivalent to 9.67 TECU along the line-of-sight. Here I assume the filters correct perfectly for such a dispersion – see Ref. (99) for bandpass measurement of our preliminary design: variations from the ideal case are small. Since the pulses are symmetric, an over-correction

in the dedispersion filters for a low STEC ('Slant' TEC – the TEC along the line of sight) is equivalent to an under-correction for a large value. Therefore the zero on the x-axis in Fig. 6.7 represents a real STEC of 9.67 TECU.

Using the measured values of VTEC, the lunar elevation, and the results given above, the loss of sensitivity due to sampling and dispersive effects can be calculated. The results are given in Fig. 6.8. The VTEC measures are given by the dashed lines – minima come just after sunset and just before sunrise (local times 10pm and 6am). Interestingly, there appears to be a secondary peak around midnight, which while much lower than the solar peak at mid-day, still causes a significant increase in the TEC of 2 – 3 TECU above the minima. Using a linear interpolation between these points and the known lunar elevation gives the line-of-sight (LOS) TEC measure (dotted lines). Since the Moon transited almost overhead around midnight, this secondary peak actually helped to keep the expected line-of-sight TEC stable, since the VTEC is high when the distance travelled through the ionosphere is low and vice versa. At low elevations, the LOS will probe a large horizontal distance, so using a constant VTEC measure may not be appropriate. However, since the TEC goes as $\sin^{-1}(\text{elev})$ and consequently blows up at low elevations, the sensitivity in this regime will be low in any case. Combined with the mean losses for the two spectra in Fig. 6.7, the range of losses for the experimental periods is calculated as the shaded regions, limited by the extreme coherent/incoherent spectra. Also shown in Fig. 6.8 are the observation times on each night.

The effect of the approximated de-dispersion is estimated by taking the mean loss over both observation time and spectral type, giving equal weighting to both the coherent and incoherent spectra (a full calculation of the time-domain signature of each event generated by the Monte Carlo simulation would take an exorbitant amount of computer time). Though it is safe to assume that near the threshold neutrino energy for detection, the only detectable events will be fully coherent pulses from near the Cherenkov angle, much above this threshold there is no easy method to estimate the relative contributions of each spectral type, though qualitatively one would expect an increasing contribution of incoherent events with primary neutrino energy. For simplicity of simulation I use the mean of the fractional losses for coherent and incoherent signals for all neutrino energies.

6.5 Relative Timing Calibration with Astronomical Point Sources

For the 2007 observations, our timing was given only by the computer clocks to poor (1 ms) accuracy, which generally is insufficient to either reject all false events or accept a real one. This section deals exclusively with our 2008 observations, in which we had clocks accurate to 0.5 ns running at each antenna, the value of which was returned with each buffer. However, these clocks had unknown timing offsets between them, which had to be determined to allow a sufficiently rigorous pulse search. In this section I describe my efforts

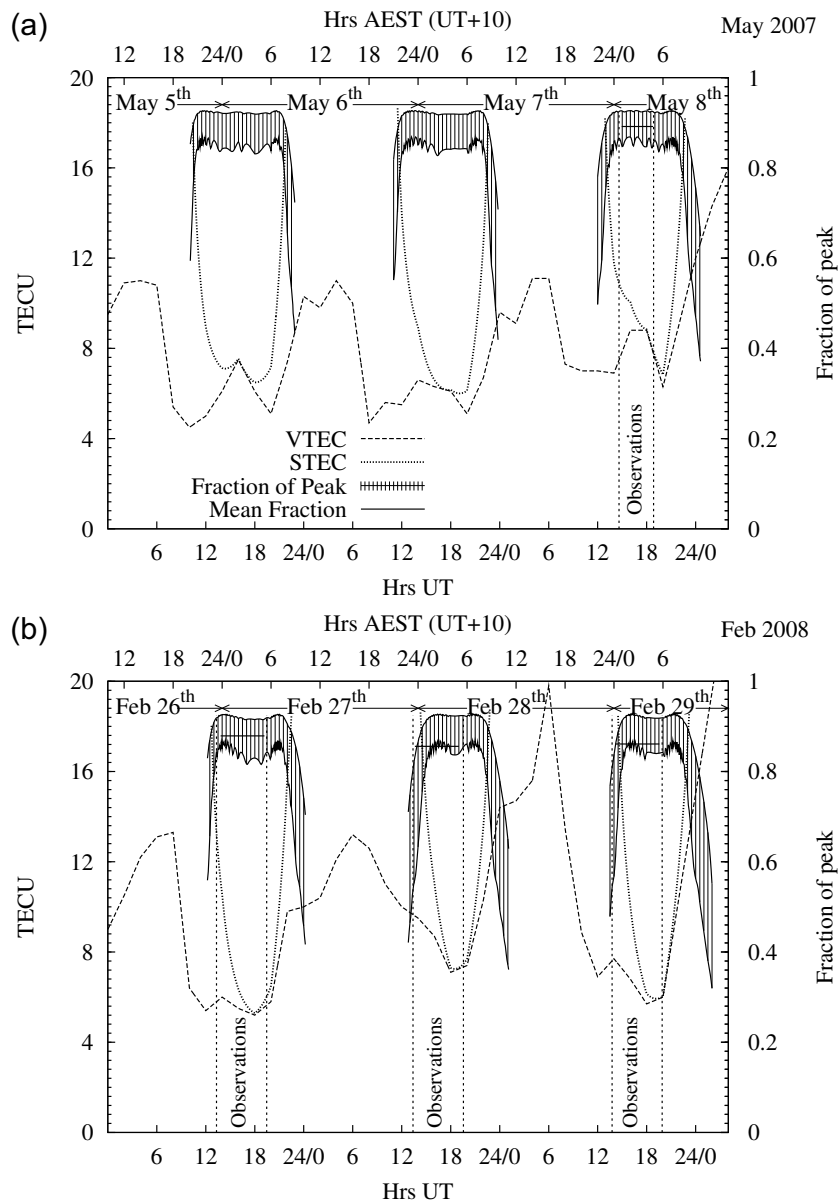


Figure 6.8: (continued over page)




Figure 6.8: Ionospheric influence over our observation periods. The key of (a) applies to all three plots. Dashed lines: measured VTEC (TECU) from Ref. (128). Dotted lines: STEC (slant TEC along the line of sight). Solid lines: mean fractional loss for coherent (lower line) and incoherent (upper line) pulses due to ionospheric effects only: the shading gives the range. Vertical solid lines indicate the times of the observations, while horizontal solid lines give the mean fractional loss during the ‘observations’.

Period Date	May '07	Feb '08				May '08			
	7 th	26 th	27 th	28 th	Mean	17 th	18 th	19 th	Mean
Best (%)	93	92	91	91	91	91	91	91	91
Worst (%)	86	84	80	81	82	82	80	81	81
Mean (%)	90	88	86	86	87	87	86	86	86

Table 6.6: Estimated nightly average fractions of the peak signal detected (%) for the observations periods indicated. The best case corresponds to incoherent signals, the worst case to completely coherent signals, as discussed in text.

Symbols $[F/M]_n^{[C/M]}$ describing continuous observations.	
$[F/M]$	Indicates F: February, M: March 2009.
$[C/M]$	C : time calibration period, M : on-Moon detection mode.
n	n^{th} observation of that type in the observation period.
Symbols used in calculating calibration times (Sec. 6.5.1).	
b_o	‘Buffer overlap’: length of common captured signal.
$t_i^g, \Delta t_{ij}^g$	Geometrical time delay from a source position.
Δt_{ij}^t	Trigger time for a noise cal pulse.
$t_i^c, \Delta t_{ij}^c$	Cable delay to an antenna.
Symbols used to calculate the time offsets (Sec. 6.5.3).	
l_b	Buffer length.
Δt_{ij}^p	Position of the peak in the correlation function.
Δt_{ij}^b	Time offset between the centres of returned buffers.
Δt_{ij}^0	Absolute time offsets between the clocks.
Symbols pertaining to individual candidate events (Sec. 6.5.4).	
$t_i (\Delta t_{ij})$	Buffer trigger times (difference between).
$\Delta' t_{ij}$	Adjusted time between returned buffers.
$\delta \Delta t_{ij}$	Magnitude of adjustment: $\Delta' t_{ij} - \Delta t_{ij}$.

Table 6.7: Description of symbols used in the time alignment – units alternate between seconds and samples ($2.048 \cdot 10^9$ per second) in the text. Symbols with an ‘ i ’ subscript apply to each antenna, those with a ‘ ij ’ give the difference between the values for each antenna; the convention is that quantity i is subtracted from quantity j when $j > i$ (the i, j indicate antenna number, e.g. ‘3’ for CA03).

towards determining these offsets and calibrating the clocks. Symbols used in this section are defined in Table 6.7.

Regarding contributions from other members of the collaboration, the method described in Sec. 6.5.1 below was developed predominantly by R. Ekers, while the software methods described in Sec. 6.5.3 were largely based upon those developed by C. Phillips. An independent analysis using Phillips’ software – similar to my work from Sec. 6.5.1 to Sec. 6.5.3 – was performed for the February 2008 data by R. Robinson and C. Phillips, with similar results.

6.5.1 Method

To calibrate the times, we required both a common signal in each antenna with some known time-delay, and a method to trigger the buffers with sufficient simultaneity that enough of the common signal seen by all three antennas would be captured to produce a significant correlation. For the trigger, we had the choice of using either a portable pulse generator or the noise-diode pulse, while for a common signal, we had either the pulse generator with a well-measured physical position, or a bright (and preferably point-like) astronomical

source. In testing both before and during the observation, we failed to generate a signal with the pulse generator strong enough to trigger all three antennas at once. Thus we used the noise-cal pulses as our trigger, and the very bright sources 3C273 and M87 as our correlated signal.

3C273 was chosen since it is the brightest (47 Jy at 1.4 GHz) point-like source near the Moon at the time of the observations, and would thus give a strong correlation over all baselines. M87 is brighter (215 Jy at 1.4 GHz) (130) but resolved on our long baselines, and was chosen to maximise the correlated signal over our shortest (CA01-CA03) baseline only. Another consideration was the length of correlated signal between buffers at each antenna (the ‘buffer overlap’), which is a function of buffer length, the relative trigger times of the noise cal pulse, and the geometrical delay due to the position of the source.

Required buffer overlap

The calculations in this portion of Sec. 6.5.1 are independent from, but motivated by, those of Ron Ekers, who attained similar results.

For each individual antenna, the signal-to-noise ratio per independent sample (σ_s) from a source of flux \mathcal{S} (W/m²/Hz) viewed by an antenna of effective area A_{eff} and system temperature T_{sys} is given by:

$$\sigma_s = \mathcal{S} \frac{A_{\text{eff}}}{2k_b T_{\text{sys}}} \quad (6.14)$$

where k_b is the Boltzmann constant. The gain from bandwidth comes because samples taken once every $\Delta\nu^{-1}$ seconds are independent. Therefore, the strength of a correlation σ_c over time t is given by:

$$\sigma_c = \sigma_s \sqrt{t \Delta\nu}. \quad (6.15)$$

Alternatively, the required time length of the correlated signal captured by a buffer to achieve a given s/n strength in the correlation σ_c is given by:

$$t = \frac{1}{\Delta\nu} \left(\frac{\sigma_c}{\mathcal{S}} \frac{2k_b T_{\text{sys}}}{A_{\text{eff}}} \right)^2. \quad (6.16)$$

For ATCA parameters ($A_{\text{eff}} = 266 \text{ m}^2$ for an aperture efficiency of 0.7 on a 22 m antenna, $T_{\text{sys}} \approx 31^\circ \text{ K}$ (130), $\Delta\nu = 600 \text{ MHz}$), we find $t_{3\text{C}273} = 2.8 \mu\text{s}$ for $\sigma_c = 6$ on all baselines.

Buffer overlap - geometrical

The buffer overlap b_o is given by $b_o = l_b - |\Delta t^g - \Delta t^t|$, where the terms are the buffer length l_b , the geometrical time-offset Δt^g between antennas due to the source position as viewed by the baselines, and the difference in triggering time Δt^t . It is most useful to calculate the factor $|\Delta t^g - \Delta t^t|$ (the relative offset of buffer centres with respect to a source signal) for

t^c (μs)	CA01	CA03	CA05
	3.39	1.73	2.91
Δt^c (μs)	Δt_{31}^c	Δt_{51}^c	Δt_{53}^c
	1.66	0.485	-1.17

Table 6.8: Cable delays t^c and resulting time offsets Δt^c for the February observations (see Table 6.1) as measured by P. Roberts. The cable refractive index is specified as 1.47.

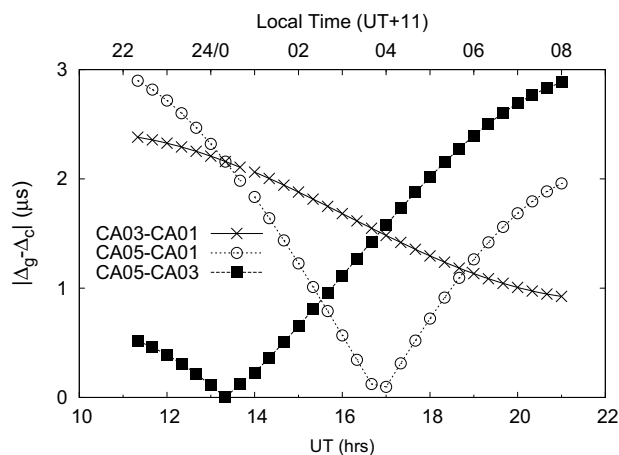


Figure 6.9: Expected offsets of buffer centres with respect to signals arriving from 3C273 when triggered by the noise cal pulse for Feb. 27th, 2008.

all three baselines, and hence determine the optimal observation time and required buffer length.

The factor Δt^t can be assumed to be a constant function of cable-delays which take the signal to turn the noise-cal on/off from the control room to the antennas. The cable delay times t^c and corresponding time offsets $\Delta t^c \sim \Delta t^t$ to the antennas were measured by P. Roberts, and are given in Table 6.8 for the baselines of our February observations.

Assuming the in-antenna delays are constant (or at least small compared to the control-room–antenna delays), the offsets Δt^c in Table 6.8 give the offsets in trigger times Δt^t . The geometrical (or astronomical) delays are simply given by the source positions, which vary with time. The factors $|\Delta t^g - \Delta t^t|$ between the antennas for the nights of our February 2008 observations are plotted in Fig. 6.9 for 3C273 – for M87, the plot is almost exactly the same, since they have similar right ascensions. The variation is minor (± 4 minutes) over a three-day observation period.

From Fig. 6.9, note that we could accept times when only two of the three baselines had a significant buffer overlap, since we could still solve for the time offsets – however, we preferred to have all three baseline correlations to be visible to give us some redundancy. Therefore times around 17:00 UT were ideal for calibration both on 3C273, and also for M87, since in the latter case the weakest (longest) baseline would capture the most signal.

Since we require a minimum of $\sim 3 \mu\text{s}$ of common signal, and at the most ideal moment there is approximately a $2 \mu\text{s}$ offset between buffer centres, the minimum required buffer length was determined to be $5 \mu\text{s}$. The maximum available was 16,324 samples ($\sim 8 \mu\text{s}$), which should therefore have been more than sufficient to get a 6σ spike on the correlation. However, the 25% of the buffers with no common signal provided correlation ‘noise’, and the 6σ correlation over the common part only was expected to be reduced by a factor of $\sqrt{0.75}$ to 5.2σ , which should still have been significant over the $< 4\sigma$ expected random maximum. Nonetheless, since the noise cal itself introduces a significant amount of power into even a long buffer, we expected to add the correlations over multiple noise-cal triggers to ensure the correct correlation peak was identified.

6.5.2 Observations

To observe in timing calibration mode, we pointed the antennas at either 3C273 or M87, set the buffer length to maximum, switched the noise cal on, and set the trigger thresholds such that we were triggering off all, and only, the noise cal pulses, at roughly 8 Hz. This could be done because the pulses were far stronger than most RFI signals, although occasionally we observed equally strong RFI. Triggering off an RFI signal could result in the triggers missing a noise cal pulse due to the dead-time associated with long buffers, though this was never a significant effect.

Although our calculations indicated that we only required a single coincident trigger once over our entire observation period in order to align the clocks, we typically observed in calibration mode for a few minutes at a time and thus took of order 2000 pulses, repeating this procedure a few times each night. Table C.1 in Appendix C gives the time spent observing in calibration mode for both the February and May observations. Since during the experiment we succeeded in obtaining correlation peaks only on the second night, we cautiously spent more time observing the calibration source than our calculations indicated was necessary.

6.5.3 Analysis

To extract a timing offset required knowing exactly which recorded buffers to correlate, whereas we had no explicit way of knowing which returned buffer on antenna X corresponded with which on Y and Z. However, the accuracy of the control-room computer’s time stamp ($O \sim 30 \mu\text{s}$) was much better than the typical times between triggers (0.128 s). Therefore I wrote a simple routine to match buffers between antennas, which checked the computer time stamps on all files from all three antennas during the calibration periods and selected those with matching times on all three antennas and which were separated by the characteristic 0.128 s interval from another three-fold match. Setting a tolerance of $30 \mu\text{s}$ for times between successive triggers on a given antenna, and $300 \mu\text{s}$ to match coincidences between antennas, typically accepted ~ 440 coincidences a minute from a maximum of 469, giving an efficiency of 94%. A greater tolerance for the timing between antennas is allowed

because only one buffer can be recorded at once, while the lower tolerance for successive triggers on a given antenna reflects the stability of the system.

The correlation program

I wrote a program to sum the correlations between coincident buffers, as output by the alignment procedure described above. The basic steps involved are as follows:

1. Initialise using the first coincident buffer set:
 - a) Calculate the base astronomical delays from the source position.
 - b) Note the time offsets between buffers – label these the ‘base offsets’.
2. Loop through and sum the correlations:
 - a) Extract absolute and relative timing information from the buffers.
 - b) Calculate the astronomical delays from the average computer clock time of the coincident buffers.
 - c) Determine the time offsets between buffers relative to the base offsets.
 - d) Zero buffer entries –150 to 50 to eliminate the noise cal pulse.
 - e) Fit the buffers and clean their spectra using a simple algorithm.
 - f) Multiply and reverse-fft to obtain the individual correlation functions.
 - g) Add to a running total, with offsets as given previously.
3. Clean the summed correlation functions:
 - a) Fast-Fourier-transform (fft) the correlation functions.
 - b) Eliminate spikes in the spectra with a simple cleaning algorithm.
 - c) Reverse fft.
4. Run the above simultaneously over all three baselines and both polarisations.
5. Plot the correlation functions and extract the offsets.

It was found that the relative times between triggers on different antennas varied by approximately ± 200 ns. This ‘jitter’ needed to be corrected for by a ‘de-jittering’ algorithm when adding multiple correlations. This is performed easily in step 2(c), since the relative times for any given buffer are known precisely. Note that the astronomical delays only need to be compensated for relative to the original delay, since we want the ‘zero’ to correspond not to a simultaneous arrival of an astronomical signal, but to ‘absolute’ simultaneity (or for relativistic purists, simultaneity in the control-room frame).

An example of the process for Pol A of CA01/CA03 is given in Fig. 6.10. The most important step is the cleaning of each individual spectrum prior to correlation, which is

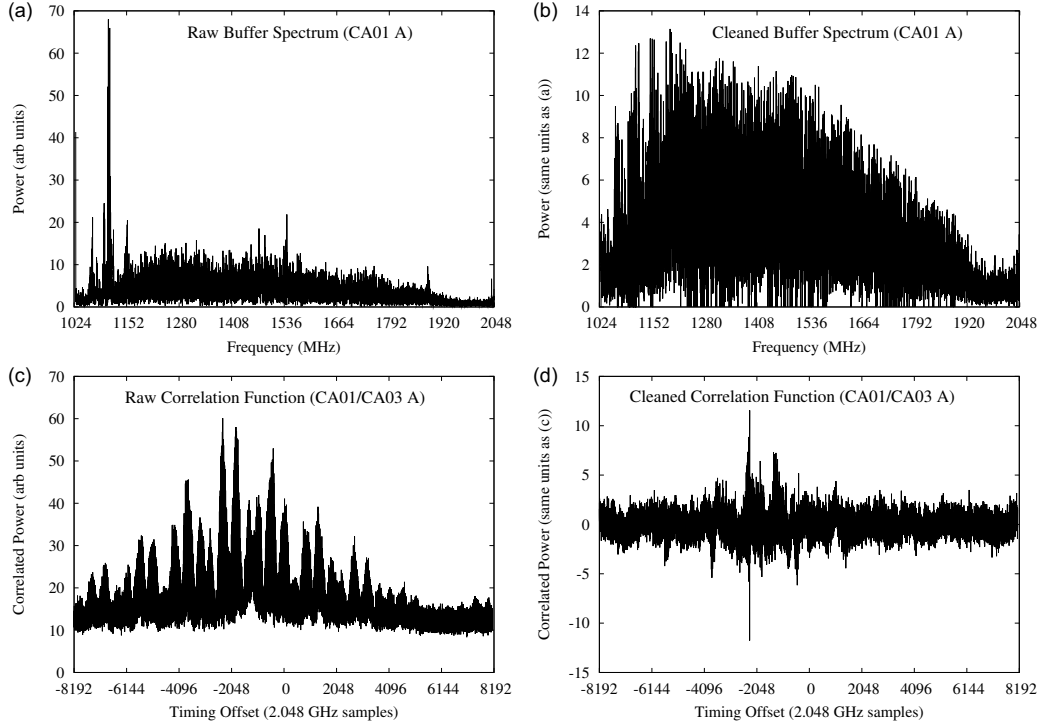


Figure 6.10: (Continued over page)

done by using a running window (typical length 1000 samples) and setting each spectral point with power more than 4 times the mean in the window to zero. For a white noise spectrum, this would (falsely) reject approximately 2.5% of the available data, which is a small price to pay for greatly increased quality. Despite this, at times the RFI environment was sufficiently bad that no correlation could be extracted.

Relating the correlation peak to the timing offset

The values of the base timing offsets Δt_{ij}^0 (given by $t_j - t_i$ for $j > i$ and $i, j \in \{1, 3, 5\}$) are derived by adding the time offsets between buffer centres Δt_{ij}^b and the positions of the cross-correlation peaks Δt_{ij}^p relative to a zero offset. Labelling the cross-correlation function between buffers from antennas i and j as $C_{ij}(\Delta t)$, the multiplication in Fourier space is chosen such that if $B(t) \leftrightarrow \mathcal{B}(\nu)$, then the cross-correlation theorem gives:

$$C_{ij}(\Delta t) \leftrightarrow C_{ij}(\nu) \quad (6.17)$$

$$\equiv \mathcal{B}_i(\nu) \mathcal{B}_j^*(\nu) \quad (6.18)$$

$$\leftrightarrow \sum B_i(t) B_j(t - \Delta t) \quad (6.19)$$

i.e. a peak in the correlation function C_{ij} corresponding to a time delay $\Delta t_{ij} = \Delta t_{ij}^p$ means that the buffer sample corresponding to the peak on CA0*i* matches in true time sample 0 on

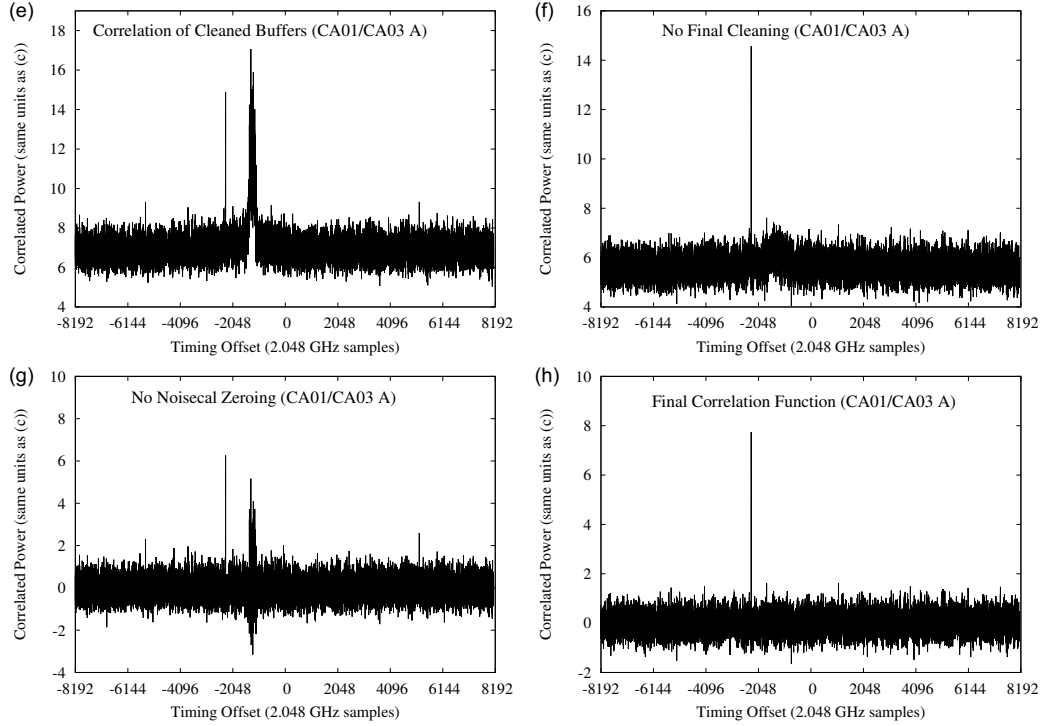


Figure 6.10: Effects of data manipulation in the correlation program. (a): raw spectrum from an individual buffer, (b): the same spectrum after cleaning. (c–h): Correlation functions with varying degrees of data manipulation over 100 aligned buffers from CA01 and CA03 Pol A, with (c): no data manipulation, (d): cleaning only the final correlation function, (e): cleaning each spectrum individually prior to correlation, (f): cleaning both the individual and final spectra, (g) cleaning the individual spectra and zero-ing the noise-cal pulse, (h) full cleaning and zeroing as used to extract timing offsets.

$ca0j$, i.e. the clock at i is Δt_{ij}^p faster relative to the buffer centres. Therefore the total time offset $t_j - t_i = \Delta t_{ij}^0$ is given by $\Delta t_{ij}^b - \Delta t_{ij}^p$.

The pathological data of May 17th

Network errors on May 17th reduced the trigger rate to 3 – 4 Hz when the buffer length $l_b = 16320$ samples. This meant that at most a third of all the triggers could be captured with a full buffer length. It also meant that, in absence of some randomisation, the captured noise-cal triggers must fall into a repeating pattern which may or may not correspond to the correct alignment. It would be perfectly possible (actually, very likely) that triggers 1, 4, 7, ... would be captured from CA01, triggers 2, 5, 8, ... from CA03, and triggers 3, 6, 9, ... from CA05, so that while one in three triggers would be captured, *none* would align properly. Three methods were used to try to get a correct alignment: M_1^C : reducing the buffer sizes to 8192 samples, hence reducing the overlap by perhaps two thirds but increasing the trigger

rate; M_2^C : setting the buffer length on CA03 to 16320 while keeping that on CA01 and CA05 to 8192 samples; M_3^C : keeping all three buffers at length 16320 while setting the trigger rate to ~ 1.5 with a threshold such that some triggers would be detected, while others would not, thereby introducing a random factor to prevent long-lasting trigger patterns such as that described above.

The resulting analysis must be done mostly manually, and is very time-consuming. Constructing an algorithm to match corresponding buffers did not work in this instance, since the returned computer times are a function of the inherent delay in the network. Normally quite low, the network errors meant the delay could be up to half a second, so that this search had to be done manually via trial and error. Also, so few three-fold matches were found that the correlation had to be done piece-wise, treating each of 1–3, 1–5, and 3–5 separately, with no redundancy. Finally, the RFI environment for May 17th was very bad, and the only attainable correlations were 1–3 and 3–5 for the M_3^C calibration. The most important result for any further investigations is that buffers 46 in May17/CA01_138_142100.lun, 47 in May17/CA03_138_142100.lun, and 321 in May17/CA05_138_142100.lun originate from the same noise-cal trigger – from here, further matches can be found with relatively little effort.

Results of the time alignment

The full results of the timing alignment for each period are given in Appendix C, Table C.6. The output in each case consists of the base offset (number of 2.048 GHz samples between the centres of the first buffers to be correlated) and the position of the graphically-determined correlation peak. In general, there was sufficient data to enable multiple independent alignments to be derived from the same calibration pointing; in the odd instance (particularly when the Moon’s elevation was low, so that there was little correlated signal and more RFI) timing information for one or more baselines could not be extracted. The exception is of course on May 17th, where network issues caused significant problems (see Sec. 6.5.3). The estimates were made by simultaneously looking at the Pol A & B data, and the errors come from the peak having a finite width. When the correlation statistics were low, this reflected our over-sampling (600 MHz b/w at 2.048 GHz).

A consistency check of the Δt is given in Figs. 6.11 for both the February (left) and May (right) observation periods. For the May observations, the clocks were reset inbetween each day, less calibrations were performed, and on May 17th only one calibration produced results. Therefore only the February data gives a good check of consistency, for which the clocks continued over the entire period. In the top row of Figs. 6.11, the vertical axis shows the calibrations Δt_{ij}^0 relative to the first calibration after each clock reset – i.e. the first data points in February, and the first on each day in May, have been normalised to have y-values of 0. The expectation was that all data would therefore have y-values within a small ($O \sim 1$ sample) error range of 0, reflecting uninterrupted clocks ticking at the same rate. Obviously this is not the case: a cursory glance – particularly at the Feb. graphs – shows that the time

offsets Δt_{ij}^0 jump around by up to ~ 800 samples. A less-cursory glance however shows that the deviations are not random, with most being in multiples of 192 samples (93.75 ns). This is illustrated in the middle row, where the offsets from the top row have been divided by 192. Most points take integer values. Suspiciously, calibrations within a period are always consistent to within a few samples (a more extensive search than is evident here proved this), indicating that between calibrations, something in the hardware sometimes caused one or more of the clocks to jump in random multiples of 192. As it turned out, this was indeed a hardware issue, which is discussed (and mostly solved) in Sec. 6.5.4. The effect is less obvious in the May data where less correlations per clock reset were performed.

The bottom row of Fig. 6.11 shows the relative alignments after artificially adjusting for the offsets in units of 192 samples, and dividing by the baseline length in units of ~ 15 m, plotted against hours UT. The variation (of order a few samples) with UT is obvious, indicating a systematic error in either the source or antenna positions. Neglecting the curvature of the Earth over ~ 750 m in calculating antenna positions implies an error of ~ 9 cm (0.3 ns), or perhaps one sample. More likely would be offsets in the positions of the antennas themselves, the coordinates of the sources used, or the assumption of a pure E-W baseline. That the data between 3C273 and M87 shows the same variation indicates that it is probably the former.

The few calibrations during February which produced deviations from the 192-sample intervals thus aroused suspicion. Using 3C273 as the source for F_6^C produced a much better alignment, consistent with an offset in integer multiples of 192 samples. This demonstrates the power of precise timing for correcting observer mistakes!

Similar calculations were performed by R. Robinson and C. Phillips, using completely independent software and analysis. Importantly, the antenna positions used were those utilised for VLBI, although the source positions were exactly the same. Both the 192 sample jumps, and a (much larger) variation with UT, were confirmed. While the UT variation introduces only a small error which could readily be corrected for by fitting a sinusoid, the random jumps in units of 192 ns obeyed no specific behaviour. Since the error is extremely large, considerable effort (and the next section) is devoted to correcting for these.

6.5.4 Aligning times in a world of 192-sample offsets

In this section I describe my efforts to account for the apparently random 192-sample offsets in the timing calibration observed in Sec. 6.5.3. At a collaboration meeting in mid-2008 at the ATNF in Sydney, it was found that the clocks were not updated during the ‘set’ operation, which was needed to recalculate the buffer centre after the buffer length was changed, and also upon initialisation. The length of this operation was 192 samples. This explained why times during calibrations were consistent (buffer lengths do not change), but why sometimes the calibration times between different periods did change, since in trigger-mode the buffers were much shorter than during calibrations. If everything ran smoothly, set-ing all three buffers would result in the same time loss between them, so that their

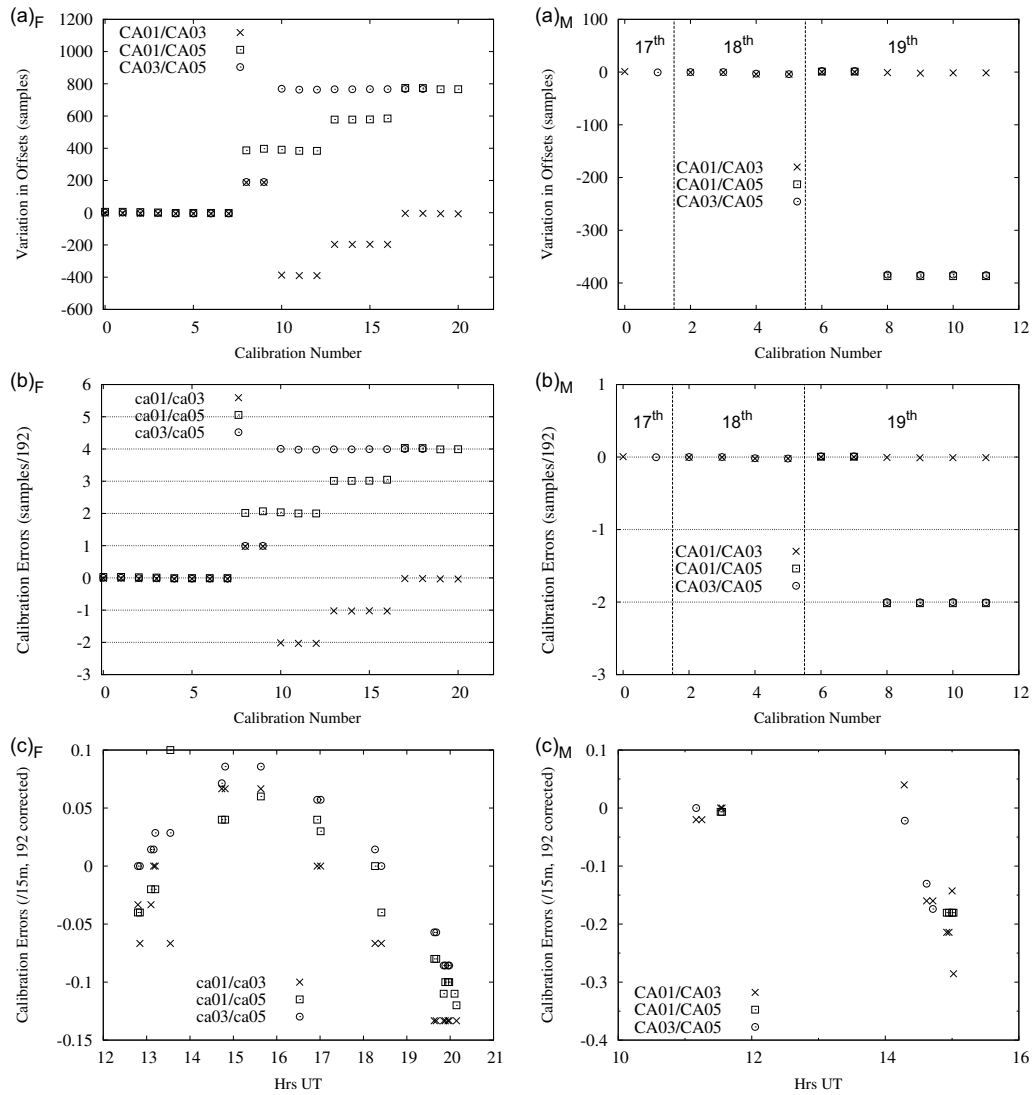


Figure 6.11: Consistency check of correlation times in the February (left) and May (right) data. Top row: alignment relative to the first calibration of the February period (left), and to the first of each day in May (right), in (x-axis) chronological order. Middle row: as top row, but with the y-axis divided by 192. Bottom row: Alignments corrected for 192-sample offsets, with the result divided by the baseline length in units of 15.306 m, plotted vs hrs UT on each day.

alignment would be preserved. However, occasionally a buffer would be set twice, or the set command would be used when re-adjusting trigger levels on one antenna only. In general, every time a calibration was performed or trigger levels adjusted, the relative timing could move by intervals of 192 samples. Additionally, sometimes a buffer computer would suffer an error and need to be restarted, further changing the alignment with that particular buffer. Since each of these occasions were recorded in the observation logs (though not how often the set command was used), it is possible to break the observation period down into intervals where the time alignment will at least be consistent, if not known.

There are two basic methods of accounting for the 192-sample offsets. One is to use the log-book to determine what a *likely* alignment would have been, which is essentially making an educated guess. The second is to use the data itself. In actuality the first method is rather quick and painless, while the second can involve numerous types of analysis with increasingly greater levels of complexity which are worth describing separately. To summarise, using the log-book proved very useful, while only in one instance did the analysis produce a useful result, and then only for the 1–3 alignment.

Educated guesses at the alignment

In Tables 6.9 and 6.10, I summarise the intervals spent observing the Moon between which the alignments could have changed, and the periods of known calibration (columns 3–6 will be explained in the next section). For each calibration period, offsets are given as integer intervals of 192 samples relative to this (the first) calibration after the clock reset (those marked with a *) in the usual 1–3, 1–5, 3–5 order, so that the value of +1/+2/+1 for F_5^C implies that $t_3 - t_1$ for that calibration period was 192 samples greater than it was for F_1^C , $t_5 - t_1$ was $2 \times 192 = 384$ samples greater etc. Unless otherwise specified, the changes which separated the periods observing the Moon were due to an adjustment of the thresholds on all three antennas.

Working through the list from start to finish, the alignment appears to be stable over February 26th. The change from 0/0/0 to +1/+2/+1 from the 26th to 27th probably resulted from differing reset commands being sent to the buffers upon starting the triggering software. The shift from +1/+2/+1 to -2/+2/+4 can easily be explained by the failure and subsequent re-initialisation of the CA03 system, since the alignment between CA01 and CA05 stays unchanged. Again, a change between the nights of the 27th and 28th is to be expected, while the shift from -1/+3/+4 to +0/+4/+4 can also be explained by the failure of CA01.

For the May observations, only one (marginally successful) calibration was performed on the 17th. However, there were also no ‘strange’ events (such as control computers failing) which would be likely cause a change in the offsets; also, each time the levels were re-adjusted, all three were, so likely the relative timing was kept constant. This was probably the case for the first four periods of the 18th (between M_4^C and M_5^C), despite only the levels on CA01 being adjusted at 14:07 UT (the alternative is that an error occurred here which

ID	Start (UT)	(mins)
Feb 26 th		
F ₁ ^M	13:19	34
F ₂ ^M	14:02	36
F ₁ ^C :	+0/+0/+0*	
F ₃ ^M	14:55	52
F ₄ ^M	16:04	94
F ₅ ^M	17:50	21
F ₂ ^C :	+0/+0/+0	
F ₆ ^M	18:30	59
F _{3,4} ^C :	+0/+0/+0	
Feb 27 th		
F ₅ ^C :	+1/+2/+1	
F ₇ ^M	13:23	51
F ₈ ^M	14:27	41
CA03 Failure - Reinitialising		
F ₉ ^M	15:23	10
F ₆ ^C :	-2/+2/+4	
F ₁₀ ^M	15:53	221
F ₇ ^C :	-2/+2/+4	
Feb 28 th		
F ₉ ^C :	-1/+3/+4	
F ₁₁ ^M	13:46	89
CA01 Failure - Reinitialising		
F ₁₂ ^M	15:37	73
F ₁₀ ^C :	+0/+4/+4	
F ₁₃ ^M	17:11	80
F ₁₄ ^M	18:42	73
F ₁₁ ^C :	+0/+4/+4	

Table 6.9: Start times (UT) and duration (minutes) for each lunar observation period (continuous period in detect mode) for the February 2008 run.

was coincidentally undone during the calibration M₅^C, which is very unlikely); thus I assume that these times are OK. For the latter half of May 18th, twice were individual trigger levels were adjusted, with the possibility of a change in timing. Finally, for the 19th, the changes observed between M₆^C and M_{8,9}^C are consistent with the adjustments to CA05 only, though it is ambiguous as to where exactly the changes took place. It could be assumed that for M₁₈^M-M₂₀^M the alignment is consistent.

The outcome is that in general, the changes between calibrations can be readily explained. However, there are some periods (particularly in May, when less calibrations were performed) for which the alignment must be extrapolated rather than interpolated. Even for the periods about which we are very sure of the alignment, we would not have sufficient proof to claim (for instance) a positive detection. The important outcome though is that

ID	Start (UT)	(mins)
May 17 th		
M ₁ ^M	11:00	58
M ₂ ^M	12:00	60
M ₃ ^M	13:06	63
	M ₃ ^C :	+0/+0/+0*
M ₄ ^M	14:47	7
M ₅ ^M	15:17	70
M ₆ ^M	16:33	65
May 18 th : Resetting All Clocks		
	M ₄ ^C :	+0/+0/+0*
M ₇ ^M	11:30	111
M ₈ ^M	13:22	39
M ₉ ^M	14:03	4
	Adjusting CA01 only	
M ₁₀ ^M	14:08	24
	M ₅ ^C :	+0/+0/+0
M ₁₁ ^M	14:55	106
M ₁₂ ^M	16:48	58
	Adjusting CA05 only	
M ₁₃ ^M	18:09	16
	Adjusting CA01 only	
M ₁₄ ^M	18:26	18
May 19 th : Resetting All Clocks		
	M ₆ ^C :	+0/+0/+0*
M ₁₅ ^M	11:41	106
	Adjusting CA05 only	
M ₁₆ ^M	13:30	15
	Adjusting CA05 only	
M ₁₇ ^M	13:46	66
	M _{8,9} ^C :	+0/-2/-2
M ₁₈ ^M	15:06	55
M ₁₉ ^M	16:13	179
M ₂₀ ^M	19:26	9

Table 6.10: Start times (UT) and duration (minutes) for each lunar observation period (continuous period in detect mode) for the May 2008 run.

over the entire time, we have narrowed the calibration times down to a small range, which should prove sufficient both to guide the re-alignment attempts described in the following sections, and also enable a pulse search with few enough candidates that any positive detection would be extremely suggestive, while a lack of candidates would be a certain null result.

Approximate coincidences and two-fold analysis

The ‘log-book’ alignment method described above allows the times to be known to within a few intervals of 192 samples. For the maximum baseline of 765.3 m, the largest true offset possible is ± 5225 samples, corresponding to signals arriving from due East/West, i.e. along the length of the baseline. Cautiously allowing for up to $\pm 10 \times 192$ sample offsets, this gives a total window of ± 7145 samples, which I round up to 8000 for extra caution; alternatively, this allows a path difference due to reflections (e.g. off the structure of the antennas) of up to 125 m. We can be confident that a signal from any source (as opposed to random noise) triggering two or more antennas will appear in this range. In particular, a strong source of RFI might appear repeatedly, producing a number of coincident triggers in this interval with the same relative times. If such a source can be found over more than one continuous run on the Moon, then it should be possible to adjust the offsets in units of 192 samples accordingly, thereby recalibrating the times. Of course, this requires a large number of RFI-sourced coincidences. A first step was to use the calibrations from Sec. 6.5.4 to search over the entire observation period for times when the triggers from two or more antennas fell within this 8000 sample ($\sim 4 \mu\text{s}$) window. The results are given in Tbls. 6.11 (Feb) and 6.12 (May), where the values N_{ij}^8 are the number of such coincidences in each interval between antennas CA0i and CA0j (including three-fold coincidences), while N_{135}^8 is the number of three-fold coincidences. The counts of two-fold events include the three-fold events, so by definition $N_{ij}^8 \geq N_{135}^8$, with equality only if there are no two-fold events without a matching event on the third antenna.

The most obvious result is the extremely large number of such coincidences, and the large number of three-fold coincidences compared to two-fold coincidences. The expected rate of two-fold and three-fold coincidences from purely random arrival times is given below in Eq. (6.22), where R_i (R_{ij}) [R_{135}] is the rate (Hz) of single (two-fold) [three-fold] triggers in antenna(s) i (i,j) [1,3,5], and W_t is the time window (seconds) required for a ‘coincidence’:

$$R_{ij} = R_i R_j W_t \quad (6.20)$$

$$R_{135} = R_i R_j R_k W_t^2 \quad (6.21)$$

$$R_{135}/R_{ij} = R_k W_t. \quad (6.22)$$

Therefore the ratio between two-fold and two-fold coincidences increases with the trigger rate. For a maximum trigger rate R_i of 3 kHz and time window $W_t = \pm 3.906 \cdot 10^{-6}$ s

ID	N_{13}^8	N_{15}^8	N_{35}^8	N_{135}^8
Feb 26 th				
F_1^M	1	0	0	0
F_2^M	145	25	17	10
F_3^M	109	120	235	54
F_4^M	5374	657	681	213
F_5^M	0	0	0	0
F_6^M	816	484	600	172
Feb 27 th				
F_7^M	45	41	47	0
F_8^M	1298	566	900	329
F_9^M	0	0	0	0
F_{10}^M	67551	39291	42277	29722
Feb 27 th				
F_{11}^M	10481	6336	6752	5289
F_{12}^M	1869	905	1004	527
F_{13}^M	2807	1142	2821	913
F_{14}^M	8139	207	495	52

Table 6.11: N_{ij}^8 : Number of two-fold coincidences between antennas i and j , and N_{135}^8 : number of three-fold coincidences, within an 8000 sample ($\sim 4 \mu\text{s}$) window for each lunar observation period in February 2008. The count of two-fold events also includes the three-fold events.

(8000 samples), the two-fold rate is 70 Hz. However, the three-fold rate is only 1.6 Hz, i.e. only $\sim 2\%$. While no period has a 70 Hz two-fold rate over the entire observation time (the highest average is M_{17}^M at 25 Hz), periods such as M_4^M and M_{11}^M which exhibit the correct ratio R_{135}/R_{ij} can be explained by assuming the triggers all came in a brief time due to some strong narrow-band RFI source increasing the trigger rate. Clearly this explanation breaks down for periods with R_{135}/R_{ij} over 2%. Extreme examples are F_{11}^M , F_{13}^M , and M_{17}^M , where approximately 90% of 1–5 coincidences are also coincident with a CA03 event, indicating a very strong RFI source. The obvious conclusion therefore is that the vast majority of observed three-fold coincidences do not occur purely by coincidence, but rather are triggered from a common event with significant time-structure. By extension, there will be many such events seen only in two antennas, and the same must therefore apply to the two-fold coincidences, of which there are (generally) many more.

In Figs. 6.12 and 6.13, I have plotted all two-fold coincidences for each of the three combinations over the entire observing period for both runs – over 5×10^5 data points. On the x-axis is the time of observation (UT), while the y-axis gives the time offset $t_j - t_i$ ($j > i$) between each coincidence. Upwards is East, since a large value of $t_j - t_i$ indicates the antenna with the higher number (which is West-most) received the trigger at a ‘larger’ (i.e. later) time; thus downwards is West. Time increases continuously except in breaks between days as indicated by the thick vertical lines. Marked are the intervals between

ID	N_{13}^8	N_{15}^8	N_{35}^8	N_{135}^8
May 17 th				
M_1^M	98	91	124	29
M_2^M	159	56	91	19
M_3^M	59	59	193	13
M_4^M	237	311	198	1
M_5^M	1093	642	617	34
M_6^M	698	766	771	0
May 18 th				
M_7^M	9455	8161	9118	2953
M_8^M	140	114	148	15
M_9^M	10	10	8	0
M_{10}^M	43	35	33	0
M_{11}^M	11407	10446	10732	435
M_{12}^M	416	390	375	33
M_{13}^M	72	65	48	0
M_{14}^M	231	224	173	1
May 19 th				
M_{15}^M	1100	1714	1734	272
M_{16}^M	364	499	634	56
M_{17}^M	103683	63352	59728	55737
M_{18}^M	820	760	889	142
M_{19}^M	7600	7120	7706	886
M_{20}^M	816	866	622	400

Table 6.12: N_{ij}^8 : Number of two-fold coincidences between antennas i and j , and N_{135}^8 : number of three-fold coincidences, within an 8000 sample ($\sim 4 \mu\text{s}$) window for each lunar observation period in May 2008. The count of two-fold events also includes the three-fold events.

each continuous run, and also the assumed offsets (units of 192 samples) and the calibration times at which each offset has been measured. In all figures, a stationary source of pulse-like RFI producing triggers over a long period of time will show up as a horizontal line, while a very brief period of strong, narrow-band RFI will be observed as a very large number of coincidences over a small time range (small x-width) but large vertical extent (since the times will be random), i.e. a vertical line. Fig. 6.14 shows a similar plot for both February and May, but with two exceptions: only triple coincidences (those with a trigger on all three antennas within the time window) are plotted, and the delays on the y -axis have been divided by the baseline length in units of 15.306 m and plotted for all baselines on the one graph, effectively measuring wavefront curvature.

The resulting graphs provide an amazing amount of information. The features can be approximately classed as below, essentially from the least to the most interesting:

1. Short time periods exhibiting a high rate of coincident triggers at all offsets, mostly

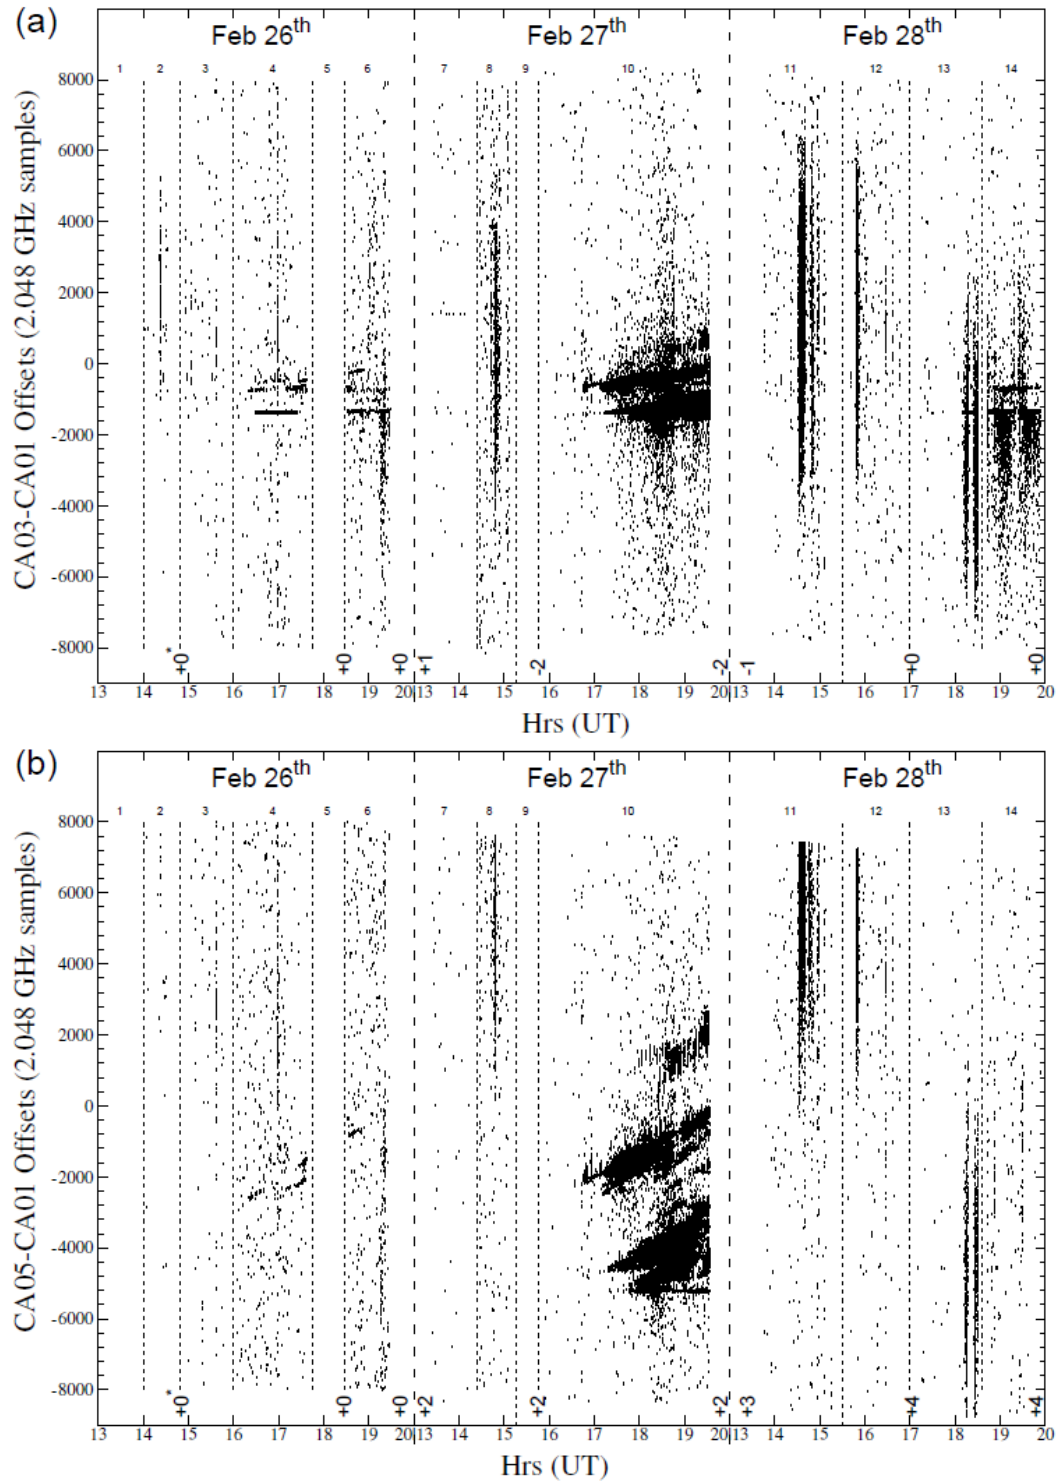


Figure 6.12: (continued next page)




Figure 6.12: Doubly-coincident triggers observed during the February run between antennas (a) 1 and 3, (b) 1 and 5, and (c) 3 and 5. See text.

in May (e.g. 15:35 May 17th, 12:50 May 18th, 10:00-17:20 May 19th). These appear as vertical features which extend uniformly over the full range of offsets.

2. Short time periods exhibiting a high rate of coincident triggers for a broad range of offsets, mostly in February (e.g. 14:25 UT Feb 26th, 14:40 and 14:50 28th, and 19:40 UT May 19th). These appear as vertical features which do *not* extend over the full range of offsets.
3. Purely horizontal, typically thin features occurring at a characteristic offset, sometimes over many days (e.g. the lines $t_3 - t_1 \approx -1300$ for February, $t_5 - t_3 \approx 1600$ on May 19th, and the sometimes-periodic feature near $t_3 - t_1 = -400$ throughout February and (at other offsets) on all baselines for May 19th.)
4. Sloped, typically broad features occurring most obviously around 17:00–19:00 UT on Feb. 27th, but also on Feb. 26th.

Type 1 features are exactly what would be expected from a high random trigger rate, with triggers evenly spread in time-offsets. The cause of the increased trigger rate must be a lowering of the effective threshold by an increase in the background containing no timing information, probably from narrow-band RFI – an increase due to ground temperature or the galactic background would be unlikely to produce such short-duration bursts. As would

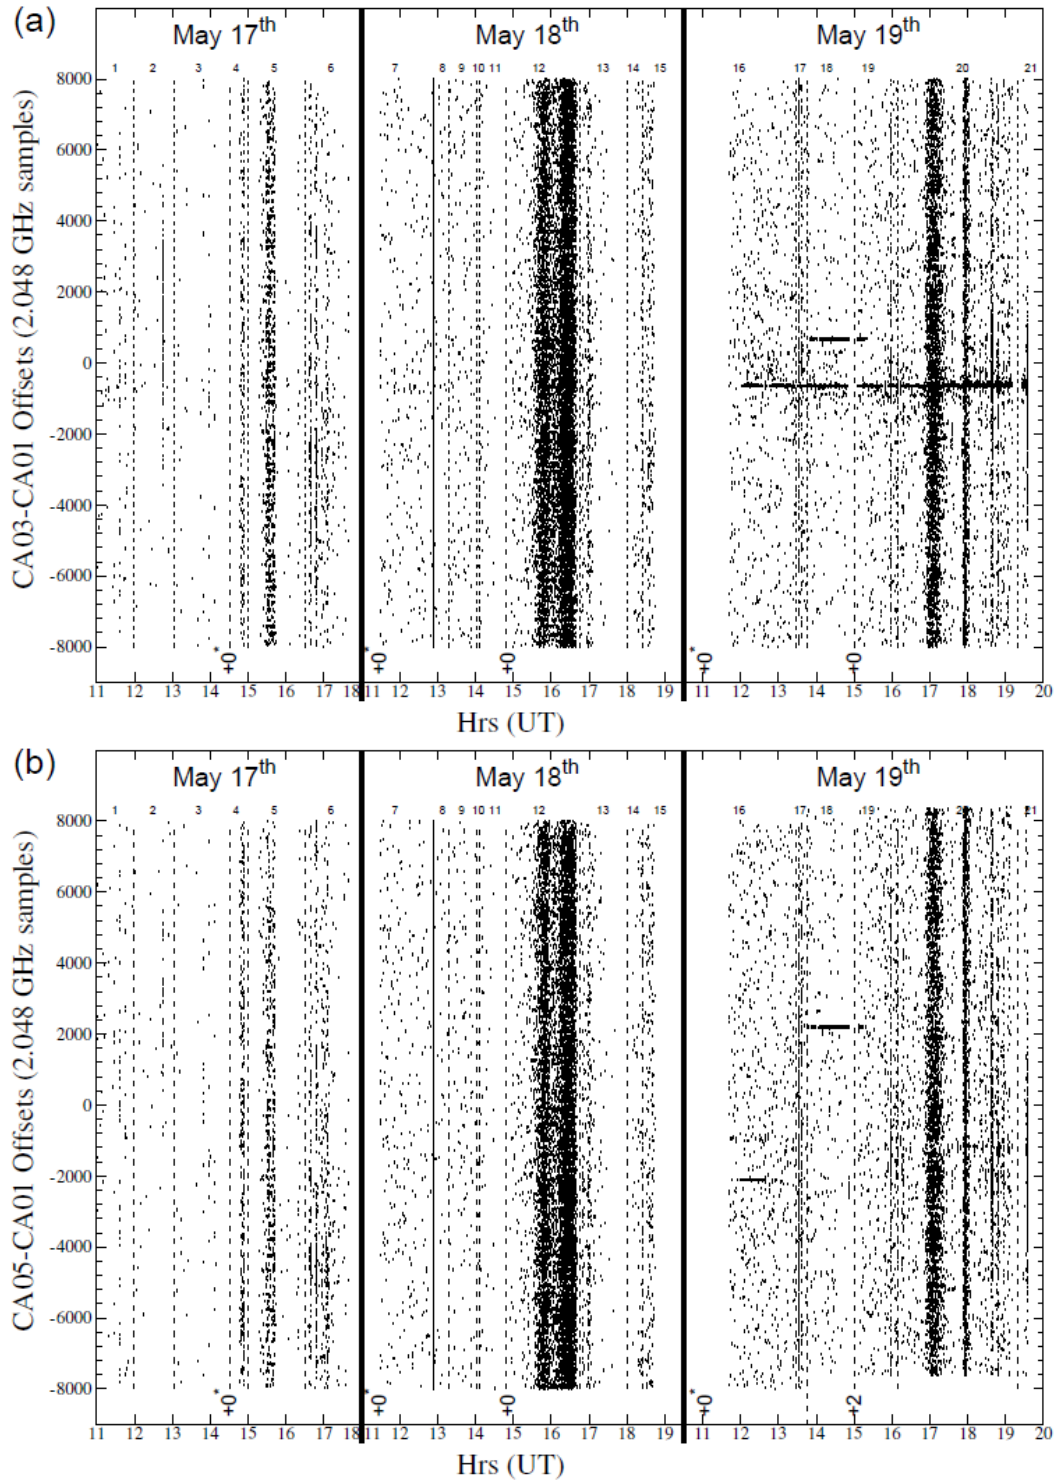


Figure 6.13: (continued next page)

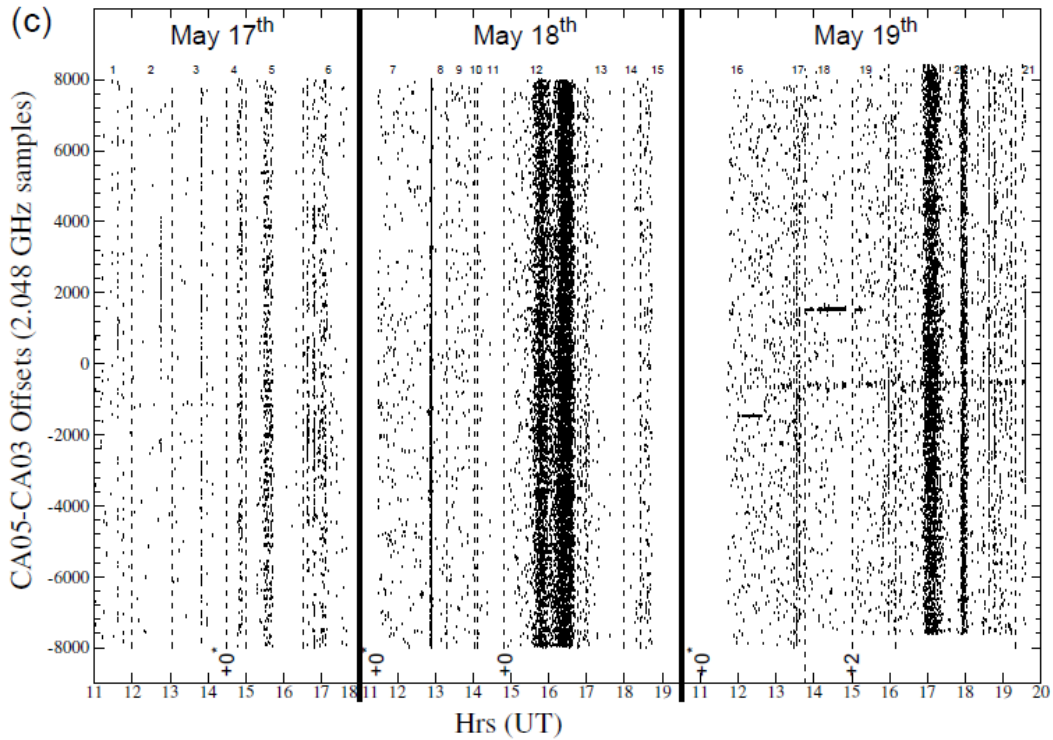


Figure 6.13: Doubly-coincident triggers observed during the May run between antennas (a) 1 and 3, (b) 1 and 5, and (c) 3 and 5. See text.

be expected, the triples rate is much lower than the doubles rate during these times, due to the random nature of the trigger times.

Type 2 features show both a small increase in triggers at all offsets corresponding to a random component from an increased background, and a large increase in triggers spread about a specific offset. Some bands likely correspond to the nearby towns of Narrabri (East) and Wee Waa (West-North-West) which present potentially large, extended sources of RFI. While these may be too weak to be detected under normal conditions, a decrease in the effective threshold due to the presence of narrow-band RFI could make such sources detectable. Another explanation is a single source of RFI with a broad time-structure, with a high likelihood of different antennas triggering off different parts of the signal, thereby adding a random component with a preferred direction. These features appear strongly for both doubles and triples, indicating the non-randomness of the timing.

Type 3 features act exactly as fixed sources of short-duration RFI. The fixed geometrical delay results in a common time-offset regardless of antenna pointing position (and hence time). However, a jump in the clock alignment will also show a break and vertical jump in the horizontal feature, although here these have mostly been corrected for already. These give us our first result on alignment from the data. The continuity of the $3 - 1/5 - 3 \approx -600$ line over different lunar pointings on May 19th indicates that the relative times for both these combinations of antennas (and therefore for all three) are correct for this date. Also, the par-

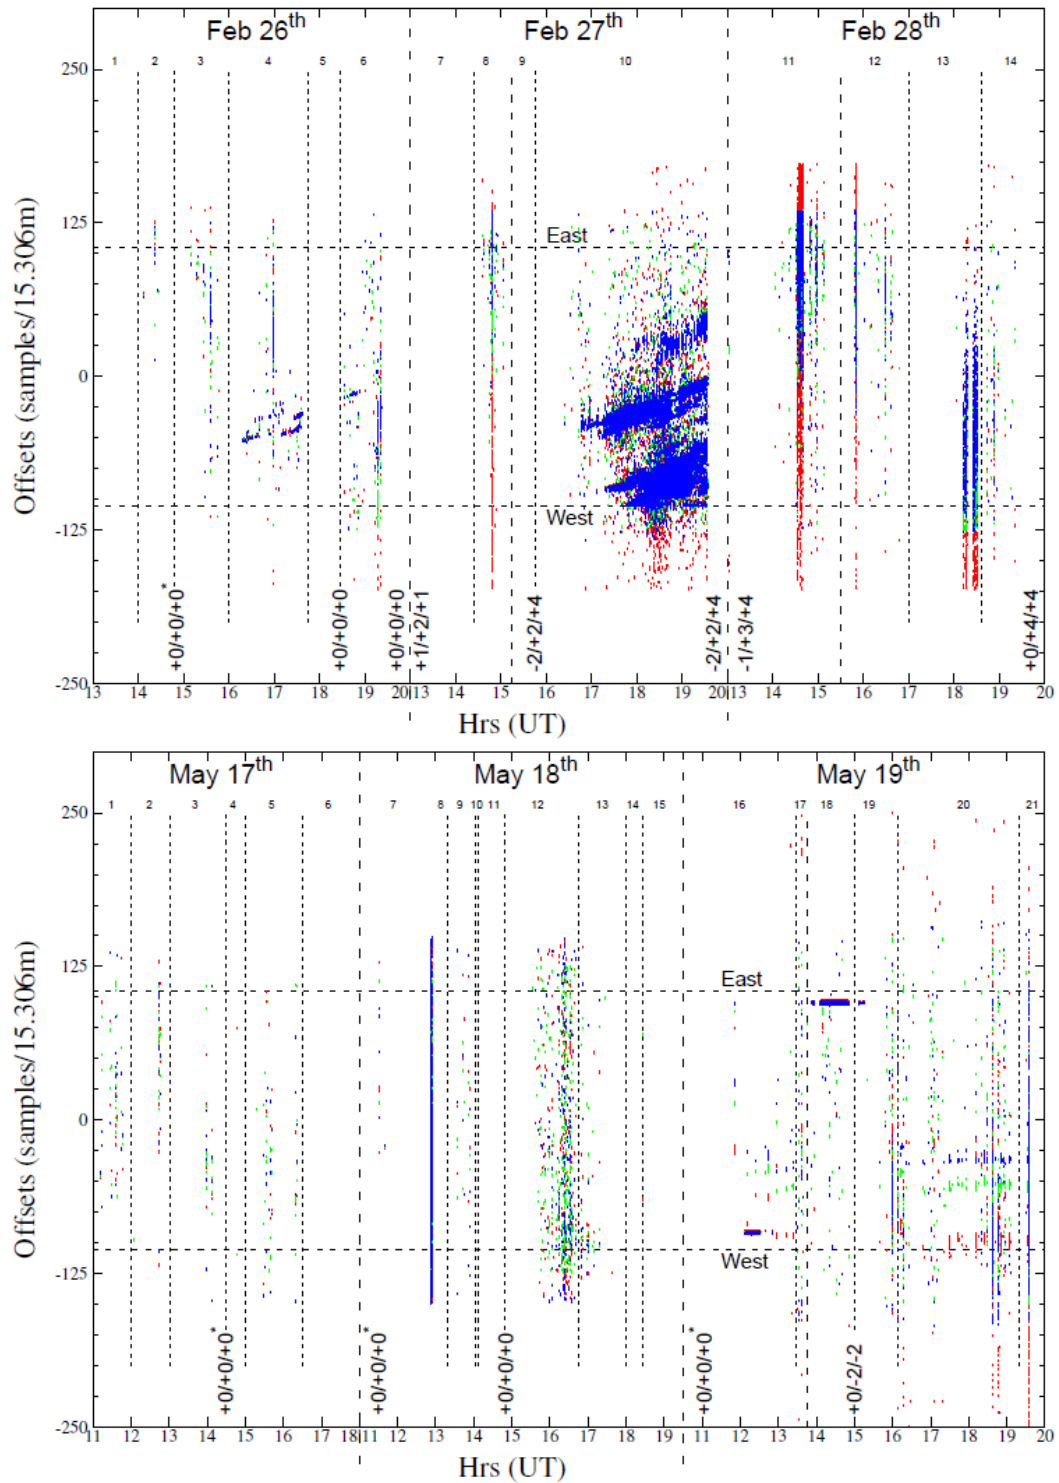


Figure 6.14: Three-fold coincident triggers within an 8000-sample ($\sim 4 \mu\text{s}$) window, plotted with UT (x-axis) vs $j - i$ delay ($j > i$, y-axis) in 2.048 GHz samples per 15.306 m baseline, for both Feb. (top) and May (bottom). The order for the offsets along the bottom are 3-1/5-1/5-3 – the corrections are thus the negative of these values multiplied by 192 – which correspond to red, green, and blue points respectively. Where points overlap, only the 5 – 3 (blue) points show.

Known Δt	Feb	May
All	4,6,10	16–21
Δt_{13} Only	2,3,11–14	
None	1,5,7–9	1–15

Table 6.13: Calibration information gleaned from Figs. 6.12, 6.13, 6.14 only.

allel horizontal features at $3-1 \approx -700, -1300$ appear at $F_{M4}, F_{M6}, F_{M6}, F_{M10}, F_{M13}$, while the intermittent, periodic ‘blips’ at $31 \approx -400$ span F_{2-4}^M, F_{11-13}^M , and perhaps F_6^M and F_{14}^M also, so that we can be confident of having a consistent alignment over all but a few periods here. Disappointingly, there are no such features over either May 17th or 18th. Interestingly, during February no such features were detected in three-fold coincidence. However, on May 19th, perhaps as many as five were. Two of these had approximately the same delay/m over all baselines, one being almost West at $\sim 12:20$ UT, the other almost East at $\sim 14:00$ UT. Though in these directions the apparent baselines are small, nonetheless the sources must have been rather distant – and therefore quite powerful – to produce the observed triggers. The other three – 12:00 UT, 12:40 UT, and $\gtrsim 17:30$ UT – must have originated in the very near-field, since the apparent angle to the source between each baseline is very different. In each case, it should therefore be possible to locate the actual source, perhaps with a view to removing its influence in further observations.

Type 4 features remain largely unexplained. The locus of the coincident triggers (lines with constant slope) are consistent with some RFI source moving with constant speed in $\cos \theta$ (θ the angle w.r.s.t. the baseline), which is strong enough to give a high rate of triples-coincidences, and in the far-field, since the delays per unit baselines match. Also, since multiple features are seen at once, there must be either many such sources all moving in unison, or many multiple reflections keeping the same (and extremely large) angular offsets over a broad period of time. Also, the apparent motion is at the sidereal rate, but in the opposite direction. One suggestion is that a far-field RFI source is being observed over multiple signal paths due to atmospheric distortion. However, these would cause a ‘wobbly’ curve about a constant delay. Reflections off the antennas themselves cannot explain the rate of change of delay being proportional to the baseline length, nor is the antenna size of 22 m sufficient to produce more than a ~ 70 ns change in delays. We can be certain however it is not an equipment fault due to the presence of the aforementioned type 3 feature during this time period. Other than a military ‘black-ops’ conspiracy or little green men, the most plausible explanation is a series of reflections off an extended object, allowing each reflection point to move smoothly with time. The most likely such object is the railway track that the antennas run on, though further investigation would be needed to determine if the geometry allows this. Despite the origin and cause for this strange phenomena remaining unidentified, the original source of the signal is evidently in the far-field, so that the constant delay/baseline of F_4^M, F_6^M , and F_9^M confirms our timing is correct here.

Source identification with three-fold triggers

The next level of sophistication is to search for actual RFI source positions. Identifying a particular RFI source during a period when the relative timing is known precisely allows a semi-random search for that source during periods when the timing is unknown by using trial offsets. Using the timing from only two antennas at once doesn't give enough discriminatory power during periods when there are no obvious type 3 or 4 features, especially over multiple trials. Also, if any on-site RFI sources are found, they could be deactivated and/or shielded in time for follow-up experiments, since in many cases these events dominate our trigger rates and limit sensitivity.

Given three antenna positions on an East-West baseline, we can solve for the source position to within a North-South ambiguity, since an event some distance North of the baseline would produce exactly the same time structure if its location were directly South of the baseline by the same distance. In the following analysis, I break the three-fold coincidences into two types of events: near-field and far-field. Far-field events can only be characterised by their apparent direction, in θ degrees from East. Technically, solutions fall along the arc of a cone of opening angle 2θ , with a tip pointing West for $\theta < \pi/2$ and East for $\theta > \pi/2$. For simplicity I assume that these events come from either the North or South the intersection points of this cone with the horizon. For near-field events I use as the x -coordinate the East-West distance from CA01, and the y -coordinate the North-South distance to the East-West baseline. Again, I assume events originate on the ground.

For a 'far-field' event, the angle θ to the baseline is given by averaging the three fitted values θ_{13} , θ_{15} , and θ_{35} given by Eq. (6.23):

$$\cos \theta_{ij} = \frac{c\Delta t_{ij}}{x_i - x_j}. \quad (6.23)$$

An error margin of ± 2 samples on each of the baselines is allowed, so that the 'far-field' begins approximately 100 km directly North/South of the antenna, but is much closer in more Easterly and Westerly directions, since the apparent baselines here are shorter. For an event inside the far-field, the time-delay Δt_{13} between the signal arrival times at antenna 3 at $(x_3, 0)$ and antenna 1 at $(0, 0)$ from a point source at position (x, y) is given by Eq. (6.24):

$$c\Delta t_{13} = \sqrt{(x - x_3)^2 + y^2} - \sqrt{x^2 + y^2} \quad (6.24)$$

where CA01 is at the origin, and the coordinates of CA03 are $(x_3, 0)$. Here I approximate $n_{\text{air}} = 1$. Adding the distance to both sides, squaring and cancelling gives:

$$c^2\Delta t_{13}^2 + 2c\Delta t_{13}\sqrt{x^2 + y^2} + x^2 + y^2 = x^2 - 2xx_3 + x_3^2 + y^2 \quad (6.25)$$

which can be reduced to:

$$\sqrt{x^2 + y^2} = \frac{x_3^2 - 2xx_3 - c^2\Delta t_{13}^2}{2c\Delta t_{13}} \quad (6.26)$$

Note that squaring both sides of Eq. (6.25) would give an equation similar to that for a hyperbola, being the solution set for a two-fold coincidence only. Performing the same operations for Δt_{15} produces the same LHS as Eq. (6.26) – equating the RHSs eliminates y to give:

$$\frac{x_3^2 - 2xx_3 - c^2\Delta t_{13}^2}{2c\Delta t_{13}} = \frac{x_5^2 - 2xx_5 - c^2\Delta t_{15}^2}{2c\Delta t_{15}} \quad (6.27)$$

which solves for x as:

$$x = \frac{\Delta t_{13} (x_5^2 - c^2\Delta t_{15}^2) - \Delta t_{15} (x_3^2 - c^2\Delta t_{13}^2)}{2x_5\Delta t_{13} - 2x_3\Delta t_{15}} \quad (6.28)$$

The solution for y can then be found by squaring Eq. (6.25) and substituting for x as:

$$y = \left(\frac{x_3^2 - 2xx_3 - c^2\Delta t_{13}^2}{2c\Delta t_{13}} - x^2 \right)^{0.5} \quad (6.29)$$

While Eq. (6.29) gives the set of solutions for real event, the positional errors for distant events can be large.

Correlation correction

To use the above equations to extract positional information, the Δt_{ij} themselves must be reliable. This will not be the case if the antennas trigger off different parts of an incoming signal, which will be quite common for low-power, long-duration RFI. Therefore, for each pair-wise coincidence, the buffers are cross-correlated, and the peak used to correct for this offset by an amount $\delta\Delta t_{ij}$, to produce a new $\Delta' t_{ij} = \Delta t_{ij} + \delta\Delta t_{ij}$. Since this is done pair-wise, there is no guarantee that the three adjustments $\delta\Delta t$ will add to zero even for a strong signal, especially since the sampling rate is above the Nyquist rate. For weaker signals, especially in cases where there is strong narrow-band RFI, the correlation function can produce many peaks of equal height. The procedure for picking the correct alignment was to identify all correlation points within 90% of the peak value for each of the three cross-correlations, choose from these candidates the combinations which gives the smallest error, and if this produced multiple solutions, choose the combination of the three which resulted in the smallest adjustment. In Fig. 6.15 are plotted histograms of the individual adjustments $\delta\Delta t_{ij}$, and the resulting ‘closure’ error $|\delta\Delta t_{13} + \delta\Delta t_{35} - \delta\Delta t_{15}|$

From Fig. 6.15, we can see that a large fraction (of order half) of modifications satisfied closure, with the next most significant fraction having errors of between 1 and 3




Figure 6.15: Histograms of the change in offsets $\delta\Delta t_{ij}$ for the February run, and a histogram of the resulting error in closure.

samples, which falls within the $\sim \pm 1$ samples errors expected from the sampling rate being approximately twice the Nyquist rate. Away from zero, the error count is falling rapidly. Importantly, while the bulk of adjustments are also centred at 0, the distribution is much broader than that for the errors, so there must be a significant number of events where the raw trigger times did not reflect the arrival direction accurately.

Results using three-fold coincidences

Using the offsets shown in Fig. 6.14, a search for both far-field and near-field events was performed for each block of data in each of the February and May observation periods. The majority of near-field solutions occur in the very near-field (within a km of the antennas), and the rates of both near- and far-field events are highly variable. In many cases point-like sources of RFI are seen, both in the near- and far-field, and it makes an interesting game trying to align the positions of possible RFI sources with those detected. However, as previously mentioned, this analysis has not been successful in shedding light on the correct offsets, for reasons which are discussed below. Therefore here I present a sample of results which, while they do nothing to assist in reconstructing the times, are still of casual interest.

In Fig. 6.16 are shown the fits for the far-field solutions based on ‘best guess’ offsets (those shown in Fig. 6.14) and including correlation corrections with an allowable closure error of ± 3 samples. The bins have 0.25° spacing, corresponding to the resolution of a 33 m apparent baseline at a 2.048 GHz sampling rate. For one definite case only (M_{18}^M and M_{19}^M), and one less definite (F_6^M and M_8^M), is there any consistency in the number and direction of far-field events. This is likely due to such events entering through far side-lobes, so that the sensitivity to these RFI sources is non-negligible only for a few brief time periods. Also, if a strong RFI source was detected, the thresholds were adjusted to bring the trigger rate under control - which necessarily meant adjusting the thresholds until the source was no




Figure 6.16: Histograms of the arrival directions of ‘far-field’ (\sim planar wave-front) events for February (top) and May (bottom) observation runs, using correlation-corrected data (see text). Each continuous pointing period is binned separately as a function of the angle from the North-South line, and individually normalised (the bin with the most counts in each period is black). On the bottom is the sum before normalisation, typically reflecting the periods with the highest count rate.

longer visible.

Plots of the near-field events are similar – some individual plots show strong point-like near-field sources, others a broad smear of apparent origins, and others have no fitted events at all. The localisation of near-field events is expected to be less accurate than that for far-field events, since all near-field RFI must come in a side-lobe of the antennas, or some more complex path involving perhaps multiple reflections off the antenna superstructure. While this is also the case for far-field events, the angle of each antenna to a near-field source (and hence the signal path) will vary between antennas, leading to a potential path differences of order the size of each antenna – ± 22 m, or ± 150 samples.

Near-field sources include RFI generated on-site, and also reflections off buildings, some of which are quite close to antenna locations (see the site-map of the central area shown in Fig. 6.17). The potential spread resulting from a ± 30 sample error for a range of these buildings is shown for all three antenna configurations in Fig. 6.18. The ‘blurring’ of the reconstruction due to these errors depends greatly upon the relative positions of the source and the antennas – when the source is seen at similar angles by all antennas (e.g. most central buildings in the May 17th – 18th configuration), a small error leads to a large offset in the distance. In these cases a bright, near-field event will appear ‘comet-like’, with a concentration of events about the true position, and an extended tail out to the far-field. For sources seen at sufficiently different angles (e.g. those $x \in \{x_5, x_1\}$), the effect of errors is smaller.

Comparing the expectations of Fig. 6.18(c) with the actual reconstructions of Fig. 6.19, we have almost certainly detected the residence, the control-building/lab, and one or more of the solar observatory and the lodge during the period M_{15}^M . For period M_{17}^M , either the visitor’s centre or the IPS act as a strong source, while for periods M_{19}^M and M_{20}^M we have again detected the control building/lab. In the case of F_{10}^M , there is obviously a strong near-field source present in the (possibly reflected) vicinity of the IPS or visitor centre, though it does not correspond exactly with either location.

Interesting as all this is, it doesn’t help with the calibration, since few sources appear over multiple periods, and those that do can be aligned equally well with a 2-D analysis only.

6.5.5 Conclusion on the time alignment

The ‘best-guess’ timing solutions using all reconstruction methods are shown in Table 6.14. We can be 100% confident that for a slight majority of the observation periods, these ‘guesses’ are correct, so that we can correctly align the data over all three antennas, allowing a search for pulses from the direction of the Moon to be conducted with a very narrow search window. For the remainder of the data, although the timing is not known exactly, it is implausible that the alignment could vary by more than a few factors of 192 samples – these periods, along with the range of plausible variations, are given in Table 6.14. Note that while these figures appear somewhat arbitrary, they are based on recollections of what

NOTE:
This figure is included on page 158 of the print copy of
the thesis held in the University of Adelaide Library.

Figure 6.17: Plan of the central area of ATCA. From www.atnf.csiro.au

was done at the time of the observations, and on known alignments for preceding and subsequent observation periods. For these cases, a search can be performed using each possible alignment, with one window existing for each possibility. Together, this gives sufficient discriminating power to eliminate most candidates, but for only half the data could we be confident that a pulse-like event falling within the search window would genuinely have timing consistent with a lunar origin.

6.6 Sensitivity Calibration

To simulate the sensitivity of this experiment and place limits on a flux of UHE neutrinos, trigger levels at each antenna in terms of real quantities must be known. The usual specification for an experiment such as this is the detection threshold E^* in V/m/MHz (a V/m threshold divided by the bandwidth) in a given polarisation just prior to being received by the antenna. For instance, the GLUE threshold was a maximum field strength per bandwidth of $E_{\text{GLUE}}^* = 1.23 \cdot 10^{-8}$ V/m/MHz, calculated by taking the threshold of $6.46 \cdot 10^{-9}$ in each circularly polarised data channel, and accounting for vacuum–receiver transmission and the splitting of power between polarisations (121). Although the normal radio-astronomical

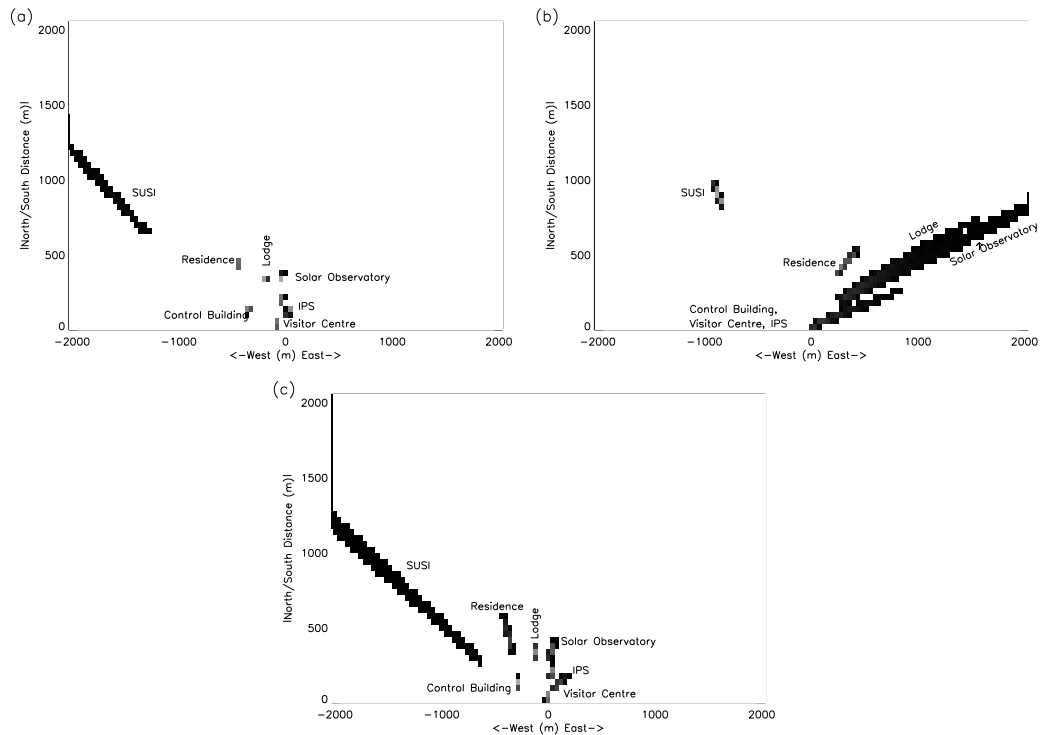


Figure 6.18: Predicted reconstruction of RFI positions – assuming a ± 30 -sample error at each antenna – for RFI originating from various potential on-site locations. The effect of the antenna positions w.r.s.t. each source is obvious. Most buildings appear in Fig. 6.17: the ‘residence’ is that labelled ‘(4)’ in Fig. 6.17; ‘SUSI’ is the Sydney University Stellar Interferometer, located approximately 1.5 km East and 1 km South of the control building. Each building has been specified with a point-like location only, when in fact some are quite extended.

measure of sensitivity is the antenna effective area A_{eff} divided by the system temperature T_{sys} (measured automatically via the noise diode), the conversion between the two measures is simple. Of greater concern is that even when we had measurements of T_{sys} , they only applied to the standard signal path and over a small frequency range. Therefore another method had to be used, which is described in this section. Also dealt with here are the effects of ionospheric dispersion, which while approximately corrected for, still reduced our sensitivity to some degree.

Note that unless otherwise stated, $\Delta\nu$ refers to the total bandwidth from 1.024-2.048 GHz, although the sensitivity outside the range 1.2-1.8 GHz will be minimal.

6.6.1 The calibration function $k(\nu)$

For a wide-band experiment such as this one, the signal is expected to change significantly in strength over the band. Also, the de-dispersion filters have a variable bandpass. Therefore, our sensitivity will change as a function of frequency. Given that our threshold is set in terms

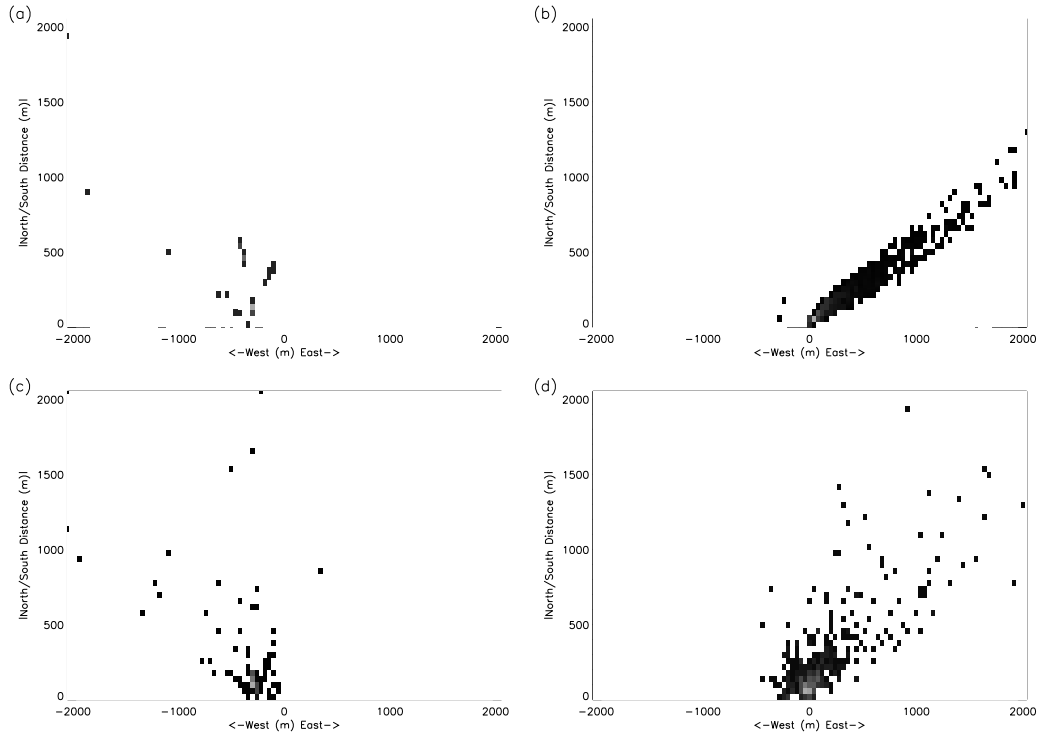


Figure 6.19: Reconstructed positions of near-field events for three periods on May 19th. These are (a): M_{15}^M , (b): M_{17}^M , (c): the sum from two like periods M_{19+20}^M , and (d): F_{10}^M .

ID	Alternative Offsets
M_{7-9}^F	[+1/ + 2/ + 1]-[-2/ + 2/ + 4]
M_{1-6}^M	± 5 on all
M_{7-15}^M	± 1 on all

Table 6.14: Alternative trial offsets from shown in Fig. 6.14. For F_{M7-9} , I try the range between the two calibrations, while the timing is expected (or known to be stable) for the rest of February. The network errors of May 17th meant there were a lot of adjustments and thus potential for offsets, while May 18th was both more stable and had two calibrations in agreement. The times of May 19th are completely aligned.

of our 8-bit sampling, we need to be able to convert from a coherent signal $E(\nu)$ (V/m/MHz) to the received values in the buffer $b(t)$, the Fourier-transform of which I denote $b(\nu)$. From here on, I simply call the values of b ‘sampling units’, or ‘s.u.’. The relationship between $E(\nu)$ and b involves an unknown (but determinable) function $k(\nu)$, defined as below:

$$b(\nu) \text{ (s.u.)} = \int_{\nu_1}^{\nu_2} k(\nu)E(\nu)d\nu. \quad (6.30)$$

The units of k are thus s.u./(V/m). It gives the conversion between real field strength at the antennas and the measured units in the CABB digitiser boards. For simplicity’s and comparability’s sake (and because the simulation program is not yet able to accept a generalised frequency-dependent sensitivity function), a more useful measure is \bar{k} , being the mean over a nominal bandwidth $\Delta\nu$ between ν_1 and ν_2 :

$$\bar{k} = \frac{1}{\Delta\nu} \int_{\nu_1}^{\nu_2} d\nu k(\nu). \quad (6.31)$$

The sensitivity of the experiment E_{thresh} , defined in terms of a threshold electric field strength per unit bandwidth, can then be calculated by knowing the trigger threshold b_{thresh} using equation 6.32:

$$E_{\text{thresh}}(\text{V/m/MHz}) \approx \frac{b_{\text{thresh}}}{k\Delta\nu} \quad (6.32)$$

for $\Delta\nu$ in MHz. Therefore, this section calculates separately $k(\nu)$, \bar{k} , and hence E_{thresh} for each data channel (antenna and polarisation band) over the entire observation period.

In order to calculate $k(\nu)$, a measurement of a known flux $F(\nu)$ (W/m²/Hz) is required. For an incoherent signal (random phases), the relationship between $F(\nu)$ and the electric field over a given bandwidth is given by:

$$\int_{\nu_1}^{\nu_2} F(\nu)d\nu = \mathcal{E}_{\text{rms}}^2/Z_0. \quad (6.33)$$

where I use \mathcal{E}_{rms} for the RMS field-strength in V/m; Z_0 is the impedance of free space, which I approximate to 377 Ω . This quantity is the same whether it is calculated in the time-domain or the frequency-domain (Parseval’s Theorem), i.e.:

$$\mathcal{E}_{\text{rms}}^2 \equiv \frac{1}{\Delta t} \int_{t_1}^{t_2} \mathcal{E}(t)^2 \quad (6.34)$$

$$= \frac{1}{\Delta\nu} \int_{\nu_1}^{\nu_2} \mathcal{E}(\nu)^2 \quad (6.35)$$

where $\mathcal{E}(t) \leftrightarrow \mathcal{E}(\nu)$ are the time- and frequency-domain expressions of the electric field

respectively. In the case of quantised data with n data points, Eq. (6.35) becomes:

$$\mathcal{E}_{\text{rms}}^2 = \frac{1}{n} \sum_{\nu_1}^{\nu_2} \mathcal{E}(\nu)^2. \quad (6.36)$$

Using the frequency-domain calculation of \mathcal{E}_{rms} in Eq. (6.33) and substituting Eq. 6.35 gives Eq. (6.37):

$$\int_{\nu_1}^{\nu_2} F(\nu) d\nu = \frac{1}{\Delta\nu Z_0} \int_{\nu_1}^{\nu_2} \mathcal{E}^2(\nu) d\nu. \quad (6.37)$$

Since this relationship holds for an arbitrary bandwidth, the integration can be eliminated (e.g. let $\nu_2 \rightarrow \nu_1$), giving:

$$F(\nu) = \frac{1}{\Delta\nu Z_0} \mathcal{E}^2(\nu). \quad (6.38)$$

Using Eqs. (6.30) and (6.38), the flux $F(\nu)$ as seen by the relevant data channel can then be related to the required calibration constant $k(\nu)$:

$$F(\nu) = \frac{1}{\Delta\nu Z_0} \left(\frac{b(\nu)}{k(\nu)} \right)^2 \quad (6.39)$$

which can be rearranged to give $k(\nu)$:

$$k(\nu) = \frac{b(\nu)}{\sqrt{F(\nu)\Delta\nu Z_0}}. \quad (6.40)$$

Eq. (6.40) states that we require a known flux $F(\nu)$, and the Fourier-transform of the corresponding sampled output $b(\nu)$, to determine $k(\nu)$. Note that for incoherent thermal emission, we do not need to track the phase change in each antenna over the bandwidth, i.e. we are interested only in the magnitude of $k(\nu)$.

The Moon as a flux calibrator

For our calibrator, we chose the Moon. The lunar temperature T_M is stable to within a few degrees over the lunar cycle at approximately 225° K in the 1–2 GHz range (see (115)). The amplitude variation with the lunar cycle is of the order of 1–2%, with comparable variation in the effective temperature across the bandwidth (the Moon is hotter at 1.8 GHz than at 1.2 GHz). Both errors could be reduced by using a frequency-dependent effective temperature ($T_M \rightarrow T_M(\nu)$) and determining the phase of the Moon over the observation period. The apparent lunar temperature will vary over the visible lunar disk, while throughout we use the mean temperature; also, the emitted strength will be polarisation-dependent, particularly at the limb. Combining these errors, this method should be accurate to within 5% or better, which is acceptable.

λ	a_1	a_2	a_3	a_4
20cm	$8.99 \cdot 10^{-4}$	$2.15 \cdot 10^{-6}$	$-2.23 \cdot 10^{-9}$	$1.56 \cdot 10^{-12}$
13cm	$1.02 \cdot 10^{-3}$	$9.48 \cdot 10^{-7}$	$-3.68 \cdot 10^{-10}$	$4.88 \cdot 10^{-13}$

Table 6.15: Parameters for the beam-shape fit – see Eq. (6.42).

Under these approximations, the lunar flux $F_M(\nu)$ (Jy) captured by the beam (it will be half this in any given polarisation channel) is given by:

$$F_M(\nu) = 2 k_b T_M \nu^2 / c^2 \int_{\Omega_M} \mathcal{B}(\Omega, \nu) d\Omega \quad (6.41)$$

where $\mathcal{B}(\Omega, \nu)$ is the beam-power pattern at a given celestial coordinate $\Omega(\alpha, \delta)$ at frequency ν , and c and k_b are the speed of light and Boltzmann constant respectively. We have two methods to approximate the beam-power-pattern: as an airy-pattern, using the antenna diameter of 22 m, and also from the ATNF memo in Ref. (119), which gives $\mathcal{B}(\Omega, \nu)$ for small offset angles, appropriate at these wavelengths to the diameter of the lunar disk. The fit is rotationally symmetric about the beam centre, so that $\mathcal{B}(\Omega, \nu) \equiv \mathcal{B}(\theta, \nu)$ where θ is the angular distance to the beam centre. The fit is given by the function:

$$\mathcal{B}(\theta, \nu) = \frac{1}{1 + a_1 x^2 + a_2 x^4 + a_3 x^6 + a_4 x^8} \quad (6.42)$$

where x is the distance in arcminutes times the frequency in GHz, i.e. $x = (\theta/1 \text{ arcmin}) \times (\nu/1 \text{ GHz})$. Coefficients for the fit out to $x = 50$, where $\mathcal{B} \sim 3\%$, are given in Table 6.15 – note that the x -width of the Moon at 1.8 GHz is 56.3, so these fits should be suitable. Both the Airy pattern from a uniform circular aperture and the fitted beam are compared with the size of the lunar disk in Fig. 6.20 (a), along with the error $|2 \text{ (fitted-airy)/(fitted + airy)}|$. Unsurprisingly, the error becomes large near the Airy nulls. Fig. 6.20 (b) shows the received fraction of lunar thermal emission in both centre-pointing and limb-pointing configuration, again for both beam patterns. In both configurations, it appears to matter little which beam pattern is used. In this section from here on we exclusively use the fit of Eq. (6.42).

6.6.2 Measurements

To measure $b(\nu)$ required the spectrum of sampled signal due to the lunar flux $F_M(\nu)$. To obtain this, we took an unbiased (trigger level of 0, i.e. maximally triggering) sample of data pointing both on and off the Moon. The received flux $F_M(\nu)$ from the Moon can be detected by subtracting the measured bandpass $b_{\text{off}}(\nu)$ when pointing away from the Moon from the bandpass $b_{\text{on}}(\nu)$ when pointing at the Moon's centre. The pointing-position for the off-Moon data was 'nothingness', i.e. a position at similar galactic latitude far from any strong sources, which we checked with the ATCA catalogue. We set the buffer lengths to maximum for this procedure, since then the product of trigger rate and buffer length is

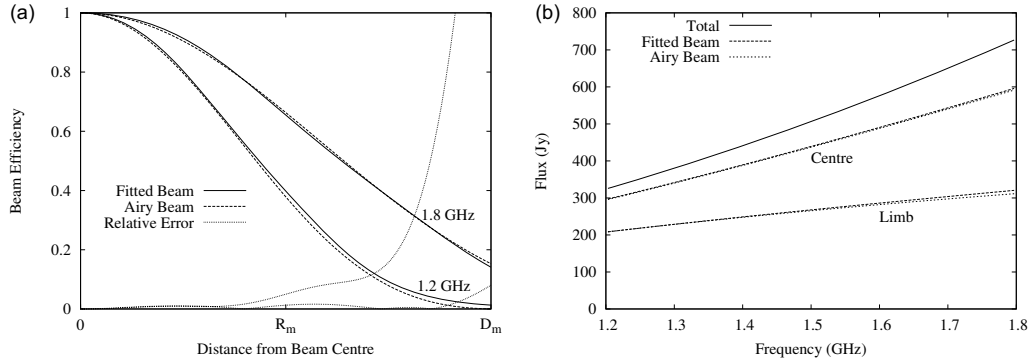


Figure 6.20: (left) Beam power patterns at 1.2 and 1.8 GHz using both an Airy beam with a uniform 22 m circular aperture, and the fitted beams of Eq. (6.42) and Table 6.15. (right) Total and received lunar thermal flux per polarisation channel as a function of frequency for both beam-shapes (almost indistinguishable) and also both centre-pointing and limb-pointing modes.

largest, and also we obtain the best spectral resolution. This was done once every time the configuration was changed.

For the 2007 observations, we also used a ‘nothingness’ pointing far from the galactic plane, so that the effect of ignoring that part of the background blocked by the Moon could be estimated. The difference between the two ‘nothingness’ pointings turned out to be, appropriately, nothing, so I ignore this effect from here on.

Each recorded buffer $b(t)$ was discrete-Fourier-transformed (DFT’d) to produce $b(\nu)$. Some care must be taken with the normalisation of the DFT - for a buffer length of n_s samples, the DFT must be normalised by $1/\sqrt{n_s}$, so that Eq. (6.40) holds for each individual sample. The resulting spectra are squared and the average over all the buffers (typically 5000) recorded for each calibration period/target taken, giving $|b_{\text{rms}}|^2(\nu)$. Each averaged spectrum is then cleaned with a very simple cleaning algorithm to remove the worst of the RFI, which simply sets the power of all RFI spikes above a running threshold to zero, and the subsequent analysis ignores them. An example of the raw and cleaned spectra is given in Fig. 6.21.

The squaring, summing, and cleaning process was repeated for both the nothingness and on-Moon (centre and limb) spectra for each antenna/polarisation, and the nothingness power-spectrum is subtracted from the corresponding on-moon spectra. Taking the square root gave the required $|b_{\text{rms}}|(\nu)$ corresponding to the lunar contribution as required for Eq. (6.40). This was then divided by the product $\sqrt{Z_0 F_M}$ to give $|k(\nu)|$.

Fig. 6.22 compares the resulting function $k(\nu)$ between different antennas/polarisations for the centre-nothingness calibration of May 18th. Note that the large values ($O \sim 10^7$ s.u./(V/m)) are not unexpected, since the induced voltage of radio signals is very small. While the shapes of the spectra are approximately the same, especially over the nominal bandwidth, the most obvious effect is the different amplification (height) of the signals,




Figure 6.21: Raw (top) and cleaned (bottom) spectra for the ‘nothingness’ pointing of May 18th on CA05.




Figure 6.22: Comparison of $k(\nu)$ between antenna and polarisation channel for May 18th. The vertical dashed lines indicate the nominal 1.2–1.8 GHz bandwidth.




Figure 6.23: All measurements of $k(\nu)$ for the CA01 A channel ($k_{1A}(\nu)$) over the 2009 observation periods. The vertical dashed lines indicate the nominal 1.2 – 1.8 GHz bandwidth.

which vary over both polarisation and antenna. This is neither surprising, since the amplification was set ‘by hand’ for each data channel, nor of concern, since trigger levels were individually set to match the data rate, and hence real threshold.

Fig. 6.23 compares the calibrations for CA01 (pol A) obtained at different periods/pointing positions. The effect of re-scaling the trigger levels between February and May is immediately obvious, while the similarity of the plots in May essentially states that the procedure and the antenna pointing solutions are correct. Note that there is no apparent reduction in the visibility of the Moon for the May 19th calibration, as might be expected if CA02 obstructed CA01 (see Sec. 6.2.3). The reduced RFI at the bottom of the band in the May data is due to the addition of a filter.

A frequency offset from the CA04 spectra from May 2007

The raw spectra (arb. units) obtained while pointing at the Moon in May 2007 are plotted in Fig. 6.24. Even a quick glance shows that the spectra for CA04 are anomalous – the noise spike at approximately 1.1 GHz in CA03 and CA05 appears just above 1 GHz (shifted down approximately 60 MHz), while all other features (i.e. the bandpass) appear shifted up by a similar amount. As explained by Paul Roberts, the likely cause was that the synthesiser reference (32 MHz signal) in the digitiser for CA04 failed, causing the synthesiser VCO (voltage-controlled oscillator) to drift to one end of its frequency range. Because the sampled band drifted lower, the sampled frequencies appear higher. The reason that the noise spike has shifted lower is that this is probably aliased during down-conversion from its true band at ~ 900 MHz, which is in the mobile phone frequency range. Performing a very basic by-eye fit indicates that the true frequency range was approximately 960-1920 MHz, which




Figure 6.24: Power spectra (arbitrary units) for all data channels for May 2007, obtained while pointing at the Moon – see Sec. 6.6.2 for the method.

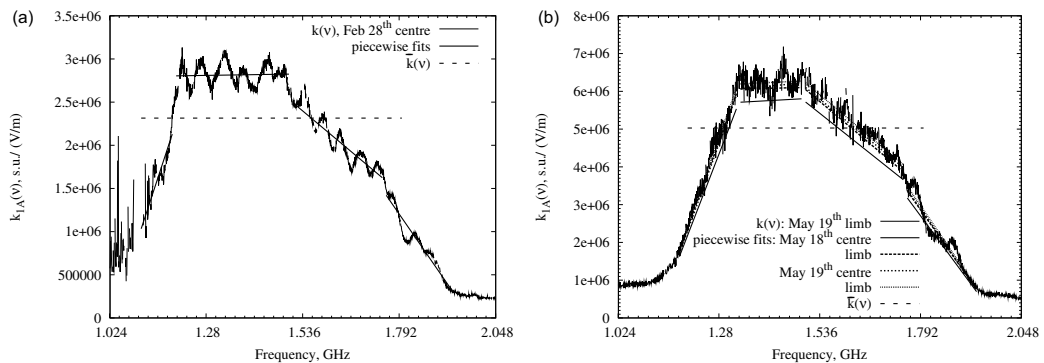


Figure 6.25: $k(\nu)$: (a) as measured for February CA01 A centre-nothingness, piecewise fits to $k(\nu)$, and the mean $\bar{k}(\nu)$; (b) as measured for May 19th CA01 A limb-nothingness, piecewise fits for $k(\nu)$ to all CA01 A calibration data in May, and the mean $\bar{k}(\nu)$ averaged over all data.

still includes the entire nominal bandwidth. No sensitivity will have been lost, although matching the CA04 buffers with the baseband data is thus more difficult.

6.6.3 Results of the calibration

Fitting for $k(\nu)$

In order to characterise $k(\nu)$ in a meaningful way, a piece-wise linear fit to $k(\nu) = a + b\nu$ was performed. Fig. 6.25 shows the fits for CA01 A in February (a) and May (b) 2008 –

four fits have been used, with different frequency ranges for each of the February and May periods. The frequencies were chosen simply to give ‘good’ fits, and the process was not rigorous. Also, a mean $\bar{k}(\nu)$ has been fitted to the bandwidth 1.1-1.8 GHz for February, and 1.2-1.8 GHz for May. While some sensitivity is not included by limiting the range of the fitted bandwidth (in particular the sensitivity ‘bump’ at 1.8-1.9 GHz), including this range in the fit would artificially reduce $\bar{k}(\nu)$ at lower frequencies where a signal is more likely to be observed. Conversely, taking the fit below the low-frequency cut-off gives undue sensitivity where a signal is more likely. Unsurprisingly, the ranges which gave a good trade-off between these effects and artificially reducing the effective bandwidth $\Delta\nu$ were close to/exactly the nominal bandwidth of 1.2-1.8 GHz.

Only one calibration was performed in February, and no cross-checking was possible. In May, it was found that while fits from both May 19th measures and the limb-nothingness fit of May 18th were always in good agreement ($\pm 3\%$), the centre-nothingness fits of May 18th were consistently low, as evident in Fig. 6.25(b). Therefore these were excluded, and the fits averaged over the remaining data. The mean fitted values, and frequency ranges, for all periods are given in Table C.7 in the Appendix.

Conversion to meaningful units

Using the values of $\bar{k}(\nu)$ given in Table C.7, it is possible to calculate the thresholds in meaningful units – here V/m/MHz – given the thresholds in s.u. These were constantly altered throughout the experiment, and are summarised in Table C.3. The thresholds vary significantly between data channels and over time, since the thresholds were chosen to keep the trigger rate on each channel (rather than the thresholds themselves) constant.

Since a positive detection requires a three-fold trigger, it is useful to define an effective threshold over all antennas. A good measure is to choose a mean V/m/MHz signal strength that over the bandwidth will trigger all three antennas with probability 50%. The random noise component can either increase or decrease the measured signal, so we require sufficient intrinsic signal strength above threshold that the chance of the random component pushing the signal below threshold is small. Thus the effective threshold is dependent upon this random component. These were calculated from the mean recorded spectra from the relevant limb/centre pointing calibrations by averaging the RMS over the entire 1.024 GHz bandwidth, and converting the measured RMS in s.u. into V/m/MHz using \bar{k} . During times of significant out-of-band RFI, the effective thresholds will vary due to a greater RMS, but since these occasions are both rare and have a low effective efficiency, their contribution to the average threshold will be negligible and is neglected here. The resulting RMS voltages are reported in Table 6.16. Note that these quantities have no real meaning (RMS voltage obviously does not add linearly over a bandwidth), but are very convenient for comparative purposes.

Assuming a normally-distributed RMS field strength, the probability of the total of signal plus noise falling above threshold for any given signal strength can be readily calcu-

Period	CA01(03)		CA03(04)		CA05	
	A	B	A	B	A	B
May07	2.11 10 ⁻⁹	2.99 10 ⁻⁹	2.17 10 ⁻⁹	3.06 10 ⁻⁹	2.12 10 ⁻⁹	2.88 10 ⁻⁹
Feb08	2.12 10 ⁻⁹	2.15 10 ⁻⁹	2.15 10 ⁻⁹	2.08 10 ⁻⁹	2.15 10 ⁻⁹	2.20 10 ⁻⁹
May08	2.12 10 ⁻⁹	2.15 10 ⁻⁹	2.15 10 ⁻⁹	2.20 10 ⁻⁹	2.12 10 ⁻⁹	2.14 10 ⁻⁹

Table 6.16: Effective RMS electric field strengths per unit bandwidth E_{rms} ($= \mathcal{E}_{\text{rms}}/\Delta\nu$; V/m/MHz) for each polarisation channel. E_{rms} is the field strength equivalent at the antennas to the sampled RMS. The bracketed antenna IDs apply to the May07 observations.

lated, and thus the probability that the global condition (CA01A OR CA01B) AND (CA03A OR CA03B) AND (CA05A OR CA05B) will be met. The detection probability is dependent on the alignment of the field vector with the A and B receivers – since the trigger condition is A OR B, the probability is highest when the field vector is parallel with either A or B, and lowest when it is 45° from both. Therefore the ‘effective threshold’ is defined for a signal polarised at 22.5° to either A or B. Since in general both the thresholds and RMS values are different for each polarisation, this threshold is the average of that calculated for a signal polarised 22.5° (67.5°) from A (B) and B (A). Using Microsoft Excel’s ‘goalseek’ function to find the value of the field strengths when the global detection probability is 50% produced the values found in the right-most column of Table C.8, given in the Appendix. While the individual thresholds are of order 1.6-1.9 10⁻⁸ V/m/MHz, the effective thresholds are 2.0–2.2 × 10⁻⁸ V/m/MHz, showing the decrease in sensitivity from requiring a three-fold coincidence to confirm detection.

6.7 Results

6.7.1 Search for Lunar Pulses

May 2007 pulse search

The search for candidate events for May 2007 involved using four criteria. Beginning with approximately 100, 150, and 120 thousand candidates on CA03, CA04, and CA05 respectively, the criteria reduced this number as follows:

1. **Coincident times** A script to search through the computer-recorded times for each event was used to generate a list of possible three-fold coincidences. The window used was $\pm 500 \mu\text{s}$, which was the assumed accuracy of the computer clock (approximately coincident noise-diode pulses had times matching within this window). This left 41240 three-way coincidences, with many events belonging to multiple coincidences.

2. **Pulse shape** A simple classification program written by C. Phillips was used to class each event according to the number of strong peaks in the data. Whereas a true event should only have one or two strong peaks, the noise-cal triggers would contain of order half a dozen, and RFI events – usually being narrow-band sources – would tend to contain many more. A by-eye check of this procedure for a few hundred events showed it to work well. In two stages, this program was used to eliminate:
 - a) all coincidences containing one or more events classified as a noise-cal pulse (2054 candidate threefold coincidences remained).
 - b) all events classified as wide rfi (382 remained).

3. **RFI bursts** Many of the resulting candidates came from periods with a maximal trigger rate, probably caused by an increase in the rms due to narrow-band RFI. Eliminating all events with triggers both immediately before and after reduced the number of candidates to 110.

4. **By-eye RFI removal** The pulse-shape classification program was intentionally written to err on the side of accepting too many RFI-like events to avoid throwing out true pulses. A manual analysis of the remaining 110 candidates eliminated most as being un-pulse-like, leaving 10.

5. **Consistent polarisation** A lunar pulse will be strongly linearly polarised, so the relative pulse heights in each of the two polarisation channels on each antenna should be the same for all antenna. Requiring a consistent polarisation left four candidate events.

Without further information, the four events (one of which is shown in Fig. 6.26) must remain unclassified. However, an important question is: how many false events do we expect? Purely random (non-RFI) events will almost certainly pass criteria 2, 3, and 4 above (it would be extremely rare to have multiple 5σ spikes in the one buffer). The rate of CA04 triggers was once per 98 ms, and for CA05, once per 125 ms. The search program first took each CA03 trigger, and looked for a CA04 event within 0.5 ms. If it found one, it then looked for a CA05 event within 0.5 ms of either. Given that the average CA04 trigger will be 0.25 ms from the CA03 time, this gives a search window for the CA05 event of $2 \times 0.5 + 0.25 = 1.25$ ms. There is thus a $1/98 \times 1.25/125 = 1/9800$ chance of getting a false detection for each CA03 candidate. Since there were 99556 such CA03 triggers, we expect approximately 10 false events. With the correct polarisation requirements included (essentially, each buffer could have triggered off either A or B), this becomes 2.5 events – this is not significantly different from 4, so we have no reason to expect these pulses to be of lunar origin.

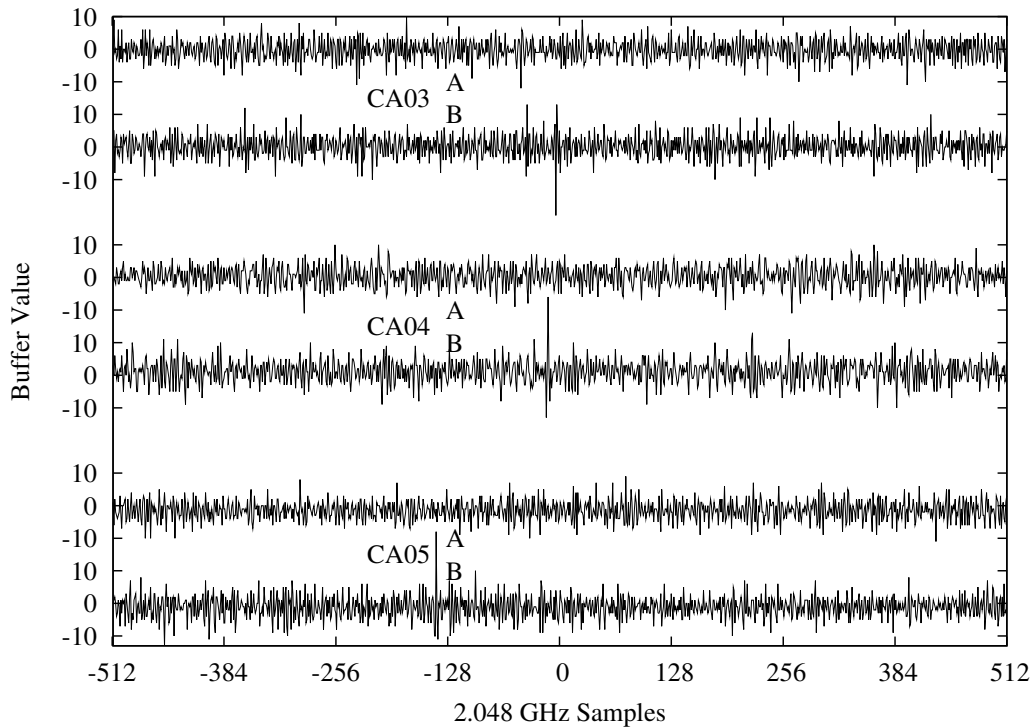


Figure 6.26: One of the four remaining candidate events from May 2007.

2008 pulse search

The main search criteria used for eliminating false events for the 2008 observations were the timing requirements. The search window is given by the apparent angular width of the Moon in the East-West direction, since the North-South component is unresolved by the East-West baseline. This is $\sim 0.5^\circ$ at transit (when the Moon achieves its greatest elevation), and considerably less near Moon-rise/set. This gives an intrinsic time-window of up to 23 ns (46 samples) over the maximum baseline of 765 m. While neither the raw nor the correlation-corrected times can be completely trusted, true lunar pulses will have a sharp time structure which will allow only a small variation in trigger times. Alternatively, those with extended structure (i.e. multi-peaked electromagnetic showers viewed away from the Cherenkov angle – see Section 5.6) will be strong enough to give a correct correlation. Nonetheless, the search uses *both* the corrected and raw times, and allows for a ± 4 -sample error on either. The small number of candidates thus produced are indeed returned by both searches with only a few exceptions, and can readily undergo a manual (‘by-eye’) examination.

Performing the search over both observation periods resulted in approximately 60 candidates. Fig. 6.27 effectively replots Fig. 6.14, but with both the expected lunar offsets (orange line), and the locations of time-consistent events (orange crossed circles), plotted.

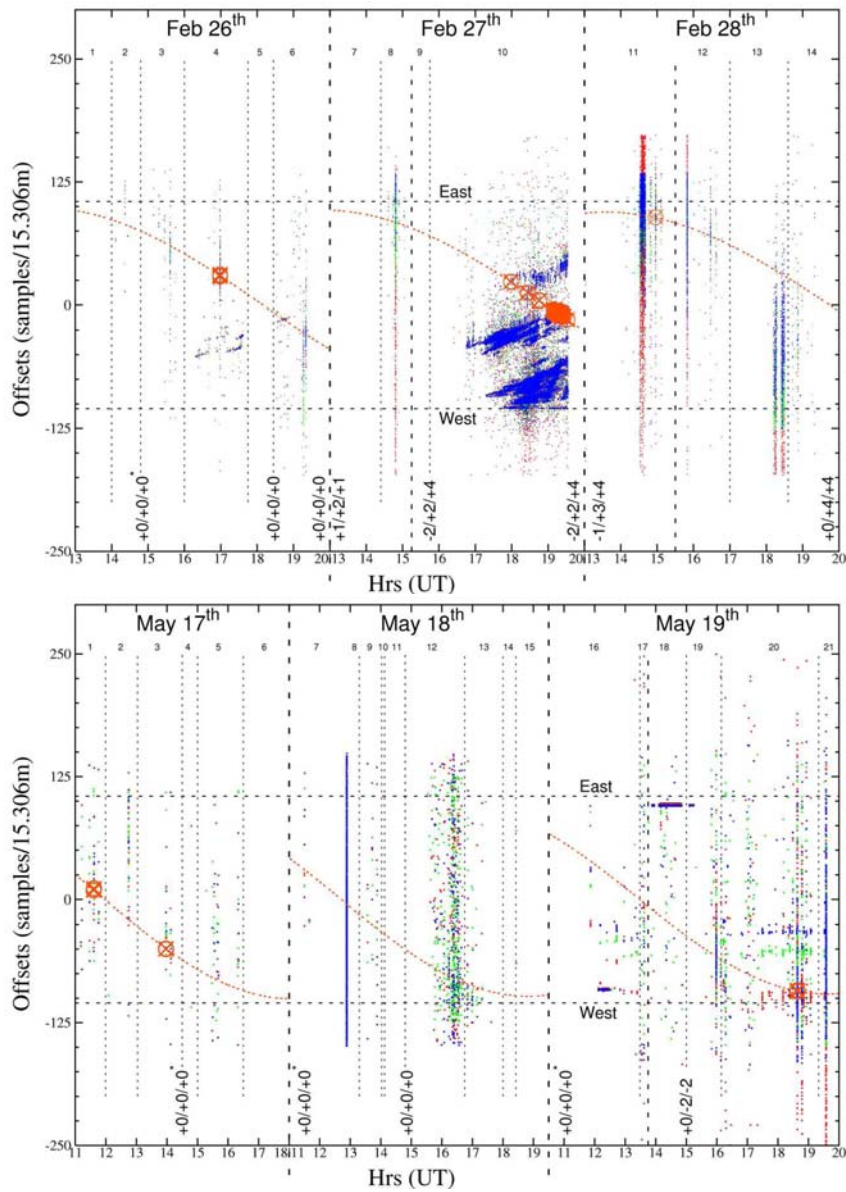


Figure 6.27: Antenna delay per (15.306 m) unit baseline against UT for three-fold coincident events (as per Fig. 6.14), with both the expected lunar delay (orange dotted line) and positions of candidate events (crossed circles) marked, for (top) February, and (bottom) May 2008.




Figure 6.28: Example of a narrow-band-RFI-dominated three-fold trigger satisfying the timing criterion.

All are detected during periods of intense RFI. These events were then searched through by eye for pulse-like events.

Conducting this search, a majority of the candidates were dominated by a narrow-band RFI signature, with the recorded time-domain signals being strong over the entire buffer length. A typical example is given in Fig. 6.28 – obviously, the trigger time here is somewhat arbitrary, since the buffer contains many samples over the threshold of ~ 30 -50. Such events would be excluded at stage 2.2 using the 2007 criteria in Sec. 6.7.1, whereas here our timing information alone is sufficient to reduce the (much greater) number of candidates to a manageable size.

A minority of events – 16 in total – had a narrow time-structure, an example of which is given in Fig. 6.29. All came within a two-hour period on February 27th, which was one of the most RFI-intense periods of all the observations. These could not be immediately excluded by eye, and had to pass more stringent tests. These are described below.

Ensuring possible origin to within sampling accuracy.

The search algorithm allowed both small deviations from a far-field event, and small offsets in origin from the Moon, to account for potential errors in the automated alignment process. A by-eye check of the corrected alignment times (described in Sec. 6.5.4), if necessary including further adjustments, would be expected to yield accurate timing information in cases where the detected event has significant time-structure, as is the case with all 16 candidates. This proved correct, with the correlation-corrected times at most needing further




Figure 6.29: A narrow time structure event from February 27th.

adjustment by one sample on a single baseline, and generally needing none. For each, the resulting alignment was compared visually to that required for the event to have a far-field origin; this was done quantitatively by comparing the CA03 times t_3 to that expected (t'_3) from t_5 and t_1 :

$$t'_3 = ((t_5 - t_1)(b_{13}/b_{15}) + t_1). \quad (6.43)$$

An example of this procedure is given in Fig. 6.30. In no case did any of these events appear to be a far-field event. In most cases, structure was evident in both polarisations, but was so weak in one that the other only could be used for determining the alignment – however, an alignment to within sampling accuracy could always be obtained. Since at these times the Moon made an angle of nearly 90° to the ATCA baseline, events within 360 km would result in a wavefront curvature with measureable differences between t_3 and t'_3 – in the case shown in Fig. 6.30, the distance is of order 100 km. Likely this contributed somewhat to the ability to exclude events using timing only.

Correct dispersion

RFI signals are expected to be of terrestrial origin and arrive undispersed by the Ionosphere, while (needless to say) lunar pulses will not be and suffer the usual dispersive effects. Terrestrial RFI will therefore be dedispersed without suffering any actual dispersion; reversing the dedispersion to produce an ‘undispersed’ signal should give a greater pulse height. In the case of a lunar event however, the correct dispersion measure will differ from that built




Figure 6.30: Comparison of best farfield-fitted alignment (top) with the best unrestricted fit (bottom) for polarisation B on one of the sixteen narrow-time-structure candidates on February 27th, for the values of the offsets (Δt_{31} , Δt_{51} , Δt_{53}) shown. The ‘wavefront curvature’ is given by subtracting the measured t_3 the expected value t'_3 given by Eq. 6.43.

into the dedispersion filter, and recovering the true pulse will involve a small correction. By performing both procedures on the detected pulses, and comparing which one sharpens the signal to which one spreads the signal out, the nature (terrestrial or extraterrestrial) of each candidate should be determinable.

To make this procedure rigorous would require two more bodies of work. Firstly, an analysis of the effects of dispersion (really, undispersion) on the pulse height of a ‘true’ (i.e. simulated) lunar pulse similar to that in Sec. 6.4, but with a random noise component included, in order to determine the possibility of a true lunar pulse having a greater pulse height when undispersed than the intrinsic (correctly dedispersed) pulse. Secondly, the reverse: an analysis of the effects of dedispersion on an intrinsically undispersed RFI signal, as measured in the hundreds of thousands by the antenna. Using these results, it would be possible to build up a statistical criteria to identify definite terrestrial, definite extraterrestrial, and unknown-origin pulses. However, since all candidates are already ruled out from Sec. 6.7.1 above, I did not perform such an analysis, and simply compared the results naively.




Figure 6.31: The ratio of pulse heights for 15 of the 16 short-time-duration candidates assuming a lunar-origin (correctly dedispersed) signal for the numerator, and a terrestrial-origin (undispersed) signal for the denominator.

Fig. 6.31 plots the ratio h_l/h_t , where h_l and h_t are respectively the lunar-origin (correctly dedispersed) and terrestrial-origin (undispersed) maximum pulse heights. Note that in each case the arbitrary phase has been varied to ensure the intrinsic height is captured, as would be the case for an infinite sampling rate. Values above 1 are evidence for an origin above the ionosphere, while those below 1 are evidence for a terrestrial origin; there is thus some evidence to suggest that many of these events had an origin above the Earth's atmosphere.

Each polarisation channel contains both some signal (be it RFI or a lunar Cherenkov pulse) and a random noise component. Whereas the signal will behave in one of two ways depending on its origin (terrestrial or above the ionosphere), the noise component will behave randomly regardless of its origin. Those polarisation channels in which the signal is stronger will give a better indication as to the signal's origin, whereas a noise-dominated channel will have a random response to any dedispersion. Therefore I fitted the mean deviation from 1 in Fig. 6.31 for both the dominant polarisation channels (whichever out of A or B had the greatest measured pulse height), and the secondary (the other) polarisation channels, averaged over all three antennas. That the dominant polarisation channels show a greater mean deviation from 1 than the secondary channels indicates there is indeed some information contained in the measured signal, since the response of the signal to dispersion is stronger than that of noise.

It is also useful to measure the effect on peak signal height of varying the assumed STEC over a large range. This is done for one event in Fig. 6.32, where the peak pulse height (averaged over all three antenna outputs) has been plotted against the assumed STEC value. The dedispersed pulse heights on polarisation A mean little, since the low values indicate




Figure 6.32: Maximum intrinsic height of a pulse as measured (red, green); assuming a terrestrial (dark blue, pink) origin; and assuming different values of the slant total electron content (STEC) for a lunar-origin pulse (light blue, yellow). The peak of the lunar pulse lies near the value of the STEC (9.67 TECU) used for the analogue dedispersion filters, which was also close to the actual value (~ 10 TECU).

the Pol A channels were noise-dominated. However, for sum of the Pol B channels, there is a well-defined peak about the true STEC value of ~ 10 . Although this particular event gives the best-defined peak, other events have similar behaviour, though not always about the true STEC value. The only conclusion that can be drawn here is that *some* information on the signal origin is contained in a dispersion analysis; however, *what* that information is must be left to another project. For now it is enough to use only the results of the previous section to conclude that these events, due to being in the antenna near-field, are not of lunar origin.

6.7.2 Effective Apertures to an Isotropic Flux

The simulation program described in Chapters 2 to 4 was modified to accept a weighted bandwidth according to a general piece-wise linear fit to $k(\nu)$. Since the thresholds are calculated relative to \bar{k} , the weightings assigned to each frequency in the calculation become multiplied by the fitted $k(\nu)/\bar{k}$. Another complicating factor was that the sensitivity of the experiment did not remain constant. Most obviously, modifications to the hardware between each observation period should have improved the sensitivity; also, every time the threshold was changed, the sensitivity would also have shifted. These effects were taken into account by running simulations using both the most- and least-sensitive thresholds (defined by the lowest and highest values of E_{thresh} in Table C.8 respectively) for each of the three observation periods.

The deviation of true TEC compared to that designed into our dedispersion filters changed our sensitivity continuously. The average peak recorded signal strength as a fraction of intrinsic peak strength (Fig. 6.8 and Table 6.6) for each of the three periods was folded into the simulations: for the ‘best’ (most sensitive) case, using the results for incoherent pulses; for the ‘worst’ (least sensitive) case, using that for coherent pulses; and for the average, using the mean effect for both pulse types and averaging the results from both the most- and least-sensitive thresholds.

Finally, the total sensitivity in limb-pointing mode will also depend on the unknown orientation of the polarised receivers with respect to the Moon. While it is possible to extract the alignment of the polarised receivers from the ATCA observing information, the effect of the orientation is expected to be small, and it is easier to simulate the value of the sensitivity averaged over all possible orientations. For the limb-pointing (May 2008) observations, three pointings were simulated, one with each polarisation aligned perpendicular to the targeted part of the lunar limb (so that a given polarisation is at 0° and 90°), and one with them both at 45° to the limb. The mean isotropic aperture for May can be estimated to be half of the 45° result plus a quarter of each aligned result, since over the full 360° , the 45° result is applicable also at 135° , 225° , and 315° , while the 0° result is only also equivalent to 180° , and the 90° result only to 270° . Using this methodology allowed the calculations presented below.

Apertures and Limits to an Isotropic Flux

The resulting range of effective apertures to an isotropic flux from each period is given in Fig. 6.33, assuming the existence of the megaregolith (see Sec. 3.3.2). Also plotted are the effective apertures from prior experiments as calculated in Chapter 3. For our ATCA observations, the threshold is lower than in past experiments (since the high bandwidth has compensated for smaller dishes), while the effective aperture at high energies is greater (due to increased coverage of the lunar limb and lower frequencies). Unlike the previous large-dish experiments, the sensitivity to UHE ν is here greater in centre-pointing mode (May 2007, Feb 2008) than in limb-pointing mode (May 2008), since the beam-width of ATCA near 1–2 GHz is comparable to the apparent diameter of the Moon. Note also that previous calculations (Refs. (121; 25; 104)) have not included the loss from a non-infinite sampling rate.

The corresponding limit from the combined February and May 2008 observations on a UHE ν flux is given in Fig. 6.34(a) – for comparative purposes, the scale is the same as in Fig. 6.34(b), which is a reproduction of Fig. 3.8.

Effect of the orientation of the polarised receivers

The effects of the orientation of the polarised receivers on this experiment are worth investigating for their own sake, since it tells us how important it is to know this orientation in determining the apertures and limits from the experiment. In Fig. 6.35 I plot the three




Figure 6.33: The range of effective apertures (see text) for the three LUNASKA ATCA observations, compared to that from previous experiments (see Chapter 3), assuming the existence of a sub-regolith layer of comparable dielectric properties to the regolith itself.




Figure 6.34: ‘Model-independent’ $(2.3(t_{\text{eff}}A_{\text{eff}})^{-1})$ limit on a flux of UHE neutrinos, from (left) our 2008 ATCA observations, and (right, from Ref. (62); also Chap. 3) from previous experiments. The range on the LHS reflects experimental uncertainties, while on the right the range (where applicable) reflects the inclusion or otherwise of a sub-regolith layer, which is always assumed on the left.




Figure 6.35: Effective apertures as fractions of the average for each orientation of the polarised receivers w.r.s.t. the lunar limb, for the best, worst, and mean sensitivity cases.

apertures (0° , 45° , 90°) normalised by the average ($0.25 \times 0^\circ + 0.5 \times 45^\circ + 0.25 \times 90^\circ$) for each of the best, worst, and mean sensitivity cases. The effect of the orientation is negligible except very near the detection thresholds, so that even for the ‘worst’ case (with the highest threshold), the effective apertures of each orientation are effectively equal above 10^{21} eV.

Near the threshold energies however, the effect of orientation is large, with the (evidently more sensitive) 0° orientation being up to twice as sensitive as the 90° orientation. Since more neutrinos are expected with energies nearer to 10^{20} eV than to 10^{21} eV (downward-sloping spectrum), the effective aperture near the threshold energy will be more relevant for the detection rate, with increasingly steep spectra making the importance of the orientation greater. This is investigated in Fig. 6.36, where I have multiplied the effective apertures for the best- and worst-sensitivity cases by factors of E_ν^{-2} and E_ν^{-3} (and an arbitrary constant C for ease of plotting). Thus the area beneath each graph corresponds to the expected event rate in that energy range for neutrino flux spectra of $dN/dE = E^{-3}$ and $dN/dE = E^{-4}$ respectively. Even for such steep spectra, the increase in the number of events detected by going from the least optimal (90°) orientation to the most (0°) is small. Therefore I conclude that our lack of data on this orientation can in no way have affected our calculations.

6.7.3 Directional-Dependence of the Aperture

Since the observations were targeting specific regions of the sky, the directional dependence of the instantaneous aperture (the ATCA ‘moon-beam’ as it were) is also of interest. Using the same definitions, methodology, etc., as in Chapter 4, the effective area to UHE neutrinos




Figure 6.36: Effective apertures multiplied by factors of $CE_v^{-\alpha}$ for the best- and worst-sensitivity cases and each of the limb orientations. The event rates for $dN/dE = E^{-3}$ and $dN/dE = E^{-4}$ are thus given by the area beneath the plots.

as a function of primary particle arrival direction at given energies is shown for both the February and May periods in Fig. 6.37.

The shape of the instantaneous apertures in Fig. 6.37 are as expected. For the centre-pointing configuration, the ATCA-Moon ‘beam-shape’ forms an annulus, with peak exposure around 15° – 20° degrees from the Moon, while for the limb-pointing configuration, I find the characteristic ‘kidney’ shape, similar to that found for the Parkes experiment (Chap. 4). As expected, the shapes broaden with primary particle energy (the contours in each plot are approximately 10% of the peak value), as the increased strength of the pulses produced allows the instrument to be sensitive to a wider range of interaction geometries.

Taking a one-dimensional ‘slice’ through the figures in Fig. 6.37 allows the effective areas to be plotted and compared as a function of distance along the slice. This is done in Fig. 6.38 for each of the plots in Fig. 6.37 – for the February observations, the slice is through the lunar centre (arbitrarily) along the η -axis, while for May it is through the centre in the $\xi = \eta$ direction, to pass through the peak sensitivity.

Several features are apparent in Fig. 6.38. The effective area of the limb-pointing May configuration in the direction of the beam-pointing position (positive direction) is almost identical to that of the centre-pointing February configuration, though in the opposite (negative) direction, the difference is large, particularly at low energies. The February slices show that the effects of lunar ‘shadowing’ (the decrease in effective area about the direction of the Moon) are small at very high energies, with a decrease in area of only 35% relative to the peak sensitivity at 10^{23} eV; however, the effects becomes drastic at low energies, where there is almost no sensitivity to any events behind the Moon at 10^{21} eV. For both configurations, the peak sensitivity shifts marginally closer to the lunar centre at higher energies,




Figure 6.37: Effective area (km^2) as a function of UHE neutrino arrival direction in the coordinate system defined in Chapter 4. Neutrino energies are (top to bottom) 10^{21} eV, 10^{22} eV, and 10^{23} eV, while the left-hand column is for the centre-pointing February 2008 observations, and the right-hand column for the May 2008 limb-pointing observations. The ‘+’ marks the position of peak effective area; the Moon is at the centre; for the May observations, the pointing position is at the top-right part of the limb, with the polarisation axes aligned such that Pol A $\parallel \hat{\eta}$, B $\parallel \hat{\xi}$.




Figure 6.38: Effective area (km^2) as a function of neutrino arrival direction (degrees from the lunar centre) for both the February and May 2008 observations (see text). The neutrino energies are (top to bottom plots) 10^{23} , 10^{22} , and 10^{21} eV.

though the effect is not great.

6.7.4 Directional-dependence of the limits and exposure to Cen A and Sgr A

Combining the instantaneous sensitivities as shown in Fig. 6.37 with the known beam-pointing positions on the Moon, and the Moon's position itself at the time of observing, allows a calculation of the exposure (area-time product) as a function of celestial coordinates, and also of the exposure to particular objects, such as the nearby active galaxy Centaurus A (Cen A), and our own galactic centre, Sagittarius A (Sgr A). The resulting exposure maps from the 2008 observations only (I do not include the 2007 observations, since these could not rule out false events) to UHE neutrinos of energy 10^{21} , 10^{22} , and 10^{23} eV are shown in Fig. 6.39.

The concentration of exposure about Cen A, and the broad galactic centre region (nominally Sgr A), indicates that the targeting of these regions was successful. Approximately half of the observed > 56 EeV UHE CR arrival directions are also covered by the exposure, and taking these as indicative of the true distribution of UHE CR arrival directions indicates that the LUNASKA observations had a greater sensitivity to these particles than if their flux had been isotropic. Without a more accurate estimate of the sensitivity of the lunar Cherenkov technique to UHE CR however, no meaningful sensitivity or limit can be calculated.

The experiment with the greatest exposure to UHE neutrinos in the 10^{21} to 10^{23} eV range is the Antarctic balloon-born experiment ANITA. The most recent flight of this instrument (ANITA-2, during the 2008/09 Antarctic summer) has only just finished, while an analysis for the first flight (ANITA-1, during the 2006/07 Antarctic summer) is only now

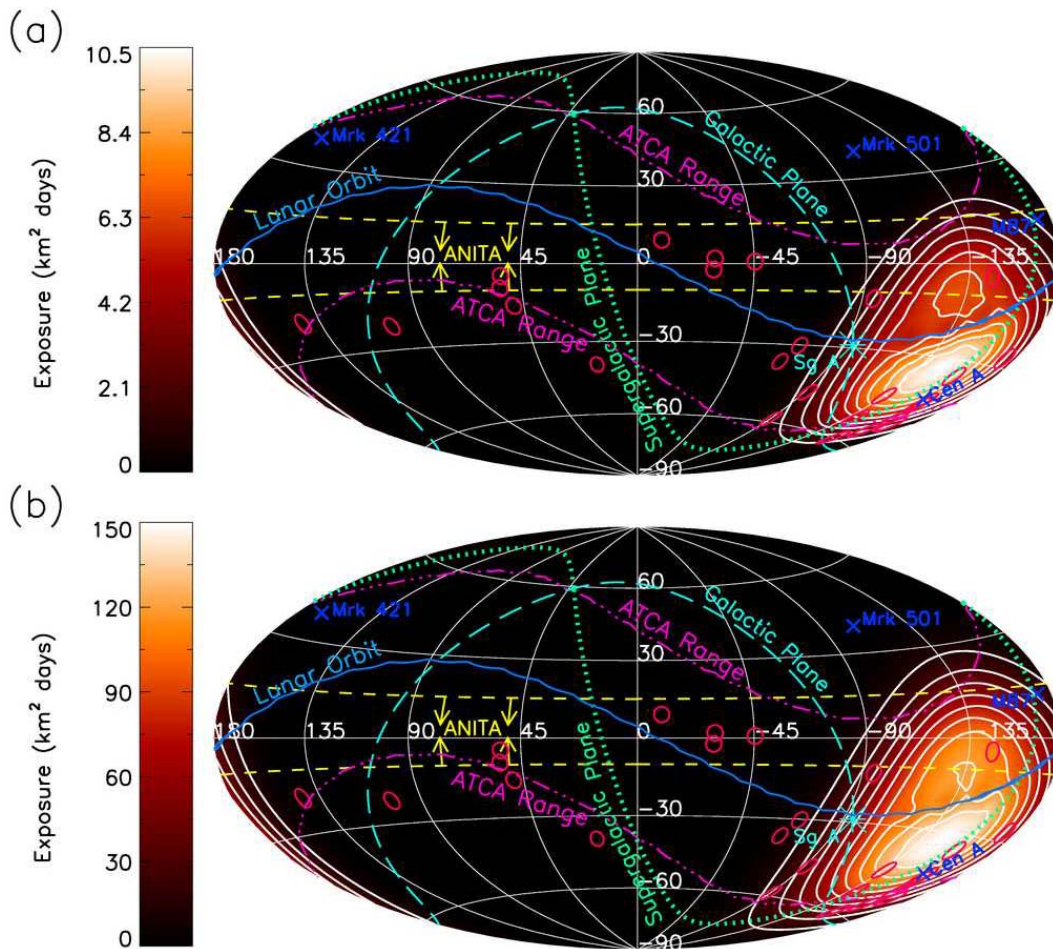


Figure 6.39: (continued next page)

being completed. Results from Gorham *et al.* (50) for the ANITA-1 exposure to 10^{20} eV neutrinos indicate that estimates of the expected exposure range (Ref. (82)) are at least approximately accurate, with the majority confined to the declination range $15^\circ > \delta > -10^\circ$. Even if the true exposure from these observations is half of that predicted by Barwick *et al.* (23), the exposure of the ANITA flights will dominate over other experiments in this declination range (see Sec. 3.8.1).

Other experiments with a significant limit on UHE ν in the $E \gtrsim 10^{21}$ eV range are RICE, GLUE, Kalyazin, and Parkes – see Fig. 6.34. RICE was an in-ice implemented-volume experiment in Antarctica – it thus had a much lower threshold than the other three (lunar Cherenkov) experiments, albeit with only a slowly-increasing exposure with neutrino energy. However, it had a very long observation time of several years as compared to several days for the lunar Cherenkov experiments, which makes up for the lower instantaneous effective aperture.

The directional-dependence of the GLUE and Parkes exposure has been calculated in




Figure 6.39: Total exposure ($\text{km}^2 \text{ days}$) of the 2008 LUNASKA ATCA observations to UHE neutrinos (Top: 10^{21} , Middle: 10^{22} , Bottom: 10^{23} eV) as a function of celestial coordinates. The lunar orbit is from May 2008; the Auger events are those used to deduce a correlation with AGN (1). For a description of other features, see Fig. 4.7.

Chapter 4, though sufficient information on the times of the Kalyazin observations is unavailable to allow an accurate calculation of the exposure of this experiment. RICE was mostly sensitive to down-going neutrino events (those appearing to arrive from above the local horizon), which combined with RICE's position effectively at the South Pole, means the exposure of RICE will be chiefly concentrated in the Southern Hemisphere. Taking the instantaneous effective volume as a function of neutrino arrival direction at 10^{22} eV (Ref. (29)) as indicative of the relative sensitivity at all energies, the exposure in celestial coordinates can be calculated using the known observation time of $7.4 \cdot 10^7$ s, using the UHE neutrino interaction cross-section from Gandhi *et al.* (43) to convert between effective volume and effective area. The combined exposure from Parkes, GLUE, RICE, and LUNASKA, to UHE neutrinos of 10^{22} and 10^{23} eV is shown in Fig. 6.40, while the individual and total exposures of these experiments to UHE ν from Cen A and Sgr A are given in Table 6.17. The exposure to 10^{21} eV neutrinos is not plotted, since at this energy the RICE exposure dominates. The nominal range of the ANITA exposure is also shown.

The contributions from Parkes and GLUE to either the total exposure, or the exposure to Cen A or Sgr A, are negligible except at the very highest energies, of order 10^{23} eV. The Parkes experiment had no exposure to either source, and its total exposure was so low as to not produce a visible shading even at 10^{23} eV in Fig. 6.40(c). For this figure, GLUE contributes at up to the 5% level, which barely shows in an arc from approximately $(\alpha, \delta) = (0^\circ, 15^\circ) - (135^\circ, 35^\circ)$. The peak GLUE sensitivity immediately to the North of Sgr A is barely visible, since it lies next to the peak exposures of both RICE (to more Southern declinations) and ATCA (to lesser values of R.A.). Nonetheless, GLUE had a

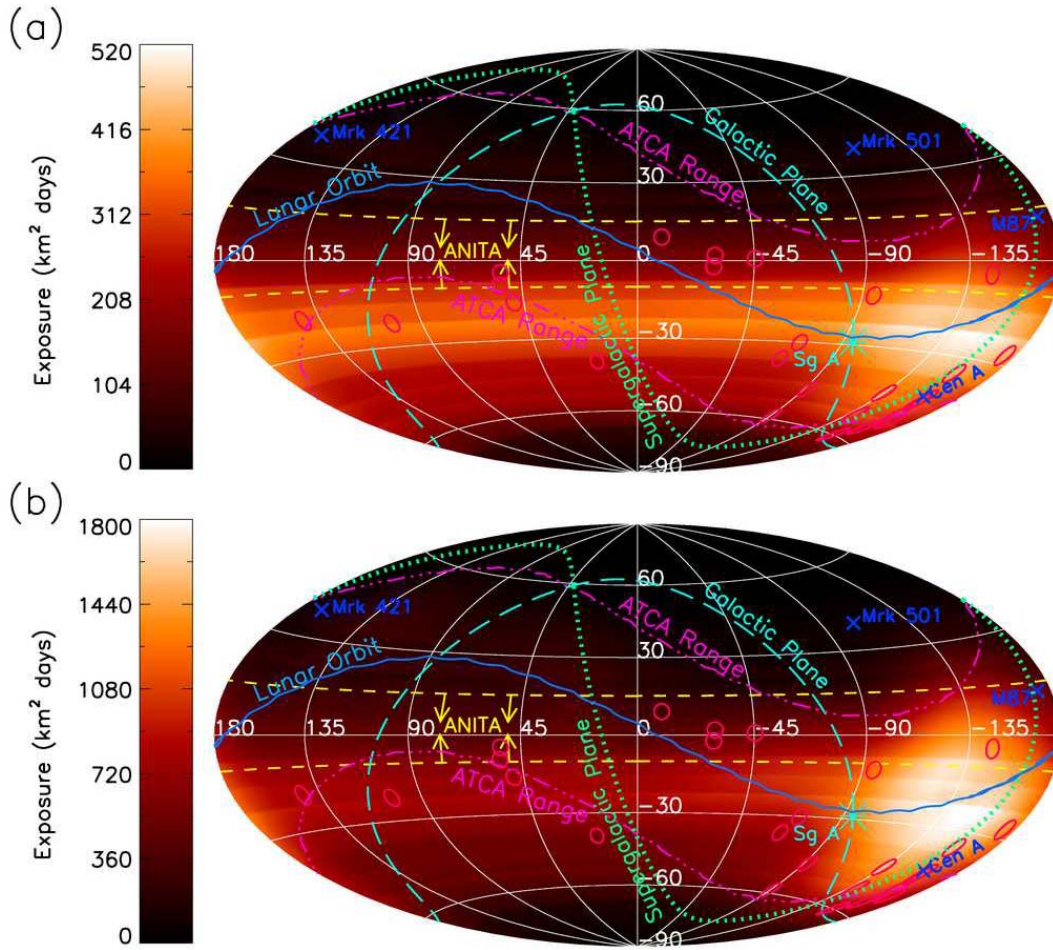


Figure 6.40: The combined exposure of the RICE, GLUE, Parkes, and LUNASKA 2008 UHE ν detection experiments to (a) 10^{22} eV and (b) 10^{23} eV neutrinos. Other features are as per Figs. 6.39 and 4.7.

E_ν	Sgr A				Cen A			
	GLUE	RICE	ATCA	Total	GLUE	RICE	ATCA	Total
10^{21} eV	0.5	195	2.9	198	0.015	242	6.9	249
10^{22} eV	14	333	54	401	2.1	242	111	355
10^{23} eV	175	706	409	1290	43	512	745	1300

Table 6.17: Experimental exposures ($\text{km}^2\text{-days}$) of GLUE, RICE, and the LUNASKA ATCA observations to UHE neutrinos at discrete energies from Cen A and Sgr A.

significant exposure to Sgr A at almost one half of that of RICE, and $\sim 15\%$ of the total averaged over all directions. The LUNASKA ATCA observations exhibit a meaningful contribution at both 10^{23} and 10^{22} eV, particularly in the point-source exposures of Table 6.17, where the ATCA exposure surpasses that of RICE. Importantly, observe that had the ATCA observations been a serious attempt lasting a year rather than the actual six nights, our experiment would produce a sensitivity to both sources comparable to that of RICE at 10^{21} eV, and a vastly superior one at higher energies. Finally, note that the Northern region of ATCA exposure is coincident with that of ANITA, which will dominate in this range.

The low peak sensitivity of the May 2008 Observations

That the peak sensitivity for the (limb-pointing) May observations was found to be lower than that for the (centre-pointing) February observations – except to neutrinos near the detection threshold ($E_\nu \lesssim 10^{21}$ eV) – is completely unexpected. While a lower isotropic aperture is both expected and found (since only a fraction of the Moon is seen at once), the peak sensitivity should be greater, both because the reduced lunar thermal emission is less, and because the beam will have full sensitivity to one part of the limb. Indeed, the limb-pointing mode was chosen specifically to increase the peak sensitivity, and an initial interpretation would be that either this methodology or the simulation results have failed.

There are two more likely causes. Firstly, a filter to reduce sensitivity to low frequencies was implemented between the two observations, which increased the lower edge of the bandpass from 1.1 GHz to 1.2 GHz. The resulting reduction in effective bandwidth decreased the sensitivity. The expectation was that there should have been a corresponding reduction in false trigger rates; combined with the reduced lunar thermal noise received by moving to the limb in May, this should have allowed a lower trigger threshold to be set, compensating for the reduced sensitivity. This does not appear to have been the case, or at least only marginally so – this is the second cause. The trigger rate is a combination of the thermal lunar noise received by the antenna, receiver noise, and RFI rate. If the rate of strong RFI pulses being received is high, then the trigger threshold will reflect this, rather than the other two factors. Hence moving to the limb may not result in a significant reduction in threshold. A quick check of the trigger levels relative to the rms noise levels confirms that the May observations triggered at a higher ratio of $V_{\text{thresh}}/V_{\text{rms}}$ than those in February.

6.8 Summary and Conclusion

This chapter effectively performs the entire post-experimental data analysis of the LUNASKA Collaboration's 2007 and 2008 observations of the Moon at the Australia Telescope Compact Array (ATCA), using the lunar Cherenkov technique to search for the signatures of UHE particles. Though no lunar pulses of any source were positively identified, the instantaneous apertures for all the observations were the most sensitive yet achieved using the technique. Although the corresponding limits on an isotropic flux are not signif-

icant in comparison with experiments such as ANITA, our methods to improve sensitivity to certain patches of the sky were a success, so that the limit from these LUNASKA ATCA observations is dominant about Centaurus A at primary neutrino energies of 10^{22} eV and above. Whether topological defect models – the only models of UHE CR production to predict a flux of neutrinos at these energies – might predict an excess from either the galactic centre region, or a nearby massive galaxy such as Cen A, is unknown.

The non-observation of nanosecond lunar pulses reported here was the expected result. This is because these observations represent the developmental stage of the project, and their success should be evaluated in this context. Most importantly, for the 2008 observations, I can be 100% sure that no true lunar pulses were detected on all three antenna. This demonstrates primarily the power of nanosecond timing over a significant baseline, which was the main criterion used to discriminate against false events. Importantly, it was found in Sec. 6.7.1 that for events of reasonable time-structure (the category into which all candidates must fall), the automatic procedures produced times accurate to better than 1 ns. Therefore, relying on such a procedure to search for candidate pulses in the future is justified, making a comparable analysis for a serious observation run of a month's length or more feasible.

An unresolved flaw in the use of timing to eliminate RFI events was found in the sheer prevalence of coincident RFI triggers (as opposed to those from thermal noise). While many of these came from near-field events and should be eliminated in real-time in future experiments, those coming from the far-field or near far-field will not be. Importantly, some of these – those on February 27th – exhibited both a short time-structure and a dispersion signature apparently consistent with an extra-terrestrial, in-solar-system origin. If this is the case, then a move from a 2D to a 3D array – which will eliminate RFI from the horizon – will not eliminate these events. Further investigation is therefore recommended, both as to the source of these events (many of which were seen), and ways of eliminating them; such an investigation should be feasible with the existing data.

The method of determining the sensitivity worked very well, with results from different calibrations being in good agreement. Whether this method remains appropriate if greater accuracy is required is beside the point, since future observations will automatically include sensitivity measurements using automatic, in-built calibration procedures. It was only for our rather unique set-up, where using the noise diode was inappropriate, that this procedure was needed. It is good to know however that should automatic procedures fail to be built into the system for future radio telescopes (e.g. if all noise diodes produce false pulses) then a workable backup has been developed.

For our observations, the use of an analogue dedispersion filter proved highly successful. The dispersion measure assumed in the filters' construction turned out to be very close to the actual values during observation periods, so that incorrect dedispersion lost less sensitivity than our non-infinite sampling rate. Though such filters must inevitably be superseded by a digital method, their continued use in the meantime will be valuable. Conversely, the finding that the loss from non-infinite sampling was greater than that from incorrect dedispersion is extremely important, and that in fact our 'over'-sampling was an important factor

in increasing (or rather, reducing the loss of) sensitivity. An increase in the sampling rate for future observations would therefore be recommended, and perhaps should take as high a priority as digital dedispersion – at least until the Sun becomes sufficiently active that the analogue filters are no longer appropriate.

Further improvements, such as real-time coincidence logic between three or more antennas, or even the ultimate goal of a coherent addition of the signals, would also improve the sensitivity. Without a further analysis of the typical RFI structure, it is not possible to determine the utility of real-time anti-RFI logic, though given the prevalence of RFI-triggered pulses, this too should be considered.

The lessons learnt above should in all cases be applicable to any use of the lunar Cherenkov technique with an array of radio antennas. The advantage of using a giant radio array such as the Square Kilometre Array to search for lunar pulses has only been highlighted by these observations, especially since it will be placed in a low-RFI environment.



Comparison of Simulation Results using the Same Methods

Four other simulations (to my knowledge) have been developed to model the lunar Cherenkov technique. The first two were associated with the GLUE experiment at Goldstone – one, the ‘UCLA’ monte-carlo by D. Williams, is well-documented in Ref. (121), while the other, known as the JPL/UH monte-carlo, has no documentation that I am aware of, though its results are compared to the UCLA monte-carlo in Ref. (121). At high energies, the results of these two simulations converge, although they produced detection thresholds differing by a factor of approximately two.

The next simulation was developed A. Beresnyak to model the Kalyazin experiment. The initial calculations of the effective experimental aperture were presented in Ref. (25), using detection thresholds similar to that of the GLUE experiment; the final results for the actual experiment obfuscated the effective aperture by publishing a limit assuming a particular neutrino spectrum, so it is the former result only which is of use in comparing simulations. As it turns out however, the initial calculations are *more* useful than those of the second experiment, since the (fictitious) experiment they model is very similar to that of GLUE, allowing the results to be usefully compared. While this was done graphically by Beresnyak in Ref. (25), little was made of the fact that at high energies, the effective apertures for essentially the same experiment differed by an order of magnitude. In the published articles (Chapters 2 and 3 – Refs. (59) and (62) respectively), I therefore discussed this apparent discrepancy between the two simulation results (really between three, since the two Goldstone monte-carlos agree at the highest energies). The original argument for the discrepancy is presented in Sec. 2.4 of the first chapter, while its resolution – in favour of the calculations for Kalyazin of A. Beresnyak – is presented in Sec. 3.6 of the second.

In the aforementioned work, I deliberately avoided comparing the calculations too closely, and never attempt to reproduce the results of Williams (for Goldstone) or those of Beresnyak (for Kalyazin) using the same methods. In fact I had done so extensively before the publication of either of these works, as a natural step for checking the veracity of my calculations. Such a reproduction requires knowing the simulation methods in great detail, whereas typically only some of the equations and methods used are specified, and substantial time was spent reproducing the techniques faithfully. Despite an exact match




Figure A.1: Comparison between published effective apertures and my attempted reproduction using ‘similar methods’, for (a): GLUE, taking results for the limb-pointing configuration from Williams (121), and (b) Kalyazin, using the results of Beresnyak (25). Also plotted are the estimates generated by using the standard simulation methods of Chap. 3 (‘my calculations’).

between simulation results being very unlikely, the discrepancy between the Goldstone and Kalyazin results is a factor of approximately 30, while the uncertainties in methods tend to be of order 2 at most. The original attempts at reproducing the results of Williams and Beresnyak are given in Fig. A.1.

From Fig. A.1, it is immediately apparent that my results are in closer agreement with those of Beresnyak than of Williams, although the greater uncertainty in reproducing the Kalyazin methods makes it hard to predict just how closely the simulations agree. Certainly, my simulations differ from the existing results less than the existing results differ from each other. Importantly, whereas the results of Beresnyak appear consistent with the documented methods used, those of Williams do not. Therefore, even taking the position that all methods and assumed lunar parameters are equal leads to the conclusion that the effective apertures (and hence limits) of the Goldstone experiment have been overestimated by up to an order of magnitude at the highest energies.

A further test was enabled by the publication of estimated apertures for LOFAR, the low-frequency array (104). These simulations were somewhat simpler, and the use of low frequencies made them less subject to particulars of the geometry (as did the exclusion of surface roughness). A comparison between my simulation results and those of Scholten *et al.* is given in Fig. A.2, for three of the five frequency ranges modelled. The results for 1 and 2.2 GHz are effectively identical, while for the 100 MHz results, I estimate a lower aperture by a factor of order 2 at 10^{23} eV and above, but also a lower detection threshold. Importantly, the discrepancy between the effective aperture at the highest energies for LOFAR is similar to that between LOFAR and the SKA, as noted in Chapter 3. Note that the maximum achievable aperture is $2\pi^2 R_m^2 \approx 6 \cdot 10^7$ km² sr, which requires every primary particle to interact in the detectable regolith volume. The estimates of Scholten *et al.* reach half this at 10^{24} eV. A quick check of the relative interaction length of UHE neutrinos at




Figure A.2: Comparison between published effective apertures for LOFAR by Scholten *et al.* (104) (‘Scholten’; crosses) and my attempted reproduction using similar methods (‘James’; open squares).

10^{24} eV (~ 6 km) and the depth of the regolith (500 m) used in the LOFAR simulations readily shows that only half or more of all neutrinos will be detected for impact angles at less than 10° incidence to the regolith (so that $500 \text{ m} / \sin 10^\circ \approx 6 \text{ km}/2$). The corresponding range of impact parameters covers only 3% of all incoming neutrinos. Obviously, there is some inconsistency between the published methods and results of Scholten *et al.*, though I know not what.

The comparisons above cannot verify explicitly the accuracy of this simulation, since after all in no case were the methods identical. However, it does demonstrate that at least the results produced by this simulation are ‘reasonable’. All I can do is to make my methods transparent in the publications in order to facilitate the ease of reproduction of my results.

Further Computation of Secondary Muon and Tau Interactions

In Chapter 3, I present a case for ignoring secondary μ and τ interactions based on a limited set of simulations. A significant limiting factor in the calculations of the contribution of these secondary particles to the detection rate for the experiments concerned was available computer time. I subsequently ran more detailed simulations over the course of some months in order to produce an updated version.

The goal was to get more accurate results at more primary energies. However, I also tested the effects of my choice of $y_{\min} = 0.001$ as the energy fraction separating discrete interactions from ‘continuous’ energy losses. Larger values of y_{\min} reduce the computer time, but run the risk of ignoring potentially detectable interactions. One might expect that for primary neutrinos of energy E_ν that a value of $y_{\min} = E_{\text{thresh}}/E_\nu$ would be sensible. However, the effective energy detection threshold E_{thresh} gained from simulations is generally higher than the ‘true’ threshold calculated by artificially setting all variables to their most optimal values ($\theta = \theta_C$ etc.), since the probability of such interactions is negligible. The sheer number of secondary low-energy interactions however makes it much more likely to get the optimal geometry corresponding to the theoretical limit, so that the effects of a high y_{\min} become significant at an energy lower than anticipated.

Countering this is the fact that at sufficiently low values of y_{\min} , detectable discrete interactions are so common that reducing y_{\min} any further, regardless of the probability of detection, will simply result in multiple pulses being detectable, rather than increasing the probability of a detection itself, which is the sole quantity of interest here.

To test the effect of changing y_{\min} , I ran simulations for the SKA dishes in the energy-range 10^{22} eV to 10^{23} eV, and varied y_{\min} over the values $= 10^{-3,-4,-5}$. The results are plotted in Fig. B.1.

Fig. B.1(a) gives the total effective aperture of the SKA dishes under the three approximations. Note first the large error ($O \sim 6\%$) in the quick calculation at 10^{22} eV presented in Chapter 3 (also Ref. (62)), which used a much shorter calculation (the estimate at 10^{23} eV has, for these purposes, negligible deviation) While the new estimates at each energy appear to be accurate to 1% only, there appears to be a systematic trend for the total effective aperture under smaller values of y_{\min} to be lower – which is physically impossible, since

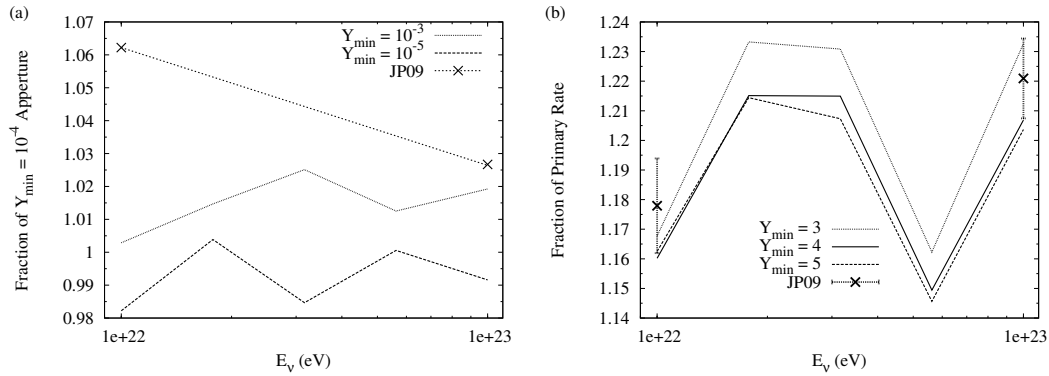


Figure B.1: (a) Total aperture of the SKA dishes as a function of neutrino energy E_ν for $y_{\min} = 10^{-3}$, 10^{-5} , normalised by the apertures at $y_{\min} = 10^{-4}$. (b) Total rate of secondary interactions for the three values of y_{\min} . Also included are the results from Chap. 3 (Ref. (62)).

more interactions cannot mean a lower detection rate. The actual reason for this is the surface correlation routine, which only generates a new surface if the distance to the previous interaction is greater than 10 m. Having a low value of y_{\min} means that there will be a much greater chance of an interaction occurring in this distance, resulting in less surface regenerations, which is the dominant effect (see Sec. 3.4). Importantly though, the difference between using $y_{\min} = 10^{-3}$ and $y_{\min} = 10^{-5}$ is only 2% on average in this energy range.

A comparison of the estimates of the relative contributions of secondary interactions derived from Fig. B.1(b) shows that my choice of $y_{\min} = 10^{-3}$ resulted in a larger estimate of the importance of such interactions – by up to 2% at the highest energies – than would have been found from using smaller values of y_{\min} . The relative rates from Ref. (62) are within the predicted error range.

In Fig. B.2 are plotted the fractional increases in effective aperture due to the inclusions of secondary interactions, using a value of $y_{\min} = 10^{-4}$, for (a): the SKA dishes, and (b): the low-frequency aperture array. Essentially, the information contained is the same as Fig. 3.3 but in more detail; the lack of error bars for the new estimates indicates that the size of the errors is ‘small’, i.e. less than 1%.

Fig. B.2(a) reproduces the trend of growing importance of secondary interactions with primary energy at high frequencies, and also their lessened importance at low frequencies. The earlier estimates appear to be accurate to within the error ranges given for the SKA dishes, but as shown in Fig. B.2(b), are significantly out for the aperture array. Partly, this could be explained by the use of $y_{\min} = 10^{-4}$ for the updated calculations, compared to $y_{\min} = 10^{-3}$ for the original calculations. At low frequencies, the importance of ‘regenerating’ the lunar surface is less important, so the gain at high energies from including more interactions (due to a smaller y_{\min}) will be greater than the loss due to less surface regenerations. Though this explains the greater contribution from secondary interactions in the new results in the $E_\nu > 10^{21}$ eV range, it does not explain the reduced significance found in the

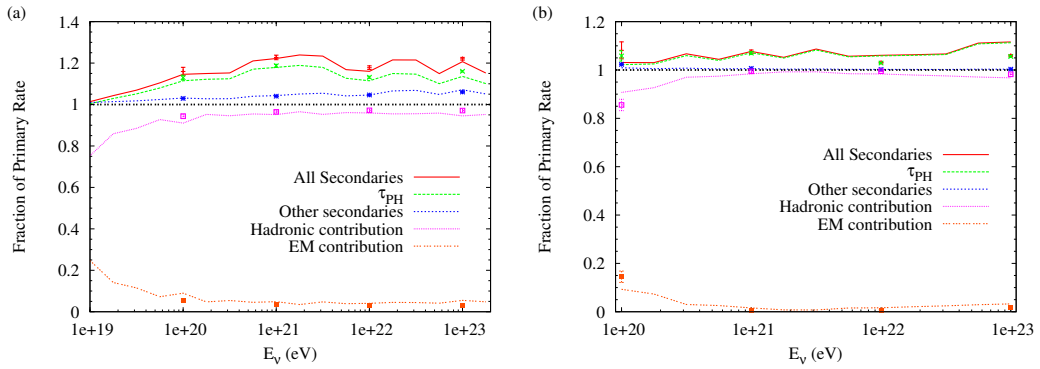


Figure B.2: Fractional contributions to the effective aperture of the SKA dishes (including the megaregolith) normalised by the detection rate for neutrino interactions only (thick black dotted line at $y = 1$), as a function of primary neutrino energy. Points from Fig. 3.3 are also included.

$E_\nu < 10^{21}$ eV range. There appears to be no obvious explanation for this – in absence of an explanation, I can only comment that this latter result should be more trustworthy, since it is based upon more statistics.

The extra calculations interpolating between energies mean that the variation from the trend with energy is now partially resolved, so that the fractions change gradually between data points, rather than apparently randomly as in the original calculations. Much of this ‘wobble’ appears to come from the contribution of τ photo-hadronic interactions, since this is the dominant contribution at all frequencies. Likely it represents changing relative cross-sections interacting in a complex manner.

The key result of the original analysis was that while the contribution of secondary interactions would be significant for the SKA dishes (of order 20%) in the high-energy range, it would be small in the energy range at which any given detector would dominate the sensitivity. The relevant ranges are high energies for low-frequency experiments such as the SKA aperture arrays, and low energies for high-frequency experiments such as the SKA dishes. Thus the greatest effect of secondary interactions on the estimates is of the order of 10%, being that found for the aperture array in the high-energy range. These conclusions are definitely supported in the results presented here, so I conclude therefore that the exclusion of secondary interactions in the majority of the analysis presented in this thesis remains justified.

Data Tables from Chapter 6

In this appendix I include some of the larger data tables from Chapter 6. All of these were essential in the analysis and calculations, but not in the understanding thereof. Their content is explained below.

Calibration times (Table C.1)

This table lists the times spent in ‘calibration’ mode, pointing at either M87 or 3C273 in order to obtain a common signal which would produce a significant correlation in the data between antennas. The labels for each period are used throughout Chapter 6.

Trigger thresholds (Tables C.2 and C.3)

In Table C.2 are plotted the raw (hexadecimal) inputs into the computer controlling the trigger levels, recorded here for sake of convenience. The trigger is defined as a threshold on the square of the buffer values, with the threshold set via a four-digit hexadecimal input. The most significant two digits (having units of 256) for the A (B) polarisation are given by input 51 (53), while the least-significant digits are given by input 50 (52). Thus for CA01 polarisation A, period 1, May 2007, the threshold is $1 \times 256 + 8 \times 16 = 384 \text{ s.u.}^2$. Since the sampled values are integers ranging from 0 to 128, the 256^2 threshold settings map to only 129 effective thresholds. Therefore, in Table C.3, I have converted these threshold into (decimal) thresholds on the buffer values, in sampler units (s.u.).

Noise levels (Tables C.4 and C.5)

The measured RMS noise levels over the entire bandwidth are given in sampling units (s.u.) in Table C.4. Note that the calibration function $k(\nu)$ is different each observation period and data channel. The bandpass for each observation period was not limited to the nominal bandwidth, and some fraction of the noise came from ‘useless’ parts of the spectrum, i.e. those which did not contribute to the sensitivity to a coherent pulse. The ‘necessary’ RMS noise levels (that coming from the range of the nominal bandwidth) are given in Fig. C.5.

Date	ID	Start Time (UT)	Source	On-time (mins)
February 2008				
Feb 26 th	F ₁ ^C	14:42	3C273	11
	F ₂ ^C	18:15	M87	12
	F ₃ ^C	19:50	M87	6
	F ₄ ^C	19:56	3C273	4
Feb 27 th	F ₅ ^C	13:05	3C273	6
	F ₆ ^C	15:37	M87	10
	F ₇ ^C	19:38	3C273	6
Feb 28 th	F ₈ ^C	12:45	3C273	10
	F ₉ ^C	13:05	3C273	28
	F ₁₀ ^C	16:55	3C273	7
	F ₁₁ ^C	20:05	3C273	4
May 2008				
May 17 th †	M ₁ ^C	10:31	3C273	11
	M ₂ ^C	14:11	3C273	5
	M ₃ ^C	14:17	3C273	25
May 18 th	M ₄ ^C	11:10	3C273	9
	M ₅ ^C	14:35	3C273	13
May 19 th	M ₆ ^C	11:32	3C273	4
	M ₇ ^C	14:55	M87	5
	M ₈ ^C	15:00	3C273	4

Table C.1: Calibration times for the 2008 observations. † Network errors – see Sec. 6.5.3.

Comparing the values shows that having an imperfect filter has resulted in a $\sim 15\%$ increase in the RMS noise levels, and hence approximately a 15% loss of sensitivity.

Calibration data (Table C.6)

Table C.6 gives the results of the timing correlation analysis for the 2008 February and May observation periods, which aim to derive the absolute time offsets Δt_{ij}^0 between antennas. The Δt_{ij}^b ($= t_j - t_i$, where $j > i$) give the timing offsets between the centres of buffers i and j . These offsets are very large numbers since the clocks run at 2.048 GHz, so the base offsets are given in units of thousands at the beginning of each period. Thus for the calibration pointing F₉^C, $\Delta t_{15}^b = -161065121490 \times 10^3 - 002 = -161065121490002$. The Δt_{ij}^p give the positions of the peaks in the calibration data relative to the buffer centre. The absolute timing offsets Δt_{ij}^0 are given by $\Delta t_{ij}^0 = \Delta t_{ij}^b - \Delta t_{ij}^c$. † M₃^C was the only calibration period producing a detectable correlation on May 17th, and this could only be performed using buffers aligned two-at-a-time.

May 2008													
Day	Time	CA03 A		CA03 B		CA04 A		CA04 B		CA05 A		CA05 B	
		50	51	52	53	50	51	52	53	50	51	52	53
6 th	17:19	80	01	80	01	c0	01	c0	01	96	01	96	01
	17:30	80	01	80	01	c0	01	c0	01	8e	01	8e	01
	17:36	6e	01	6e	01	b4	01	b4	01	8e	01	8e	01
7 th	14:37	a0	01	a0	01	c0	01	c0	01	a0	01	a0	01
February 2008													
		CA01 A		CA01 B		CA03 A		CA03 B		CA05 A		CA05 B	
26 th	12:46	80	03	00	07	00	06	00	05	84	03	00	04
	13:57	80	03	00	07	00	06	00	05	60	03	00	04
	15:46	09	03	c0	05	fb	05	f0	04	20	03	c0	03
27 th	13:23	80	03	00	07	fb	05	b0	05	49	03	e9	03
	14:13	20	03	ff	05	03	a0	06	50	05	49	03	e9
	15:23	20	03	ff	fb	05	50	05	05	49	03	e9	03
28 th	13:46	80	03	00	07	a0	06	ff	05	84	03	43	04
	15:37	40	03	70	06	20	06	99	05	70	03	25	04
	17:11	40	03	30	11	20	06	52	05	70	03	25	04
May 2008													
17 th	11:00	aa	0a	99	05	88	0c	99	0f	aa	07	77	07
	12:00	88	0a	99	05	33	0c	99	0f	77	07	33	07
	13:06	aa	09	33	05	55	0c	99	0f	30	07	aa	06
	15:17	ff	0c	cc	06	ff	10	ff	0f	05	0a	33	09
	16:26	05	0c	01	06	cc	0f	ff	0f	33	08	77	07
18 th	11:30	44	0a	99	05	ff	0a	77	0f	10	07	88	06
	13:22	11	0a	44	05	dd	0a	77	0f	dd	06	88	06
	14:03	dd	09	11	05	77	0a	77	0f	88	06	33	06
	14:08	aa	09	ff	04	33	0a	77	0f	88	06	33	06
	14:55	55	0a	88	05	ff	0a	77	0f	ee	06	88	06
	16:48	dd	0a	88	05	dd	0b	77	0f	66	07	99	06
	18:09	dd	0b	dd	06	dd	0d	aa	0f	88	08	aa	07
18:26	88	0c	dd	06	dd	0d	aa	0f	88	08	aa	07	
19 th	11:41	33	0a	88	05	cc	0a	77	0f	dd	06	88	06
	17:50	00	09	00	05	88	0a	33	0f	01	06	dd	05
	18:06	ff	0a	ff	05	88	0c	aa	0f	aa	07	66	07
	19:01	88	0b	00	0f	ff	0c	aa	0f	ff	07	66	07

Table C.2: The raw hexadecimal input into the trigger levels.

May 2007							
Day	Time	CA03		CA04		CA05	
		A	B	A	B	A	B
6 th	17:19	19	19	21	21	20	20
	17:30	19	19	21	21	19	19
	17:36	19	19	20	20	19	19
7 th	14:37	20	20	21	21	20	20
February 2008							
Day	Time	CA01		CA03		CA05	
		A	B	A	B	A	B
26 th	12:46	29	42	39	35	30	32
	13:57	29	42	39	35	29	32
	15:46	27	38	39	35	28	30
27 th	13:23	29	42	39	38	29	31
	14:13	28	39	41	36	29	31
	15:23	28	39	39	36	29	31
28 th	13:46	29	42	41	39	30	33
	15:37	28	40	39	37	29	32
	17:11	28	66	39	36	29	32
May 2008							
Day	Time	CA01		CA03		CA05	
		A	B	A	B	A	B
17 th	11:00	52	37	56	63	44	43
	12:00	51	37	55	63	43	42
	13:06	49	36	56	63	42	41
	15:17	57	41	65	63	50	48
	16:26	55	39	63	63	45	43
18 th	11:30	51	37	53	62	42	40
	13:22	50	36	52	62	41	40
	14:03	50	36	51	62	40	39
	14:08	49	35	51	62	40	39
	14:55	51	37	53	62	42	40
	16:48	52	37	55	62	43	41
	18:09	55	41	59	63	46	44
18:26	56	41	59	63	46	44	
19 th	11:41	51	37	52	62	41	40
	17:50	48	35	51	62	39	38
	18:06	53	39	56	63	44	43
	19:01	54	61	57	63	45	43

Table C.3: Trigger levels (s.u.) for all three observation runs.

2007	CA03		CA04		CA05	
	A	B	A	B	A	B
May	3.159	4.542	3.130	4.586	3.157	4.467
2008	CA01		CA03		CA05	
	A	B	A	B	A	B
Feb	4.869	6.976	7.042	6.265	4.853	5.421
May	9.031	7.116	9.766	10.016	7.750	7.495

Table C.4: Buffer RMS (s.u.) whilst pointing at the Moon (May 07, Feb 08) and the lunar limb (May08) over the entire bandwidth.

2007	CA03		CA04		CA05	
	A	B	A	B	A	B
May	2.582	4.220	2.649	4.217	2.596	4.219
2008	CA01		CA03		CA05	
	A	B	A	B	A	B
Feb	3.886	5.513	5.631	5.074	3.901	4.277
May	7.859	6.095	8.464	8.776	6.792	6.505

Table C.5: Buffer RMS (s.u.) after band-limiting to the frequency range 1.1-1.8 GHz (May 07, Feb 08) and 1.2-1.8 GHz (May 08).

Piecewise fitted values to the function $k(\nu)$ (Table C.7).

The function $k(\nu)$ is described in Section 6.6, and gives the relationship between the field-strength in frequency-space at the antennas (V/m), and the observed values of the digitised output (in sampling units. s.u.). The function $k(\nu)$ (s.u./(V/m)) was fitted (as described in Sec. 6.6.3) for all observation periods. Table C.7 gives the results of the fits. Piecewise fits of the form $k(\nu) = a + b\nu$ are given for ν in GHz, along with the fitted ranges. Outside these ranges it can be assumed that $k(\nu) \approx 0$. Also given are fits for \bar{k} , over the range 1.1 – 1.8 GHz (May 2007; Feb 2008), and 1.2 – 1.8 GHz (May 2007).

Effective thresholds in each data channel (Table C.8)

Table C.8 plots the thresholds (10^{-8} V/m/MHz) of each data channel for each lunar pointing over all three observation periods. The final column gives the effective thresholds corresponding to a signal which would be detected with 50% probability (see Sec. 6.6.3). For the effective thresholds, polarisation has already been taken into account, while it has not been (nor can it be) accounted for in the threshold values for the individual data channels. The effective thresholds are generated mostly for comparative purposes, e.g. with the GLUE threshold of $1.23 \cdot 10^{-8}$ V/m/MHz (121) – however, this was estimated assuming a system temperature of 50° K (unlikely, even in limb-pointing mode) and assuming a 6σ threshold on the main antenna only (i.e. excluding possible vetoes from the 35 m dish not triggering),

ID	Δt_{13}^b	Δt_{15}^b	$\Delta^b t_{35}$	Δt_{13}^p	Δt_{13}^p	Δt_{13}^p
February 2008						
Base/ 10^3 :	-23551890720			-161065121490		-137513230769
F_1^C	-224	-544	-1320	-2357 ± 1	1112 ± 0.5	3467.5 ± 0.5
	-264	-704	-1440	-2397 ± 0.5	952 ± 1	3347 ± 0.5
F_2^C	-024	-528	-1504	-2155 ± 1	1130 ± 2	3285.5 ± 1
	-224	-672	-1448	-2355 ± 1	988 ± 2	3342 ± 1.5
F_3^C	-136	-554	-1418	-2266 ± 0.5	1109.5 ± 0.5	3375 ± 1
	-096	-624	-1528	-2226 ± 0.5	1039 ± 0.5	3265 ± 0.5
F_4^C	-064	-538	-1474	-2194 ± 1	1125 ± 1	???
	-264	-512	-1248	-2394 ± 1	1151 ± 1	???
F_5^C	-032	-139	-1107	-2355.5 ± 0.5	1136 ± 1	3490.5 ± 0.5
	-040	-296	-1256	-2364 ± 0.5	969 ± 0.5	3341.5 ± 0.5
F_6^C ‡	-504	-147	-643	-2253 ± 1	1124 ± 1	3376 ± 2
	-560	-002	-442	-2306 ± 1	1276 ± 0.5	3582 ± 1
F_7^C	-640	-136	-496	-2386 ± 1	1142 ± 0.5	3528 ± 1
F_8^C	-424	-107	-683	-2363.5 ± 0.5	977 ± 0.5	3339 ± 1
	-400	-088	-688	-2339 ± 0.5	996 ± 0.5	3334 ± 1
F_9^C	-464	+8	-528	-2404 ± 1	1091 ± 0.5	3493 ± 0.5
	-256	+56	-688	-2195 ± 1	1133 ± 1	3333 ± 1
F_{10}^C	-048	+157	-795	-2180 ± 0.5	1045 ± 0.5	3225 ± 1
	-184	+232	-584	-2316 ± 0.5	1120.5 ± 0.5	3436 ± 1
F_{11}^C	-048	+312	-640	-2178 ± 1	1208 ± 5	???
	-296	+149	-555	-2425.5 ± 1	1044.5 ± 5	???
May 2008						
Base/ 10^3 :	8915974519			???	-178041183664	
M_3^C †	384	???	???	3950 ± 1	???	???
	???	???	-360	???	???	3660.5 ± 1
Base/ 10^3 :	8424267630			7051057573		-1373210056
M_4^C	216	653	-563	3899.5 ± 0.5	???	3799 ± 1
	368	472	-896	4051.5 ± 0.5	???	3466 ± 1
M_5^C	160	457	-703	3847 ± 0.5	???	3662 ± 0.5
	112	504	-608	3799 ± 0.5	???	3758 ± 0.5
Base/ 10^3 :	151274996611			-15903724926		-167178721537
M_6^C	840	-131	-971	-1116 ± 1	-2353 ± 1	-1237 ± 1
	992	-000	-992	-964 ± 1	-2222 ± 1	-1258 ± 1
M_7^C	864	-451	-1315	-1091 ± 1	-2285 ± 0	-1195 ± 1
	840	-328	-1168	-1114 ± 1	-2162 ± 1	-1047.5 ± 1
M_8^C	992	-256	-1248	-962.5 ± 1	-2090 ± 1	-1128 ± 1
	992	-432	-1424	-962.5 ± 1	-2266 ± 1	-1303 ± 1

Table C.6: Calibration data for the observation periods in February and May 2008 (see text). ‡ Using 3C273 as the source (M87 was listed in the notes).

Fit Range (GHz):	1.1-1.2 GHz		1.2-1.5 GHz		1.5-1.75 GHz		1.75-1.94 GHz		1.1-1.8 GHz
$k(\nu) = a + b(\nu)$:	a	b	a	b	a	b	a	b	\bar{k}
ca03 A	$-1.09 \cdot 10^7$	$1.02 \cdot 10^7$	$1.70 \cdot 10^6$	$1.41 \cdot 10^5$	$3.59 \cdot 10^6$	$-1.36 \cdot 10^6$	$9.81 \cdot 10^6$	$-5.01 \cdot 10^6$	$1.53 \cdot 10^6$
ca03 B	$-9.48 \cdot 10^6$	$9.09 \cdot 10^6$	$3.13 \cdot 10^6$	$-8.81 \cdot 10^5$	$4.44 \cdot 10^6$	$-1.90 \cdot 10^6$	$5.94 \cdot 10^6$	$-2.93 \cdot 10^6$	$1.53 \cdot 10^6$
ca04 A	$-1.34 \cdot 10^7$	$1.25 \cdot 10^7$	$3.72 \cdot 10^6$	$-1.41 \cdot 10^6$	$3.41 \cdot 10^6$	$-1.30 \cdot 10^6$	$8.41 \cdot 10^6$	$-4.26 \cdot 10^6$	$1.46 \cdot 10^6$
ca04 B	$-1.21 \cdot 10^7$	$1.13 \cdot 10^7$	$1.82 \cdot 10^6$	$1.30 \cdot 10^4$	$3.97 \cdot 10^6$	$-1.58 \cdot 10^6$	$4.21 \cdot 10^6$	$-1.82 \cdot 10^6$	$1.53 \cdot 10^6$
ca05 A	$-1.04 \cdot 10^7$	$9.66 \cdot 10^6$	$9.91 \cdot 10^5$	$6.10 \cdot 10^5$	$4.31 \cdot 10^6$	$-1.73 \cdot 10^6$	$6.02 \cdot 10^6$	$-2.88 \cdot 10^6$	$1.51 \cdot 10^6$
ca05 B	$-1.09 \cdot 10^7$	$1.03 \cdot 10^7$	$1.56 \cdot 10^6$	$2.69 \cdot 10^5$	$6.48 \cdot 10^6$	$-3.04 \cdot 10^6$	$5.44 \cdot 10^6$	$-2.66 \cdot 10^6$	$1.57 \cdot 10^6$
Fit Range (GHz):	1.1-1.2 GHz		1.2-1.5 GHz		1.5-1.75 GHz		1.75-1.94 GHz		1.1-1.8 GHz
ca01 A	$-1.37 \cdot 10^7$	$1.33 \cdot 10^7$	$2.73 \cdot 10^6$	$6.72 \cdot 10^4$	$8.03 \cdot 10^6$	$-3.66 \cdot 10^6$	$1.26 \cdot 10^7$	$6.34 \cdot 10^6$	$2.32 \cdot 10^6$
ca01 B	$-1.54 \cdot 10^7$	$1.53 \cdot 10^7$	$9.72 \cdot 10^5$	$2.18 \cdot 10^6$	$1.26 \cdot 10^7$	$-5.84 \cdot 10^6$	$1.77 \cdot 10^7$	$-8.84 \cdot 10^6$	$3.28 \cdot 10^6$
ca03 A	$-2.08 \cdot 10^7$	$2.02 \cdot 10^7$	$6.25 \cdot 10^6$	$-1.60 \cdot 10^6$	$1.04 \cdot 10^7$	$-4.67 \cdot 10^6$	$1.71 \cdot 10^7$	$-8.70 \cdot 10^6$	$3.31 \cdot 10^6$
ca03 B	$-1.43 \cdot 10^7$	$1.41 \cdot 10^7$	$2.64 \cdot 10^6$	$7.67 \cdot 10^5$	$8.03 \cdot 10^6$	$-3.28 \cdot 10^6$	$1.77 \cdot 10^7$	$-8.99 \cdot 10^6$	$3.04 \cdot 10^6$
ca05 A	$-1.78 \cdot 10^7$	$1.65 \cdot 10^7$	$1.47 \cdot 10^6$	$8.91 \cdot 10^5$	$8.84 \cdot 10^6$	$-4.09 \cdot 10^6$	$9.14 \cdot 10^6$	$-4.47 \cdot 10^6$	$2.28 \cdot 10^6$
ca05 B	$-8.39 \cdot 10^6$	$8.58 \cdot 10^6$	$9.84 \cdot 10^5$	$1.41 \cdot 10^6$	$9.07 \cdot 10^6$	$-4.06 \cdot 10^6$	$1.05 \cdot 10^7$	$-5.09 \cdot 10^6$	$2.49 \cdot 10^6$
Fit Range (GHz):	1.17-1.325 GHz		1.325-1.5 GHz		1.5-1.75 GHz		1.75-1.94 GHz		1.2-1.8 GHz
ca01 A	$-3.25 \cdot 10^7$	$2.90 \cdot 10^7$	$5.35 \cdot 10^6$	$5.73 \cdot 10^5$	$1.92 \cdot 10^7$	$-8.70 \cdot 10^6$	$2.98 \cdot 10^7$	$-1.50 \cdot 10^7$	$5.03 \cdot 10^6$
ca01 B	$-2.74 \cdot 10^7$	$2.45 \cdot 10^7$	$2.86 \cdot 10^6$	$1.38 \cdot 10^6$	$1.48 \cdot 10^7$	$-6.77 \cdot 10^6$	$2.29 \cdot 10^7$	$-1.16 \cdot 10^7$	$3.92 \cdot 10^6$
ca03 A	$-3.95 \cdot 10^7$	$3.54 \cdot 10^7$	$8.94 \cdot 10^6$	$-1.47 \cdot 10^6$	$2.31 \cdot 10^7$	$-1.11 \cdot 10^7$	$1.56 \cdot 10^7$	$-7.64 \cdot 10^6$	$5.35 \cdot 10^6$
ca03 B	$-3.72 \cdot 10^7$	$3.31 \cdot 10^7$	$6.63 \cdot 10^6$	$2.11 \cdot 10^4$	$1.84 \cdot 10^7$	$-8.19 \cdot 10^6$	$3.08 \cdot 10^7$	$-1.56 \cdot 10^7$	$5.37 \cdot 10^6$
ca05 A	$-2.88 \cdot 10^7$	$2.59 \cdot 10^7$	$6.17 \cdot 10^6$	$-8.22 \cdot 10^5$	$1.33 \cdot 10^7$	$-5.50 \cdot 10^6$	$1.94 \cdot 10^7$	$-9.38 \cdot 10^6$	$4.31 \cdot 10^6$
ca05 B	$-2.69 \cdot 10^7$	$2.40 \cdot 10^7$	$2.80 \cdot 10^6$	$1.48 \cdot 10^6$	$1.71 \cdot 10^7$	$-7.95 \cdot 10^6$	$1.73 \cdot 10^7$	$-8.44 \cdot 10^6$	$4.13 \cdot 10^6$

Table C.7: Fitted values of $k(\nu)$, in s.u./(V/m) (see text).

May 2007								
Day		CA03A	CA03B	CA04A	CA04B	CA05A	CA05B	E_{thresh}
6 th	17:19	1.27	1.25	1.45	1.40	1.34	1.29	1.68
	17:30	1.27	1.25	1.45	1.40	1.27	1.22	1.66
	17:36	1.27	1.25	1.38	1.33	1.27	1.22	1.63
7 th	1:33	1.32	1.45	1.40	1.34	1.29	1.700	
February 2008								
Day		CA01A	CA01B	CA03A	CA03B	CA05A	CA05B	E_{thresh}
26 th	12:46	1.27	1.29	1.19	1.16	1.33	1.30	1.56
	13:57	1.27	1.29	1.19	1.16	1.28	1.30	1.55
	15:46	1.18	1.17	1.19	1.16	1.24	1.22	1.48
27 th	13:23	1.27	1.29	1.19	1.26	1.28	1.26	1.55
	14:13	1.22	1.20	1.25	1.20	1.28	1.26	1.53
	15:23	1.22	1.20	1.19	1.20	1.28	1.26	1.52
28 th	13:23	1.27	1.29	1.25	1.30	1.33	1.34	1.59
	15:37	1.22	1.23	1.19	1.23	1.28	1.30	1.54
	17:11	1.22	2.03	1.19	1.20	1.28	1.30	1.86
May 2008								
17 th	11:00	1.22	1.12	1.24	1.38	1.20	1.23	1.54
	12:00	1.19	1.12	1.21	1.38	1.18	1.20	1.52
	13:06	1.15	1.09	1.24	1.38	1.15	1.17	1.51
	15:17	1.34	1.24	1.43	1.38	1.37	1.37	1.66
	17:11	1.29	1.18	1.39	1.38	1.23	1.23	1.60
18 th	11:30	1.19	1.12	1.17	1.36	1.15	1.14	1.49
	13:22	1.17	1.09	1.15	1.36	1.12	1.14	1.48
	14:03	1.17	1.09	1.12	1.36	1.09	1.11	1.47
	14:08	1.15	1.06	1.12	1.36	1.09	1.11	1.46
	14:55	1.19	1.12	1.17	1.36	1.15	1.14	1.49
	16:48	1.22	1.12	1.21	1.36	1.18	1.17	1.52
	18:09	1.29	1.24	1.30	1.38	1.26	1.26	1.59
18:26	1.31	1.24	1.30	1.38	1.26	1.26	1.59	
19 th	11:41	1.19	1.12	1.15	1.36	1.12	1.14	1.49
	17:50	1.12	1.06	1.12	1.36	1.07	1.09	1.45
	18:06	1.24	1.18	1.24	1.38	1.20	1.23	1.55
	19:01	1.26	1.84	1.26	1.38	1.23	1.23	1.77

Table C.8: Effective thresholds (10^{-8} V/m/MHz) for all data channels over the specified frequency ranges. See text.

so probably the thresholds are equivalent.

Bibliography

- [1] The Pierre Auger Collaboration (J. Abraham *et al.*), *Science* **318**, 938 (2007).
- [2] The Pierre Auger Collaboration (J. Abraham *et al.*), arXiv:0712.1909v1 [astro-ph] (2007).
- [3] The Pierre Auger Collaboration (J. Abraham *et al.*), *Phys. Rev. Lett.* **101**, 061101 (2008).
- [4] H. Abramowicz, A. Levy, arXiv:hep-ph/9712415v2 (1997).
- [5] T. Abu-Zayyad *et al.*, *Astropart. Phys.* **23**, 157 (2005).
- [6] H.R. Aggarwal, V.R. Oberbeck, *Lunar and Planetary Science Conference* **10**, 2689 (1979).
- [7] D. Allard, M. Ave, N. Busca, M.A. Malkan, A.V. Olinto, E. Parizot, F.W. Stecker, T. Yamamoto, *JCAP* **9**, 5 (2005).
- [8] R. Aloisio, V. Berezhinsky, M. Kachelrieß, *Nuc. Phys. B Proc. Supp.* **136**, 319 (2004).
- [9] J. Álvarez-Muñiz, *E. Zas, Phys. Lett. B* **434**, 396 (1998).
- [10] J. Álvarez-Muñiz, R. Engel, T.K. Gaisser, J.A. Ortiz, T. Stanev, *Phys. Rev. D* **69**, 103003 (2004).
- [11] J. Álvarez-Muñiz, *E. Zas, Phys. Lett. B* **411**, 218 (1997).
- [12] J. Álvarez-Muñiz, *E. Zas, Proceedings of the XXV International Cosmic Ray Conference, Durban, South Africa*, **7**, 309 (1997).
- [13] J. Álvarez-Muñiz, R.A. Vázquez, *E. Zas, Phys. Rev. D* **62**, 063001 (2000).
- [14] J. Alvarez-Muñiz, *E. Zas*, in: D. Saltzberg, P.W. Gorham (eds), *Radio Detection of High Energy Particles — RADHEP 2000*, AIP Conf. Proc. **579**, 117 (AIP, New York, 2001).
- [15] J. Álvarez-Muñiz *et al.*, *Phys. Rev. D* **68**, 043001 (2003).
- [16] J. Álvarez-Muñiz, E. Marqués, R.A. Vázquez, *E. Zas, Phys. Rev. D* **67**, 101303 (2003).
- [17] J. Álvarez-Muñiz, *Phys. Rev. D* **74**, 023007 (2006).
- [18] J. Álvarez-Muñiz, C.W. James, *E. Zas*, R.J. Protheroe, *Thinned simulations of extremely energetic showers in dense media for radio applications*, in preparation (2008).

- [19] C. Amsler *et al.*, 2008 Review of Particle Physics, Phys. Lett. B **667**, 1 (2008).
- [20] L.A. Anchordoqui, H. Goldberg, D. Hooper, S. Sarkar, A.M. Taylor, Phys. Rev. D **76**, 123008 (2007).
- [21] N. Armesto, C. Merino, G. Parente, E. Zas, Phys. Rev. D **77**, 013001 (2008).
- [22] G.A. Askar'yan, Sov. Phys. JETP, **14**, 441 (1962); **48**, 988 (1965).
- [23] S.W. Barwick *et al.*, Phys. Rev. Lett. **96**, 171101 (2006).
- [24] R. Beck, Astronomische Nachrichten **326**, 608 (2005).
- [25] A.R. Beresnyak, arXiv:astro-ph/0310295v2 (2004).
- [26] A.R. Beresnyak, R.D. Dagkesamanskii, I.M. Zheleznykh, A.V. Kovalenko, V.V. Oreshko, Astronomy Reports **49**, 127 (2005).
- [27] V.S. Berezinskij, G.T. Zatsepin, Phys. Lett. B, **28**, 423 (1969).
- [28] V. Berezinsky, arXiv:astro-ph/0509675 (2005).
- [29] D. Besson, private communication (2008).
- [30] L.B. Bezrukov, E.V. Bugaev, Sov. J. Nucl. Phys. **33**, 635 (1981).
- [31] P. Billoir, O.B. Bigas, The Pierre Auger Collaboration, in D. Barret, F. Casoli, G. Lagache, A. Lecavelier, eds, SF2A-2006: Semaine de l' Astrophysique Francaise, 115 (2006).
- [32] D.J. Bird *et al.*, ApJ **441**, 144 (1995).
- [33] R.V. Buniy, J.P. Ralston, Phys. Rev. D **65**, 016003 (2002).
- [34] L. Cazon, *The Northern Auger Observatory*, AAS, High Energy Astrophysics Division meeting **10**, 18.05 (2008).
- [35] B.M. Connolly, S.Y. BenZvi, C.B. Finley, A.C. O'Neill, S. Westerhoff, Phys. Rev. D **74**, 043001 (2006).
- [36] R.M. Crocker, F. Melia, R.R. Volkas, ApJ Lett. **622**, L37 (2005).
- [37] R.D. Dagkesamanskii, I.M. Zheleznykh, Sov. Phys. JETP Lett. **50**, 233 (1989).
- [38] S.I. Dutta, M.H. Reno, I. Sarcevic, D. Seckel, Phys. Rev. D **63**, 094020 (2001).
- [39] T. Ebisuzaki *et al.*, Nuc. Phys. B Proc. Supp. **175**, 237 (2008).
- [40] R.D. Ekers, C.W. James, R.J. Protheroe, R.A. McFadden, presented at Acoustic and Radio EeV Neutrino detection Activities (ARENA), Rome, Italy (2008).
- [41] R. Engel, D. Seckel, T. Stanev, Phys. Rev. D **64**, 093010 (2001).
- [42] R. Gandhi, C. Quigg, M.H. Reno, I. Sarcevic, Phys. Rev. D, **58**, 093009 (1998).
- [43] R. Gandhi, C. Quigg, M.H. Reno, I. Sarcevic, Phys. Rev. D **58**, 093009 (1998).

- [44] D.S. Gorbunov, G.I. Rubstov, S.V. Troitsky, *Phys. Rev. D* **76**, 043004 (2007).
- [45] H. Falcke, P. Gorham, R.J. Protheroe, *New Astron. Rev.* **48**, 1487 (2004).
- [46] P.W. Gorham, K.M. Liewer, C.J. Naudet., D.P. Saltzberg, D. Williams, in: D. Saltzberg, P.W. Gorham (eds), *Radio Detection of High Energy Particles — RADHEP 2000*, AIP Conf. Proc. **579**, 177 (AIP, New York, 2001).
- [47] P.W. Gorham *et al.*, *Phys. Rev. Lett.* **93**, 041101 (2004).
- [48] P.W. Gorham *et al.*, *Phys. Rev. D* **72**, 023002 (2005).
- [49] P.W. Gorham *et al.*, *Phys. Rev. Lett.* **99**, 171101 (2007).
- [50] P.W. Gorham *et al.*, arXiv:0812.1920v1 [astro-ph] (2008).
- [51] K. Greisen, in: *Prog. of Cosmic Ray Phys.*, ed. J.G. Wilson, Vol. III (North Holland Publ. Co., Amsterdam, 1956), 1.
- [52] K. Greisen, *Phys. Rev. Lett.* **16**, 748 (1966).
- [53] T.H. Hankins, R.D. Ekers, J.D. O’Sullivan, *MNRAS* **283**, 1027 (1996).
- [54] T.H. Hankins, R.D. Ekers, J. . O’Sullivan, in: D. Saltzberg, P.W. Gorham (eds), *Radio Detection of High Energy Particles — RADHEP 2000*, AIP Conf. Proc. **579**, 168 (AIP, New York, 2001).
- [55] W. Heitler, ‘*Quantum Theory of Radiation*’, Oxford University Press (1944).
- [56] A.M. Hillas, *Nucl. Phys. B Proc. Supp.* **52**, 29 (1997).
- [57] A. Ishihara for the IceCube Collaboration, arXiv:0711.0353v1 [astro-ph] (2007).
- [58] C.W. James, *Simulation of the Radio-Detection of UHE Neutrino Interactions in the Lunar Regolith* (honours thesis), University of Adelaide (2006).
- [59] C.W. James, R.M. Crocker, R.D. Ekers, T.H. Hankins, J.D. O’Sullivan, R.J. Protheroe, *MNRAS* **379**, 3 (2007).
- [60] C.W. James, R.D. Ekers, R.A. McFadden, R.J. Protheroe, presented at the 30th ICRC, Merida, Mexico (2007); arXiv:0709.0160v1 [astro-ph].
- [61] C.W. James, *Status Report and Future Prospects on LUNASKA Lunar Observations with ATCA*, presented at Acoustic and Radio EeV Neutrino detection Activities (ARENA), Rome, Italy (2008).
- [62] C.W. James, R.J. Protheroe, *Astropart. Phys.* **30**, 318 (2009).
- [63] S. Jester, H. Falcke, *Science with a lunar low-frequency array: From the dark ages of the Universe to nearby exoplanets* (private communication).
- [64] S. Johnston *et al.*, *Experimental Astronomy* **22** 3, 151 (2008); S. Johnston *et al.*, *PASA* **24**, 174 (2007); www.atnf.csiro.au/projects/askap/

- [65] P. Kalberla, *Search for UHE neutrinos using a refurbished 25-m telescope*, presented at Acoustic and Radio EeV Neutrino detection Activities (ARENA), Rome, Italy (2008).
- [66] K. Kamata, J. Nishimura, Prog. Theor. Phys. (Kyoto) Suppl. **6**, 93 (1958).
- [67] S. Klein *et al.*, Presented at Particles & Fields 92, Batavia, IL, USA; SLAC-PUB 5969 (1992).
- [68] S.R. Klein, 44th Workshop on QCD at Cosmic Energies: The Highest Energy Cosmic Rays and QCD, Erice, Italy, 2004; arXiv:astro-ph/0412546 (2004).
- [69] M. Kobal, Astropart. Phys. **15**, 259 (2001).
- [70] E. Konishi, A. Adachi, N. Takahashi, A. Misaki, J. Phys. G: Nucl. Part. **17**, 719 (1991).
- [71] I. Kravchenko *et al.*, Phys. Rev. Lett. **93**, 041101 (2004).
- [72] I. Kravchenko *et al.*, Phys. Rev. D **73**, 082002 (2006).
- [73] L. Landau, I. Pomeranchuk, Dokl. Akad. Nauk SSSR **92**, 735 (1935); **92**, 535 (1953).
- [74] G. Langston, *GLINT: Green bank Lunar Interferometer for Neutrino Transients*, presented at Acoustic and Radio EeV Neutrino detection Activities (ARENA), Rome, Italy (2008).
- [75] N.G. Lehtinen, P.W. Gorham, Phys. Rev. D **69**, 013008 (2004).
- [76] M.S. Longair, *High Energy Astrophysics*, Vol. I, 2nd Ed., Cambridge University Press, UK (1992).
- [77] P.G. Lucey, G.J. Taylor, E. Malaret, Science **268**, 1150 (1995).
- [78] K. Mannheim, R.J. Protheroe, J.P. Rachen, Phys. Rev. D **63**, 023003 (2001).
- [79] R.A. McFadden, R.D. Ekers, C.W. James, D. Jones, P. Roberts, R.J. Protheroe, presented at the 30th ICRC, Merida, Mexico (2007); arXiv:0801.3304v1 [astro-ph].
- [80] A.B. Migdal, Phys. Rev. **103**, 1811 (1956); Zh. Eksp. Teor. Fiz. **32**, 633 (1957) [Sov. Phys. JETP **5**, 527 (1957)].
- [81] G. Miele, S. Pastor, O. Pisanti, Phys. Lett. B **634**, 137 (2006).
- [82] P. Miočinović *et al.*, arXiv:astro-ph/0503304v1 (2005).
- [83] P. Miočinović, R.C. Field, P.W. Gorham, E. Guillian, R. Milinčić, D. Saltzberg, D. Walz, D. Williams, Phys. Rev. D, **74**, 043002 (2006).
- [84] R. Nahnauer, A.A. Rostovtsev, D. Tosi, Nuc. Instruments and Methods in Phys. Res. Sec. A, **587** 1, 29 (2008).
- [85] G. Olhoeft, D. Strangway, Earth Planet Sci Lett. **24**, 394 (1975).
- [86] T. Ono, H. Oya, Earth, Planets, and Space **52**, 629 (2000).
- [87] A.A. Petrukhin, V.V. Shestakov, Can. J. Phys. **46**, S377 (1968).

- [88] R.J. Protheroe, P.A. Johnson, *Astropart. Phys.* **4**, 253 (1996).
- [89] R.J. Protheroe, T. Stanev, *Phys. Rev. Lett.* **77**, 3708 (1996).
- [90] R.J. Prothro, *Astropart. Phys.* **21**, 415 (2004).
- [91] R.J. Protheroe, R.W. Clay, *PASA* **21**, 1 (2004).
- [92] J.P. Rachen, P.L. Biermann, *Astron. Astroph.* **272**, 161 (1993).
- [93] J.P. Ralston, D.W. McKay, in: *Proc. of High Energy Gamma-Ray Astronomy Conference (Ann Arbor, Michigan, 1990)*, J. Matthews (ed), *AIP Conf. Proc.* **220**, 295.
- [94] K.L. Rasmussen, P.H. Warres, *Nature* **313**, 121 (1985).
- [95] R.A. Robinson, private communication (2009).
- [96] S. Razzaque *et al.* *Phys. Rev. D* **65**, 103002 (2002).
- [97] S. Razzaque *et al.* *Phys. Rev. D* **69**, 047101 (2004).
- [98] M. Risse *et al.*, arXiv:astro-ph/0512434 (2001).
- [99] P. Roberts, *First measured results of a microwave de-dispersion filter for Neutrino detection experiment*, ATNF internal communication (2006).
- [100] D. Saltzberg *et al.*, *Phys. Rev. Lett.* **86**, 2802 (2001).
- [101] D. Saltzberg *et al.*, *International J. Modern Phys. A* **21**, 252 (2006).
- [102] O. Scholten *et al.*, arXiv:astro-ph/0508580v2 (2005).
- [103] O. Scholten, *Ultra-High Energy Cosmic ray and Neutrino Physics using the Moon*, presented at Acoustic and Radio EeV Neutrino detection Activities (ARENA), Rome, Italy (2008).
- [104] O. Scholten *et al.*, *Astropart. Phys.* **26**, 219 (2006).
- [105] D. Seckel, T. Stanev, *Phys. Rev. Lett.* **95**, 041101 (2005).
- [106] M.K. Shepard, R.A. Brackett, R.E. Arvidson, *JGR* **100** (1996) 11709.
- [107] Y.G. Shkuratov, N.V. Bondarenko, *Icarus* **149**, 329 (2001).
- [108] N.M. Short, M.L. Forman, *Modern Geology* **3**, 69 (1972).
- [109] T. Stanev *et al.*, *Phys. Rev. D* **25**, 1291 (1982).
- [110] R.M. Sternheimer *et al.*, *At. Data Nucl. Data Tables* **30**, 261 (1984).
- [111] T. Takahashi, Y. Shibata, K. Ishi, M. Ikezawa, M. Oyamada, Y. Kondo, *Phys. Rev. E* **62**, 8606 (2000).
- [112] Y. Takano, Y. Takizawa, S. Sasaki, *Acta Astronautica* **57**, 112 (2005).
- [113] M. Takeda *et al.*, *Astropart. Phys.* **19**, 447 (2003).

- [114] Committee On The Physics Of The Universe, Board On Physics, Astronomy, Division On Engineering, Physical Sciences, National Research Council Of The National Academies, *Connecting Quarks with the Cosmos: Eleven Science Questions for the New Century*, National Academies Press, Washington D.C. (2003).
- [115] V.S. Troitskij, T.V. Tikhonova, *Izv. Vyssh. Uchebn. Zaved. Radiofiz.* **13**, 1272 (1970).
- [116] R. Ulrich, *Z. Phys.* **194**, 180 (1966).
- [117] R. Ulrich, *Z. Phys.* **199**, 171 (1967).
- [118] M.A. Wieczorek, R.J. Phillips, *JGR* **103**, 1715 (1998).
- [119] M. H. Wieringa & M. J. Kesteven, *Measurements of the ATCA primary beam*, ATNF Technical Memo Series, d96b7e1 (1992).
- [120] B.B. Wilcox, M.S. Robinson, P.C. Thomas, B.R. Hawke, *Meteoritics & Planetary Sci.* **40**, 695 (2005).
- [121] D.R. Williams, *The Askaryan Effect and Detection of Extremely High Energy Neutrinos in the Lunar Regolith and Salt*, Dissertation, University of California (2004).
- [122] T. Yamamoto (for the Auger Collaboration), arXiv:0707.2638v1 [astro-ph] (2007).
- [123] E. Zas, F. Halzen, T. Stanev, *Phys. Rev. D* **45**, 362 (1992).
- [124] G.T. Zatsepin, V.A. Kuzmin, *JETP Lett.* **4**, 78 (1966).
- [125] ATCA Observing Characteristics Calculator, www.atnf.csiro.au/observers/docs/at_sens/
- [126] SKA Project Website: www.skatelescope.org
- [127] SKADS website: www.skads-eu.org
- [128] Crustal Dynamics Data Information System (CDDIS), NASA, 2007; http://cddisa.gsfc.nasa.gov/gnss_datasum.html
- [129] <http://ssd.jpl.nasa.gov/?horizons>
- [130] Australia Telescope National Facility: <http://www.atnf.csiro.au/> and pages contained therein.
- [131] <http://antares.in2p3.fr/>
- [132] <http://nemo.in2p3.fr/>
- [133] <http://icecube.wisc.edu/>
- [134] <http://www.ps.uci.edu/~anita/>