

Detecting High-Energy Cosmic Ray Electrons With CREST (the Cosmic Ray Electron Synchrotron Telescope)


by

Joseph C. Gennaro

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Physics)
in the University of Michigan
2015

Doctoral Committee:

Professor Gregory Tarlé, Chair
Professor Carl W. Akerlof
Professor Lennard A. Fisk
Professor Jeffrey J. McMahon
Research Scientist Michael Schubnell



©Joseph C. Gennaro

2015

This dissertation is dedicated to my mother, father, sister and brother, whose constant love and support make anything possible, and to Jana Wilbricht, for reminding me of the Grace of God.

A C K N O W L E D G M E N T S

I would especially like to thank: my advisor, Professor Greg Tarlé, for his constant confidence, advice and patience, without which this dissertation would not be; Michael Schubnell for his expert guidance and assistance; Jon Ameel, for his many contributions to the designing, building, launching and recovery of CREST, and who facilitated my early forays into the world of firmware programming and hardware testing; CREST P.I. Professor James Musser, who provided strong leadership by word and example throughout our mission; Professor Scott Wakely, who provided invaluable advice and guided my efforts in many ways, especially in the area of software development; Professor Dietrich Müller for his sagely advice throughout our mission; Dr. Nahee Park, whose significant contributions to CREST testing, calibration, simulation and analysis enabled much of my work; Rich Northrop for his exceptionally successful channel and chassis designs; Professor Stéphane Coutu, Dr. Matt Geske and Professor Scott Nutter, whose design, construction and integration of the CREST veto system made our detection technique possible (Professor Nutter also contributed much of the simulation work that guided my data analysis techniques and enabled accurate characterization of CREST's end result); Professor Wolfgang Lorenzon, Professor Carl Akerlof (who performed important analysis of CREST's signal source) and Professor Jeff McMahon each provided me with critical assistance and advice, without which this work would have been severely impoverished; Professor Lennard Fisk; Mark Gebhard, Mike Lang and Alex Shroyer at Indiana University for their many and varied contributions; the University of Michigan Physics Department, especially Christina Zigulis and Christine Bolang for their unwavering confidence and support; the scientists, engineers and staff at the various federal agencies and stations who made CREST possible, especially those at the Glenn Research Center in Plum Brook, OH, at the Columbia Scientific Balloon Facility in Palestine, TX and at McMurdo Station and the Long Duration Balloon facility in Antarctica; and Elaine F. Gennaro and Jana Wilbricht, for their valuable editing contributions, which greatly enhanced the quality of this work.

TABLE OF CONTENTS

Dedication	ii
Acknowledgments	iii
List of Figures	vii
List of Tables	xli
List of Abbreviations	xliv
Abstract	xlvi
Chapter	
1 Production, Propagation and Current Measurements of High Energy Cosmic Ray Electrons	1
1.1 Ultra-High-Energy Cosmic Ray Protons: The GZK Limit	2
1.2 Existing Measurements of the High-Energy Electron Spectrum	6
1.3 High Energy Cosmic Ray Electrons: Production and Propagation	6
1.3.1 Cosmic Ray Electron Energy Loss in the ISM	7
1.3.2 Galactic Sources of High Energy Cosmic Ray Electrons	9
1.4 Cosmic Ray Electrons: From Source to Detection	39
2 Synchrotron and Bremsstrahlung Signal Photon Production	41
2.1 Synchrotron Radiation by the Primary Electron	41
2.1.1 Synchrotron Power and Energy Spectrum	42
2.1.2 Calculating The Illuminating Pathlength	45
2.1.3 The Expected Number of Photons	48
2.1.4 Integration Limits: Extreme Trajectories	49
2.2 Bremsstrahlung Emission by the Primary Electron	51
2.2.1 Complete Screening and Radiation Length	53
2.3 Summary and Next Steps	54
3 The CREST Detector: Design and Implementation	57
3.1 Overview: Meeting Physical and Operational Constraints	57
3.2 Instrument Design	62
3.2.1 Hierarchical Design	63
3.2.2 BaF ₂ Scintillating Crystals	64

3.2.3	Photomultiplier Tubes	66
3.2.4	Veto System	70
3.2.5	Pulsar System	73
3.3	Readout Electronics	76
3.3.1	Hardware Components and Communication Scheme	77
3.3.2	CRESTBus and VOLTBus	79
3.3.3	Channel Endcap Boards	79
3.3.4	Digitize Signaling	80
3.3.5	Front-End Electronics Readout	80
3.3.6	PMT Voltage Control: Sedecim Voltage Interpreter (SVI) Boards	86
3.3.7	Digitize Decision: Development, Testing and Flight Operation	86
3.4	The Science Flight Computer (SFC) and the Flight Control Software	100
3.4.1	Pedestal Data Taking	103
3.5	Magnetic Field Sensor	104
3.6	Summary, and a Brief Look Ahead	104
4	Photon and Charged Particle Energy Loss in CREST	108
4.1	Charged Particle Energy Loss	109
4.1.1	Stopping Power for Protons	109
4.1.2	Fast Electron Energy Loss	117
4.2	Photon Energy Loss	120
4.2.1	Interaction Processes	120
4.2.2	Mass Attenuation Coefficient	120
4.2.3	Estimated Interaction Probabilities	125
4.2.4	Effect of the Lead Shield	126
4.2.5	Energy Deposition and Collected Light for Photon Interactions	130
4.3	Energy Calibration	135
4.3.1	Ground Calibration and Gain Matching	135
4.3.2	In-Flight Energy Calibration and Linearization	143
4.4	Summary	144
5	Antarctic Expedition: Preparation and Flight	146
5.1	Thermal Shield Assembly and Instrument Testing	146
5.1.1	The Thermal Plan	147
5.1.2	Plum Brook Vacuum Test	148
5.1.3	Arrival and Preparation	148
5.1.4	First Test Procedure: Two Critical Failures	152
5.1.5	Diagnosis and Adaptation	153
5.1.6	T/V Test: Take Two	158
5.2	Integration Test at CSBF: Palestine	159
5.2.1	Power System Integration	161
5.2.2	Delivery to Palestine	161
5.2.3	CSBF Integration	162
5.3	Pre-Antarctic Preparation	173
5.3.1	Instrument Operation Training	173

5.3.2	Ground Station Maintenance at the POCC	175
5.4	Antarctic Expedition	176
5.4.1	Living in McMurdo	176
5.4.2	Working at LDB	177
5.4.3	First Launch Attempt: Bitter Cold	178
5.4.4	Second Launch Attempt: Battery and Power Issues	179
5.4.5	Third Launch Attempt: Success	181
5.4.6	The Incident	183
5.4.7	Recovery	188
5.4.8	Flight Profile	191
5.5	Summary, and A Look Back	195
5.5.1	Pictures from McMurdo and LDB	196
6	Data Analysis	253
6.1	Signal Event Morphology	254
6.1.1	Anatomy of a Signal Event	255
6.1.2	Identifying Signal Events: Two Key Features	259
6.1.3	Estimating Electron Momentum Direction	261
6.2	<i>EventDisplay</i> : Events in Three Dimensions	262
6.3	The Raw Data	267
6.3.1	Note on Monte Carlo Simulations	267
6.4	Analysis Software Chain	269
6.4.1	<i>Raw Reader</i> : Freeze Cycles	270
6.4.2	<i>Calibrator</i> : Time and Energy Calibration	271
6.4.3	<i>EventFinder</i> : Splitting Freeze Cycles into Events	272
6.4.4	<i>EventCutter</i> : Data Reduction Using Simple Parameters	274
6.4.5	<i>Lineal</i> : Clustering, Linear Analysis and Crystallography	275
6.5	Summary of Linear Analysis Parameters	290
6.6	<i>Geometizer</i> : Geometric Analysis	292
6.6.1	Procedure for Veto Hit Location Estimation	293
6.6.2	<i>Geometizer</i> : Characterizing Veto Reponse to Signal Events	298
6.7	<i>MomentumTracker</i> : Pointing at the Primary Electron	301
6.7.1	An Additional Signal-Seeking Cut?	302
6.7.2	Estimating \vec{p} : An Example Event	303
6.8	Outline of a Technique for Positron Detection	311
6.9	Analysis Summary	314
7	Conclusions	318
7.1	The Mainstream Analysis: Preliminary Results	318
7.2	Extending the Mainstream Analysis	319
7.3	Ideas for a Future CREST-Like Detector: Defeating the Proton Background	322
7.4	Final Thoughts: Looking Ahead, by Looking Nearby	324
	Bibliography	326

LIST OF FIGURES

1.1	Figure 1 (from [48]) showing the possible acceleration sites for ultra-high energy protons. As Hillas notes, “Objects below the diagonal line cannot accelerate protons up to 10^{20} TeV.”	4
1.2	Figure 2 (from [13]) showing the steep cutoff in the ultra-high energy cosmic ray proton spectrum at the energy predicted by Greisen, Zatsepin and Kuz’min in 1966.	4
1.3	Existing all-electron flux ($\times E^3$) measurements. We are just beginning to probe energies above one TeV. The Pamela spectrum is e^- only. Error bars are systematic+statistical, except for HESS which reported the shown systematic error band (green shaded area). HESS points also suffer from an undrawn 15% energy scale systematic uncertainty. Undrawn error bars indicate errors smaller than the marker size. [2][21][25] [27][32][34]	5
1.4	Figure 1 (from [38]) showing the energy loss coefficient from Equation 1.6. As energy increases, this coefficient drops, but not faster than E^2 rises; thus the overall rate at which the electron loses energy increases with energy.	8
1.5	The average electron diffusion range as a function of electron energy. As energy increases, average range decreases. The vertical envelope at each energy represents the range of choice of diffusion constant D_{xx} as in Equation 1.11. Also shown are the distances of several Supernova Remnants (SNR)s identified by [38] as near and young enough to produce observable electron fluxes under various production assumptions. By confining one’s self to energies above a certain value on the X-axis, one detects flux from no further away than the corresponding value on the Y-axis. For example if we were to observe a large amount of 10 TeV electrons, we could exclude them as being from SN 185, S147, HB 21, and G65.3+5.7, but not from the Cygnus Loop, Geminga, Vela, Monogem or Loop 1 remnants (not on energy loss considerations alone - Loop 1 for example has a maximum theoretical energy of 1.2 Tev. See Table 1.3.2.5 for details.). I generated this plot using Equations 1.10, 1.11, the energy loss coefficient in Figure 1.4 and the distances of supernova remnants in Table 1.3.2.5 (all from [38]).	10

1.6	Higher energy electrons must come from nearer and younger sources. From [38], shows flux contours as a function of distance and age of theoretical source SNR, as compared to the measured source and age of various actual SNR. Each group of three lines represents a flux contour with the given number of electrons (in units of $(E/GeV)^3 J GeV^2 m^{-2} s^{-1} sr^{-1}$). Within each group of three lines, the solid line represents no energy cutoff (that is, the spectrum is a pure power law), the long-dashed line represents $E_c = 20$ TeV, and the short-dashed line represents $E_c = 10$ TeV. As noted in the figure, this assumes a diffusion coefficient on the higher end of the usual range; for lower values of the diffusion coefficient, a given flux contour would move down in distance (but not in age).	20
1.7	From [60], a spatial map of non-thermal x-ray emission between 2 and 4.5 keV from SN1006 produced from a compilation of data collected by the <i>XMM-Newton</i> instrument. The energy of the x-rays combined with estimates of the magnetic field in the acceleration region implies a maximum electron energy of about 25 TeV. Color image retrieved from http://xmm.esac.esa.int/external/xmm_science/gallery/images/sn1006_4images.jpg on 2/22/2015.	22
1.8	Figure 4 from [24], showing reconstructed synchrotron emission spectra (dashed lines at low energy) for three magnetic field strength possibilities given the observed inverse Compton emission spectrum (solid line and data points at high energies) for electrons with energy between 5 and 100 TeV accelerated by a γ -ray source near the Vela pulsar.	23
1.9	From [38], shows potential contributions to the high energy all-electron spectrum from various SNR overlaid on the state of the measured spectrum in 2004.	23
1.10	Recent measurements of the positron fraction $e^+/(e^+ + e^-)$. HEAT e^\pm from [37]. HEAT \bar{p} from [22]. PAMELA from [35]. Fermi from [28]. AMS-02 from [26]. Errors are systematic plus statistical and are drawn within the marker when smaller than the size of the marker.	24
1.11	Figure 9 from [18]. Some typical channel-dependent boost factors B_e amplifying the WIMP thermal relic interaction cross section required to reproduce the positron fraction (as measured by PAMELA) ranged from 10 to 10^6	28
1.12	Figure 1 (bottom panel) from [31]. A high-resolution (sub-200 pc) model of the Milky Way DM halo predicts that with very high probability the Solar System does not reside in a region over-dense enough with WIMPs to produce an enhanced indirect electron or positron signal.	29
1.13	Figure 9 from [30] showing the positron fraction resulting from two separate good-fit regions of WIMP parameter space, a high mass, mixed quark/lepton model (left) which has increasing positron fraction with energy and a low mass four-lepton (though non-muonic) model (right) which displays a sharp cutoff past the WIMP mass m_χ . The dotted green line shows their estimate of the astrophysical background.	31

1.14	Figure 4 from [50] showing the positron flux and fraction resulting from a combination of the discrete pulsar sources Geminga and B0656+14 (Monogem) and a Monte Carlo ensemble of pulsars at distance greater than 500 pc and a creation frequency of four per century. The ensemble has an injection spectral index $k = -1.6$ and injection energy cutoff at $E_{max} = 80$ GeV while the discrete sources have injection spectral index $k = -1.5$ and injection cut-off energies of $E_{max} = 600$ GeV. The dotted lines indicate the contribution from a GALPROP-estimated diffuse Galactic background.	34
1.15	Figure 4 from [53] showing the positron fraction and electron-plus-positron flux resulting from the discrete pulsar sources Geminga and B0656+14 (Monogem). The best-fit injection spectral indices are $\gamma = 1.9(1.95)$ for Geminga (Monogem) and injection energy cutoffs are $E_{max} = 2$ TeV for both pulsars. The dotted line is the GALPROP-estimated diffuse Galactic background.	35
1.16	(a) Figure 1 from [4]. Comparison of recent measurements (from AMS-01 and AMS-02) showing the seeming coincidence that the proton and positron fluxes vary with the same spectral index above approximately 40 GeV. (b) Figure 4 from [4] demonstrates the agreement of the primary electron and secondary electron/positron model with AMS-02 and HESS data over the measured energy range. As the primary electron flux falls off due to lack of local sources, the nearby diffuse secondary electron/positron flux takes over and would be characterized by an isotropic flux with positron fraction asymptotically approaching 60% (the fraction of positrons produced by the maximally polarized decay products of primary cosmic ray protons).	37
2.1	Critical synchrotron photon energy envelope (in MeV) as a function of primary electron energy (in TeV). The upper edge of the envelope corresponds to an average B_{\perp} of 0.9 Gauss, the center to 0.6 Gauss, and the bottom edge to 0.3 Gauss.	43
2.2	Analytical approximation (red dashed line) to the function $S(x)$ (black solid line), necessary for computing the synchrotron spectral distribution [5]. The most probable energy, $E_{\gamma}/E_c = 0.29$ is marked by the vertical dashed blue line.	44
2.3	From [5], courtesy of Professor C. Akerlof. Portions of circular primary electron trajectories are shown (online, left in blue and right in red). The portion of the track during which each electron illuminates the CREST detector is bolded. The left (blue) track has a longer illuminating pathlength and therefore a higher probability of producing a large number of synchrotron photons on the detector plane with width w	46

2.4	Illustration of the relative probability of signal photon detection over the coincidence threshold of 3 as a function of trajectory parameter. The effective width of the detector at a specific primary electron energy and magnetic field strength (in this case, 10 TeV and 0.5 G) is proportional to the integral of this quantity. The dotted lines delineate two classes of signal events. First, events where the electron does not traverse the CREST detector plane (dubbed Termination Condition 1, to the left of the red line), and second, events where it does (dubbed Termination Condition 2, between the two dotted lines). To the right of the blue dotted line, the trajectory produces a synchrotron beam path in the detector shorter than 1m and is regarded as undetectable. At 10 TeV primary electron energy approximately 60% of CREST’s detection power derives from events where the electron traverses the detector. This fraction increases at lower energy and decreases at higher energy.	50
2.5	Simulated Bremsstrahlung photon relative energy spectra as a function of the primary electron energy illustrating three key features. First, the spectrum falls off rapidly at low energies. Second, the spectrum is nearly uniform at higher energies. Third, the relative spectrum is nearly completely independent of primary electron energy. The blue plot shows the spectra of photons emitted by 10 TeV primary electrons traversing Earth’s atmosphere. The red plot shows the same for a electron energies ranging from 500 to 1.25 TeV. Plot generated from simulated CREST signal events using CRESTMag8, developed by Scott Nutter.	55
3.1	Overview schematic of the CREST detector. Visible are 1,024 of the cylindrical crystal/PMT assemblies (in grey), partially obscured by the transparent green veto plastic scintillator paddles. Not visible are those paddles on the bottom and rear two sides of the instrument; all together the paddles provide approximately 100% geometrical coverage of the detector plane defined by the crystals. Courtesy of R. Northrop, CREST lead mechanical engineer (University of Chicago). The area below the bottom veto would later house the CREST electronics and flight computer, the power conversion and storage system, and the NASA-provided flight integration package.	58
3.2	An Opal Kelly XEM-3050 board, integrating a Xilinx Spartan-3 4,000,000 logical gate FPGA with various additional features. CREST utilized four such boards, along with 72 Opal Kelly XEM-3010 boards (lower capacity chips with 1,500,000 gates each). From www.opalkelly.com.	62
3.3	Lead mechanical engineer Rich Northrop and CREST’s aluminum gondola which he designed and built.	63
3.4	Mechanical design of a crystal C-channel (also referred to as a bus) containing one row of eight Sedecim modules. Units are in inches. Courtesy of R. Northrop, CREST lead mechanical engineer, University of Chicago.	64
3.5	Chuck Boyer and Alex Shroyer assembling a crystal C-channel at the Indiana University high bay in Fall 2010. On the blackboard to Alex’s right the hexadecimal numbering scheme for tubes within a Sedecim has been drawn.	65

3.6	One fully assembled crystal C-channel rests on top of CREST’s aluminum chassis. Picture courtesy of R. Northrop. Each silver tube is one crystal/PMT assembly.	65
3.7	Picture of two 1 cm thick and one 2 cm thick BaF ₂ crystals supplied by Proteus. Picture by A. Yagi [66].	66
3.8	A Hamamatsu R7724CW photomultiplier tube, like those reading out CREST’s 1,024 crystal/PMT assemblies (and its top veto paddles).	68
3.9	Schematic of the crystal/PMT assembly. Not shown is the lead shielding subsequently installed to reduce secondary hits migrating from crystal to crystal. Reproduced from [66].	69
3.10	Initial tests of the Hamamatsu PMT’s timing capability when used to detect scintillation photons from BGO and BaF ₂ crystals performed by Atsushi Yagi. Picture reproduced from [66].	70
3.11	Mechanical design of a 4x4 Sedecim module. Units are in inches. Courtesy of R. Northrop, CREST lead mechanical engineer, University of Chicago.	71
3.12	Detailed side elevation of a 4x4 sedecim module. Units are in millimeters. The grey outlines of the crystal/PMT assemblies show their foam-ensconced positions relative to the readout electronics. The red rectangle is an edge on view of a Analog Front End (AFE) board, which houses the discriminator, charge sampling and charge amplification circuitry. Just below that is the blue edge-on view of a Sedecim Timing And Control (STAC) board, which interfaces with upstream electronics, generates a clock capable of time-stamping hits with 1ns precision and digitizes analog voltage levels on the AFE corresponding to charge signals. Just to the right, the shorter blue rectangle is an edge-on view of an SVI board, which sets PMT control voltages and provides the input to the Cockcroft-Walton high voltage generators located in the base of each PMT. Schematic courtesy of J. Ameal, CREST lead electrical design engineer, University of Michigan.	72
3.13	Veto paddle showing paddle base with grooved channels containing fiber light guides. Courtesy of M. Geske, PSU.	74
3.14	Hamamatsu PMT and attached Fiber Optic Light Guide Array (FOLGA), designed to mate with the intra-paddle fibers and provide an efficient pathway for photons trapped in those fibers to the PMT. Each paddle had one of these assemblies on each end. Courtesy of M. Geske, PSU.	74
3.15	Hamamatsu PMTs wrapped in black Tedlar tape reading out bottom veto paddles.	75
3.16	Firmware architecture for encoding the 8ns STAC FPGA clock signal into 8 1ns blocks, thereby enabling signal readout with ~1 ns time resolution. Architecture by J. Musser and C. Smith. Diagram by J. Musser.	82

3.17	Simplified schematic of the readout architecture integrated with the Digitize decision. Each Sedecim module's 16 PMT anode outputs feed into the 16 independent discriminators on the AFE. Small boxes represent firmware modules or physical board components. Large dashed boxes indicate the carrier board in which the boxed modules reside. Heavy lines (red online) indicate signals that only proceed if the DIGOL has initiated an instrument-wide Digitize command (see Section 3.3.7 for details). Latched discriminator signals are timestamped (with a TDC word) by the STAC's CPLD array with 1ns time resolution. These TDC words are sent to a delay line (the TDC Holdpipe) to await their fate. The STAC also reports to the DIGOL that it has received a hit. The DIGOL integrates this SH with all other SH signals from all STACs and over the last 100 ns. If this integrated signal meets the coincidence trigger, the DIGOL issues a Digitize command. If the VA Control module receives this command <i>and</i> a TDC word from the hold pipe, it will then sample the charge accumulated on the VA0 (high gain) and VA1 (low gain) charge collection ASICs on the AFE. It digitizes these analog charge estimates into the ADC0 and ADC1 words, pairs them with the TDC word on the readout buffer. The next time the CROL polls this STAC for data the readout buffer will transmit all the TDC/ADC data words built in this manner via CRESTBus. Veto Sedecims read out in an identical manner, except that the <i>single hit</i> signal is not sent to the DIGOL and therefore veto hits do not contribute to the DIGOL's coincidence trigger.	83
3.18	Diagram by M. Schubnell showing the schematic relationship between hardware components of the CREST detector. Note that the dynode signals from the PMTs were not used as they did not meet the specified noise requirements and that the number of veto scintillator panels is 27 (for 54 veto PMT channels) and not 21 as shown. Also not shown is the Digitize sub-system (at the time the diagram was made the hardware pathway existed but was not being utilized) which represents an additional independent two-way signaling pathway between the Digitize Decision Overlord (DIGOL) (mounted on the XEM Carrier Board) and the 72 Sedecim modules.	85
3.19	A channel Endcap board with metal support plate custom designed by the author. 16-pin CRESTBus and VOLTBus connections to the XEM Carrier are on the far left, along with CRESTBus and VOLTBus chain connectors. Power is input via the left facing connector, converted to the 5V and 3.3V required by the SVI and STAC/AFE boards, respectively, then output along the Sedecim chain. The power converter chips are hidden beneath the board; one such chip is shown in Figure 3.20. The 64-pin facing connector interfaces with the DIGOL FPGA on-board the CROL, and the 8 rightmost connectors send Digitize commands and receive Single/Multi-Hit signals to/from each STAC board. These cables were of matched length to minimize variations in trigger behavior between STACs.	87
3.20	A DC-DC converter chip identical to those used on the channel Endcap boards.	87

3.21	The crimping jig used to attach the drop connectors mid-strand of the CREST-Bus and VOLTBUS cables. All such cables were purchased as a single strand, cut to length, then fitted with two endpoint connectors (one for each Endcap) and 8 drop connectors (one for each Sedecim module in each C-channel) at the Michigan E-Shop.	88
3.22	The UM Test Stand used to test a single C-channel's worth of electronics. Essentially a 1-channel version of the XEM Carrier it houses an Opal Kelly XEM-3050 FPGA (lower left) with USB connection to a controlling PC. In middle front the CRESTBus and VOLTBUS cables for the single attached C-channel depart the test stand.	88
3.23	Lower front: one AFE(red)/STAC (blue, underneath the AFE) mated combination. The 16 thin grey cables provide the PMT anode signal to the 16 discriminator channels on board the AFE. Upper rear: an SVI board with 16 voltage control and readback channels. This is the test mode in which the UM E-Shop confirmed the read-out functionality of each STAC, AFE and SVI board sent to Indiana University for instrument integration. In total UM provided 100 functioning copies of each board.	89
3.24	Another view of the test stand pictured in Figure 3.23 also showing the terminating channel Endcap board (far right) and the connection to the neighboring STAC/AFE/SVI Sedecim unit (far left).	89
3.25	Another view of the test setup in Figures 3.23 and 3.24 showing a series of STAC/AFE/SVI Sedecim modules connected in daisy chain fashion.	90
3.26	Another view of the test setup in Figures 3.23-3.25 showing 7 Sedecim modules chained together. The angular red device in the left foreground is the Xilinx tool used to program every FPGA on board each STAC in the C-channel via the JTAG chaining scheme devised by J. Ameel. This enabled us to rectify STAC firmware issues after instrument integration by interfacing only with the first STAC in each C-channel, saving us from weeks of delay in re-assembling the instrument that would have otherwise been necessary.	90
3.27	CREST's readout electronics mounted in the instrument below the crystal C-channels. Each C-channel's boards are linked in a fashion identical to those in the test stand pictured above. Visible are the green stickers on each board indicating that they had been tested and received Michigan E-Shop Seal of Approval. There were a total of 64 crystal PMT readout assemblies (one for each Sedecim) and an additional 8 dedicated to veto PMT readout. Picture by R. Northrop.	91
3.28	The veto "honeycomb" before mounting to the instrument, upside-down, supporting four of nine bottom veto paddles with FOLGA's mated to their PMTs. The white veto boxes are the veto system's equivalent of the crystal C-channels; each would later hold four STAC/AFE/SVI electronics modules. Picture by R. Northrop.	92

3.29	Oscilloscope-generated timing diagrams of digitize signaling within and between the DIGOL on board the XEM Carrier and a STAC board's XEM-3010 FPGA. Horizontal divisions represent 40ns. Each digital signal is represented by an analog voltage level between 0 and 3.3V. See discussion in Section 3.3.7 for details.	95
3.30	How coincidence is defined in the context of the Digitize system. The Digitize FPGA receives 64 singlehit signals (one from each crystal STAC). It also remembers for 20 5ns clock cycles the Single Hit history for each crystal STAC. To count the number of coincident single hits from unique STACs, it logical OR's together each truth value in each STAC's signal history, then counts the number of 1's in the resulting 64-bit array. If this number meets or exceeds the coincidence threshold (set to 3 for flight) then the Digitize FPGA will instruct every instrument STAC (crystal and veto) to digitize (read the charge resulting from) any hits it receives for a brief period.	99
3.31	Effect of modifying the digitize pre-scale value. On the x-axis is STAC ID number (for these runs crystal C-channels 2 and 8 were not operational). On the y-axis is the number of hits reported by each STAC in a data run taken over several minutes. The legend shows the pre-scale value used for each line; from highest to lowest rate the values were 1, 3, 65 and 257.	99
3.32	Effect of modifying the digitize coincidence threshold. On the x-axis is STAC ID number (for these runs crystal C-channels 2 and 8 were not operational). On the y-axis is the number of hits reported by each STAC in a data run taken over several minutes. The legend shows the coincidence threshold used for each line; from highest to lowest rate the values were 1, 2, 3 and 4.	100
3.33	Histogram of the number of hits in individual Freeze/Accumulate cycles (periods of 131.074 μ s) from a pulser calibration data run with STACs operating in two distinct modes. In the left plot STACs digitize and report every hit to the CROL ('internal commanding'). On the right STACs are set to the mode used in flight where they only digitize and report hits when instructed to by the Digitize FPGA ('external commanding'). The DIGOL was set to a coincidence threshold of 7 unique STACs. Doing so eliminates events with fewer than 7 participating STAC boards.	101
3.34	<i>Left:</i> Comparison of data taken with STACs in internal commanding (black line) and in external commanding with a coincidence threshold of 1 (red dashed line). "Cluster size" is the number of STACs with signals in an individual freeze cycle. <i>Right:</i> Ratio of plots in left pane. An ideal distribution would be identically 1.	101
3.35	Identical plots to those in Figure 3.34 but with a coincidence threshold of three STACs instead of one STAC reporting a single hit. Freeze cycles with one or two STACs are greatly reduced in number. The remaining cases are likely straddling consecutive freeze cycles (so that hits on STACs contributing to the digitize decision were read out as part of the previous or next freeze cycle). The clusters with size one result from the single hit pre-scale set to its maximum value of 32000 (so that one out of every 32000 clock cycles that have one or more STACs active are read out as well).	102

3.36	Identical plots to those in Figures 3.34 and 3.35 but with a coincidence threshold of seven STACs reporting a single hit, the strictest setting possible. Freeze cycles with fewer than 7 STACs participating are nearly completely eliminated. The remaining cases are likely straddling consecutive freeze cycles (so that hits on STACs contributing to the digitize decision were read out as part of the previous or next freeze cycle).	102
3.37	The Science Flight Computer assembled by M. Schubnell. On the left are visible the four 0.5 TB solid state disk drives used to store data during flight. Picture by J. Gennaro, McMurdo, Antarctica.	103
3.38	Control flow of the flight software.	104
3.39	The magnetic field sensor provided by the University of Michigan Atmospheric Oceanic and Space Sciences department. It flew on board CREST and recorded the three local components of magnetic field strength.	105
4.1	Stopping power as a function of energy for protons, muons and electrons in BaF ₂ calculated from Equation 4.1 and $I_{BaF_2} = 375.9$ eV. Note that this model is inapplicable at the low end of the plotted energy range for muons and protons because the energy is below the location of the Bragg peak, below which particle velocities become comparable to that of shell electrons and protons can capture electrons. At the high end of the range electron interaction becomes dominated by radiative interactions (i.e. Bremsstrahlung) and the model is also invalid. Particles in the intermediate range near the minimum of these curves (at around 1.326 MeV / (g/cm ³)) are said to be minimum ionizing particles, or ‘mips’.	112
4.2	Stopping power as a function of energy for protons, muons and electrons in polyvinyltoluene (the polymer base for the EJ-200 veto scintillator) calculated from Equation 4.1 and $I_{EJ-200} = 64.7$ eV. The average energy loss of minimum ionizing particles is approximately 2.022 MeV / (g/cm ³).	113
4.3	Distribution of ADC1 counts for signals in the PMTs at the -X (top) and +X (bottom) ends of top veto paddle 4 for ground events that met a muonic cut set. According to the estimates in Table 4.2 the peaks correspond to approximately 0.85 MeV deposited energy. The veto PMT gains were adjusted only coarsely to ensure the Landau most probable energies were visible in such data runs.	115
4.4	The probability that both ends of a top (<i>a</i>) or non-top (that is, bottom, side or inside slanted) (<i>b</i>) veto paddle detect a signal given that either end of the paddle detects a signal as a function of combined ADC channel for background Monte Carlo (red online) and flight data (black). For most of the interesting ADC range the flight data displays a two-sided inefficiency of approximately 86%.	116

4.5	Mass attenuation coefficients for the three major interaction processes (and their summed effect) for photons in Barium Fluoride (BaF_2) (plot produced from data provided by the XCOM database, courtesy of NIST [33]). Also shown by the vertical line at $E_\gamma = 0.511$ MeV is the electron rest mass. Visible at low energies are the discontinuities in the photoelectric cross section arising from the discrete energies associated with atomic transition edges. These edges are listed in Table 4.4. The highest energy edge (the 56K edge at around 37 keV) is quite close to the potential critical synchrotron energies for electrons at around 1 TeV, as shown in Figure 2.1. Signal synchrotron photons between ~ 100 keV and \sim a few MeV are very likely to interact via Compton scattering.	122
4.6	Cross section against the three major interaction processes for photons in EJ-200 plastic scintillator (plot produced from data provided by the XCOM database, courtesy of NIST [33]). Also shown by the vertical line at $E_\gamma = 0.511$ MeV is the electron rest mass. Note that the absorption edges relevant for this polymer (those of Carbon and Hydrogen) occur at energies which are below the minimum photon energy tabulated here.	123
4.7	Cross section against the three major interaction processes for photons in the atmosphere (plot produced from data provided by the XCOM database, courtesy of NIST [33]). Also shown by the vertical line at $E_\gamma = 0.511$ MeV is the electron rest mass.	124
4.8	Mass attenuation coefficients in Air, BaF_2 , EJ-200 and Lead (data from [33]). Between $E_\gamma \sim 600$ keV and ~ 5 MeV BaF_2 actually has a smaller mass attenuation coefficient than both air and EJ-200. This is due to the Compton mass attenuation length scaling like Z/A ; EJ-200 is composed mainly of Carbon and Hydrogen and therefore has a higher Z/A ratio than heavier elements such as Barium. In terms of interaction probability this effect is more than compensated for by the increased mass thickness of the BaF_2 crystals compared to the veto paddles.	126
4.9	Interaction probabilities for photons in air, BaF_2 crystals and EJ-200 veto paddles computed using Equation 4.11, the total mass attenuation coefficients in Figure 4.8 and the various Geometric Cases listed above.	127
4.10	Cross section against the three major interaction processes for photons in lead (plot produced from data provided by the XCOM database, courtesy of NIST [33]). Also shown by the vertical line at $E_\gamma = 0.511$ MeV is the electron rest mass. Note the absorption edge at just over 100 keV, in the center of the signal synchrotron photon typical energy range.	128
4.11	Energy deposited by simulated Bremsstrahlung photons and their daughter products in the BaF_2 crystals. The dominant feature at the electron rest mass energy is primarily caused by pair conversion in the lead shield surrounding each crystal/PMT assembly. The blue plot shows the energy deposition resulting from Bremsstrahlung photons emitted by 10 TeV primary electrons traversing Earth's atmosphere. The red plot shows the same for a electron energies ranging from 500 GeV to 1.25 TeV. Plot generated from simulated CREST signal events using CRESTMag8, written by Scott Nutter.	129

4.12	Spectrum from crystal/PMT assembly number 56 while exposed to 662 keV photons from a Cesium-137 source during a ground calibration run. After subtracting pedestal noise (the green dotted line) from the raw spectrum (black line) two features are clearly evident: the Compton continuum (from approximately channels 100-350) and the 662 keV photopeak at approximately ADC1 (low gain) channel 500.	132
4.13	The Klein-Nishina differential cross section per unit solid angle for outgoing photons and recoil electrons following a Compton interaction. <i>Top</i> : Cross section per unit solid angle as a function of <i>outgoing photon polar angle</i> (Equation 4.15) normalized to unit probability in the forward direction. The polar angle is the angle off horizontal at which the outgoing photon departs relative to the direction of the incident photon. <i>Bottom</i> : Cross section per unit solid angle as a function of <i>recoil electron polar angle</i> (Equation 4.18). The angle on the x-axis, in degrees, is the angle relative to the incident photon direction with which the recoil electron departs. As E_γ increases, the probability of a forward scattering of the electron increases drastically, making the recoil direction of the Compton electron a good estimator for the momentum direction of the incident photon. Polarization effects are neglected in both plots, resulting in an azimuthally symmetric distribution.	134
4.14	Time progression of the hit rate resulting from a Cesium-137 source transiting across the center (between the 2nd and 3rd rows of assemblies) crystal c-channel 4 during energy calibration on the ground at Indiana University (May 2011). The x-axis is in units of Freeze Cycles (each lasting 131.074 μ s). Each row of the y-axis corresponds to an individual crystal/PMT assembly. Increased box areas correspond to more hits per Freeze Cycle. As the source passes above the crystal the hit rate reaches a maximum. Each run along a c-channel lasted approximately two minutes. For this mode of data taking the Digitize system was set to instruct the instrument to collect charge with a coincidence threshold of 1 STAC (in other words, STACs read out every hit).	136
4.15	Low-gain (ADC1) channel spectra from crystal/PMT assembly number 638 (which happens to be the 15th PMT on C-channel 5, STAC 8) without (<i>left</i>) and with (<i>right</i>) exposure to a Cesium-137 662 keV γ -ray photon source. . . .	137
4.16	Example of the Gaussian fit performed for each 1,024 crystal/PMT assemblies using a Cesium-137 662 keV γ -ray source. The location of the peak was determined as in Figure 4.12. Its width is found here as the σ of the best-fit Gaussian function. Using this procedure, I measured a FWHM peak channel resolution of approximately 20-25% for all 1,024 independent instrument crystal channels. In this case, the FWHM resolution = $2.355 \times \sigma/\bar{X} \times 100 = 23.5\%$	138
4.17	Gaussian mean (filled circle) and sigma values (width of error bars) for each of the 128 tubes in crystal c-channel 1 using the procedure outlined in Section 4.3.1. The four left plots used the old high voltage values, and the four right plots use the new high voltage values.	139

4.18	Gaussian mean (filled circle) and sigma values (width of error bars) for each of the 128 tubes in crystal c-channel 2 using the procedure outlined in Section 4.3.1. The four left plots used the old high voltage values, and the four right plots use the new high voltage values.	140
4.19	Gaussian mean (filled circle) and sigma values (width of error bars) for each of the 128 tubes in crystal c-channel 3 using the procedure outlined in Section 4.3.1. The four left plots used the old high voltage values, and the four right plots use the new high voltage values.	140
4.20	Gaussian mean (filled circle) and sigma values (width of error bars) for each of the 128 tubes in crystal c-channel 4 using the procedure outlined in Section 4.3.1. The four left plots used the old high voltage values, and the four right plots use the new high voltage values.	141
4.21	Gaussian mean (filled circle) and sigma values (width of error bars) for each of the 128 tubes in crystal c-channel 5 using the procedure outlined in Section 4.3.1. The four left plots used the old high voltage values, and the four right plots use the new high voltage values.	141
4.22	Gaussian mean (filled circle) and sigma values (width of error bars) for each of the 128 tubes in crystal c-channel 6 using the procedure outlined in Section 4.3.1. The four left plots used the old high voltage values, and the four right plots use the new high voltage values.	142
4.23	Gaussian mean (filled circle) and sigma values (width of error bars) for each of the 128 tubes in crystal c-channel 7 using the procedure outlined in Section 4.3.1. The four left plots used the old high voltage values, and the four right plots use the new high voltage values.	142
4.24	Gaussian mean (filled circle) and sigma values (width of error bars) for each of the 128 tubes in crystal c-channel 8 using the procedure outlined in Section 4.3.1. The four left plots used the old high voltage values, and the four right plots use the new high voltage values.	143
4.25	Example of parameters used in Nahee Park's gain-stitching procedure. On the left she plots ADC1 (low-gain) channel on the vertical axis against ADC0 (high-gain) channel on the x-axis for crystal/PMT assembly 268. Below the break channel around ADC1 channel 500 ADC1 displays linear behavior with respect to energy. Above the break the relationship between ADC1 channel and energy is fitted by a second-order polynomial. ADC1 saturates at channel 2048, at which point ADC0 is relied upon to provide a coarse energy reading. Plots courtesy of Nahee Park.	144

5.1	The CREST instrument on wheels in the Indiana University high bay. All C-channels have been installed, but the thermal shield has not. The black tedlar-wrapped veto paddles surround the crystal plane with triangular FOLGA's connecting the paddles to their readout PMTs. The edges of the bottom veto paddles peek out from beneath the crystal array. The coiled fiber optic cables are pulser system fan-outs waiting to be connected to the pulser box. In the lower-near corner the power panel accepts the bench power supply. In the lower-far right corner the SFC is visible, and to the left of that (blue board) is the XEM Carrier. Yet to be installed in the lower portion are the batteries, the rest of the power system and the CSBF SIP module (later integrated at Palestine).	149
5.2	One of the 2"-thick foam panels (of the second batch, laminated in McMurdo) covering the dark side of the top portion of the instrument.	150
5.3	A ruined panel not used in the Plum Brook test constructed by laminating aluminized mylar sheeting onto foam backing. When applied prematurely (before allowing the glue to dry) a reaction between the aluminum, glue and/or foam occurs resulting in the formation of an intriguing green pattern. J. Musser dubbed this beautiful mistake "RTFM."	150
5.4	Various tapes used as part of CREST's thermal shield. From left to right: aluminized mylar, perforated silverized teflon, conventional packing tape, Nashua duct tape. None of these tapes' adhesive functioned at all during our first launch attempt (scrubbed to due surface wind conditions), which took place during the coldest weather we experienced in McMurdo.	151
5.5	CREST is delivered via flatbed truck to the Plum Brook testing site on 6/14/2011. The white box is the aluminum shipping container constructed and painted at Indiana University. Visible in the background is the upper section of the vacuum chamber with the mushroom hatch in the open state.	154
5.6	Nahee Park, Stéphane Coutu, Scott Nutter and Matt Geske examine the detector's condition following delivery by flatbed truck to Plum Brook in the B-2 facility's loading dock.	154
5.7	CREST is lowered via crane (a Critical Lift Procedure) into the testing chamber at Plum Brook on 6/15/2011.	155
5.8	A view from the ground level of the test chamber. Michael Lang, IU engineer oversees the installation of skin temperature sensors. The cylindrical wall of the chamber consists of densely packed fine copper tubing through which would later flow liquid N ₂ gas to cool the chamber to a minimum environmental temperature of about -40 °C.	156
5.9	Below the instrument inside the Plum Brook B-2 facility test chamber. The hexagonal rods with orange tape on the end are heat lamps. These lamps would simulate the albedo heat flux resulting from sunlight reflecting upwards from the ocean and ice shelf during the actual flight.	156
5.10	Deep in the bowels of Plum Brook' B-2 facility, CREST members watch over the performance of the instrument. The LN ₂ line which would later rupture during the first cold fill bides its time in the background. From left: Stéphane Coutu, Matt Geske, Michael Schubnell and Nahee Park.	157

5.11	The LN ₂ pipe which ruptured during the first cold fill. Due to Plum Brook's strict safety protocol (and the grace of God) no one was injured.	157
5.12	Schematic showing the locations of external temperature and heat flux sensors during the test initiated on 6/27/2011, courtesy of NASA Plum Brook staff. The sun-side (+X direction) is to the left and the dark-side (-X) is to the right. The green tubes indicate the location and orientation of the heat lamps used to simulate the heat produced by the solar and albedo radiation flux that would be experienced during flight.	158
5.13	Plot showing the pressure over time measured by a vacuum ion gauge (sensor IG0121) during the second Thermal/Vacuum test at Plum Brook. The minimum pressure achieved was approximately 0.035 torr. The time axis shows seconds since 6/27/2011 at 9:00 AM.	159
5.14	Plot showing the measurements over time of two temperature sensors (CRST-103, the sun-side skin temperature and CRST-109, the sun-side environment temperature) (left axis) and of one heat flux sensor (HFS-03, the sun-side heat flux sensor, right axis) during the second test procedure. The time axis shows seconds since 9:00 AM on 6/27/2011. While the skin temperature was sensitive to both the incoming flux and environment temperature, at the simulated flight conditions (-30 °C and the hot flux case detailed in Table 5.1.3) the sun-side skin temperature was stable at approximately 40 °C, well within stable operating range.	160
5.15	Plot showing the measurements over time of three temperature sensors (CRST-103, the sun-side skin temperature, CRST-101, the dark side sensor, and CRST-102, the bottom sensor) during the second test procedure. Temperatures were well within the nominal range for the hot and cold cases in all regions of the detector.	160
5.16	Scott Wakely, Nahee Park and Michael Schubnell install the power panel inside the CREST instrument. Photo courtesy of R. Northrop.	162
5.17	Full-sunlight veto tube hit rates (in hz) during the power system integration test at Indiana University on 7/13/2011. Tubes 17, 20 and 33 displayed uncharacteristically high rates on only one side of a veto paddle, indicative of light leaks in the FOLGA mating the PMT to the paddle. Tubes 53 and 54 comprise two ends of the same (small) paddle, indicating either a light leak on the paddle itself or light propagating out of the folga and back into the paddle, thus lighting up both ends. Plots of the veto tube rates such as this one formed an important part of shift operators' sense of whether the instrument was operating normally or not during the Antarctic flight. Ignoring the abnormal rates of the outlier tubes, the normal rates are largely determined by paddle size (top, bottom and side paddles are larger than the inside slanted tubes), proximity to the crystals (top paddles are closest) and geometry (side paddles have their surface normal facing sideways and thus have less flux from down-going particles).	163
5.18	Schematic overview of CREST's power system. From [64], courtesy of S.P. Wakely (University of Chicago).	163

5.19	Assembly, with handles, of CREST’s two Nickel Metal Hydride battery banks. The grey modules on top are the charge controllers.	164
5.20	From left to right, Jim Musser, Scott Nutter, Alex Shroyer, Scott Wakely and a helpful IU faculty member whose name escapes me help perform initial integration of the power system with the CREST instrument outside the IU high bay prior to mounting the solar panels. Supporting the panels here are four wooden ‘veto boxes’ which were used to transport the veto paddles to Plum Brook, Palestine and McMurdo. They were also used to retrieve the veto paddles during recovery.	164
5.21	(a) Plot of the voltage and current provided by one nickel metal hydride battery bank to the instrument running in normal operation from an unknown charge level to discharge. During flight two such batteries would be active at once. (b) Plot of the power supplied to the instrument over the same time period showing that the total instrument power consumption is between 800 and 850 watts.	165
5.22	The Ryder truck full of tools and supplies Alex Shroyer and I drove from Bloomington, IN to Palestine, TX.	166
5.23	CREST’s shipping container, prior to being unpacked in Hangar 2 at CSBF: Palestine.	166
5.24	CREST in CSBF Hangar 2 prior to SIP integration. The power panel (middle left), power control unit (middle right) and relay box (slightly further right, just to the left of the SFC) are visible. Photo by M. Geske.	167
5.25	One of CREST’s solar panel wings awaits attachment to the instrument in its shipping container. Photo by M. Geske.	167
5.26	One of CREST’s solar panel wings (on the +Y side) mounted on the instrument in Hangar 2 at CSBF: Palestine. The vector normal to the solar panel face points in the +X direction. Photo by M. Geske.	168
5.27	Alex Shroyer examines CSBF’s SIP (Systems Integration Package) prior to assisting CSBF install it onboard CREST. Photo by M. Geske.	168
5.28	CSBF personnel fit CREST with crush pads (stacks of cardboard tubes), the ‘skirt’ of solar panels that will power the SIP and the ballast container (center). The SIP is shown here mounted in the lower right portion of the instrument. The white-bearded fellow at far left facing the camera is CSBF’s Dave Sullivan, who ran LDB during most of our time in McMurdo and was of immense help to CREST in many ways during our time in Palestine and in Antarctica. Photo by M. Geske.	169
5.29	CREST crystal rates during the final hang test at Palestine. All tubes had acceptable rates except for Bus 5, STAC 1, Tube 13. The problem was later diagnosed as a control voltage provision issue on the associated SVI card, which was replaced in McMurdo.	169
5.30	Alex Shroyer (left) and yours truly with CREST during its final hang test aboard ‘Tiny Tim’ at CSBF: Palestine. Although Alex was unable to travel to Antarctica, his efforts were instrumental to CREST’s successful completion of the Plum Brook and Palestine tests and the Antarctic expedition. Photo by M. Geske.	170

5.31 Preparing CREST for a hang test in Hangar 2 at CSBF: Palestine. CSBF had just finished attaching this metal superstructure, to which all antenna equipment would be mounted in McMurdo. Photo by M. Geske. 170

5.32 Schematic of information flow from CSBF ground station in Palestine, TX, to the airborne CREST instrument via satellite array. Courtesy of staff at CSBF: Palestine. 171

5.33 CREST on the flight pad at CSBF following our successful hang test. Clearly visible on the lower lefthand thermal panel is a circular region of the aluminized mylar sheeting which bubbled off of the panel’s surface during the low pressure portion of the Plum Brook thermal/vacuum test. These panels were replaced in McMurdo, along with most of the rest of CREST’s thermal shield. The new panels were laminated with perforated silverized teflon tape rather than aluminized mylar sheeting, both to avoid this failure mode and for improved thermal performance. 174

5.34 Surface temperatures of CREST’s two NiMH battery banks shortly before and after bank 1’s apparent charge controller failure. Bank 1 and its charge controller were replaced with spares hours before the second launch attempt. Bank 2’s temperature and operating characteristics remained stable throughout. . . . 180

5.35 Plot showing the disk space remaining on SFC disk 2 (in GB) and the current drawn by the SFC from the power system. At the same moment disk space stopped decreasing and current draw increased by a measurable amount. It is believed that at this moment SFC’s connection to disk 2 was lost. 185

5.36 A housekeeping plot generated during the last few hours of flight, showing crystal STAC (top right) and veto tube (bottom right) rates in hz. With the veto high voltage provision at 30% of normal the full instrument could be operated in a stable manner. With the veto high voltage at normal values, the CROL would crash within minutes, reducing crystal STAC rates to zero. . . . 186

5.37 A composite graphic overlaying electron and proton flux data from NOAA’s Polar-orbiting Operational Environment Satellite (POES) (yellow shading) with CREST’s flight path up to that point (red squiggly line). This was subsequently referred to as the “Yellow Ring of Doom” as a possible explanation for CREST’s degraded performance beginning the second day of flight. Areas of intense yellow indicate high power of protons and electrons of solar origin during an M2-class solar flare which occurred on 12/26/2011 at 8:00 PM UT. The onset of instrument performance degradation (on 12/26/2011 at 6:00 UT) coincides with CREST leaving the safety of the eye in the center of the solar storm. Whether or not this abnormally high solar activity in fact had anything to do with CREST’s degraded performance is unknown. The instrument’s behavior is consistent, however, with the veto system detecting an abnormally high flux of low energy charged particles as would be expected from this kind of situation. 189

5.38 CREST’s flight trajectory, starting from McMurdo on 12/25/2011 NZDT (12/14/11 EDT) and ending in nearby Victoria Land ten days later. 191

5.39	The environmental pressure and atmospheric overburden at CREST’s flight altitude. The pressure was measured by onboard instruments while the overburden was calculated from Antarctic atmospheric models. CREST’s overburden was below 6 g/cm ² for only a short duration which roughly matches up with data taken when the full (rather than half) detector plane was operational. Plot by S. Coutu.	192
5.40	The magnetic field sensor’s measurements during the Antarctic flight. <i>Top, a</i> : The Z-component of the field was approximately constant at 0.6 G (downwards), with the X and Y components oscillating daily with a maximum transverse magnitude of approximately 0.12 G. Error estimates for each data point were not available. <i>Bottom, b</i> : The verticality of the magnetic field at CREST’s location ranges between approximately 75% and 90%.	193
5.41	CREST’s raw event rate as a function of time (blue, left axis) along with its flight altitude (red, right axis). Major events at times shown by the large red numbers are detailed in Table 5.4.8.	194
5.42	CREST’s ‘No Top Or Side’ event rate as a function of time (blue, left axis) along with its flight altitude (red, right axis). Major events at times shown by the large red numbers are detailed in the text.	194
5.43	CREST’s ‘Two Side’ event rate as a function of time (blue, left axis) along with its flight altitude (red, right axis). Major events at times shown by the large red numbers are detailed in the text.	195
5.44	We received our cold weather gear at the CSBF departure facility in Christchurch. This was my first meeting with my ‘Big Red’, the standard issue coat issued to every McMurdo resident, which would serve me well over the next several months.	197
5.45	The USAF C-17 which transported the initial CREST ice team (and several hundred other passengers) from Christchurch to McMurdo.	197
5.46	A view of McMurdo with Scott’s Hut in the foreground. Between the Hut and McMurdo is Winter Quarters Bay, where when the ice is broken in summer resupply ships will anchor.	198
5.47	Our first glimpse of Ivan the Terra Bus, our principal mode of transport between McMurdo proper and the Long Duration Balloon (LDB) facility.	198
5.48	The first elements of the CREST team arrived in McMurdo on 11/2/2011.	199
5.49	The exterior of ‘HQ’ at LDB, housing the conference room, mess hall and kitchen. LDB’s social hub also includes a foosball table.	199
5.50	CREST PI Jim Musser stands in front of CREST as it arrives at LDB inside of its shipping container.	200
5.51	Lifting CREST into Hangar 096 via sliding ceiling crane.	200
5.52	Home, sweet home. When CREST next leaves this hangar it will be by balloon.	201
5.53	Professor Scott Wakely of University of Chicago pauses while examining a CREST C-channel on the ground level of Hangar 096 at the LDB facility near McMurdo Station, Antarctica. The aluminum jig used to lift the channel was designed by Alex Shroyer. This allowed us to safely and quickly remove the C-channels from the instrument for maintenance.	201

5.54	Nahee Park inspects CREST’s thermal outer covering during a rollout test at LDB. Visible in the foreground is CREST’s flight computer’s umbilical (ethernet) cord, which we had to subsequently wean ourselves off of prior to flight.	202
5.55	A view of Mt. Erebus from the patio of Hangar 096. Mt. Erebus, approximately 30 miles away in the photo, is the southernmost known active volcano in the world. Clouds such as this unfortunate one often cling to its slopes, sometimes in stacks of three or four at a time. The puff of smoke at the summit is volcanic in origin and varies significantly in size and direction, at times a useful indicator of the wind pattern at higher altitudes.	202
5.56	New Zealand’s Scott Base is far smaller than McMurdo population-wise but more than makes it up for it in their zealous hospitality (and with the quality of their gift shop). It lies directly between McMurdo and LDB and we passed it twice each day on our way to and from Hangar 096.	203
5.57	Michael Schubnell poses in front of McMurdo on our way up Observation Hill (referred to locally as Ob Hill). Fuel tanks dominate the foreground, with McMurdo’s organic architectural style evident behind them.	203
5.58	Michael Schubnell, myself and Scott Nutter in front of Mt. Erebus from the slope of Observation Hill.	204
5.59	A flock of four antarctic terns near the summit of Castle Rock. Picture taken by Stéphane Coutu during our duo climb of said rock while awaiting our second launch attempt.	204
5.60	The trail back to McMurdo from Castle Rock. Orange flags mark the safe path; by deviating from the flags one risks falling into an unmarked cravasse, likely never to escape. Photo by S. Coutu.	205
5.61	Cutting ice blocks at Happy Camper training session. To survive in the Antarctic waste, construct a wall out of such blocks to block the incessantly heat-thieving wind, pitch a tent, then radio for help.	206
5.62	These handy ice blocks can also serve as solid foundations for kitchen appliances.	206
5.63	A ‘Foremost Delta Two’ transport vehicle. Supposedly the safest transportation vehicle for ice roads, its light weight and immense tires make it susceptible to getting stuck in ice potholes in warmer temperatures, or whenever the driver spins the tires too quickly (see the next Figure). Riding in the undulating toolshed hitched back of one is a rapid way to lose your lunch. If one had to sprint 100 meters to steal your buddy’s seat aboard Ivan the Terra Bus thereby dooming them to a ride aboard the Delta, one did so without regret. Overall, McMurdoans least-favorite mode of transportation vehicle, barely beating out “on foot.”	207
5.64	A heavily laden Delta stuck in the snow, awaiting rescuing.	207

5.65	CSBF’s rotator, responsible for pointing CREST at the sun during flight to heroic effect. CSBF claimed it had pointing accuracy of a few degrees and precision of less than one degree. Without this device maintaining a 10-30° aspect towards the sun our batteries would have overcharged mid-flight due to higher than expected solar power production. Micro-managing this device to avoid over- or under-charging our battery banks proved to be the chief occupation of shift operators during otherwise quiet periods.	208
5.66	The crate used to transport the solar panels from Chicago to Indiana, from Indiana to Palestine and from Palestine to McMurdo. Once there, its foam blocks and cozy frame served mainly as a convenient resting place. Note: physicist not included.	208
5.67	The magnetic field sensor supplied by the University of Michigan Department of Space Science. Data taken by this device is presented in Section 5.4.8. . . .	209
5.68	Before using the snowmobiles we were given snowmobile training. Snowmobile training started off slowly for some of us.	209
5.69	One of the two snowmobiles assigned to CREST for use during launch procedures and for ‘training’. This particular vehicle left me stranded on the launch pad during a recreational jaunt after its carburetor detached during a routine jump.	210
5.70	Thanksgiving dinner at the main McMurdo mess hall (and probable breeding ground for the Crud super-virus). From left: Jim Musser, Stéphane Coutu, Michael Schubnell, Joseph Gennaro, Scott Wakely, Nahee Park and Scott Nutter.	210
5.71	Thanksgiving dessert at LDB mess hall. All the lunches there were truly extraordinary. Completely unrelated to this is the fact that all McMurdo personnel tours of LDB occur just before lunch time.	211
5.72	CREST’s first of many roll-outs onto the balcony of Hangar 096. Here we were testing re-integration of the solar panels.	211
5.73	Our neighbor, STO. Due to damage suffered by their cryogenic system during shipment, their launch was delayed until after CREST’s. During our launch they served nobly as filmtakers.	212
5.74	CREST PI Jim Musser skis out towards the launch pad at LDB. Skiing was an essential recreational component of our time at LDB, providing much-needed respite from the stress of hangar life.	212
5.75	It turns out the best way to smooth an ice road is to roll a very, very heavy wagon on top of it.	213
5.76	‘The Boss’, CREST’s launch vehicle and for many years the largest vehicle in Antarctica by mass.	213

5.77	The sun-side of the CREST instrument (facing the +X direction). The top portion of the detector is shielded by 2"-thick foam panels assembled into a hermetic box surrounding the veto system. The lower portion shows the layers of perforated silverized teflon tape covering sheet metal panels. The same tape was also used to cover all metal surfaces of and attached to the instrument chassis (e.g. the bar installed by CSBF supporting their communications array) with the exception of the solar panel mounting brackets. Both solar wings mounted on the -Y and +Y instrument sides are visible.	214
5.78	CSBF personnel testing attachment of the skirt solar panels (that power the SIP) and the high-tech crush pads (made of layered cardboard). The white oblong shape at center bottom is the ballast container, which can release up to several hundred pounds of glass beads during flight. Notice that the side panels are unattached; CSBF required access to the interior of the instrument to attach the skirt panels. This required the instrument to be lifted out of the hangar doors before closing it up (as the doors were not tall enough to extract CREST with the ballast container attached). This caused an enormous amount of trouble and nearly scrubbed our first launch attempt before it began. This was amended for the next two launch attempts, enabling us to button up the instrument inside rather than outside the hangar.	215
5.79	The Boss taking CREST out for its first hang test at LDB.	215
5.80	CREST and our two assigned snowmobiles during its first hang test aboard The Boss at LDB.	216
5.81	CREST's bottom surface. The sheet metal panels are laminated with perforated silverized teflon tape. Exposed steel boltheads were covered with aluminized mylar tape to prevent them from compromising the thermal scheme. Visible near top left is some tape that has been blown loose by the wind. The cold temperatures and persistently vigorous breeze vastly reduced the effectiveness of all of our tapes' adhesives.	216
5.82	Close-up of a solar panel showing the amount of damage sustained in and between Palestine and McMurdo. Despite the scars the panels performed above specifications during flight.	217
5.83	In order to maintain our position in the launch queue, CREST formed a vicious bicycle gang known as the CRESTaceans. From left to right: Scott Wakely, Jon Ameel, Michael Schubnell and Scott Nutter.	217
5.84	In this photo I am precisely at the minimum legal distance from the seal specified in the Antarctic Treaty.	218
5.85	CREST declared flight readiness on 12/6/2011. From left to right, top: Michael Lang, Jon Ameel, Scott Nutter, Scott Wakely, Jim Musser, Michael Schubnell. Bottom: myself, Nahee Park, Stéphane Coutu.	218
5.86	The CREST group has generously supported me in many ways and I am truly grateful for the opportunity to have worked with and gotten to know each of them.	219

5.87	A Kress tractor hauling a massive load, stuck in a pothole on the ice road to the Pegasus airfield. Yours truly formed part of the rescue team. For those of you asking which vehicle displaced The Boss as the heaviest on the continent, this is it.	219
5.88	The rescue vehicle for the heaviest vehicle on the continent: another Kress tractor. Jules, fearless leader of the ice road maintenance team, coordinates a rescue effort via radio. This was the first ever rescue of a fully laden Kress truck on the continent.	220
5.89	Cool physicists don't look at explosions...	220
5.90	This seal had nothing whatsoever to do with the scientific aims of the CREST expedition. Or did it?	221
5.91	Celebrating declaration of flight readiness with an improvised musical instrument. From left: Jon Ameel, Stéphane Coutu, myself and Nahee Park.	221
5.92	The CREST Ground Station inside Hangar 096. Front left: GSE3. Facing away: GSE4. Jon Ameel is visible working at his engineer's station near the rear of the hangar.	222
5.93	Sheet metal panel showing cutout region for power cables that will remain connected to CREST until moments before launch to ensure full battery charge during ascent. As the instrument will likely be spinning, pointing (and thus reliable solar power) will be impossible until it stabilizes. The region was to be blocked by a foam block covered in aluminized mylar after disconnecting those cables, secured by tape. It is doubtful that said block remained in place given that we were typically unable to adhere tape securely outside the hangar (not to mention at flight altitude).	222
5.94	The foam block used to cover the power cable access did in fact shrink during flight. Photo by J. Ameel (during recovery).	223
5.95	CRESTaceans plot their next move. From top left, clockwise: Scott Nutter, Nahee Park, Jim Musser, Stéphane Coutu.	223
5.96	<i>'Standing on a hill in my mountain of dreams, telling myself it's not as hard, hard, hard as it seems'</i> (Led Zeppelin, Going to California). Michael Schubnell during our approach on foot to Castle Rock. The foggy weather precluded an ascent on that day.	224
5.97	Ivan the Terra Bus waiting to take us back to McMurdo from LDB. Though comfortable the ride was also very slow, taking approximately 50 minutes to travel between LDB and McMurdo (vans did the same trip in approximately 10 minutes). While climbing the switchback up the hill past Scott Base the bus travels more slowly than a jogging CSBF engineer.	224
5.98	View of LDB from the launch pad. All non-tent buildings are deployed on skis and moved to higher ground every winter. Otherwise, non-mobile buildings eventually sink below the ice surface.	225
5.99	Inside a Delta transport desperately trying to avoid losing today's lunch. It is customary for experiments, unions and military outfits to paste their sticker all over the inside of the vehicle, lending it an decor eerily reminiscent to that of a men's bathroom. CREST had no sticker.	225

5.100	View of the dorm buildings in McMurdo where we spent the ‘nights’. We were two to a room. Bathrooms were shared and we were encouraged not to shower every day to save clean water. Each building also had laundry facilities (each room was assigned one laundry day per week) and common living areas with seating, board games and televisions.	226
5.101	(a): Detailed map of the road system for McMurdo and the surrounding ice shelf. The only solid land here is colored in green. The distance from McMurdo to LDB is approximately three miles following the road. We arrived on the Ice Runway (top left) and departed months later from the Pegasus Runway (bottom left). An arrow near the top middle states that the ice shelf flows to the left at a rate of approximately 385 feet per year. Closer to Pegasus Runway the rate is only 95 feet/year. (b): Closer view of the LDB surroundings, which abuts against Williams Field (called Willy Field). Clearly visible is the circular launch pad, around which we skied many times. (c): Closer view of McMurdo, Hut Point and the Ice Runway. (d): The edge of the permanent ice shelf lies just beyond the main Ice Road between LDB and Pegasus Runway. It was on this stretch that we rescued the stranded Kress truck.	227
5.102	IU Electrical Engineer Michael Lang. All of the tools transported by Alex Shroyer and myself to Palestine were then shipped to McMurdo, enabling significant work to be done on CREST in Hangar 096 prior to flight readiness.	228
5.103	An early prototype for the GoPro mobile camera, used by myself to make videos of launch procedures, ski trips, snowmobile rides and kite flying. . . .	228
5.104	During our first launch attempt on 12/18/2011, CSBF’s skirt and crush pad attachment procedure called for us to button up the instruments bottom thermal panels outside the hangar. Temps of below -10°F and a bitterly frigid wind made this an incredibly difficult task. Without mittens in addition to thinner gloves my hands froze; with mittens I was unable to screw anything in. Furthermore, the tension induced in the chassis by the crane lift misaligned some of the screw holes, making the task even more difficult. Consequently a procedure which normally took a few minutes inside took nearly two hours. . . .	229
5.105	CREST, finally buttoned up for its first launch attempt.	229
5.106	CREST, finally buttoned up for its first launch attempt.	230
5.107	One example of the <i>fata morgana</i> (Morgana’s Curse) mirage, caused by layers of air with differing temperature near the ice surface.	230
5.108	A slightly more striking example of the <i>fata morgana</i> mirage.	231
5.109	The foam block covering the power cable pass-through hole during our first launch attempt.	231

5.110 Stéphane Coutu attempts to activate Nashua duct tape’s adhesive with a heat gun while I stand ready to catch him should the ladder tip over. The low temperatures during our first launch attempt prevented our silverized teflon tape’s adhesive from adhering. At first we attempted to tape over it with packing tape, but it also would not stick. Neither would Nashua duct tape (which CSBF staff claimed had never happened to them before). I hit upon the idea of blasting the tape with a heat gun in order to activate the adhesive. This worked but was very time consuming. Pressed for time CSBF assisted us in securing the foam panels by tying ripcord around the instrument.	232
5.111 M. Lang, N. Park, J. Musser and J. Ameel (ladder) attaching the -X, +Y side panel in the bitter cold.	232
5.112 Our bench power supply, “Maverick,” braved the elements along with us on the failed first launch attempt. Not pictured is the previous power supply, “Viper,” which powered CREST at Indiana and in Plum Brook. No “Top Gun” references were harmed during the creation of this instrument.	233
5.113 One of CREST’s two Nickel Metal Hydride battery banks. The wires marked ‘CUT’ are for the benefit of recovery personnel.	234
5.114 Jim Musser, Nahee Park and Myself after replacing one of CREST’s two battery banks, which had been overcharged by a faulty charge controller.	234
5.115 The power cables which would remain attached to CREST until moments before launch. Michael Lang bravely performed this task, thereafter transported away from the launch pad by a snowmobile driven by CREST’s most elite snowmobile pilot.	235
5.116 To assuage our pain at the failure of our first launch attempt we fled to the comfort of the neighborhood wine house.	235
5.117 View of Mt. Erebus from atop Castle Rock.	236
5.118 View northwest from midway up Castle Rock.	236
5.119 The same view northwest from atop Castle Rock.	237
5.120 View south from Castle Rock looking over the ice shelf. The tiny, barely visible buildings are LDB.	237
5.121 A Skua. These hardy residents of the icy coast feed on anything to survive, including each other’s young and the contents of cafeteria trays of unwitting McMurdoans. Protected by the Antarctic Treaty and free from retribution, they grow bolder by the day.	238
5.122 Life in McMurdo was not without its comforts. Jon Ameel, Jim Musser, myself and Nahee Park at the Café, one of three establishments at McMurdo serving alcoholic beverages. I cannot recall the photographer; judging by the rapidly emptied wine glass, neither can he/she.	238
5.123 CREST being picked up by The Boss the morning of our third and final launch attempt. It would not touch ground again until the end of its flight.	239
5.124 Mt. Erebus’ plume the morning of our third and final launch attempt.	239
5.125 Overcome by happiness at the prospect of our imminent launch, Jon Ameel flies a kite.	240
5.126 A massive spool unwinds the tether that will connect CREST to its host balloon.	240

5.127	Prior to launch dozens of CSBF engineers and support staff expertly prepared the balloon and tether. (a): Helium tanks that will inflate CREST's host balloon. (b): A view down the flightline: preparing CREST's tether. Visible in the distance is The Boss with CREST hanging behind. (c): Moving down the flightline towards The Boss. (d): The rotator and ring harness.	241
5.128	CREST hanging from The Boss minutes before launch. The power cord linking the Maverick generator-fed power supply aboard The Boss snakes in from the lower left and would need to be detached shortly.	242
5.129	CREST hangs majestically from The Boss minutes before launch.	242
5.130	Erebus' plume shows evidence of calm winds minutes before launch. CSBF used an actual weather balloon to determine this, but I prefer to use the sulphurous emission of a volcano.	243
5.131	Initiating helium fill of the balloon.	243
5.132	The balloon begins to inflate.	244
5.133	The balloon continues to inflate.	244
5.134	The ballon continues to inflate via the fill tubes (white wisps coming out either side of the inflated portion of the balloon). The loud 'whoosh' of rushing helium dominates the scene. At left, CREST's orange and white parachute lies in wait.	245
5.135	Shortly before launch.	245
5.136	A minute or so after launch one could barely make out the solar wings.	246
5.137	Several minutes after launch all that was visible of CREST to the naked eye was periodic flashes of reflected sunlight as CREST spun every few seconds.	246
5.138	Though it was hard to say goodbye, I would much prefer CREST continue to rise than the alternative.	247
5.139	The remaining ice team celebrates with champagne in LDB-issued spill-proof coffee cups. A well-deserved celebration.	247
5.140	A skua, likely fleeing the scene of an unimaginably horrible crime. Even they appear majestic from the right angle.	248
5.141	Bassler aircraft shrouded by fog at the Pegasus airfield. The fog cancelled inbound and outbound flights that day.	248
5.142	While awaiting the recovery mission, Jon Ameel encountered an amiable LDB visitor, separated from his tour group and asking for directions to the kitchen tent.	249
5.143	Pictures of and from aboard the LC-130 emergency medevac aircraft that also transported Stéphane Coutu and myself from McMurdo to Christchurch fuels up while awaiting takeoff at the Pegasus Runway on 12/28/2011. (b): Three phases of water: fragmented ice sheets float on a blue ocean beneath the clouds. (c): The melting, fragmented edge of the Ross Ice Shelf. (d): Rugged splinters of the Transantarctic Mountain Range.	249
5.144	The shifting snows digest CREST as it awaits recovery. The +Y side solar wing (far side) was martyred by the falling rotator attachment - ironic given that the rotator moderated the solar panels' power production all flight long by pointing them up to 30° off sun. Photo by J. Ameel.	250

5.145	The rotator crashed right through the +Y-side solar wing. We were fortunate it did not crush the crystal system, the batteries or the SFC. The +Y-side solar wing was not. Photo by J. Ameel.	250
5.146	The SFC during recovery. Disk 2 (the one that stopped working) is second from the front. Photos by J. Ameel.	251
5.147	The recovery team finishes picking CREST’s bones clean. Only a battery bank, the SIP and some odds and ends remain. Photo by J. Ameel.	251
5.148	CREST awaits recovery in a desert of windblown snow, casting a shadow in the downwind ice pattern. Photo by J. Ameel.	252
6.1	Geometry archetype for signal events. A primary electron’s trajectory illuminates the CREST detector plane. The plane containing the electron’s motion intersects the detector plane on a line herein referred to as the “main event line.” The angle between these two planes θ is equal to the electron’s zenith angle against the horizon when CREST’s detector plane is parallel to the ground. I refer to this angle as the polar angle. To characterize the trajectory with respect to detector axes a second angle, ϕ is required, which characterizes rotation of the electron’s plane of motion around the detector’s z axis. I refer to this angle as the azimuthal angle. My analysis efforts revolve around reconstructing these two angles. Note also that a positron with the same initial momentum would behave identically to the electron with the exception that it would be pushed in the opposite direction by the geo-magnetic force.	256
6.2	A negative electron illuminates CREST while interacting with the perpendicular component of the geo-magnetic field, $B_{geo}^{\vec{}}$ which points in the $-\hat{z}$ direction (this is close to the actual orientation for the initial segment of CREST’s flight). The resulting Lorentz force on the negatively charged electron is directed as shown by $F_B^{\vec{}}$. When illumination starts, the electron has initial momentum \vec{p}_1 and its extended trajectory intersects the CREST detector plane at point P_1 . When illumination stops, the electron has final momentum \vec{p}_2 and its extended trajectory intersects CREST at point P_2 . θ measures the polar angle between detector horizontal and \vec{p}_1 . Synchrotron and Bremsstrahlung photons emitted by the electron will have nearly-parallel trajectories such as those marked by gray arrows; the spacing between photon emissions is random. These photon trajectories will intersect CREST’s detector plane somewhere along \vec{L} , the main event line, marked by the teal line between points P_1 and P_2 . Note that a positron would have an identical diagram, except that the directions of $F_B^{\vec{}}$ and \vec{L} would be reversed, and the roles of initial and final momentum and intersection points would be exchanged.	257
6.3	A top-down view (we are looking in the $-\hat{z}$ direction) of the same event as in Figure 6.2. The initial and final electron momentum vectors, \vec{p}_1 and \vec{p}_2 are foreshortened due to our perspective. Also foreshortened are the trajectories of photons emitted by the primary electron (gray arrows). In the projection of the detector (\vec{x}, \vec{y}) plane, the electron momentum is perpendicular to the main event line \vec{L} when the magnetic field $B_{geo}^{\vec{}}$ points in the $-\hat{z}$ direction.	258

6.4	Normalized veto activity as measured by the number of paddles with pairs (hits in PMTs on both ends) in simulated 10 TeV signal events, by subsystem. Cutting out events with any top veto activity eliminates over 90% of these events. The bottom veto has high average multiplicity which suggests using it as part of the reconstruction process.	260
6.5	Plot by Jim Musser of the inverse track velocity distribution for flight data events. I discuss the meaning of and a procedure for calculating the inverse track velocity in Section 6.4.5.2. The cuts applied to generate this plot are listed in Table 6.1. Out of the entire flight, a few hundred events remain in the signal region of this plot (near zero inverse track velocity). Whether this represents an excess over the amount of background events expected to bleed into this region via random coincidence or event mis-characterization will depend on the results of additional background Monte Carlo event production (which, to my knowledge, is still in progress at the time of this writing).	260
6.6	CREST as visualized in “EventDisplay,” the three-dimensional display tool I developed to help visualize simulated and flight events. The veto paddles are represented by transparent green rectangular solids. Beneath the top veto lies the crystal plane, each crystal depicted as a solid yellow disk. The user can interact with the display by rotating, zooming and panning, as well as inspecting instrument sub-systems (e.g., the bottom veto system) or individual detector components (e.g., bottom veto paddle 3). On the left is the control pane, allowing the user to navigate between events, view event metadata (upper-left) and view and/or set event parameters (middle left). Clicking on a detector module in the display window allows further interaction with it in the lower left (e.g. to view hit information or change its appearance).	263
6.7	Another view from EventDisplay making clear the location of paddles belonging to various veto subsystems, delineated by color. For visibility the facing two side panels and several top panels have been made invisible. The eight top paddles lying immediately above the crystals with long dimension along the X direction are colored blue. The four side paddles are green (the facing two are invisible here). The nine bottom paddles running along the Y dimension are colored red. One of the bottom paddles is slightly offset below the rest; this is to accomodate a structural beam. The six smaller, inside slanted vetos are colored magenta. The same beam that displaces the center bottom paddle causes a significant gap between four of the inside slanted paddles. Particles that sneak through these gaps and interact with the crystals have a relatively large zenith angle, and can not make tracks through the crystals longer than a few tens of cm (and can therefore be eliminated in detailed event analysis). . .	264

6.8	A template “traversal”-type event in Event Display. The primary electron follows the (very exaggerated) curved teal trajectory, crossing detector altitude inside the veto box and crystal plane (qualifying this as a traversal event). The thin teal line within the crystal plane is the intersection of the electron’s extended trajectory with the crystal plane. The magenta-colored veto paddles have been activated on both ends and thus have a pair. The relative time difference between these matched-end PMT signals provides an estimate of the location along the paddle where the hit occurred; these locations are marked by red elliptical cones.	266
6.9	Fraction of simulated signal events in which the primary electron passes through CREST’s veto system. As I predicted in Section 2.1, the fraction is approximately 60% for 10 TeV primary electrons, and rapidly increases as primary energy decreases, until at 500 GeV the fraction is over 80%.	269
6.10	Scott Nutter’s plots which demonstrate utilizing the pulser information to determine timing offsets between crystal tubes. (a): Gaussian fit to the distribution of timing differences between tube 2 and an arbitrary reference tube in a single pulser run used to determine the mean t_0 . The x axis is measured in nanoseconds. (b): Plot of the time variation of tube 2’s mean timing offset over the course of the flight. The x axis is time relative to the what the collaboration referred to as “Joe Time,” which was merely a round number near midnight of launch day. Variation over the course of the flight is a fraction of a nanosecond.	271
6.11	Two parameters used to guide my choice of 100 ns for the minimum time delay between events inside of freeze cycles. (a): Distribution of Freeze Ticks for events in a subset of flight run 03499. Bins on the x axis are Freeze Tick with bin widths of 12000. Each bin contains approximately 12000 events, showing that each freeze cycle in this particular run contained, on average, approximately one event. (b): Distributions of event durations in 10 TeV signal Monte Carlo simulation. Event duration is defined as the time in nanoseconds between the first and last TDC values of hits in the event.	273
6.12	Hit times (compared to the time of the earliest hit) in simulated signal events for primary Bremsstrahlung photons and their secondary products (green online), primary synchrotron photons and their products (blue online) and the primary electron and its secondary products (red online). Three time scales are apparent. First, the simultaneous arrival (at 0 ns) of the primary electron and photons. Second, intra-detector scattering takes place within one light crossing time of the detector (between two and six ns). Third, showers formed in the air above the detector can result in the delayed arrival of shower products, ranging from eight ns all the way up to hundreds of ns. Synchrotron photons contribute relatively few hits to the late (air shower) regime because of their lower average energy; they more often interact via the photoelectric effect or Compton scattering than pair production (see Figure 4.7).	278

6.13	Electron signal candidate event from run 03488 showing cluster formation consistent with all three time regimes shown in Figure 6.12. The two clusters towards the right (lower right in the top-down view) happen at time 0, indicating they are probably primary photons. The left cluster (upper left in the top-down plot) has a primary photon with two slightly later hits that are most likely secondary Compton scatters. The center cluster is likely a Bremsstrahlung primary photon that air showered above the detector. These hits arrived 10-15 ns later and belong to the latest time regime.	279
6.14	Distance, in units of crystal separation parameter (1 separation = 7.5 cm), between the location where a primary photon pierced the crystal plane, and the location assigned to the cluster containing its secondary hits. These distributions resulted from applying my clustering algorithm to simulated 10 TeV signal events while varying the temporal FoF window parameter Δ_t , and the method used to determine the cluster center. I plot the largest error within each event. Using a shorter time for Δ_t reduces the frequency of catastrophically mistaken cluster locations. “Earliest” and “heaviest” refer to whether I take the earliest hit or most energetic hit as the cluster center; for most clusters, there is little practical difference between the two.	280
6.15	Plot of the distribution of distances between the crystal struck by a primary particle or photon and the crystals struck by its secondary products with a full lead shield configuration. The x axis is in units of crystal separations (one center-to-center crystal separation is equal to 7.5 cm). Plot produced by Scott Nutter using the signal Monte Carlo he developed. Over 90% of secondary hits occur in crystals within two crystal separations of the primary crystal leading to my choice for maximum separation between hits in the same cluster of 15 cm (2 crystal separations).	281
6.16	Distribution of the statistic $\Delta_\phi \equiv \bar{\phi} - \phi$ in simulated 10 TeV signal events when using the main event line estimating algorithm outline in Section 6.4.5.2. Main event lines in events with very high crystal multiplicity are mis-characterized more often than those in events with low crystal multiplicity.	282
6.17	Distribution of the inverse track velocity in simulated 10 TeV signal events of varying maximum extent (event length is defined as the length of the fitted main event line). There are two main populations. First, the “simultaneous” population at zero inverse velocity (corresponding to simultaneous line endpoints). This value occurs when the main event line endpoints happen to be correctly chosen as primary photons. Second, the “scatter” population at a wide distribution centered approximately at $v/c = 1$. This value occurs when one endpoint of the line is chosen to be a cluster that results from intra-detector scattering. Since <i>Lineal</i> used a time window of 6 ns to generate this analysis, hits from the second regime in Figure 6.12 become valid cluster locations and times. Evidently <i>Lineal</i> is more likely to discover the correct line endpoints in longer events.	283

6.18	Schematic summarizing how Lineal dissects crystal hit morphology in an event. The two most distant clusters (at lower left and upper right) define the endpoints of Lineal's estimate for the main event line (solid teal line). The length of this line defines the analysis parameter "event extent." All of the other clusters in the event are sorted according to their projected coordinate along the main event line (that is, along the rotated axis \vec{x}'). This enables generation of a plethora of analysis parameters. The lengths of gaps between clusters (Sep_{i-1} , Sep_i , etc.) are measured as shown for use in the crystallography cut (discussed in Section 6.4.5.3). Second, the projected coordinates orthogonal to the main event line (that is, along the rotated axis \vec{y}') are measured by the lengths of the dotted lines. This provides an estimate of the "event width" and the starting point for calculating the goodness-of-fit of the line (as characterized by the reduced χ^2 variable). The last parameter extracted via this process is the ratio of the spatial distance and time difference between the line endpoints. This statistic (the "inverse track velocity") measures the inverse of the implied velocity of the particle that made this track in units of $1/c$. If the line endpoints are simultaneous then the inverse track velocity is zero, implying that if a side-going particle created this track it must have had infinite velocity (in other words, a side-going particle likely did not make this track). If the line endpoints are on a null interval (as would be the case for hits produced by side-going charged particles) then the inverse track velocity is approximately 1. Figure 6.17 shows the distribution of this statistic in Monte Carlo signal events.	284
6.19	The two-point function for distances between crystals as a function of the azimuthal orientation of the vector between the two crystals. Orientations of 0 and 90 degrees both correspond to alignment with a principal axis of the crystal lattice. Gaps along the main event line in signal events should be characterized by this spectrum.	285
6.20	The un-occluded two-point function for distances between crystals. Generated by removing those distances from the plot in Figure 6.19 spanned by vectors which intersect a third crystal at some intermediate distance. Proton tracks at a given angle can leave gaps of these sizes between consecutive hits, but not larger.	286
6.21	Subset of the plot shown in Figure 6.20 focusing on tracks aligned with the crystal lattice's principal axes. Such tracks should be characterized by large distances between long runs consecutive hits. At these angles, occluded spectra (from proton tracks) and un-occluded spectra (from signal electron events) are not easily distinguishable.	287
6.22	A linear event with no veto activity from run 03488. Since it is not aligned with the crystal lattice principal axes it has characteristically short gaps between hits.	287
6.23	Another linear event with no veto activity from run 03488. Near alignment with a crystal lattice principal axis leads to gaps and runs of hits with a characteristic size as the proton enters and exits rows of crystals.	288

6.24	Another linear event with no veto activity from run 03488. Even closer alignment with a principal axis has enabled the proton to leave a gap (and a string of consecutive hits) nearly the size of the instrument.	288
6.25	Visual representation of the definition of an “illegal hit,” forming the basis of my crystallography cut. Gaps between hits that reside inside the red box in this plot are “illegal” for side-going protons: in order to create a gap between hits of those distances and at those track angles, they would first have to pass through another crystal. Events that create gaps inside the red box are likely not the result of side-going charged particles, independent of veto activity. Requiring multiple illegal separations in an event helps eliminate side-going charged particles that evaded detection by the veto system.	289
6.26	Event display with overlay describing the meaning of several key linear analysis parameters. The event length is measured as the length of the teal main event line. The event width is measured as the maximum orthogonal distance between the main event line and hits in the event. The veto activity in this event has been suppressed in this display for demonstrative purposes. The number of illegal separations between hits is calculated as in Section 6.4.5.3.	291
6.27	The same event from flight data shown in Figure 6.26 with some veto information made visible. The red ellipses correspond to estimated locations of hits within veto paddles. The cones with the sharp ends upward represent hits in the top veto system; the cones with the sharp ends downward represent hits in the bottom veto system. Here I have highlighted those hits that are consistent with the proposed traversal location of the primary electron. The size of the cones reflects the uncertainty in the position estimate: approximately the width of the paddles in the short dimension, and approximately 30 cm (σ_{veto} in the long dimension.	292
6.28	Comparing two methods for estimating hit locations along top veto paddle 0. On the left, I plot the average location estimate and error as a function of relative response in the PMTs on either end (using the formula shown). On the right, I plot the average location estimate and error as a function of relative timing in the PMTs on either end. The timing method provides smaller errors and more uniform results across the entire paddle.	295
6.29	Hit location accuracy of the four side veto paddles. Plotted is the distribution of the distance, in cm, between the intersection of the extended track of crystal hits with a paddle and the location generated using the end-to-end relative timing method detailed in Section 6.6.1. The distributions have been corrected for the unknown timing offset between the PMTs on either end (x_{offset} in Equation 6.4. All four paddles achieve hit location estimation with error of order ~ 30 cm (a distance I call $1 \sigma_{veto}$).	296

6.30	A flight event from run 03488, likely the event of a side-going proton which evaded detection in the far-side paddle. The large red elliptical cones indicate the estimate location within the veto paddle where the particle or photon traversed the paddle. This location is based on the relative timing between the PMTs on either end of each paddle (and as you can see this occasionally results in the ludicrous location estimate outside of the paddle itself, indicating a timing error). The gradual increase of the purple crosses indicating the hit time as the crystal track progresses is consistent with a particle moving across the detector with velocity c . The uniformly short distance between hits is consistent with a side-going particle as well. The lack of veto activity beneath the far side of the event provides strong evidence against the possibility that the particle inducing the event had momentum in the vertical direction.	297
6.31	Another flight event from run 03488, also likely the event of a side-going proton. This time however the veto system has properly flagged the event with entrance- and exit-locations. The blue cone in the facing side paddle indicates <i>Geometizer's</i> best guess as to where it believes the crystal track indicates the particle exited the detector. The close agreement with the veto hit location estimate provides strong evidence, <i>independent of the timing of the crystal hits</i> , that this was indeed a side-going proton.	298
6.32	Protons with more vertical trajectories leave shorter tracks in the crystal system. I plot here a geometrical estimate of the maximum length crystal track a proton with the given angle could make given that the crystals are only 2 cm tall.	299
6.33	A top-down view of the event in Figure 6.34 clearly showing how the location of the blue cone is established by simply extending the main event line until it intersects with a veto paddle. The way the bottom veto hits shadow the clusters in the main event line suggested to me that measuring the primary electron momentum utilizing the bottom veto signals might be possible. . . .	299
6.34	Event with overlay showing how I define “pierce error,” the three-dimensional distance between the crystal-track generated particle exit location and the nearest estimated veto hit location. The minimum pierce error is the smallest of these distances. A small minimum pierce error indicates that the veto and crystal systems agree that on the trajectory of the particle that made the crystal track and veto hits. A large minimum pierce error indicates disagreement between the two systems and is a clue that the crystal track is the result of signal electron illumination, not a background event generated in response to a side-going proton.	300
6.35	Distribution of minimum pierce errors in background Monte Carlo, low- and high-energy signal Monte Carlo and flight data sets. Minimum pierce error is defined in Figure 6.34 and is measured in units of 30 cm, the RMS error for veto hit location estimates. The data sets used here are discussed in Section 6.6.2.	300

6.36	Scatterplot of the projection of electron momentum \vec{p} into the detector plane (electronPhi, on the y axis) vs. ϕ , the azimuthal angle defining the orientation of the main event line \vec{L} on the crystal plane (labeled as beamPhi on the x axis). In simulated signal events the two vectors are near-orthogonal. This is because the magnetic field orientation sampled by the electron while it illuminates the detector is near-vertically downward.	303
6.37	A simulated 10 TeV signal event with five primary photons and the primary electron. Four primary synchrotron photons (blue disks) were generated, as well as one primary Bremsstrahlung photon (green disk) and the primary electron (orange disk). The primaries are drawn at the top ends of the red dashed lines indicating their trajectories through the detector. The crystal activity produced by these primaries are colored identically, so that for example all the green crystals were activated by secondary products of the Bremsstrahlung primary (or the Bremsstrahlung primary itself). The veto activity has been removed to make the primary and crystal activity more clear.	305
6.38	Another view of the same simulated signal event.	305
6.39	The event displays bottom veto activity offset from the crystal activity roughly in the direction of the primary (photon and electron) momentum.	306
6.40	Another view of the location of the bottom veto activity with respect to the crystal activity and the primary momentum.	306
6.41	To estimate the primary momentum I connect two points, the average crystal cluster location and the average non-top veto hit location, with a line. I draw that line here as the dashed teal line.	307
6.42	View showing that the estimate for the polar angle θ happened to be quite good in this event - the teal line (momentum reconstruction) and the red lines (true primary momentum) are nearly parallel in this projection.	307
6.43	Pointing resolution (when “cheating” by utilizing the truth-values for crystal cluster locations) for the polar (y-axis) and azimuthal (x-axis) angles of the primary electron momentum in the low-energy (<i>a</i>) and high-energy (<i>b</i>) simulated signal data sets. FWHM resolution is of order 20-30 degrees for θ and of order 40 degrees for ϕ	308
6.44	Angular dependence of the pointing resolution when “cheating” by utilizing the truth-values for crystal cluster locations. (<i>a</i>): The distribution of errors in ϕ -pointing as a function of the true θ value. Very vertical primary trajectories ($\theta \approx \pi$ radians, right side of the plot) are likely to induce large errors in ϕ -pointing. (<i>b</i>): The distribution of errors in θ -pointing has no clear dependence on the true ϕ value.	309
6.45	ϕ -pointing resolution in low-energy (<i>a</i>) and high-energy (<i>b</i>) signal MC data sets as a function of true θ . When using estimated crystal cluster locations (red marks) rather than the true locations (blue marks), ϕ -pointing worsens for horizontal primary trajectories by only a slight amount. This indicates that while the crystal clustering algorithm is less accurate for horizontal primary trajectories (which tends to create larger clusters), the clusters tend to elongate in such a way that the ϕ estimate is undisturbed.	310

6.46	<p>θ-pointing resolution in low-energy (a) and high-energy (b) signal MC data sets as a function of the true θ. When using estimated crystal cluster locations (red marks) rather than the true locations (blue marks), θ-pointing worsens preferentially for very horizontal particle tracks. This contrasts with ϕ-pointing, which worsened only slightly for very horizontal tracks when not cheating. When clusters elongate due to horizontal primary trajectory, evidently they do so in a way that significantly disrupts θ-pointing.</p>	312
6.47	<p>Two diagrams illustrating the regions of primary trajectory phase space that belong to signal electrons and positrons. With infinite-precision ϕ and θ pointing these diagrams show that charge sign discrimination would be possible with an improved CREST-like instrument by identifying which half of the potential signal phase space the primary lepton's momentum occupies. (a): Given a downward magnetic field and the teal main event line as shown here, any electron trajectory within the orange box would be a valid signal trajectory (that is, the magnetic force would point along the teal line. (b): Given a downward magnetic field and the teal main event line as shown here, any positron trajectory within the purple box would be a valid signal trajectory (that is, the magnetic force would point along the teal line.</p>	313
6.48	<p>Four plots showing the results of my attempts to properly orient the main event line, \vec{L}, so that it points in the same direction as the magnetic force on the primary electron. Correctly oriented and aligned main event lines have data points on the main diagonal. Data points that are off the main diagonal represent line-misalignments (or line reversals). In the upper left I show the distribution produced by <i>Lineal</i>'s output, which uses the earlier line endpoint as the start of \vec{L}. In the upper right I apply a "traversivity" reversal, which attempts to decide which end of the line looks like it more likely came from a shower-type event, such as from a nearby electron or high-energy Bremsstrahlung photon. The lower plots are identical, except that I restrict my attention to traversal-type events only. This suggests that traversal-type events are easier to properly orient, especially after applying the traversivity reversal algorithm. This could mean that charge-sign discrimination will be viable for traversal-type events, and not for non-traversal type events.</p>	315
7.1	<p>Distributions of inverse track velocities for events from three data sets, after applying all "mainstream" cuts (listed in Table 7.1), with the instrument in the "solid-half" configuration: flight data (black), Monte Carlo background (red online) and simulated signal (blue online). The Monte Carlo data sets are still in production and have statistical bin errors approaching zero. The flight data set has statistical bin errors proportional to $1/\sqrt{N}$ (not drawn), and systematic bin errors related to the frequency of errors in estimating the inverse track velocity (not yet estimated as of the time of this writing). The exceptionally close match between the background Monte Carlo data and the flight data will likely place a strong upper limit on the estimated amount of <i>observed</i> signal events.</p>	320

7.2 Four viewpoints of my proposed “Crystal Bowl” configuration for a future CREST-like detector. The bowl shape provides inherent insensitivity to side-going protons by shortening their track length by a factor proportional to the curvature of the bowl. Signal electrons illuminating the bowl from above would still be able to produce long “main event lines.” Due to the curvature of the bowl these would have a trajectory-dependent shape rather than simply being linear. To generate the displayed curvature I calculated the z-offset as $\Delta_z = \sqrt{R^2 - r^2}$, where $R = 350$ cm and r is the transverse distance, in cm, from the center of the flat crystal plane to any given crystal disk’s center (before being vertically offset). 323

LIST OF TABLES

1.1	Nearby remnants of supernovae (both shell and compact-object types), according to [38].	19
3.1	Properties of the BaF ₂ scintillator crystals used in CREST’s main detection grid. 1,024 such crystals arranged in a square 32x32 grid and coupled to custom-built Hamamatsu PMTs comprise CREST’s main detector plane. Compilation of information from CREST design and testing documents, [39] and [8]. For a calculation of photon interaction probabilities over a range of energies and for various geometries see Section 4.9.	67
3.2	Properties of Hamamatsu photomultiplier tubes used to collect scintillation light from the BaF ₂ crystals. From [8] and [66].	67
3.3	Properties of the EJ-200 plastic scintillator used to construct CREST’s charged particle veto system. The material consists of a polymer base of Polyvinyltoluene (C ₂ H ₃ 0) doped with the fluor Anthracene. From Eljen Technology EJ-200 data sheet, http://www.eljentechnology.com/images/stories/Data_Sheets/Plastic_Scintillators/EJ200%20data%20sheet.pdf , accessed on 4/5/15. The average excitation and ionization potential I is discussed and estimated in Section 4.1.	73
3.4	Chronology of signal dispatch and arrival times for a single STAC connected to the DIGOL aboard the XEM Carrier used to optimize firmware timing parameters on STAC boards. Signal transmission along cables up to two meters long would only increase certain delays (between SH signal dispatched/received and between Digitize command dispatched/received) compared to these measurements (generated using a digitize cable of ~12” in length. Therefore timing parameters were chosen to ensure STACs would hold on to hits longer than the minimum 180ns implied by these measurements.	96
4.1	Parameters of the Landau distributions for proton energy deposition in the BaF ₂ crystals for horizontal and vertical trajectories, for which the crystals present 24.45 and 9.78 g/cm ² of thickness, respectively. Δ is the Bethe average energy loss $\langle \frac{dE}{dX} \rangle$ for minimum ionizing protons in BaF ₂ . Δ_p is the most probable energy loss per g/cm ² of thickness of BaF ₂ . Δ_p/Δ and ξ are the plateau values for $\beta\gamma \gtrsim 10$; at lower energies protons will lose more energy. The widths do not include the widening effects of convolution with finite energy resolution of the crystals and variations in pathlengths through crystals.	114

4.2	Parameters of the Landau distributions for proton energy deposition in the EJ-200 veto paddles for normal incidence (0.51 g/cm ² of thickness). Parameter definitions are identical to those in Table 4.1.	115
4.3	Radiative energy loss parameters for electrons in CREST detector materials. $E_{1\%}$ is analogous to the electron critical energy and refers to the electron energy at which the collisional losses make up 1% of the total energy loss. E_{loss}/E_0 refers to the fraction of incident electron energy radiated as Bremsstrahlung photons while it traverses the given mass thickness of medium. The geometry cases listed here are identical to those detailed in Section 4.2.3. The lead geometry cases assume traversing 0.4 cm-thick layer of lead either once or twice, for 4.5 or 9.0 g/cm ² mass thickness.	119
4.4	Photoelectric absorption edges in Barium Fluoride. From [33].	122
4.5	Mean energies of background peaks used by Nahee Park to perform the in-flight gain matching and energy calibration of the ADC0 and ADC1 channels.	143
5.1	Heat plan for the test procedure in Plum Brook’s B-2 test chamber. Infra-red lamps would simulate the indicated component of solar radiation for an assumed sun elevation of 22° above the horizon. The listed flux values are a maximum; at any given time the actual flux values would be a fraction of the maximum but in the listed proportion. Note: Aluminized mylar was used to cover these surfaces for this test. The flight configuration utilized silverized teflon, as planned.	152
5.2	Chronology of significant detector status changes during CREST’s Antarctic flight. Runs between Disk 1, 3482 and Disk 1, 3513 were generated with the full detector plane active at float altitude. Various periods after that consist of runs with only half the detector plane active, either in a checkerboard or solid-half pattern, and at lower altitudes.	192
6.1	Cut set devised by Jim Musser that I refer to as the “mainstream analysis.” I discuss the definition of event length and a procedure for calculating the reduced χ^2 goodness-of-fit variable in Section 6.4.5.2. The “largest gap” and “second largest gap” are parameters generated via the crystallographic analysis technique I devised, which I present in Section 6.4.5.3.	259
6.2	Types of binary run files written to the flight data disks by the Science Flight Computer, their filename prefix and their typical size. An underscore and a five-digit number were appended to each file to generate the full name. E.g. the pedestal file for run 03482 would be “PedRun_03482.dat” (if a pedestal run had been executed for run 03482). “Golden” data refers to data with more than 3 crystal STACs activated and no veto activity.	267

6.3	Summary of the software analysis chain I developed to analyze CREST’s signal events. The input to my first module (Calibrator) is the output of RawReader (Nahee Park’s binary file converter). The parameters generated by EventCutter, Lineal, Geometizer and MomentumTracker are the basic building blocks of cut strategies designed to discriminate between signal and background events. Moving down the pipeline reduces the number of events to process but also increases the computational intensiveness; the total computation time required for each run strongly depends on the number of events that survive <i>EventCutter</i>	270
6.4	“Surface” event parameters defined in the ParFiller module describing the multiplicity of activated detector elements. The minimum and maximum values represent the sensible range for values of the cut parameter, not actual cut parameters which vary from analysis to analysis. By providing ParFiller with a configuration file consisting of lower and upper bounds for these parameters, the user can produce a culled file of Events meeting the desired cuts.	275
6.5	Summary of the four data sets used when generating the plot shown in Figure 6.35. The flight runs form a five-hour subset of flight data from the period after reaching float altitude but before “The Incident.” The signal Monte Carlo data sets were generated by me using the modified simulation code that allows the primary electron to interact with the detector and therefore has limited statistics.	301
6.6	Response of the four data sets listed in Table 6.5 to the listed cuts. The minimum pierce error cuts results in a multiplicative signal-to-noise enhancement of approximately 5 in simulated events. “Fraction of prev. surviving” indicates the fraction of events that reach a cut stage which survive that particular cut stage. Thus the “total remaining (fraction)” is equal to the product of all the individual survival fractions in each column. This particular set of cuts achieves background rejection of nearly 10^{-6} , but also suffers signal acceptance of less than 3×10^{-3}	302
7.1	“Mainstream” cuts used to produce distributions shown in Figure 7.1.	318

LIST OF ABBREVIATIONS

ADC Analog-to-Digital Conversion

AFE Analog Front End

API Application Program Interface

CPLD Clocked Programmable Logic Device

CREST Cosmic Ray Electron Synchrotron Telescope

CROL Crystal Overlord

CSBF Columbia Scientific Balloon Facility

DIGOL Digitize Decision Overlord

FOLGA Fiber Optic Light Guide Array

FPGA Field Programmable Gate Array

HEAT High Energy Antimatter Telescope

HESS High Energy Stereoscopic System

ISM Interstellar Medium

LDB Long Duration Balloon

LVDS Low Voltage Digital Signal

MVC Master Voltage Controller

NIST National Institute of Standards and Technology

PAMELA Payload for Anti-Matter Exploration and Light-nuclei Astrophysics

PDG Particle Data Group

PROM Programmable Read-Only Memory

SFC Science Flight Computer

SIP Support Instrumentation Package

SM Standard Model

SNR Supernova Remnants

STAC Sedecim Timing And Control

SVI Sedecim Voltage Interpreter

VTOL Veto Overlord

WIMP Weakly Interacting Massive Particle

ABSTRACT

Detecting High-Energy Cosmic Ray Electrons With CREST (the Cosmic Ray Electron Synchrotron Telescope)

by

Joseph C. Gennaro

Chair: Dr. Gregory Tarlé

The Cosmic Ray Electron Synchrotron Telescope (CREST) was a balloon-borne detector designed to measure the high-energy cosmic ray all-electron energy spectrum. It utilized a novel indirect technique proposed in 1983 by Stephens and Balasubrahmanyam to detect electrons by means of the synchrotron radiation they produce while interacting with the Earth's magnetic field. CREST took data during a 10-day circum-anti-polar flight originating near McMurdo Station, Antarctica from December 2011-January 2012. In this work I describe the instrument's design, assembly, operation, detection scheme, launch, flight and recovery, as well as my original work on displaying and analyzing the flight data. In particular I describe a novel detection method combining direct detection of the primary electron and indirect detection of the electron's secondary synchrotron photons. Further I characterize CREST's ability to determine the momentum direction of signal electrons, and outline a method for determining the charge of primary leptons making use of that pointing capability. I also propose an improved detector surface configuration suitable for future CREST-like detectors which greatly reduces sensitivity to the ubiquitous charged particle background.

CHAPTER 1

Production, Propagation and Current Measurements of High Energy Cosmic Ray Electrons

Cosmic rays provide a unique window for learning about the location and nature of their sources (acceleration by supernova remnants and pulsars, or via secondary production) and the type and amount of media through which they travel prior to their detection here on or near Earth (the Cosmic Microwave Background, the Inter-Galactic/-Stellar Medium (IGM/ISM) and the extra- and intra-galactic magnetic fields). Extensive measurements of the energy-dependent flux of primary and secondary protons and antiprotons, utilizing both direct and indirect detection methods, have been carried out. Those with the lowest energies, below around 100 MeV, have intra-solar system origins (by virtue of the fact that extra-solar low energy particles are deflected by the solar wind [16]). Those with the highest energies, up to approximately 10^{20} eV, may come from outside the Milky Way. This upper limit is enforced by the distant location of proton acceleration regions relative to the Earth, and the distance scale set by the dominant energy loss processes during propagation through the Cosmic Microwave Background. Although electrons and positrons (hereafter simply “electrons”) have different sources and energy loss processes compared to protons, the detectable flux of primary electrons is also theorized to be energy-limited by an analogous situation. Experimental confirmation of this limit remains out of reach, as an experiment capable of detecting electrons with energy at or beyond the cutoff region without very large amount of systematic error has not yet been conducted. While some ground-based air shower detectors (such as HESS) have probed the spectrum at around the TeV scale, low statistics, high systematic errors and disagreement between experiments only emphasize the need for a solution at TeV and higher energies. The highest energy measurements of the positron flux indicate point to the possibility of a largely, if not purely, secondary origin for electrons beyond a TeV or so (see Section 1.3.2.9).

In 1983, Stephens and Balasubrahmanyam proposed [63] a candidate method relying on indirect detection of cosmic ray electrons via the synchrotron radiation generated by their interaction with the Earth’s magnetic field. Because synchrotron radiation power increases with increasing primary electron energy, the sensitivity of an instrument utilizing this technique would (*ceteris paribus*) tend to increase with energy, rather than decrease. An experiment successfully implementing this method could probe the high energy cutoff region of the electron spectrum, and provide direction confirmation or rebuttal of current models for the production and propagation of high-energy electrons. The CREST (Cosmic Ray Electron Synchrotron Telescope) represents the first implementation of this method, attempting to measure for the first time the energy spectrum of galactic cosmic ray electrons beyond the theorized high energy cutoff.

1.1 Ultra-High-Energy Cosmic Ray Protons: The GZK Limit

In this section I review the analogous situation for cosmic ray protons, which also have a theoretical maximum detectable energy at Earth. This analogy has been noted before (for example by Müller [56]) and I review it here because it serves as an excellent model for how the high energy electron limit comes about and how it can be probed.

With the exception of several features (e.g. the ‘Knee’ around 4×10^{15} eV and the ‘Ankle’ around 3×10^{18} eV [6]), the cosmic ray proton energy spectrum falls off as a power law up to a high energy limit known as the Greisen-Zatsepin-Kuz’min (GZK) limit, posited in 1966 by Greisen [47] contemporaneously with Zatsepin and Kuz’min.

It describes the origin of the upper limit on detectable proton energies utilizing two key pieces of information: a) the onset of energy loss as protons exceed the energy threshold for pion production via interaction with the Cosmic Microwave Background (CMB) radiation, and b) the lack of intra-galactic sources capable of accelerating protons to or above this energy threshold required for such pion production.

The processes in question can be described by the following schematic equations. The Δ ‘resonances’ are baryons (composite particles composed of three quarks) with 3/2-spin rather than the 1/2-spin possessed by protons and neutrons:

$$\gamma_{CMB} + p \rightarrow \Delta^+ \rightarrow p + \pi^0 \tag{1.1}$$

or

$$\gamma_{CMB} + p \rightarrow \Delta^+ \rightarrow n + \pi^+. \quad (1.2)$$

Greisen estimated that the energy threshold to begin pion production for a proton against a photon background with mean energy 7×10^{-4} eV is 10^{20} eV. However since the photon energies are distributed according to the Planck distribution, some photons in the high energy tail would enable protons with a slightly lower energy to produce pions as well. With an average cross section of $200\mu\text{b}$ and a mean photon number density of 550 cm^{-3} , the average pathlength between such interactions is calculated as:

$$(n\sigma)^{-1} = 200\mu\text{b} \times 550\text{cm}^{-3} \approx 9 \times 10^{24}\text{cm}, \quad (1.3)$$

or about 3 Mpc (9.5 million light years). If we let E be the initial proton energy and ΔE the energy loss per interaction, the characteristic length for energy loss, L , becomes

$$L = \frac{E}{\Delta E}(n\sigma)^{-1}. \quad (1.4)$$

At an average value for $\frac{E}{\Delta E} \approx 0.22$, this gives $L \approx 2 \times 10^{24}$ cm, about 0.65 Mpc, far larger than the diameter of the Milky Way (30 kpc). For comparison the nearest major galaxy (Andromeda) is 0.78 Mpc distant. Any protons detected on Earth above this high energy limit must have been produced nearer Earth than this characteristic length to avoid losing a significant fraction of their initial kinetic energy.

While a proton accelerator within this energy loss distance could cause protons above the GZK limit to be detected here on Earth, such acceleration regions would have to either be larger than a galaxy or have very strong magnetic fields which are uniform over very large regions. Structures that satisfy these requirements, e.g. galaxy clusters or Active Galactic Nuclei, do not exist nearer to us than this energy loss distance. Thus protons above the GZK limit should not be detected near Earth. Figure 1.1 reproduces Figure 1 in [48], which summarizes some of these potential acceleration sites.

The GZK cutoff has been observed experimentally by the HiRes experiment [13] as shown in Figure 1.2, an indirect detector observing secondary air showers produced by ultra-high energy cosmic rays in the Earth's atmosphere.




Figure 1.1: Figure 1 (from [48]) showing the possible acceleration sites for ultra-high energy protons. As Hillas notes, “Objects below the diagonal line cannot accelerate protons up to 10^{20} TeV.”




Figure 1.2: Figure 2 (from [13]) showing the steep cutoff in the ultra-high energy cosmic ray proton spectrum at the energy predicted by Greisen, Zatsepin and Kuz'min in 1966.




Figure 1.3: Existing all-electron flux ($\times E^3$) measurements. We are just beginning to probe energies above one TeV. The Pamela spectrum is e^- only. Error bars are systematic+statistical, except for HESS which reported the shown systematic error band (green shaded area). HESS points also suffer from an undrawn 15% energy scale systematic uncertainty. Undrawn error bars indicate errors smaller than the marker size. [2][21][25][27][32][34]

1.2 Existing Measurements of the High-Energy Electron Spectrum

The electron flux data up to several hundred GeV displays high statistics and excellent agreement between experiments (above the low energy discrepancies, likely explained by solar modulation). This is not the case at slightly higher energies where high-statistics results from FERMI and AMS-02 display disagreement with the ATIC data. At even higher energies (approaching 1 TeV) results from ground array detectors such as HESS display a large amount of systematic error. While there is a hint of a cutoff in the HESS data, overall the existing high energy measurements are too uncertain, too sparse and too low energy to adequately test for the high energy cutoff predicted by current production and propagation models. Fully exploring the potential contribution from Vela, for example, would require more than a factor of ten increase of the maximum energy probed (see Figure 1.9).

A new technique that is more sensitive at higher energies is needed to adequately probe this energy range.

1.3 High Energy Cosmic Ray Electrons: Production and Propagation

While indirect evidence of electrons created with or accelerated to estimated energies up to 100 TeV has been detected near intra-galactic supernova remnants (see e.g. [7], [60]), such electrons rapidly lose a substantial part of their energy to interactions with the CMB as they travel through the galaxy. They can be detected at Earth only if emitted so near to us that they arrive before most of this energy loss occurs. In this section I will review the basis for the theoretical high energy cutoff for primary cosmic ray electrons. This basis consists of the limited number of nearby sources coupled with the rapid energy losses experienced by electrons in the Interstellar Medium (ISM). By comparing the (energy-dependent) energy loss scale length of cosmic ray electrons with the nearest and youngest sources, it is possible to predict the highest energy primary electrons that reach our Solar System. Beyond this effective limit the primary electron flux could fall to zero, with any remaining flux being of secondary origin (possibly as by-products of high energy protons - see e.g. [4]). Such a secondary flux would likely have a positron fraction closer to 50% than the 5-10% observed at lower energies (see Section 1.3.2.9 for details). Fully investigating the high energy cutoff requires measuring not only the all-electron spectrum, but also the separate electron and positron spectra. While CREST was originally designed to perform only the

former measurement, it may be possible to utilize magnetic field direction information in combination with an estimate of the primary electron trajectory to also extract the positron fraction. I discuss this possibility in detail in Section 6.8.

1.3.1 Cosmic Ray Electron Energy Loss in the ISM

Relativistic electrons lose energy at an average rate based on the photonic and magnetic energy densities through which they travel. Their initial energy, the distance they travel from source to observer and the particular magnetic and photonic conditions they experience during their journey affect the energy they have when observed. The following discussion on electron energy loss employed by Kobayashi *et. al.* to model the energy loss of primary electrons produced in supernova remnants will prove equally useful when evaluating the potential flux from any hypothesized source object or process.

1.3.1.1 The Electron Energy Loss Rate

The dominant energy loss processes for high energy cosmic ray electrons in the ISM are inverse Compton scattering off of CMB photons and synchrotron emission produced during interaction with the intra-galactic magnetic field.

Kobayashi *et. al.* [38] employ the following energy loss rate equation for electrons above 10 GeV during intra-galactic propagation:

$$\frac{dE}{dt} = -bE^2, \quad (1.5)$$

with

$$b = \frac{4\sigma c}{3(mc^2)^2} \left(\frac{B^2}{8\pi} + \mathcal{W}_{ph} \right). \quad (1.6)$$

E and m are the electron energy and mass, B is the galactic magnetic field strength, \mathcal{W}_{ph} is the interstellar photon energy density, σ is the scattering cross section, and c is the speed of light. Easily identifiable are the synchrotron (magnetic field energy) and Compton (photonic field energy) contributions to the energy loss. They assume an average perpendicular magnetic field strength of $B_{\perp} = 5 \mu G$. At low energies σ can be approximated by the Thomson cross section, while at high energies this gives way to the Klein-Nishina energy-dependent Compton cross section.

These processes gradually increase in power rather than experiencing a sharp increase past a particular threshold energy (as is the case for protons producing pions against the CMB). Importantly, as is shown in Figure 1.4, the energy loss coefficient drops as energy increases, but not faster than E^2 rises; thus the overall rate at which the electron loses




Figure 1.4: Figure 1 (from [38]) showing the energy loss coefficient from Equation 1.6. As energy increases, this coefficient drops, but not faster than E^2 rises; thus the overall rate at which the electron loses energy increases with energy.

energy (the loss power) increases with energy. This means one can calculate an average electron diffusion distance as a function of energy. To put it in terms most comparable to the situation for protons, we can compute a characteristic energy loss distance L for electrons as a function of electron energy.

1.3.1.2 The Average Electron Diffusion Distance

Integration of Equation 1.5 yields ([38])

$$-\frac{dE}{E^2} = bdt, \text{ and} \quad (1.7)$$

$$\frac{1}{E_f} - \frac{1}{E_0} = b\delta t. \quad (1.8)$$

At time $\delta t = \frac{1}{bE_0}$, one finds that $E_f = E_0/2$, or that the electron's 'energy half-life' is proportional to $\frac{1}{bE}$. Encapsulating the energy dependence of the cross section inside the variable b (which varies with energy as in Figure 1.4), one has

$$\delta t = \frac{1}{bE} \approx 3.03 \times 10^5 \text{ yr}/E, \quad (1.9)$$

when E is measured in TeV. Roughly speaking, we can convert this ‘half-time’ into a ‘half-length’ using the diffusion relation:

$$R = (2D\delta t)^{1/2}. \quad (1.10)$$

Using a diffusion constant¹ of

$$D_{xx} = (2 - 5) \times 10^{29} \left(\frac{E}{\text{TeV}} \right)^{0.3} \text{ cm}^2\text{s}^{-1}, \quad (1.11)$$

we can convert a half-time into a rough approximation of the average distance traversed by an electron with energy E at Earth as

$$\bar{R} \approx \left((4 - 10) \times 10^{29} \left(\frac{E}{\text{TeV}} \right)^{0.3} (bE)^{-1} \text{ cm}^2\text{s}^{-1} \right)^{1/2} \quad (1.12)$$

I plot this function of average diffusion range against energy in Figure 1.5.

1.3.2 Galactic Sources of High Energy Cosmic Ray Electrons

Here I review some candidate sources of high energy cosmic ray electrons and positrons in our Galaxy. I first review the most widely accepted candidate source for high energy electrons: diffusive shock acceleration. CREST will contribute to testing these models by placing limits on the combined flux of high energy electrons and positrons. Recent measurements of the higher-than-expected positron fraction at high energies by the Payload for Anti-Matter Exploration and Light-nuclei Astrophysics (PAMELA) and AMS-02 have resulted in numerous models attempting to explain the increased positron flux. Measurements of the combined electron and positron flux greatly constrain these models. CREST has the capability to participate in these efforts by providing the highest-energy measurements of the positron fraction. Therefore, I also review potential explanations for this phenomenon: Dark Matter particle (WIMP) self-annihilation, pair production by nearby pulsars and a re-examination of the possibility that the positron fraction is due to the secondary flux of positrons produced by the well-measured high energy population of nearby cosmic ray protons.

¹Note that Moskalenko and Strong use a spatial diffusion coefficient of $D_{xx} = 3 - 5 \times 10^{28} \text{ cm}^2\text{s}^{-1}$ in their review on the subject [1]; however, Kobayashi *et. al.*'s estimate for D takes into account its energy dependence when estimating the number specifically for TeV electrons, while Moskalenko and Strong's value was more typical for all cosmic ray species and energies.

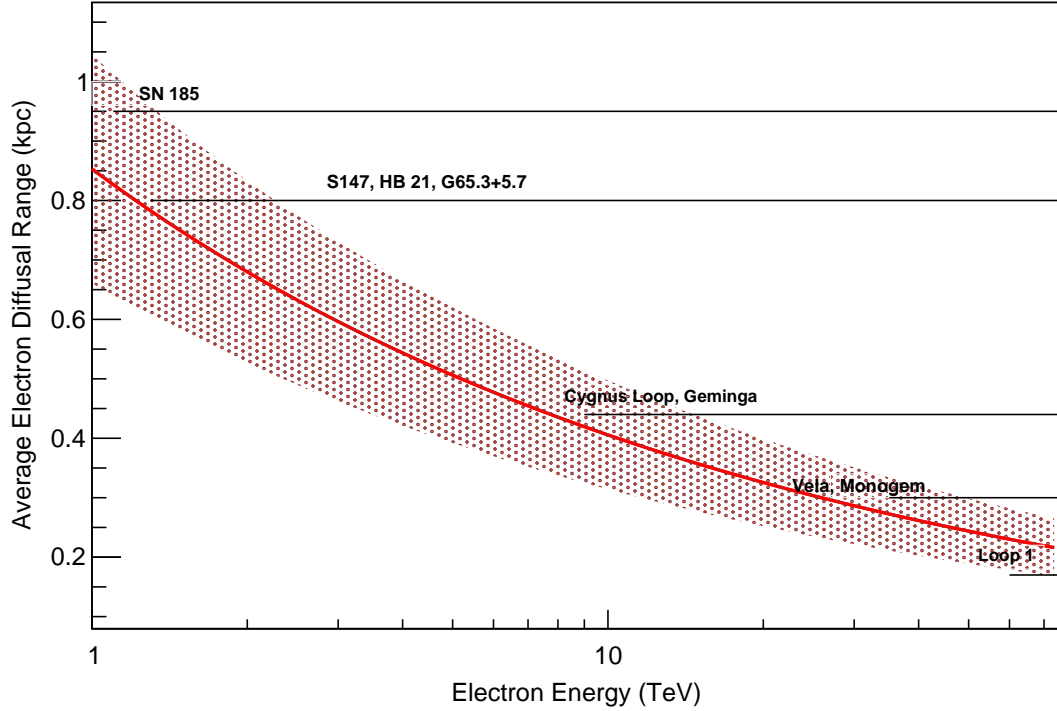


Figure 1.5: The average electron diffusional range as a function of electron energy. As energy increases, average range decreases. The vertical envelope at each energy represents the range of choice of diffusion constant D_{xx} as in Equation 1.11. Also shown are the distances of several SNRs identified by [38] as near and young enough to produce observable electron fluxes under various production assumptions. By confining one's self to energies above a certain value on the X-axis, one detects flux from no further away than the corresponding value on the Y-axis. For example if we were to observe a large amount of 10 TeV electrons, we could exclude them as being from SN 185, S147, HB 21, and G65.3+5.7, but not from the Cygnus Loop, Geminga, Vela, Monogem or Loop 1 remnants (not on energy loss considerations alone - Loop 1 for example has a maximum theoretical energy of 1.2 Tev. See Table 1.3.2.5 for details.). I generated this plot using Equations 1.10, 1.11, the energy loss coefficient in Figure 1.4 and the distances of supernova remnants in Table 1.3.2.5 (all from [38]).

1.3.2.1 Acceleration Processes: General Considerations

Many cosmic ray acceleration processes appeal to a massive (compared to an individual particle) structure imparting momentum to individual particles via repeated interactions. The exact nature of the momentum source is often abstracted away and therefore similar analysis applies to the momentum imparted to individual particles by wandering magnetic mirrors in the interstellar medium, expanding supernova shock fronts, the solar wind impinging on a planetary magneto-sphere or jets of hot material escaping an active galactic nucleus.

At least three factors strongly affect the types of particles processes can accelerate and the energies they can accelerate them to:

1. There must be a particle injection process; i.e., a means by which the to-be accelerated particles are injected into the region where the process will take place. In some cases the output energy spectrum of a process depends strongly on the spectrum of the injected particles; in others the injection is of a purely thermal nature and the output is independent of the thermally random initial particle energies.
2. The process must impart energy to the particle faster than the particle loses energy. Often this is an energy-dependent competition and the energy at which this ceases to be true enforces the upper limit of the spectrum of accelerated particles.
3. The accelerated particle must be capable of leaving the acceleration region unimpeded. For example, an acceleration region surrounded by (or consisting of) a medium with high magnetic field intensity may produce high energy electrons, but they might not escape the region of high magnetic field without severe energy loss.

Point two above is often referred to as the efficiency of the acceleration process. Processes with more frequent or more powerful acceleration events and less frequent and powerful deceleration events have a higher efficiency. Since electrons tend to lose energy faster (especially in the presence of strong magnetic fields) than protons and heavier ions, they require more efficient accelerators than do protons and heavier ions. This rapid energy loss elevates point three to critical importance for electrons escaping supernova remnants, as the same magnetic fields which contain and accelerate electrons also threaten to radiate large fractions of their energy away as synchrotron radiation.

1.3.2.2 Fermi Acceleration

In 1949, Fermi proposed [41] that particles in the ISM would occasionally encounter “wandering magnetic fields” produced by the streaming motions of the (largely ionized) inter-

stellar hydrogen gas. Kinematically, head-on collisions would be more frequent than tail-on collisions (since particles see a slightly greater flux of approaching rather than retreating mirrors) by a factor equal to $B = V/c$, where V is the peculiar velocity of the field irregularity (about 30 Km/sec). If each head-on reflection increases the fractional particle energy by a factor of B and each tail-on interaction decreases fractional particle energy by the same factor, then after many such reflections the averaged fractional energy gain would be proportional to B^2 :

$$\delta w = B^2 w, \quad (1.13)$$

where w is the total (kinetic + rest) energy of the supra-thermal particle. After N such collisions, the particle would have an energy of order

$$w = M c^2 e^{B^2 N}. \quad (1.14)$$

Considering protons in particular, Fermi postulated that the main mechanism for termination of this process would occur when the high energy proton underwent a nuclear collision, losing enough energy that the outgoing proton would be below the injection energy threshold. If the mean lifetime against such a collision for a proton in the interstellar medium is T , then if such particles are supplied at a constant rate, Fermi estimated the current lifetime of the particles observed today would be $e^{-t/T} dt/T$. Convolving the energy gained during that time via the above reflection process with this spectrum of lifetimes, Fermi found that the probability π of observing a particle with energy between w and $w + dw$ to be

$$\pi(w)dw = \frac{\tau}{B^2 T} (M c^2)^{\frac{\tau}{B^2 T}} \frac{dw}{w^{1+\frac{\tau}{B^2 T}}} \quad (1.15)$$

Of particular note is the way the power law behavior ($\pi \sim w^{-(1+\tau/B^2 T)}$) appears as a consequence of the competition between discrete accelerations (reflective scattering events with a period τ) and catastrophic termination (via nuclear interaction with a lifetime T). Matching this result with the observed spectral index (Fermi used $k = 2.9$; today the power law exponent is estimated to be closer to $k = 2.7$ over much of the all-particle cosmic ray energy spectrum) requires fine tuning of the relative lifetimes of the competing processes such that $\tau = (k - 1)B^2 T$. Such a relationship depends on the particular character of acceleration and termination processes envisioned by Fermi, the lifetimes of which would be expected to vary with energy, while the spectral index holds over a very large energy regime; in other words, the validity of this process as an explanation for the cosmic ray spectrum observed at Earth could only be, as Blandford and Eichler put it, “...some chance coincidence between the escape and acceleration timescales.” [16]

Most relevant to our present aims, Fermi noted that such a process would not accelerate primary electrons faster than they lose energy, since their energy loss from synchrotron and inverse Compton scattering would always outweigh the energy gained by random encounters with wandering magnetic mirrors. A simple calculation using our above energy loss characteristics shows his intuition to be true. The energy gain per collision is independent of the scattered particle mass. Rather it is a function of B , the scatterer velocity, β , the scattered particle velocity, and θ , the angle made between β and the orientation of the magnetic field line (see equation 13 and associated discussion in [41]). Therefore, it is valid to use the same acceleration time constant τ for electrons as for protons when making these order of magnitude estimates. Matching τ with the known spectral index for protons, Fermi calculated an average distance traveled by protons between scattering events of 10^{18} cm or about a third of a parsec. While the spectral index for electrons does vary slightly from that of other species, the difference does not make any significant impact on this estimate.

We can use equation 1.14 to estimate the number of scattering events required to accelerate a non-relativistic electron to 1 TeV:

$$e^{B^2 N} = \frac{10^{12}}{0.511 \times 10^6} \approx 2 \times 10^6. \quad (1.16)$$

Using (as Fermi did) $B = 10^{-4}$ for the wandering magnetic mirrors this implies that the number of such collisions $N \approx 1.45 \times 10^9$ and that the average total distance traveled is over 435 kpc (over ten times the diameter of the galaxy). A TeV electron loses most of its energy after diffusing a distance of about 0.85 pc (see Figure 1.5). While not all of this distance would be traveled while the electron had energy in the TeV range, the order of magnitude of energy loss is still too great for this mechanism to be a diffuse source of high energy electrons. Therefore we are motivated to seek a more efficient acceleration mechanism (that is, one with either a smaller distance between scattering events or a far higher energy gained per scattering event).

1.3.2.3 Shock Acceleration

For our purposes, a shock is any disturbance in an astrophysical gas which can accelerate a subset of the gas to a non-thermal velocity distribution. Defined this way, ‘shock’ can refer to the bow shocks produced in the solar system by the interaction between the solar wind and planetary magnetospheres, by the expanding front of a supernova explosion, the active region of matter falling into a black hole (quasars), active galactic nuclei and other sources. As we are most interested in the output of this process, that is the energy spectrum of the non-thermal particles, I shall review a simplified formulation from [59], pages 46-

56, outlining how a power law spectrum for the energies of the non-thermal particles arises with a spectral index depending on basic parameters of the shock, rather than more detailed information about acceleration event frequency or particle loss rates (as was the case for Fermi's process).

Consider a planar, thin (compared to the mean free path of gas particles), non-relativistic thermodynamic discontinuity (a plane shock) moving through an ideal gas with velocity U in the x direction. Define the region which the shock is traveling towards as Region 1, and the region from which the shock is moving away as Region 2. By requiring momentum and energy conservation across the shock front, one finds that in the limiting case of a strong shock (where the shock velocity is very high compared to the speed of sound in the gas) the ratio of the velocity of the unshocked gas (v_1) to the shocked gas (v_2) is given by $\frac{v_1}{v_2} = 4$. In the rest frame of the shock itself $v_1 = U$ and $v_2 = U/4$.

Let E and $p_x = E/c$ be the energy and momentum in the direction of U of a relativistic gas particle (perhaps from the tail of the Maxwell-Boltzmann distribution describing velocities of particles in thermal equilibrium). Further, let θ be the angle off normal of the particle's velocity relative to the direction of motion of the shock U . Suppose this particle diffuses from Region 1 into Region 2, across the shock, and interacts with a shocked particle from Region 2. Assuming U is small enough so that the shock has a Lorentz factor of approximately 1, the crossing particle will be imparted an energy $\Delta E = p_x v \cos \theta$, where $v \equiv |v_1 - v_2| = 3U/4$. This results in a fractional change in energy of

$$\frac{\Delta E}{E} = \frac{v}{c} \cos \theta, \quad (1.17)$$

which is of the order of magnitude of the velocity of the scattering mirror, hence the appellation 'first-order Fermi acceleration'. The term 'diffusive shock acceleration' is also frequently used when the mechanism for repeated exposure of the particle to the shock front is diffusion ([20]). This arises because the assumption that the particle is being swept up by the shock front guarantees a head-on collision. Contrast this to the situation described by equation 1.13, where the average fractional energy change is proportional to the square of the scatterer's velocity since some of the collisions result in an energy loss rather than gain (thus earning Fermi's original conception the appellation 'second-order Fermi acceleration'). These particles will then interact with the shocked gas in Region 2; a fraction of them will diffuse back across the shock boundary. Such re-diffusing particles will see a thermal gas approaching with relative velocity $3U/4$ and once more experience a fractional energy change $\frac{\Delta E}{E} = \frac{v}{c} \cos \theta$. Since these particles are still relativistic (having only increased energy) and a fraction of them will again diffuse across the shock boundary, this

process could, for a kinematically-determined fraction of the particles, repeat indefinitely with the same relation holding each time. Key to this step of the process is that the non-thermal particles cause no back-reaction on the shock or the material around the shock - this simplifying assumption is often referred to as the *test particle* assumption.

The particle's energy after j such crossings (beginning with energy E_0) can be written as $E = E_0\beta^j$, where $\beta \equiv 1 + v/c$, the multiplicative factor by which the energy changes upon each diffusion across the shock.

This process requires far less time and distance traveled on the part of the accelerated particle (assuming the shock boundary is very thin compared to the mean free path of gas particles) than does Fermi's process (which required protons to move a third of a parsec, on average, between interactions with wandering magnetic mirrors). In the simplest form it is assumed that catastrophic interactions (such as nuclear collisions for protons, or significant energy loss via synchrotron or inverse Compton scattering for electrons) occur on a timescale much larger than that characterizing repeated diffusion across the shock front.

Suppose that after j crossings, $N = N_0P^j$ particles remain from the initial N_0 relativistic particle population. Combining the relations for $E(j)$ and $N(j)$ yields:

$$\frac{N}{N_0} = \left(\frac{E}{E_0} \right)^{\frac{\log P}{\log \beta}}. \quad (1.18)$$

In differential form this results in a power law for the distribution of energies:

$$n(E) \propto E^{\log P / \log \beta - 1} \propto E^{-k} dE, \quad (1.19)$$

with $k = 1 - \log P / \log B$. Similar to second-order Fermi acceleration, it appears that the precise value of k will again depend on the nature of the competition between the acceleration (β) and escape (P) timescales. To compare these timescales, Rosswog and Brüggén estimate them as follows.

If the gas were at rest with respect to the shock, and the gas scatters particles isotropically, the solid angle for scattered particles to recross the shock would be half of a complete sphere, or 2π steradians. If F represents the total flux of particles recrossing the shock front, then in terms of the average velocity $\langle v \rangle$ of the diffusing particles $F = n\langle v \rangle/4$. If we also assume that these particles are relativistic, $\langle v \rangle \approx c$, so $F \approx nc/4$. In the rest frame of the shock, upstream (already shocked) particles are advected away from the shock with a bulk velocity $U/4$. The probability that the relativistic particles will avoid being carried away

by the bulk flow $P = 1 - \frac{nU}{4} / \frac{nc}{4}$, so that

$$\log P = \log \left(1 - \frac{U}{c} \right) \approx -\frac{U}{c}. \quad (1.20)$$

(The approximation in the last step is valid when the shock is non-relativistic so that $U \ll c$.) Note that in second-order acceleration the loss rate is a parameter of the process set independently from the acceleration rate; in this case the loss rate is of the order of magnitude of the relativistic velocity of the acceleration zone, a basic parameter of the shock itself.

All that is left (for this simplest of cases) is to estimate β , the factor by which energy increases after each crossing. The energy gain per shock crossing averaged over all particle crossing angles θ is given by

$$\left\langle \frac{\Delta E}{E} \right\rangle = \frac{v}{c} \int_0^{\pi/2} 2 \cos^2 \theta \sin \theta d\theta = \frac{2v}{3c}. \quad (1.21)$$

A round trip produces twice this gain, so $\beta = \frac{E}{E_0} = 1 + \frac{4v}{3c}$. Combined with the result that $v = 3U/4$ (for strong shocks), this gives

$$\log \beta = \log \left(1 + \frac{U}{c} \right) \approx \frac{U}{c}. \quad (1.22)$$

Armed with P and B we can now calculate k , the spectral index of relativistic particles produced by the shock as

$$k = 1 - \frac{\log P}{\log \beta} \approx 2. \quad (1.23)$$

For this simple case the spectral index, to first order, does not depend on the detailed nature of the acceleration and loss processes.

1.3.2.4 Extensions to the Simple Model

As my main interest in reviewing shock acceleration is to demonstrate the importance of measuring the high energy electron cutoff I conclude this section by mentioning several extensions to this most basic model and provide some useful references for further investigation by the interested reader.

The fact that the overall cosmic ray spectral index is not 2 but 2.7 indicates that we have ignored some aspects of shock acceleration and particle propagation that would tend to steepen the spectrum. The ‘leaky box’ model wherein particles escape the galaxy with a lifetime of approximately $E^{-0.7}$ is often invoked to explain the discrepancy, although it

predicts a steeper spectrum at higher energies than is observed [49].

Also interesting is modifying the assumption that the shock is so thin that the diffusion time is small compared to the timescale for other processes, such as radiative losses; the obvious candidate for electrons in particular is synchrotron losses and the relevant timescale is again the synchrotron loss time. Drury *et al* address [20] this complication by estimating the synchrotron power (in this case, momentum loss per unit time) as a function of electron momentum p :

$$\langle \dot{p} \rangle_s = -Dp^2, \quad (1.24)$$

$$D = (m_e c \tau_r)^{-1} \text{ and} \quad (1.25)$$

$$\tau_r = 6\pi m_e c / (\sigma_T B^2), \quad (1.26)$$

where m_e is the electron rest mass, σ_T is the Thomson cross section, B is the magnetic field strength and τ_r is the characteristic synchrotron loss time for electrons with momentum $m_e c$. (Note τ 's similarity with the analogous constant from [38] in Equation 1.6.)

The net energy gain per round trip is then modified to account for synchrotron losses by subtracting the momentum lost on either side of the shock front

$$(\Delta p)_s = -(D_1 p^2 \Delta t_1 + D_2 p^2 \Delta t_2), \quad (1.27)$$

where t_i is the dwell time of the electron on the i -th side. This highly energy-dependent loss process ($\propto E^2$) preferentially extinguishes high energy electrons, steepening the spectrum at higher energies and potentially providing an upper energy limit for electrons that would be different from protons and heavier ions. Interestingly, the total number of supra-thermal particles is unaffected; instead synchrotron losses “merely redistribute the electrons downward in energy” [20]. They estimate the cut-off energy (where synchrotron losses approximately balance the energy gained by crossing the shock) as approximately 10 – 100 TeV (about five to ten times higher than the cutoff estimated in [38], and roughly consistent with the maximum energies observed in SN1006 [60]).

Reynolds and Keohane make use of this formalism to explain the observed steepening in the photon spectrum from supernova remnants somewhere between the radio and x-ray components [58]. They assume that the cutoff to the power law is synchrotron-enforced, and estimate E_{max} , the maximum energy to which 14 radio-bright shell remnants could accelerate electrons. The inherent sharpness of this type of cutoff (due to the E^2 dependence of synchrotron power) allows for the highest possible E_{max} without violating the observed x-ray-to-radio flux ratio. Using observations of these ratios they estimate that the supernova remnant Kes 73 could accelerate electrons up to 200 TeV. However, Kes 73 is unusual

in that it has an x-ray bright, radio quiet central source [46]².

Another way to extend the basic shock framework is to drop the ‘test-particle’ assumption and incorporate the non-thermal component of the particle population into the model describing the behavior of the shocked and unshocked fluids. This produces a non-linear model, as the shock affects the non-thermal particle population, which then affects the shock itself, etc. Section 4 of Drury’s review on the topic [19] discusses the back-pressure induced by the accelerated particles in the strong (faster than sound speed) shock regime.

Bell [12] examined amplification of a magnetic seed field by the accelerated particles and its effect on the maximum attainable energy for particles undergoing diffusive shock acceleration in supernova remnants. Assuming the interactions are mediated by magnetic fields (‘Bohm diffusion’), the accelerated particle would have approximately unit probability of escape when its gyroradius exceeds the size of the acceleration zone itself. Since the gyroradius depends on (specifically, it is inversely proportional to) the magnitude of the perpendicular component of the magnetic field, if the supra-thermal component enhances the background magnetic field, it can effectively help contain itself to the acceleration region. Bell also demonstrated that the growth rate of the field amplification is highly dependent on the velocity of the shock, thereby raising the possibility that higher energy components of the cosmic ray spectrum may derive from faster, rather than simply younger, shocks [12] (although these two variables are in general correlated, for example supernova remnant shock waves slow down as they sweep up more matter). Of course, such a process might prove self-limiting for electrons since the same field amplifications serving to engender their production would also amplify synchrotron energy loss.

1.3.2.5 Candidates for High-Energy Electron Sources: Supernova Remnants

Supernova are “explosions of dying stars” which unleash a total of about 10^{51} ergs of energy [16], approximately 10^{48} ergs of which are estimated to be directed into particles with individual kinetic energy of above 1 GeV ([38] and references therein). The above discussion on shock acceleration is relevant to two separate acceleration zones, one or both of which could be present in various remnants depending on the supernova type and on the possible formation of a compact object (e.g. a neutron star which may or may not also be a pulsar, depending on magnetic field conditions). The first process, characteristic of shell-type remnants, consists of the shock created by the ejecta of the explosion colliding with the interstellar medium surrounding the star prior to explosion. SN 1006 is thought to be of this type (see Figure 1.7). The second process, characteristic of remnants containing a pulsar,

²At a distance of approximately 7 kpc [46] it is not a potential source of CR electrons above a few GeV at Earth.

SNR	Distance (kpc)	Age (yr)	E_{max} (TeV)
SN 185	0.95	1.8×10^3	1.7×10^2
S147	0.80	4.6×10^3	63
HB 21	0.80	1.9×10^4	14
G65.3+5.7	0.80	2.0×10^4	13
Cygnus Loop	0.44	2.0×10^4	13
Vela	0.30	1.1×10^4	25
Monogem	0.30	8.6×10^4	2.8
Loop 1	0.17	2.0×10^5	1.2
Geminga	0.4	3.4×10^5	0.67

Table 1.1: Nearby remnants of supernovae (both shell and compact-object types), according to [38].

consists of the shock that occurs as the gas excited by the pulsar’s magnetic field collides with a companion object or the interstellar medium. Vela is thought to be of this type (see Figure 1.8). The ~ 0.4 kpc distant supernova Geminga, estimated to have occurred about 300,000 years ago, endowed our neighborhood with a γ -ray emitting, 230 ms period pulsar which may be capable of producing electrons and positrons by means other than shock acceleration (see Section 1.3.2.6).

The synchrotron x-ray emission from electrons with an estimated energy of over 10 TeV detected originating from SN 1006 ([7] [60], see Figure 1.7) shows that supernovae remnants in our galaxy can accelerate electrons to very high energies. What is not known is whether these electrons escape to the galaxy at large, and once having escaped, whether they can reach Earth with their energy largely intact. While electrons from SN1006 are too distant to produce a flux at those very high energies at Earth, determining whether we can see electrons with these energies originating from closer sources would provide important clues as to the detailed structure of the expanding shock wave escaping the star’s explosion center.

The Vela supernova remnant at an estimated distance of 0.290 kpc to 0.300 kpc ([24],[38]) is likely close enough for at least some of the electrons it accelerates to TeV scales to reach Earth. Figure 1.8 shows photon spectra measured by the High Energy Stereoscopic System (HESS), a set of four ground-based Imaging Atmospheric Cherenkov Telescopes located in Namibia. Assuming these γ -rays are the result of inverse Compton-scattering, the best fit electron flux between 5 and 100 TeV is a broken power law with a pre-break index of -2.0, a break energy of 67 TeV and a post-break index of -9.0 [24]. There is some doubt as to whether these high energy photons are in fact inverse Compton scattered by electrons of these energies and not the decay products of π^0 “produced by inelastic collisions between CR protons and ambient thermal nuclei” [57].

Even if these TeV photons are in fact up-scattered by electrons and the decay products




Figure 1.6: Higher energy electrons must come from nearer and younger sources. From [38], shows flux contours as a function of distance and age of theoretical source SNR, as compared to the measured source and age of various actual SNR. Each group of three lines represents a flux contour with the given number of electrons (in units of $(E/\text{GeV})^3 J \text{ GeV}^2 \text{ m}^{-2} \text{ s}^{-1} \text{ sr}^{-1}$). Within each group of three lines, the solid line represents no energy cutoff (that is, the spectrum is a pure power law), the long-dashed line represents $E_c = 20 \text{ TeV}$, and the short-dashed line represents $E_c = 10 \text{ TeV}$. As noted in the figure, this assumes a diffusion coefficient on the higher end of the usual range; for lower values of the diffusion coefficient, a given flux contour would move down in distance (but not in age).

of protons, these electrons must still escape the remnant and propagate through the ISM before reaching us. Ohira *et al* investigated the question of whether such electrons would lose their energy to synchrotron emission *after* acceleration but *before* escape from the remnant when the diffusive shock acceleration is mediated by Bohm (magnetic interaction) Diffusion of relativistic particles. They conclude that while synchrotron cooling does limit maximum electron energy during an intermediate phase of SNR evolution, it has little effect on the maximum energy of escaped electrons, because the escape time is approximately equal to the age of the SNR ([57], Section 2.6). Rather the maximum escape energy is strongly limited by their spatial diffusion, which is also a sensitive function of the magnetic energy density (see e.g. Equation 1.6). They conclude that while electrons up to 50 TeV can be produced in SNRs, the maximum energy of electrons that escape is lower than 50 TeV. How much lower depends on the detailed time-evolution of the magnetic field, information to which we do not yet have access. In an important sense, measurements of electron flux from SNR can therefore be thought of as one of our best and only methods for discovering this kind of information. This also means that there is a vast unexplored range of interesting energies, all the way up to 50 TeV.

In this model the maximum proton and heavier nuclei energy is anti-correlated with the maximum electron energy. This occurs because nuclei acceleration time is largely limited by the SNR age, while maximum nuclei energy increases with the acceleration power, which itself is a function of magnetic field strength. Unlike electrons, their escape is not inhibited by a magnetic field that persists into the escape time scale owing to their larger mass. In this manner, magnetic fields that persist into latter stages of remnant development – perhaps including those arising from the non-linear magnetic field amplification (driven by electrons, protons and nuclei) mentioned above – preferentially prohibit the escape of electrons compared to protons and nuclei. Direct measurement of the electron spectrum near these cutoff energies, modulo propagation through the ISM, would provide valuable insight into the magnetic field evolution which is difficult to measure by other means.

There is therefore a detailed inter-dependence between mass-dependent particle escape time and magnetic field time evolution. This suggests that the sources of the highest energy cosmic ray electrons share one or more common features with regards to magnetic field behavior at late times. Specifically, these sources have strong magnetic fields that either decay rapidly after a certain time shorter than the escape time of the electrons, or have anomalous “burst release” periods. These bursts would be characterized by abnormally low magnetic energy density, or abnormally large regions of aligned magnetic fields, permitting electron escape along field lines. One example of this is the polar gap model discussed in Section 1.3.2.8. Another (as far as I can tell, unexplored) possibility for burst electron




Figure 1.7: From [60], a spatial map of non-thermal x-ray emission between 2 and 4.5 keV from SN1006 produced from a compilation of data collected by the *XMM-Newton* instrument. The energy of the x-rays combined with estimates of the magnetic field in the acceleration region implies a maximum electron energy of about 25 TeV. Color image retrieved from http://xmm.esac.esa.int/external/xmm_science/gallery/images/sn1006_4images.jpg on 2/22/2015.

escape could be pulsar “starquakes,” when rapid deviations in the star’s molecular iron crust briefly disrupt its otherwise clockwork-like rotational period.

Kobayashi *et. al.* [38] present a unified treatment of many of the above factors, combining the power law output of diffusive shock acceleration models with synchrotron cooling during acceleration and delayed escape of electrons from acceleration zones. They attempt to reflect the range of parameters in the literature for values that have uncertain theoretical values. For example [57] estimate $t_{escape,e}$ to be in the range 9.70×10^2 to 7.39×10^3 years after the initial explosion, while Kobayashi *et al* model the release as occurring in bursts at 0.5×10^3 , 1×10^4 , 5×10^4 and 1×10^5 years after the explosion. Modifications to the dwell time could either reduce or increase the maximum escape energy depending on whether the retained electrons gain or lose energy to the competing processes of heating (by magnetic diffusion) and synchrotron cooling, a competition which is inherently energy dependent. Their model uses the estimated location of known nearby SNRs (summarized in Table 1.3.2.5) to predict the flux at Earth parameterized by the source electron energy cutoffs. One figure presenting their results with $D_{xx} = 5 \times 10^{29}$ is reproduced in Figure 1.9.




Figure 1.8: Figure 4 from [24], showing reconstructed synchrotron emission spectra (dashed lines at low energy) for three magnetic field strength possibilities given the observed inverse Compton emission spectrum (solid line and data points at high energies) for electrons with energy between 5 and 100 TeV accelerated by a γ -ray source near the Vela pulsar.




Figure 1.9: From [38], shows potential contributions to the high energy all-electron spectrum from various SNR overlaid on the state of the measured spectrum in 2004.

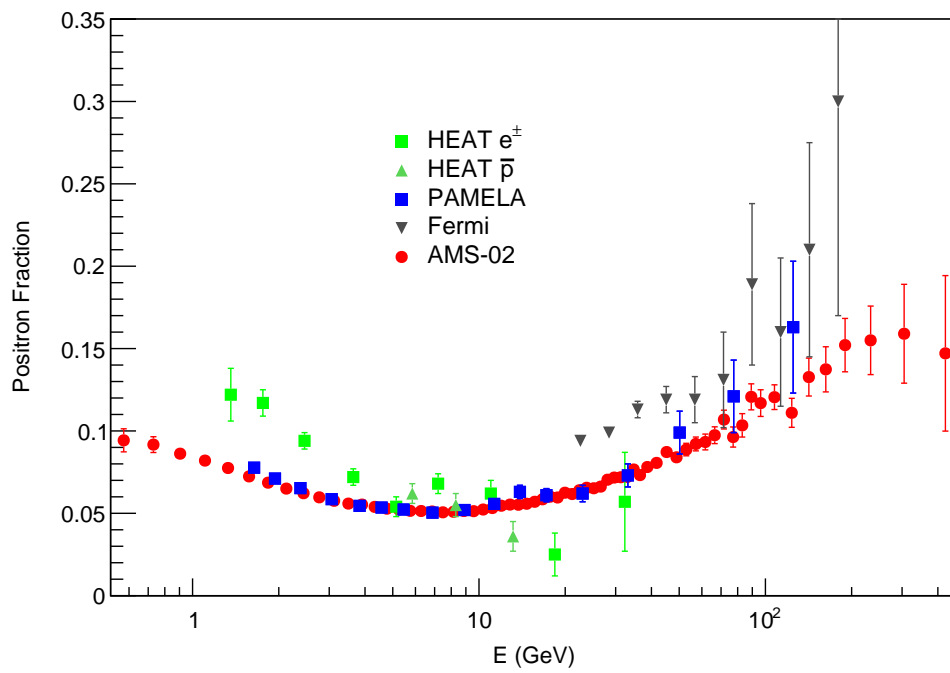


Figure 1.10: Recent measurements of the positron fraction $e^+/(e^+ + e^-)$. HEAT e^\pm from [37]. HEAT \bar{p} from [22]. PAMELA from [35]. Fermi from [28]. AMS-02 from [26]. Errors are systematic plus statistical and are drawn within the marker when smaller than the size of the marker.

1.3.2.6 Explaining the Increasing Positron Fraction: Primary or Secondary Positrons?

Data from the High Energy Antimatter Telescope (HEAT), PAMELA, Fermi-LAT and AMS-02 experiments (shown in Figure 1.10) reveal that the positron fraction rises from approximately 5% near the solar modulation region below a few GeV to approximately 15% starting at almost 100 GeV. Various candidate mechanisms have been proposed for increasing the flux of positrons relative to negative electrons. Here I briefly describe two local high energy, positron-rich production (Dark Matter particle self-annihilation) or acceleration (pulsars) mechanisms as well as the third possibility, that cosmic ray protons undergoing nuclear collisions with the ISM producing electron and positron pairs (among other products) via electromagnetic showers become the dominant flux contributor at higher energies. Though varied in nature, these mechanisms are relatively more ‘neutral’ compared to shock acceleration of negative electrons; that is, the positron fraction of particles accelerated or produced by these mechanisms would tend towards 50%³. As the share of the overall total electron flux from ‘neutral’ sources increases (as compared to the ‘negative’ sources described above which accelerate mostly negative electrons), the positron fraction should also increase from the 5-10% measured below 10 GeV.

Resolution between these candidate mechanisms may come from precise measurement of one of two features. First, Dark Matter particle self-annihilation and pulsar production should display a cutoff energy, in the former case determined by the unknown DM particle mass, and in the latter case determined by the detailed (and difficult to measure) nature of the electromagnetic processes leading to eventual production and escape of electron and positron pairs. This cutoff energy can be overshadowed by the cutoff enforced by the energy loss distance but there is significant phase space left between the maximum energy measurement of the positron fraction at 450 GeV and the energies where this transition to a propagation-energy-loss dominated regime would occur (somewhere above several TeV - see Figure 1.4). Secondary production of positrons by cosmic ray protons should extend well past the currently measured energies and should therefore only experience the energy-loss cutoff.

Second, there might be a detectable anisotropy in the arrival directions of the flux of positrons from DM/pulsars if the flux is dominated by discrete, nearby sources. To the extent that self-annihilating DM particles are evenly distributed throughout nearby space their flux should display isotropy. Alternatively, if DM particle density or interaction cross-section enhancement in certain locations is invoked (which is required in virtually

³Protonic pair production is not in fact neutral due to unequal amounts of electrons and positrons produced by the asymmetric decay of fully polarized secondary muons. Moskalenko and Strong [55] employ a ratio of 0.4 for the flux of secondary electrons to secondary positrons.

all models to achieve the required flux), one would expect an accompanying positron flux anisotropy which conflicts with the lack of anisotropy in other signal channels. The most typical such channel is γ -ray photon emission, which should be observed in the same region of the sky as enhanced DM particle density. In contrast, secondary pair production via nuclear collisions would display the same high degree of isotropy characteristic of locations of collisions between the parent protons/nuclei and the ISM target hydrogen atoms.

CREST has the capability to speak to the first point above, but probably not to the second. Although it was not designed to do so, CREST could measure the positron fraction at higher energies than AMS-02 has probed. In Section 6.8 I describe a technique capable of performing electron charge discrimination with a CREST-like instrument and its current limitations.

1.3.2.7 Pair Production via Dark Matter Annihilation

The flux of electrons and positrons resulting from the self-annihilation of Dark Matter particles depends on the density of these particles, their relative velocity and their self-annihilation cross section. The density of DM as a fraction of the total energy density of all matter and energy is estimated to be $\Omega_X \approx 0.22$ by the WMAP-7 estimates of the parameters of the standard cosmological model Λ CDM [23], according to which the DM particle is weakly but not electro-magnetically interacting: hence the term Weakly Interacting Massive Particle (WIMP). With this assumption, the interaction rate for interaction with Standard Model (SM) particles can be calculated by attributing to the WIMP all those interactions accorded to a neutral weakly interacting particle by the SM.

This results in a value for the interaction rate very close to that calculated cosmologically, roughly as follows. Assume that early in the evolution of the Universe, WIMPs are in thermal equilibrium with SM particles, meaning that the relative numbers of the DM and SM particles species are changing via interaction to ensure chemical equilibrium. At some point the falling equilibrium temperature and increasing volume of the universe means the rate of interaction between WIMPs and SM particles will become very small compared to the rate of expansion of the universe - a scenario referred to as thermodynamic de-coupling. The WIMP density at this critical time will become frozen in, as WIMPs are no longer permitted to change number appreciably by interacting with and turning into SM particles (and vice versa).

This balancing equilibrates the number of WIMPs characteristic of the Universe at some past time for the rest of cosmological evolution and is thus referred to as a thermal relic density. The exact level of this density can be matched to the cosmologically (using the best-fit Λ CDM model) determined density by choosing the appropriate cross section for

interaction with SM particles (see Equation 1.28). The fact that this best-fit cross section matches very closely to the cross section calculated by the methods natural to the Standard Model gives rise to the “WIMP miracle” and is an important argument in favor of the Dark Matter manifesting as a WIMP.

I forego a detailed derivation here; more background can be found in many reviews and the literature referenced therein, see for example [14]. The following approximation derived there is sufficient to display the dependence of the thermally averaged interaction rate $\langle\sigma v\rangle$ on two key cosmological parameters, the DM density Ω_X (expressed as a fraction of the critical total density needed for a flat universe) and the Hubble parameter $h = H_0/100\text{km s}^{-1}\text{Mpc}^{-1}$:

$$\Omega_X h^2 = \frac{3 \times 10^{-27} \text{cm}^3 \text{s}^{-1}}{\langle\sigma v\rangle} \quad (1.28)$$

With $\Omega_X \approx 0.226$ and $h \approx 0.703$ [23] this gives $\langle\sigma v\rangle \approx 2.69 \times 10^{-26} \text{cm}^3 \text{s}^{-1}$. Improved models predict a very slight dependence of this value on the WIMP mass m_X ; for example [42] estimate a maximum of $5.2 \times 10^{-26} \text{cm}^3 \text{s}^{-1}$ at $m_x \approx 0.3 \text{ GeV}$, leveling off at around $2.2 \times 10^{-26} \text{cm}^3 \text{s}^{-1}$ for $m_x > 10 \text{ GeV}$.

Although the exact relationship between interaction and self-annihilation cross sections is model dependent, this interaction cross section does determine to a large degree the rate that the theoretical relic WIMPs would self-annihilate into SM particles and therefore the flux of those particles attributable to WIMP annihilation [42]. Many calculations have been performed to determine whether flux measurements such as those in Figure 1.10 can be reproduced with such a mechanism.

In 2009 Cirelli *et al* [18] studied non-relativistic WIMP annihilation through all available SM channels and attempted to fit the resulting flux of photons, electrons, positrons, protons and anti-protons to the PAMELA measurements available at the time (Figure 1.10). In agreement with many other efforts, they found that the interaction frequency implied by the relic density is insufficient to produce the measured increase in positron flux, regardless of assumptions about the preferred SM annihilation product channels. In order to generate the requisite normalization, they (as is common in the literature) introduce a dimensionless boost factor B which multiplies the annihilation rate. Figure 1.11 shows the range of boost factors required to reproduce the PAMELA positron signal as a function of the unknown parameter m_X (WIMP mass).

The physical manifestation of the boost factor B is often left unspecified. One such possibility is Sommerfeld enhancement which amplifies the interaction rate (of any particle, not specifically WIMPS) by positing that the interaction is made more probable by




Figure 1.11: Figure 9 from [18]. Some typical channel-dependent boost factors B_e amplifying the WIMP thermal relic interaction cross section required to reproduce the positron fraction (as measured by PAMELA) ranged from 10 to 10^6 .

the intervention of some unspecified force carrier. The canonical example is the attractive electric potential between an electron and a positron enhancing the pair annihilation cross section. Both the mass of the force carrier (essentially a free parameter) and the relative velocity of the WIMPs (difficult to model for small regions) can strongly affect the level of enhancement. The nature and degree of velocity dependence often depends on the nature and mass of the force carrier and can be an important effect. Whatever level of enhancement is desired, the same amplification would affect the interaction rate that gave rise to the “WIMP Miracle.” For overly large enhancement factors the WIMPs would actually thermally recouple to the SM particles, reducing their density and thus greatly suppressing any indirect signal thereof. Feng *et al* [40] estimate that the largest Sommerfeld enhancement factors which avoid this recoupling range from $B = 7$ at $m_X = 100$ GeV to $B = 90$ at $m_X = 1$ TeV. These factors are too small to explain the PAMELA excess by factors of 25 or above.

Substructures in the WIMP density or relative velocity profile could also amplify the annihilation rate. Many such fluctuations (e.g. dwarf galaxies) are too far from the Solar System to inject high energy electron and positron fluxes at Earth; thus we are motivated to focus on fluctuations within the energy loss distance of \sim TeV electrons, that is, within a few kpc of the Solar System. This is a very small distance scale compared to most halo models of DM densities (which have a characteristic radius larger than that of the galaxies they envelop). Vogelsberger *et al* [31] attempted to model the detailed structure of DM density and velocity fluctuations against a background elliptical halo density between 6 and 12 kpc from the Galactic center (bracketing the Sun’s distance from the center at approximately




Figure 1.12: Figure 1 (bottom panel) from [31]. A high-resolution (sub-200 pc) model of the Milky Way DM halo predicts that with very high probability the Solar System does not reside in a region over-dense enough with WIMPs to produce an enhanced indirect electron or positron signal.

8.5 kpc). This is close to the distance scale (~ 1 kpc) within which such fluctuations could potentially contribute to the high energy electron and positron flux at Earth. Figure 1.12 reproduces the bottom panel of their Figure 1 which displays the resulting estimated distribution function for the pointwise density divided by the smoothed power law background density. The distribution is dominated by a lognormal peak resulting from Poisson density fluctuations (i.e., noise). Out of six resolution scales they sampled, only the two most precise (with a smoothing length of around 200 pc) are able to generate the high-density power law tail that would be characteristic of dwarf galaxies with large density-derived boost factors. The overall infrequency of such high-density regions indicates there is less than a 10^{-4} probability that the Sun is located in such a region (assuming of course that such fluctuations are uncorrelated with the location of the Solar System). They sum up: “...we can say with better than 99.9% confidence that the DM density at the Sun’s position differs by less than 15% from the average over the ellipsoidal shell on which the Sun sits” [31]. Thus the boost factor-modified electron and positron flux would vary by factors of order unity from that predicted by the canonical interaction rate, not of order $10^2 - 10^6$ as posited by most WIMP models.

The re-tuning of many models to fit the high-statistics data provided by the AMS-02 measurement in 2013 (see Figure 1.10) demonstrated the resilience of WIMP annihilation models to significantly improved positron fraction data. The new measurement undercut the sharp peak inferred by many to extend past the energies probed by PAMELA. New models arose which were able to match the steady rise in the positron fraction extending up to at least 450 GeV. The favored m_X estimates for most models were increased to at least the maximum positron energies being detected.

For example, Boudad *et al* [30] re-examined the WIMP parameter space and found two main regions displaying good fit. One model specifies $m_X \approx 20$ TeV (see Figure 1.13, left), annihilation channels comprising both lepton and quark products and an annihilation rate $\langle \sigma_{ann} v \rangle \approx 1.12 \times 10^{-21} \text{ cm}^3 \text{ s}^{-1}$. This model however is severely constrained by observations of the absence of γ -ray emission from spheroidal dwarf galaxies (see below). The other specifies $m_X \approx 600$ GeV (see Figure 1.13, right), annihilating into four-lepton states (either four electrons/positrons or four tau/anti-tau leptons, but not muons/antimuons) with annihilation rate $\langle \sigma_{ann} v \rangle \approx 7.37 \times 10^{-24} \text{ cm}^3 \text{ s}^{-1}$. The low WIMP mass model predicts that the positron fraction falls as positron energy approaches the WIMP mass. Measurement of the positron fraction at higher energy would provide an important additional test of this model. Note that the boost factors multiplying the canonical annihilation rate implied for the low mass model (about 200) and the high mass model (about 10^6) are well above the maximal Sommerfeld boost factors discussed previously and well into the high-density,




Figure 1.13: Figure 9 from [30] showing the positron fraction resulting from two separate good-fit regions of WIMP parameter space, a high mass, mixed quark/lepton model (left) which has increasing positron fraction with energy and a low mass four-lepton (though non-muonic) model (right) which displays a sharp cutoff past the WIMP mass m_χ . The dotted green line shows their estimate of the astrophysical background.

low-probability fluctuation tail shown in Figure 1.12).

In 2011, before the availability of the AMS-02 data, the Fermi collaboration investigated the amplified signal that could be produced by unusually high WIMP density in spheroidal dwarf galaxies in [29]. These ‘dSphs’ are chosen due to their very high DM content (with mass/light ratios of over 1000 [54]) and relative isolation with respect to other strong γ -ray sources (as compared to, for example, the Galactic center). Located further away from Earth than the energy loss distance for high energy electrons and positrons (their sample included dSphs between 20 and 140 kpc distant) they are unable to produce a measurable high energy positron signal at Earth. However any source of high energy positrons should also end up producing γ -rays, the flux of which would strongly depend on the same interaction rates producing the electron and positron flux. The Fermi-LAT γ -ray detector measured the γ -ray signal originating from the direction of the dSphs and found no excess over the expected background. Interpreting this as the lack of a detectable by-product of WIMP annihilation, they combined these observations with the estimated WIMP density in the dwarf galaxies to place upper limits on the annihilation cross section as a function of m_χ . They found that for some annihilation channels, the canonical value for $\langle\sigma_{ann}v\rangle$ of $3 \times 10^{-26} \text{cm}^3 \text{s}^{-1}$ is ruled out for low m_χ (below 30-40 GeV). This was not a very strong constraint as WIMP annihilation models themselves rule out the canonical value for essentially all masses.

Armed with additional Fermi observations and the higher energy AMS-02 data, Lopez *et al* ([54]) placed far stronger constraints based on the lack of a γ -ray excess in the direction of a catalog of 25 dSphs on the post-AMS-02 candidate models such as those proposed

by Boudad *et al.* Essentially the only kind of model that produces enough positrons with the right spectral shape without producing a measurable γ -ray flux is a 4- μ model with $m_x \approx 750$ TeV. They explain that the μ channel produces positrons in abundance without the accompanying γ -rays Fermi did not observe. 4- e models share those properties but fail to provide an adequate fit to the positron fraction in the parameter space explored by Lopez *et al.*⁴ The best-fit $\langle\sigma v\rangle \approx 10^{-23}$, implying a boost factor $B \approx 3000$, which according to [40] is well above the maximum Sommerfeld enhancement for preventing thermal re-coupling, indicating that at least some of this boost factor would have to be attributed to density or velocity fluctuations of order 100 – 1000 or so. As shown in Figure 1.12, such large fluctuations in our immediate vicinity are unlikely.

1.3.2.8 Pulsar Electron/Positron Pair Production

Pulsars are compact stars with magnetic fields on the order of 10^{12} G and rotation rates with periods as short as \sim ms. Various models of the interaction between a pulsar’s magnetic moment and spin angular momentum predict that they could emit high energy charged particles. Common to these models is that the particles injected into the ISM are finally produced by γ -ray induced electron and positron pair production. This naturally suppresses protons and anti-protons from being produced thereby matching the non-observation of an anti-proton fraction increase in the energy ranges measured so far. This is an attractive feature compared to WIMP annihilation models which must require that leptonic channels are favored over quark channels, largely without independent motivation in order to prevent an over-abundance of anti-protons.

One widely-used pulsar model that results in pair-produced electrons and positrons with typical energies above a GeV is the ‘polar gap’ model originally described by Ruderman and Sutherland [61]. They theorize that the trapping of positive iron nuclei by the surface of the star would tend to form a region of increasing electric potential as electrons (not bound to the surface) stream out along magnetic field lines. This forms a ‘magnetospheric gap’ approximately 10^4 cm in altitude with a voltage difference that grows up to 10^{12} V. When the potential difference becomes large enough to induce electron-positron pair avalanche breakdown occurs. Negative electrons flow back to the positively-charged surface and positrons (with up to \sim GeV energies) flow upwards out of the gap along open magnetic field lines. The positrons’ curved trajectories along these field lines produces γ -rays via curvature radiation. These photons, having sufficient energy to produce electron-positron pairs of their own, produce more electron and positron pairs and so on, leading to an elec-

⁴Recall Boudad *et al.*’s favored model was a mixed 4- e + 4- τ model with $m_x \approx 600$ GeV (but without muons).

tromagnetic cascade. Outside the gap electrons as well as positrons escape along magnetic field lines because there are no electrostatic forces pushing the gap electrons back to the star’s positively-charged surface. Thus the final leptons injected into the ISM would (without considering any follow-up interaction between these particles and the pulsar wind, the shock front between the pulsar wind and the ISM or other obstacle) be neutral, half electron and half positron.

Ruderman and Sutherland were primarily interested in explaining the coherent radio and microwave radiation observed coming from pulsars at the time and to that end estimated the maximum gap-accelerated positron energy as:

$$E_{max} = 1.7 \times 10^7 \frac{B_{12} h_4^2}{P} (m_e c^2). \quad (1.29)$$

B_{12} is the dipole component of the surface magnetic field in units of $10^{12}G$, h_4 is the gap thickness in units of 10^4 cm and P is the pulsar period in seconds. For values of parameters of $B_{12} = 1$, $h_4 = 0.1 - 1$ and $P = 0.2$, this would give a maximum energy of from approximately 400 GeV at $h_4 = 0.1$ to over 40 TeV at $h_4 = 1$. Again note that this is the energy of avalanche-produced positrons produced in the magnetospheric gap and not the energy of electrons and positrons that escape to the ISM. Positrons with near maximum energy in fact lose most of it to curvature radiation before leaving the gap and it is only “somewhat less energetic positrons... which leave the near magnetosphere without much radiative loss” ([61], pg. 64). Therefore this does not take into account any cooling or acceleration that might occur as they subsequently escape the pulsar wind and shock front at the termination of the pulsar wind.

Hooper, Blasi and Serpico ([50]) modeled the electron and positron flux at Earth due to an ensemble of pulsars within a kpc or so of the Solar System. They employed the injection spectrum

$$\frac{dN_e}{dE_e} \propto E_e^{-1.6} \exp(-E_e/80 \text{ GeV}) \text{ GeV}^{-1} \text{ s}^{-1} \quad (1.30)$$

as the input for a diffusive propagation model similar to that of Kobayashi *et al* ([38]). Two key features are evident here. First, the spectral index of $k = -1.6$ is hard compared to the canonical $k \approx 2$ (and softer) assumed for the injection spectrum for diffusive shock acceleration. Whereas the latter results naturally as a basic and general property of many kinds of shocks, this spectral index varies strongly with pulsar parameters which are not well known. Second, there is assumed to be an exponential cutoff at the energy $E_{max} = 80$ GeV for the ensemble of pulsars, while $E_{max} = 600$ GeV for the nearby Geminga and Monogem pulsars. These parameter choices result in an ensemble contribution to the positron fraction which falls off above the upper limit of the PAMELA results and a




Figure 1.14: Figure 4 from [50] showing the positron flux and fraction resulting from a combination of the discrete pulsar sources Geminga and B0656+14 (Monogem) and a Monte Carlo ensemble of pulsars at distance greater than 500 pc and a creation frequency of four per century. The ensemble has an injection spectral index $k = -1.6$ and injection energy cutoff at $E_{max} = 80$ GeV while the discrete sources have injection spectral index $k = -1.5$ and injection cutoff energies of $E_{max} = 600$ GeV. The dotted lines indicate the contribution from a GALPROP-estimated diffuse Galactic background.

discrete contribution which continues to higher energies (see Figure 1.14). No rationale was provided for the different cutoff energies of Geminga and Monogem.

Hooper *et al* acknowledge the tuning involved and cite as a basis for this choice a range of cutoff values in the literature, including the $E_{max} = 80$ GeV they chose for the ensemble pulsars and a maximum value of $E_{max} = 400$ GeV. This value was calculated using an equation derived by Chi *et al* ([17]). Chi *et al* extended the polar gap model by positing that electrons and positrons produced outside of the gap region form a relativistic plasma of e^\pm pairs. This plasma then enters thermal equilibrium with the electromagnetic waves that comprise the pulsar’s wind and are further accelerated. Assuming equipartition between the plasma and the electromagnetic waves of the pulsar wind Chi *et al* indeed find that plasma electrons and positrons will have energy *before escaping the pulsar wind* of:

$$E_{e^\pm} \approx 8.1 B_{12}^{5/7} P^{-17/7} \text{GeV}, \quad (1.31)$$

which is the equation Hooper *et al* used to estimate a cutoff energy of 400 GeV. However, Chi *et al* also take into account the fact that these electrons and positrons must still escape the pulsar wind and based on an escape-time estimate conclude that “...the particle energy ejected into the ISM is one order of magnitude lower than the equipartition energy given by [Equation 1.31].”⁵ This would indicate that had Hooper *et al* correctly incorporated the

⁵In [17] Chi *et al* refer to their equations 5 and 13.

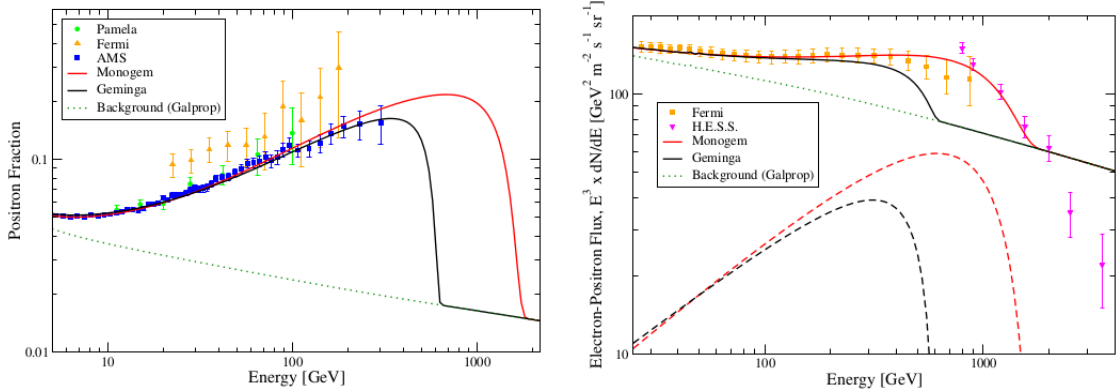


Figure 1.15: Figure 4 from [53] showing the positron fraction and electron-plus-positron flux resulting from the discrete pulsar sources Geminga and B0656+14 (Monogem). The best-fit injection spectral indices are $\gamma = 1.9(1.95)$ for Geminga (Monogem) and injection energy cutoffs are $E_{max} = 2$ TeV for both pulsars. The dotted line is the GALPROP-estimated diffuse Galactic background.

results of Chi *et al*'s work, the cutoff energy would be 40, not 400, GeV. When estimating the flux from Geminga and Monogem, Hooper *et al* use a cutoff energy of 600 GeV, even further from the values estimated in the references they used to inform their choice of E_{max} .

Because of their freedom in choosing this parameter (and the spectral index), Hooper *et al* concede that “...the results presented here are only demonstrative of the fact that a significant role of nearby pulsars, while not required to explain present data, is consistent with them... in no case the high-energy spectra presented here should be considered as a robust prediction, since they depend crucially on the detailed spectral properties of B0656+14 [Monogem], Geminga or and [*sic*] other nearby, mature pulsars that contribute significantly to the high energy positron spectrum” ([50], pp. 8-9).

Linden and Profumo [53] re-examined the flux from Geminga and Monogem in light of the AMS-02 data. They assume that the higher energy regime (up to 450 GeV) of the fitting domain rescues the model from having to rigorously specify an E_{max} for the injection spectrum since energy loss during propagation through the ISM provides its own exponential cutoff at high energies. While this may be true for their propagation model, this is not in agreement with the propagation model of Kobayashi *et al* which estimates that electrons with an energy loss distance equal to the distance to Geminga and Monogem are, respectively, approximately 10 and 25 TeV (see Figure 1.4), though these figures vary as the distance estimates to the pulsars vary. They assumed a cutoff energy of $E_{max} = 2$ TeV. The best fit spectral indices were $\gamma = 1.9$ for Geminga and $\gamma = 1.95$ for Monogem.

Independently from the energy spectra of electrons and positrons escaping such pulsars, a pulsar signal could be characterized by an anisotropy in arrival directions. Since Geminga

and B0656+14 (Monogem) are located opposite to the Galactic center, if dark matter signals are primarily directed from the Galactic center this might produce a noticeable dipole anisotropy [53]. Dark matter clumps in other regions of the sky would likely produce a γ -ray signal which (as noted in Section 1.3.2.7) is not seen by Fermi-LAT.

Interestingly, Boudad *et al* [30] (whose WIMP models were claimed to be largely ruled out by the Fermi-LAT γ -ray null result in [54]) claim that the single pulsar explanation could be ruled out if AMS-02 continues taking data, does not change the mean value of any of the bins in its positron fraction histogram, and that the statistical errors in each bin shrink with time: $\sigma_{stat} \propto t^{-1/2}$. Essentially they are pointing out that while the pulsar positron fractions match the data well at low energies the fit is less good for the high energy bins. Therefore the pulsar spectra, they claim, disproportionately benefit from the large error bars on the higher energy AMS-02 data points. Given the uncertainty in the spectral index, the cutoff energy and in the energy losses during propagation (which are much more acute at 450 GeV than at ~ 50 GeV and sensitive to uncertainties in the pulsar distance measurements as well as the propagation model parameters themselves) this may be an unwarranted criticism. In other words, the authors of a model (which suffers from extreme fine-tuning) criticize another model (which also suffers from extreme fine-tuning) for being insufficiently well-tuned.

1.3.2.9 Secondary Pair Production by Protons and Heavier Nuclei

Two-component models consisting of a distant, continuous and negative source distribution giving way to a nearby neutral component at higher energies are common in the literature. For example, Atoyan, Aharonian and Völk [10] and references therein discuss a two-component model in order to explain the rising positron fraction. Most authors of these studies conclude that the secondary flux produced in a self-consistent model of primary and secondary cosmic rays is too small to match measurements above 10 GeV (see e.g. Moskalenko and Strong [55]). Recent attempts to describe the observed positron flux with nearby discrete (spheroidal dwarf galaxies or the nearest pulsars) or distributed (WIMPs) sources do so with the assumption that the diffuse Galactic background contributes only at the background level. Here I briefly examine a forthcoming study which challenges this assumption and argues, using the most recent data from AMS-02, that the diffuse background could in fact be the dominant positron source at high energies.

Motivated by the seemingly coincidental matching of the proton and positron spectral index above around 40 GeV in the AMS-02 data, Ahlen and Tarlé [4] investigated the possibility that differing propagation histories between the standard primary shock-accelerated electrons and thought-to-be-insufficient secondary electrons and positron flux produced by

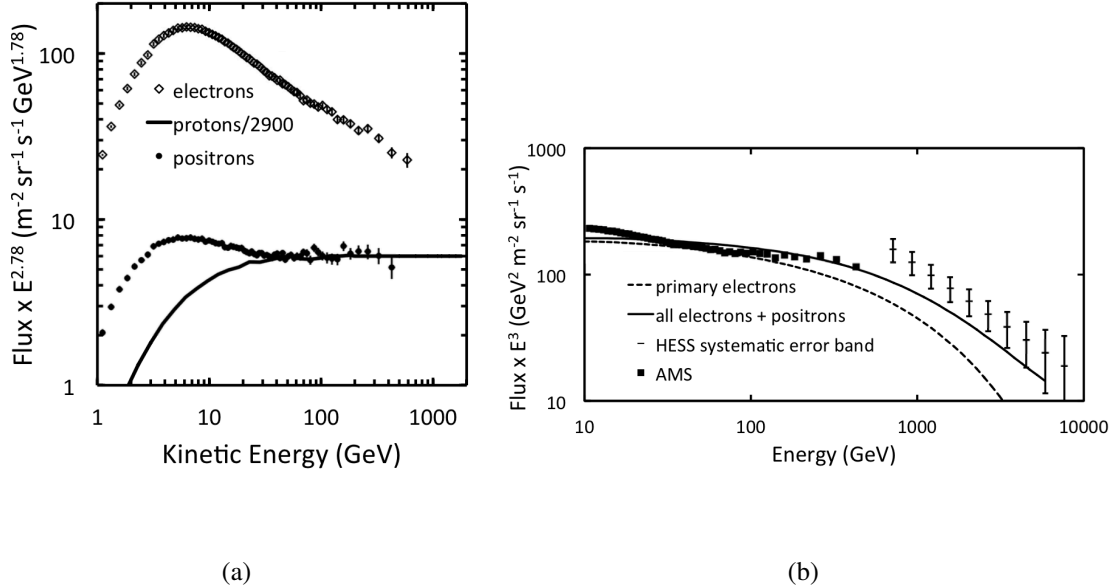


Figure 1.16: (a) Figure 1 from [4]. Comparison of recent measurements (from AMS-01 and AMS-02) showing the seeming coincidence that the proton and positron fluxes vary with the same spectral index above approximately 40 GeV. (b) Figure 4 from [4] demonstrates the agreement of the primary electron and secondary electron/positron model with AMS-02 and HESS data over the measured energy range. As the primary electron flux falls off due to lack of local sources, the nearby diffuse secondary electron/positron flux takes over and would be characterized by an isotropic flux with positron fraction asymptotically approaching 60% (the fraction of positrons produced by the maximally polarized decay products of primary cosmic ray protons).

nuclear collisions could explain the new electron and positron flux measurements. In this model the nearby neutral component is diffuse and isotropic in nature and therefore would not display the anisotropy characteristic of WIMP density fluctuations or nearby pulsars.

They propose a two-component model consisting of a primary negative source distribution outside of about 200 pc and a secondary neutral nearby source consisting of positrons produced by the nuclear collisions between cosmic ray protons and nuclei with atomic hydrogen in the ISM. This production is theorized to occur in a cylinder around Earth with a radius of approximately 5kpc and a half-height of 265 pc, the HWHM for the disk of warm neutral atomic hydrogen near the sun. They also assume that at some point in our history some kind of advective mechanism (the exact nature of which is not at issue) swept out older primary and secondary cosmic ray species, enforcing a roughly energy-independent confinement time of around 1-2 Myr. This provides the conditions for a primary electron

source which is prevented from contributing to the flux above a high energy cutoff enforced by diffusive energy losses to be overtaken by a secondary electron and positron component at high energies produced at all distances from Earth (up to 5 kpc).

Protons with a particular power law index naturally produce secondary electrons and positrons with the same index and with a fraction of the parent proton energy. Two complications specific to the energy regimes they consider (a low energy regime ranging in positron energy from 10 to 40 GeV and a high energy regime from 40 to 430 GeV) threaten to disrupt this concordance. First, the proton spectrum displays a break at around 1 TeV (possibly due to the onset of energy dependent leakage out of the galaxy decreasing the confinement times of higher energy particles), with the spectrum hardening by decreasing the spectral index from $\alpha = 2.8$ below 1 TeV to $\alpha = 2.6$ above 1 TeV⁶. Second, positrons will lose energy (in the manner described in Section 1.3.1.1) at an energy-dependent rate as they diffuse to Earth from their production point inside the cylinder.

As it turns out these two mechanisms roughly cancel out, meaning that close agreement of the spectral index above 40 GeV positron energy is somewhat of a lucky accident. The energy-dependent energy loss also accounts for the difference in flux between primary electrons and secondary electrons and positrons - although their energy loss coefficients are identical, their propagation histories are not. Primary electrons must diffuse a larger distance from their source locations in order to reach us, meaning they cut off at a lower energy than do the secondary electrons and positrons, which enjoy being produced everywhere cosmic rays interact with the ISM, including in our immediate vicinity. Therefore above the high energy cutoff for primary electrons, secondary electrons and positrons will dominate.

A criticism of the model might proceed as follows. The fortuitous cancellation brought about between energy-loss and proton spectral break indicates that perhaps the model is benefitting from fine-tuning. One free parameter used to normalize the resulting flux of positrons to measurements is the hydrogen number density n_H . Tuning this parameter is analogous to the boost factors integral to the WIMP models; however unlike those models the best-fit values are not unreasonable when taken as an average value of a quantity that varies rapidly over the productive volume. Another is the spatial diffusion constant for positrons, K , which was allowed to take on a range of values. In both the WIMP annihilation and pulsar models considered here, the propagation constants are also either marginalized over or set to best-fit values. The values of these parameters resulting in flux estimates displaying best agreement with the AMS-02 measurements are $K = 1.0 \times 10^{28} \text{cm}^2 \text{s}^{-1}$ (somewhat on the low side compared to the constant computed using Equation 1.11 from

⁶Parent proton energy of 1 TeV corresponds to average positron energy of around 40 GeV.

[38] for a 40 GeV electron of 7.6×10^{28}) and $n_H = 1.4 \text{ cm}^{-3}$. The remaining model parameters are inputs from other models or measurements (e.g. the velocity of the advective process clearing out old particles) or have minimal impact on the flux estimate (e.g. the radius of the secondary production cylinder R^7). Figure 1.16(b) compares the model's combined estimated primary electron and secondary electron/positron spectrum with the AMS-02 and HESS measurements. The small amount of free parameters and the ease with which the model matches the most recent data while utilizing largely exogenously-derived astrophysical parameters reflects favorably on its explanatory power. A measurement of the positron fraction at energies beyond those probed by AMS-02 is needed to help determine the primary or secondary nature of the flux at these energies and provide an important test of such a two-component model.

1.4 Cosmic Ray Electrons: From Source to Detection

The structures and processes capable of accelerating, or generating, high energy electrons and positrons are varied in nature. Many of them utilize magnetic fields to do so. Following acceleration or production, the electrons again must contend with magnetic field interaction, this time as they randomly diffuse through the ISM. Their uniquely large rate of energy loss make measuring their high-energy flux spectrum a key endeavor for learning about the structure of our Galactic neighborhood, within approximately 1 kpc. Known nearby structures, such as Vela, may be transmitting important information to us by means of electrons with energy above a TeV or so. Other structures, such as nearby pulsars or dwarf galaxies, may also be generating measurable signals. Detecting these electrons would provide an important confirmation of our ability to model the acceleration and propagation of cosmic ray electrons. Current experiments, utilizing both direct and indirect detection methods, are unable to probe this energy regime. By combining direct and indirect methods, CREST could represent an important step forward in detector methodology.

One upcoming experiment, CALET [36], is slated for launch within a year or so, and hopes to measure the combined electron and positron spectrum up to approximately 20 TeV. As I have related here, a great deal more information about pulsars, WIMP annihilation and secondary production of electrons by other cosmic ray species would be unlocked by measuring the individual electron and positron spectra in this energy range. Until now, no method for doing so that I am aware of has been proposed. While it was not designed to do so, I will show here a technique by which CREST (or a future, CREST-like detector) could

⁷The insensitivity to R comes about due to the positrons either being advected out of the disk by the wind or losing energy to inverse Compton and synchrotron losses before diffusing that far.

measure the positron fraction at energies above several TeV.

In order to describe CREST' detection method, it is first necessary to describe the signals produced by high-energy electrons as they interact with the geo-magnetic field, and the residual atmosphere above CREST's float altitude, which I proceed to do in the following Chapter.

CHAPTER 2

Synchrotron and Bremsstrahlung Signal Photon Production

While diffusing through the Galaxy, high energy cosmic ray electrons lose energy primarily via magnetic and electric field interactions as in Equation 1.5. These interactions also occur as they enter Earth's atmosphere and geo-magnetosphere. This results in the emission of Bremsstrahlung radiation via Coulomb interactions with the atoms in air molecules and magneto-Bremsstrahlung (synchrotron) radiation via interactions with the Earth's geomagnetic field. In this section, I review the properties of these processes and their role in shaping the signal these electrons make in the CREST detector.

2.1 Synchrotron Radiation by the Primary Electron

In this section I make use of analysis Michigan Physics Professor Carl Akerlof performed for CREST in 2009. This treatment parameterizes electron trajectories according to their distance from the detector plane at detector altitude. This allows one to estimate the expected number of detector-plane intersecting synchrotron photons each point in trajectory phase space would generate. This analysis and the discussion it generated guided the choice of the coincidence threshold value of three STACs when tuning the digitize decision trigger system (see Section 3.3.7).

I list here the relevant features of synchrotron radiation produced by ultra-relativistic electrons in a ~ 0.5 G magnetic field strength in the component perpendicular to the electron motion:

1. The differential photon energy distribution peaks at a fraction (0.29) of the critical energy E_c (which is a simple function of electron energy and the magnitude of the perpendicular component of the magnetic field along the electron's path);

2. The differential photon angular distribution is conical with an extremely narrow ($\propto \frac{1}{\gamma} \sim 10^{-6}$ radians at electron energy of 1 TeV) opening angle and cone axis centered on the primary electron's momentum at the moment of emission;
3. Since the ultra-relativistic primary electron moves at the same speed as the photons it emits and the momentum deflection from the magnetic field is in the transverse direction, the photons (and electron) comprise a simultaneous radiation front impacting on a plane perpendicular to the electron motion.

Consequently, the spatial-temporal distribution of the intersection point of these photons with an arbitrary plane in space (such as CREST's detector plane) is a nearly-straight, nearly-simultaneous line. These properties motivate the requirements for a fast (sub-nanosecond time resolution) detector with good spatial resolution outlined in Section 3.1.¹

2.1.1 Synchrotron Power and Energy Spectrum

The total power radiated by an electron of energy E_0 (measured in TeV) in a magnetic field with perpendicular component of strength B_{\perp} Gauss is given by [63]:

$$P(E_0) = 3.77 \times 10^3 B_{\perp}^2 E_0^2 \text{ MeV s}^{-1} \quad (2.1)$$

The spectral power distribution is such that half of emitted photons have frequency below the critical synchrotron frequency:

$$\nu_c = \frac{3}{2} \frac{1}{2\pi} \gamma^3 \frac{c}{\rho}, \quad (2.2)$$

where $\gamma \equiv \frac{E}{mc^2}$ and ρ is the gyroradius given by

$$\rho \equiv \frac{E}{ceB_{\perp}}. \quad (2.3)$$

This gives a critical frequency of

$$\nu_c = \frac{3}{2} \frac{\gamma^2 e B_{\perp}}{2\pi m_e}. \quad (2.4)$$

¹The assumption that the magnetic field component parallel to the detector is zero is at best only nearly true. Figure 5.40 shows the local magnetic field components during CREST's flight as measured by the on-board magnetic field sensor. It is evident that B_{xy} is non-zero while the CREST detector plane is largely parallel with the Earth's surface. This induces an acceleration in the z-component of the electron's momentum, and not that of the photons. However, the differential arrival times this would produce in the synchrotron photon wave front are very small compared to CREST's time resolution of approximately 1 ns.




Figure 2.1: Critical synchrotron photon energy envelope (in MeV) as a function of primary electron energy (in TeV). The upper edge of the envelope corresponds to an average B_{\perp} of 0.9 Gauss, the center to 0.6 Gauss, and the bottom edge to 0.3 Gauss.

If the perpendicular magnetic field is expressed in Gauss and the primary electron energy in TeV, one finds

$$E_c = 6.56 \times 10^{-2} B_{\perp} E_{TeV}^2 \text{ MeV.} \quad (2.5)$$

which I plot in Figure 2.1. Parameterizing energy as $x = \frac{E}{E_c}$ gives a spectral shape of

$$\frac{dN_{\gamma}(E_{\gamma})}{dE_{\gamma}} = \frac{3.53 \times 10^5}{E_{\gamma}} \left(\frac{E_{\gamma}}{E_c} \right) \times \int_{(E_{\gamma}/E_c)}^{\infty} K_{5/3}(z) dz \text{ photons/MeV.} \quad (2.6)$$

Professor Akerlof devised in [5] a close numerical approximation to the Bessel function of the second kind inside the integral by means of the triple-domained polynomial

$$S \left(\frac{E_{\gamma}}{E_c} \right) = S(x) = x \int_x^{\infty} K_{5/3}(z) dz \quad (2.7)$$

$$S(x) = \begin{cases} a_1 x^{\frac{1}{3}} & \text{if } 0 \leq x \leq x_b \\ S_p - a_2 \left(x^{\frac{1}{3}} - x_p^{\frac{1}{3}} \right)^2 & \text{if } x_b \leq x \leq x_c \\ a_3 e^{-b_3 x} & \text{if } x_c \leq x \leq \infty \end{cases} \quad (2.8)$$




Figure 2.2: Analytical approximation (red dashed line) to the function $S(x)$ (black solid line), necessary for computing the synchrotron spectral distribution [5]. The most probable energy, $E_\gamma/E_c = 0.29$ is marked by the vertical dashed blue line.

with the following values for the constants:

$$\begin{aligned}
 a_1 &= \frac{3}{2} 2^{2/3} \Gamma\left(\frac{5}{3}\right) = 2.1495282415 \\
 x_p &= 0.2858122477 \\
 S_p &= 0.9180123332 \\
 x_b &= 0.007465566303 \\
 x_c &= 0.9938126206 \\
 a_2 &= 2.3199736297 \\
 a_3 &= 1.4550396911 \\
 b_3 &= 0.8092153955
 \end{aligned} \tag{2.9}$$

Computing numerically the integral out to $x = 10$ and plotting against this analytical approximation demonstrates its accuracy (see Figure 2.2). Utilizing this simplification, the expression for the photon number spectral distribution becomes

$$\frac{d\dot{N}}{dx} = \frac{P_{synch}}{E_c} \frac{S(x)}{x}, \tag{2.10}$$

and by integrating over all energies, one obtains:

$$\dot{N} = \frac{P_{synch}}{E_c} \int_x^\infty \frac{S(x)}{x}. \quad (2.11)$$

This double integral (first the Bessel integral, then the integral over the Bessel integrand) can be performed analytically using the numerical approximation for $S(x)$. The result is

$$\dot{N} = 5.30 \frac{9\sqrt{3}}{8\pi} \frac{P_{synch}}{E_c}, \quad (2.12)$$

as long as E_{thresh} , CREST's minimum detectable photon energy, is sufficiently small compared to E_c , the critical synchrotron energy. Combining the above provides an estimate for the number of synchrotron photons emitted per unit pathlength of

$$\frac{dN}{ds} = \dot{N}/c \approx 0.255 \text{ photons/km}. \quad (2.13)$$

Notably this value is constant with respect to the primary electron energy.

2.1.2 Calculating The Illuminating Pathlength

Armed with an estimate for the number of photons emitted per unit of primary electron path length, to estimate the number of photons incident on the detector one must calculate how much of that path is actually illuminating our detector. The question is purely geometrical: given that the electron is moving on a circle with radius of the gyroradius ρ , over what arclength is it illuminating a detector of size w ? In Figure 2.3 I reproduce Prof. Akerlof's schematic for making this estimate. Consider the particle trajectory as fixed and the detector position as variable.² In this diagram the particle trajectory is variable, but one could imagine a similar picture where the same trajectory is used repeatedly for different positions and orientations of the detector. For example, if we keep the red trajectory but move the detector rightwards, we'd eventually have the same geometry as that of the blue trajectory and the initial detector position. With the assumption that the electron is moving on a circular trajectory with gyroradius on the order of 10^6 km (though the true trajectory is a spiral, the deflection in the transverse direction is ignorably tiny), let the angles θ_i and θ_j subtend the portion of the trajectory during which the electron points at the detector. (these segments are bolded in Figure 2.3). These arclengths are the illuminating pathlengths.

Focus on the red path. Since the detector's left edge lies along the tangent of the particle at $\theta = 0$, we have $\theta_i = 0$. Trace along the trajectory, simultaneously tracing the intersection

²This approach was also taken in some of the CREST Monte Carlo signal simulations.




Figure 2.3: From [5], courtesy of Professor C. Akerlof. Portions of circular primary electron trajectories are shown (online, left in blue and right in red). The portion of the track during which each electron illuminates the CREST detector is bolded. The left (blue) track has a longer illuminating pathlength and therefore a higher probability of producing a large number of synchrotron photons on the detector plane with width w .

of the tangent line with the $d = 1000$ plane. Stop when this intersection point reaches the other end of the detector; the corresponding angle along the trajectory is θ_j . This constitutes one iteration of a computational process for calculating the arclength: choose θ_i , calculate its point of intersection with an imaginary plane d km away (here $d = 1000$ km), then find the θ_j which has an intersection point $w = 2$ m away from θ_i 's intersection point.

Imagine increasing θ_i slightly above zero, and tracing forward to θ_j which intersects the $d = 1000$ km plane exactly 2m to the right. This arc will be slightly longer than the original arc, and starts and ends physically closer to the detector (which sits on the imaginary plane). We can continue this process, generating longer and longer arclengths from positions closer and closer to the detector, until we reach the situation depicted by the blue line - at some point the electron actually traverses the detector plane itself. This marks the point of maximal illuminating arclength, and were we to proceed in this manner the illuminating arclength would rapidly decline. During CREST's design these trajectories were considered the most extreme, as any further movement of the trajectory to the left would cause activation of the veto system, alerting the detector of the passage of a charged particle. I will return to this point shortly.

Performing this iterative process numerically provides an estimate of the illuminating pathlength and therefore (by means of Equation 2.13) the expected value of the number of synchrotron photons. It turns out that the most photon-rich geometries occur quite close to the detector, where the illuminating pathlength is maximal. Integration step sizes which calmly traverse the sparse phase space far away (about 1000-200 km) from the detector blow right by the interesting bits near (about 200-50 km) the detector. Therefore an adaptive step is required which shrinks as one gets closer to the detector. I employed a variable step size of

$$d\theta(x) = d\theta_0 e^{-(x/x_s)^2}, \quad (2.14)$$

where $d\theta_0 = \frac{\pi}{4} \frac{1}{50,000}$, x is the perpendicular distance from emission point to the imaginary plane and x_s is a scale factor which controls how quickly the step size falls (both measured in km). In the interest of performing this calculation over many energies, I kept the scale factor rather large (at about 450-500 km). This allowed the process to spend more computation time near the maximal θ_i , which region of phase space generates the longer illuminating pathlengths.

The effect of atmospheric attenuation would need to be taken into account to more closely match reality, especially for these proximal trajectories. First the attenuation of synchrotron photons by the atmosphere will be at a minimum for photons emitted from these kinds of tracks, as they will traverse the barest minimum number of air particles on their way to the detector. For photons emitted further away and at zenith angles different from

zero, this attenuation would further decrease the expected signal of distant compared to proximal trajectories. This reduced photon attenuation does not come for free, as the electron itself must now traverse the residual atmosphere instead. The rate of Bremsstrahlung emission goes up as the density of the atmosphere goes up; consequently proximal trajectories generate more Bremsstrahlung than distant ones. The increased rate of Bremsstrahlung photon emission complicates the event morphology in a non-trivial manner.

2.1.3 The Expected Number of Photons

Given the mean distance traveled between synchrotron photon creation Δl_{synch} (the constant value given by Equation 2.13) and the illuminating pathlength $\Delta s(E, x)$, one can estimate the number of photons that will interact with the detector plane for any given trajectory, with the further inclusion of F , the fraction of detector area covered by photon-sensitive BaF₂ crystal scintillators

$$\langle n \rangle = \frac{\Delta s(E, x)}{\Delta l_{synch}} F_{crystal(\phi)}. \quad (2.15)$$

The discrete random variable n , the number of photons on each trajectory, can be considered to have Poisson distribution with mean value $\langle n \rangle$ such that

$$P_n \equiv P(n; \langle n \rangle) = \frac{\langle n \rangle^n e^{-\langle n \rangle}}{n!}. \quad (2.16)$$

The probability that the number of detected synchrotron photons is greater than the detection threshold of 3 is then

$$P_{n>3} = 1 - (P_0 + P_1 + P_2 + P_3) = 1 - \left(1 + \langle n \rangle + \frac{\langle n \rangle^2}{2} + \frac{\langle n \rangle^3}{6} \right) e^{-\langle n \rangle}. \quad (2.17)$$

Define the effective width Δx_{eff} at a given primary energy and magnetic field strength as

$$\Delta x_{eff} = \int_0^{\Delta x} P_{>3}(x) dx, \quad (2.18)$$

and re-parameterizing the trajectories using θ in place of x (with $x = \rho \sin \theta$) as

$$dx = -\rho \cos \theta d\theta \quad (2.19)$$

gives the estimate effective width of the detector

$$\Delta x_{eff} = - \int_{\theta_0}^{\theta_0 + \Delta\theta} P_{>3}(\theta) \rho \cos \theta d\theta. \quad (2.20)$$

To summarize:

1. Pick an average magnetic field strength (0.5 G)
2. Pick an energy (from 0.5 to 50 TeV)
3. Calculate the gyroradius from Equation 2.3
4. Iterate over values of θ
 - (a) At each θ (that is, each trajectory), numerically develop $\Delta s(E, \theta)$
 - (b) Compute $\langle n \rangle$ using Equation 2.15
 - (c) Compute $P_{>3}$ using Equation 2.17
 - (d) Compute the integral contribution using Equation 2.20
5. Terminate either by condition 1 (more conservative) or condition 2 (less conservative; see below)

2.1.4 Integration Limits: Extreme Trajectories

The limits on the integral photon probability in Equation 2.20 represent a choice about the types of events one considers signal events. One angle at which to terminate this iteration is the angle that causes the electron’s trajectory to intersect with the detector. I refer to this as “termination condition 1” in my summary algorithm above. Figure 2.4 shows the differential probability of detecting more than three signal photons as a function of trajectory parameter. The dotted vertical red line at around 0.48 km shows the trajectory parameter at which electrons begin to intersect the detector. Events that satisfy termination condition 1 are purely indirect events: the primary electron itself would never intersect the detector plane; only synchrotron (or Bremsstrahlung) photons would. In events that violate termination condition 1 the electron would interact directly with the detector and thus manifest activity in CREST’s charged particle veto which detects the passage of the electron in one (or more - usually more) of the veto paddles that surround the crystal plane. During CREST’s pre-flight design and simulation the collaboration elected to reject any events *en masse* with any veto activity whatsoever. The same is done as part of the post-flight “mainstream” analysis. In essence CREST chose to use termination condition 1.

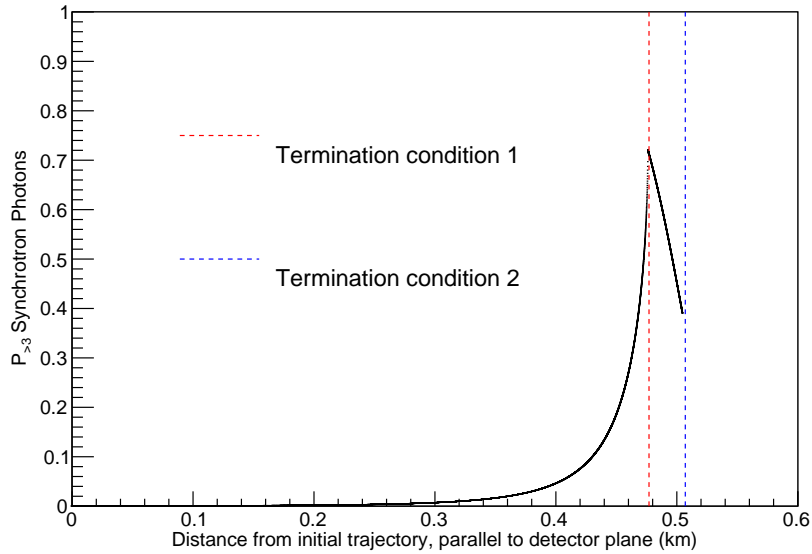


Figure 2.4: Illustration of the relative probability of signal photon detection over the coincidence threshold of 3 as a function of trajectory parameter. The effective width of the detector at a specific primary electron energy and magnetic field strength (in this case, 10 TeV and 0.5 G) is proportional to the integral of this quantity. The dotted lines delineate two classes of signal events. First, events where the electron does not traverse the CREST detector plane (dubbed Termination Condition 1, to the left of the red line), and second, events where it does (dubbed Termination Condition 2, between the two dotted lines). To the right of the blue dotted line, the trajectory produces a synchrotron beam path in the detector shorter than 1m and is regarded as undetectable. At 10 TeV primary electron energy approximately 60% of CREST’s detection power derives from events where the electron traverses the detector. This fraction increases at lower energy and decreases at higher energy.

According to this model, such a procedure systematically ignores a significant population of signal photons, as follows. Refer again to Figure 2.3 and imagine starting from the situation depicted by the blue trajectory. Picture moving the detector slightly to the right (equivalently, imagine moving the blue trajectory to the left). The illuminating pathlength (which was at a global maximum for the blue trajectory) will decrease for two reasons. First θ_i (representing the leading edge of the illuminating path length) will increase, as the photons emitted at the previous θ_i will miss the left edge of the detector. Second, θ_j will decrease as the electron will run out of the detector on the latter part of its trajectory. Now the electron will actually pass through the detector and this ‘traversal’ will cut short the illuminating path. As a consequence, rather than illuminating the entire detector, the synchrotron beam would illuminate a slightly shorter section of it. Termination condition 2 would include all such trajectories, up until the illuminated section of the detector falls below an analysis-determined threshold. A reasonable value is approximately 1 m - the exact value is not important for this argument.

This termination condition - that the intersection with the detector be at least 1m - defines what I call “termination condition 2,” depicted by the blue line in Figure 2.4. Clearly evident is the rapidly falling probability of detecting more than three photons beyond termination condition 1 (the length decreases from both sides - θ_i is increasing and θ_j is decreasing). Nonetheless, because the probability starts from the global maximum value, the total integrated number of expected synchrotron photons from such traversing trajectories is significant. In fact, at a primary electron energy of 10 TeV and a magnetic field strength of 0.5 G this geometric model predicts that 60% of all detectable signal photons derive from events that violate termination condition 1. In other words, over half of CREST’s detection power comes from traversal events.

In Figure 6.9 I plot the Monte Carlo signal simulation’s prediction for this direct detection fraction. At 10 TeV, and with a realistic geo-magnetic model, the traversal fraction at 10 TeV is estimated to be 59% (somewhat miraculously within 1% of the value predicted here using this simplified model). As primary energy decreases, the traversal fraction increases. At 500 GeV, approximately 80% of theoretically detectable events are traversal events.

2.2 Bremsstrahlung Emission by the Primary Electron

Since CREST flew below a non-negligible amount of residual atmosphere (ranging from around 6-9 g/cm², see Section 5.4.8 for CREST’s flight altitude, pressure and overburden profiles) the emission of Bremsstrahlung photons caused by the interaction of the primary

electron with atmospheric atoms must be considered. I review the properties of this emission here, largely following the canonical text *Classical Electrodynamics* by Jackson (3rd Edition) [51].

Bremsstrahlung radiation occurs when a fast charged particle collides with an atom. In our case we are concerned with ultra-relativistic electrons colliding mainly with the Nitrogen and Oxygen atoms in Earth's atmosphere. This brings about two corrections to the standard power and angular spectra:

1. The angular spectrum, normally that of standard dipole radiation in the rest frame of the colliding particle, is Doppler shifted to the rest frame of the detector. The resulting spectrum is that of a conical beam with opening angle of $1/\gamma$, and
2. The screening effect of the atom's electrons on the electric field of the nucleus becomes important when the colliding particle is very relativistic. This increases the ratio of energy lost to radiation to that lost to collisions.

The first correction is of great relevance to CREST's signal profile; Bremsstrahlung photons will lie on the primary electron illumination path since they are beamed in the forward direction by the same $1/\gamma$ factor as the synchrotron photons. The second means that the amount of bremsstrahlung radiation one would expect in any given signal event is proportional to the amount of air traversed by the primary electron while it is illuminating the detector and is not proportional to the energy of the primary electron (as long as it is in the complete-screening regime). This means the ratio of synchrotron power to Bremsstrahlung power will scale with primary electron energy (and therefore lower energy electrons will generate relatively more Bremsstrahlung).

Formula 15.36 from [51] gives the doubly differential cross section for energy radiated per unit frequency interval and per unit solid angle for $\hbar\omega \ll E$, where $\hbar\omega$ is the Bremsstrahlung photon energy and E is the energy of the incident electron, as

$$\frac{d^2\chi_R}{d\omega d\Omega_\gamma} \simeq \left[\frac{3}{2\pi} \gamma^2 \frac{1 + \gamma^4 \theta^4}{(1 + \gamma^2 \theta^2)^4} \right] \frac{d\chi_R}{d\omega}, \quad (2.21)$$

with

$$\frac{d\chi_R}{d\omega} \simeq \frac{16}{3} \frac{Z^2 e^2}{c} \left(\frac{z^2 e^2}{Mc^2} \right)^2 \ln \left(\frac{\lambda'' E E'}{Mc^2 \hbar\omega} \right). \quad (2.22)$$

Here Z is the atomic number of the scattering atoms, ze is the charge of the scattered particle (for us that's the electron, so that $z \equiv 1$), M is the electron mass, $E(E')$ is the energy of the electron before (after) interaction, and λ'' is some constant assumed to be of order unity.

Of note in Equation 2.21 is the property that angles far from zero are heavily penalized with increasing γ . Thus TeV electrons (with $\gamma \approx 10^{12}/(0.511 \times 10^6) \approx 2 \times 10^6$) emit Bremsstrahlung photons overwhelmingly in the direction they are traveling. This result can also be obtained by considering minimum and maximum momentum transfer in the frame of the scattering atom or in the frame of the relativistic particle, which illustrates the Doppler boosting of the dipole radiation spectrum into a forward-beamed cone.

2.2.1 Complete Screening and Radiation Length

Due to their low mass and high energy, electrons in the TeV regime will experience “complete” screening of the atomic nuclei by the atomic electrons ([51], pg. 722) meaning that screening effects are important for all emitted photon frequencies, rather than only below some critical frequency (as is the case for non-relativistic particles).

In this regime, at higher Bremsstrahlung photon energies the energy spectrum is approximated by

$$\frac{d\chi}{d\omega} \simeq \frac{16}{3} \frac{Z^2 e^2}{c} \left(\frac{z^2 e^2}{Mc^2} \right) \ln \left(\frac{233M}{m_e Z^{1/3}} \right), \quad (2.23)$$

Individual photon energies will thus follow a $(\hbar\omega)^{-1}$ spectrum up to the maximum allowable energy (that of the radiating electron itself). To estimate the primary electron energy from the energy of the Bremsstrahlung photons themselves a detector would thus need to be able to contain photons with up to TeV energies. Very few balloon- or space-borne detectors are capable of this as it would require a very heavy calorimeter. CREST’s crystal system is able to resolve photon energies up to only around 40 MeV (although much higher energy depositions are possible within the crystals and lead, the low-gain charge amplification channel saturates at this level).

Signal events generated by TeV electrons can include Bremsstrahlung photons of energy sufficient to themselves produce electron-positron pairs. Such photons can result in the formation of an electromagnetic shower in the atmosphere immediately *above* the detector and contribute to the late-arriving subset of signal photons (see Figure 6.12). This suggests utilizing the existence of these miniature air-showers as indicators of the passage of the electron itself in traversal-type events. It also suggests that events generated by electrons that illuminated the detector while at lower altitude (and hence denser atmosphere) will have high crystal multiplicity with some shower products deviating from the main beam path, as they diffuse in space and time before crossing the detector plane.

The rate at which these high energy photons are produced (equivalently, the rate at which the electron loses energy) as the electron travels through the atmosphere can be

expressed as an energy loss per unit length (Equation 15.48 in [51]):

$$\frac{dE_{rad}}{dx} \simeq \frac{16}{3} N \frac{Z^2 e^2}{\hbar c} \left(\frac{z^2 e^2}{Mc^2} \right) \ln \left(\frac{233M}{mZ^{1/3}} \right) \gamma Mc^2, \quad (2.24)$$

where N is the number density of the scattering atoms. This suggests re-writing Equation 2.24 as the differential equation

$$\frac{dE}{dX} = -\frac{E}{X_0}, \quad (2.25)$$

with solution $E(x) = E_0 e^{-x/X_0}$ where X_0 is the radiation length, the distance after traveling through which the electron would retain $1/e$ of its original energy (Equation 15.49 in [51]):

$$X_0 \equiv \left[4N \frac{Z^2 e^2}{\hbar c} \left(\frac{z^2 e^2}{Mc^2} \right)^2 \ln \left(\frac{233M}{mZ^{1/3}} \right) \right]^{-1} \quad (2.26)$$

Electrons in air have $X_0 \approx 37 \text{ g/cm}^2$. Given an atmospheric overburden of approximately $5\text{-}7 \text{ g/cm}^2$, this implies that electrons would radiate approximately 13% of their total energy as Bremsstrahlung if the electron illuminated the detector the entire time it was in the atmosphere. As electrons illuminate the detector for very short sections of their trajectory, typically only a few km out of the thousands of km they spend inside the atmosphere, the typical fraction for most trajectories is probably much smaller than 13%. Certain trajectories which illuminate the detector when the electron is also near detector altitude (and therefore where the electron experiences the greatest number density of air atoms) could experience this order of magnitude of Bremsstrahlung energy loss.

Figure 2.5 shows the simulated Bremsstrahlung photon energy spectra as a fraction of primary electron energy for 10 TeV and lower energy primary electrons.

2.3 Summary and Next Steps

Primary signal electrons illuminate the CREST detector plane in a unique manner. Simple geometrical models, such as that conceived of by Professor Akerlof, are of great service in illuminating our understanding of these events. In Chapter 6, I present more detailed schematics of the electron's trajectory in hopes of further clarifying the nature of its interaction with the detector plane, the atmosphere, and the geo-magnetic field. But before any detailed simulation or vector algebra, this model was of great use in helping tune the digitize system's coincidence threshold, and the expected number of signal photons. By working in the context of this model in concert with signal Monte Carlo simulations, I was

Bremsstrahlung Photon Energy Spectrum

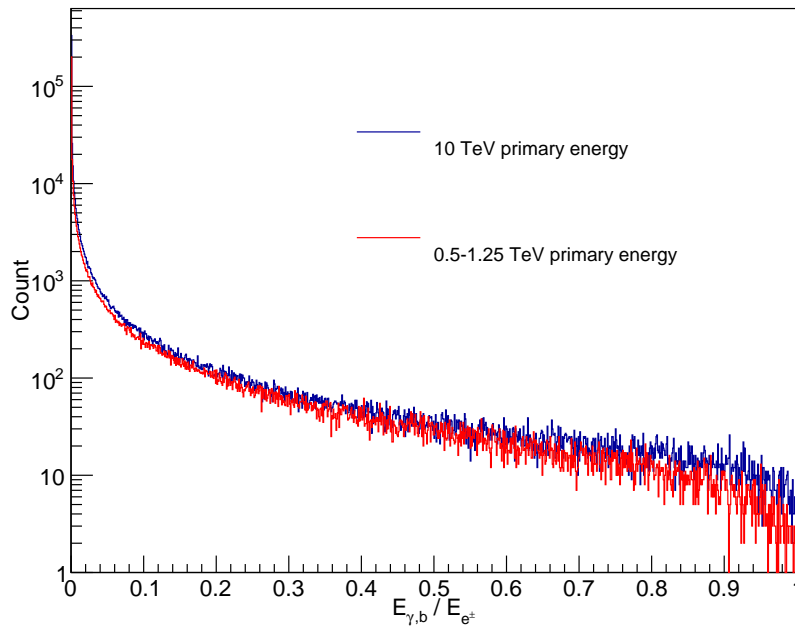


Figure 2.5: Simulated Bremsstrahlung photon relative energy spectra as a function of the primary electron energy illustrating three key features. First, the spectrum falls off rapidly at low energies. Second, the spectrum is nearly uniform at higher energies. Third, the relative spectrum is nearly completely independent of primary electron energy. The **blue** plot shows the spectra of photons emitted by 10 TeV primary electrons traversing Earth's atmosphere. The **red** plot shows the same for a electron energies ranging from 500 to 1.25 TeV. Plot generated from simulated CREST signal events using CRESTMag8, developed by Scott Nutter.

able to realize and communicate the important role traversal- and near-traversal-type events would play in CREST's attempt to observe and identify signal events.

CREST's signal event morphology consists of its response to synchrotron photons, Bremsstrahlung photons and other secondary products resulting from the electron's own interaction with the detector, in addition to the interaction of the electron itself. The relative importance of these three sources of signal hits depends intimately on the detailed properties of the signal electron's trajectory. Specifically, the strength of the perpendicular component of the magnetic field affects the power of synchrotron emission and the energy of the resulting synchrotron photons. The amount of air traversed by the electron while it illuminates the crystal plane determines the power of Bremsstrahlung emission. Trajectories that take the electron near, or even through, the detector are likely to result in many more Bremsstrahlung photons interacting with the detector. Finally, such trajectories greatly increase the probability that the electron itself will interact directly with the detector.

In Chapter 4, I address the nature of these interactions, including the amount of energy so deposited and the consequent response of the data acquisition system. Prior to that discussion, I describe the CREST detector itself, in Chapter 3. As mentioned above, I will at length return to the analysis of such signal events, and our attempts to identify them against a fierce background of side-going protons, in Chapter 6. But it should already be abundantly clear (and of little surprise) that CREST's detection method lies in a nascent twilight between direct and indirect detection methods. Our most valuable signal hits (from synchrotron photons) have indirect origin. Yet, our most frequent (and potentially, most valuable, as I suggest in Section 6.8) signal events occur when the primary electron itself passes through our detector plane. To succeed in our goal of using CREST to identify signal electrons, we should draw on the best methods from both camps.

CHAPTER 3

The CREST Detector: Design and Implementation

In this Chapter I describe the CREST detector, from mechanical, electronic, physical and thermal perspectives. The brainchild of a collaboration stocked with veterans of multiple successful balloon campaigns, CREST's design and implementation satisfies a multitude of competing constraints, ranging from high-altitude operation, to sub-nanosecond synchronization and timing resolution across a large, finely-segmented detector plane. While I joined the CREST collaboration after most of the design choices had already been made, I was afforded the opportunity to contribute to its integration, debugging and testing. Much of the information presented here derived from before my involvement is therefore necessarily largely a compilation of information made available by CREST collaboration members, including my predecessor, Dr. Atsushi Yagi.

3.1 Overview: Meeting Physical and Operational Constraints

CREST's design is motivated by the combination of constraints resulting from attempting to detect synchrotron photons using a balloon-borne instrument. Detecting the synchrotron photons emitted by electrons traveling through the Earth's atmosphere and geo-magnetic field against the entire primary and secondary background fluxes of photons and charged particles imposes the following constraints. I review the physical principles motivating these requirements separately in Section 2.1.

1. Energy resolution and containment sufficient to detect and estimate the energy of the signal synchrotron photons with energies ranging from \sim keV - \sim MeV
2. Time resolution on the order of a fraction of the light crossing time of the detector (about 6ns)




Figure 3.1: Overview schematic of the CREST detector. Visible are 1,024 of the cylindrical crystal/PMT assemblies (in grey), partially obscured by the transparent green veto plastic scintillator paddles. Not visible are those paddles on the bottom and rear two sides of the instrument; all together the paddles provide approximately 100% geometrical coverage of the detector plane defined by the crystals. Courtesy of R. Northrop, CREST lead mechanical engineer (University of Chicago). The area below the bottom veto would later house the CREST electronics and flight computer, the power conversion and storage system, and the NASA-provided flight integration package.

3. Sufficient spatial extent to detect a sufficient number of signal synchrotron photons in each event
4. Spatial resolution sufficient to resolve a linear track of activated detector elements from a signal photon
5. Discrimination of signal-like events induced by the passage of the numerous background charged particles, especially protons
6. Either a very large effective detector area or a very long exposure time to compensate for the very low expected signal flux (on order dozens of events for a flight lasting weeks, depending on the flux model assumed)

Balloon-borne instruments must operate at high altitude and therefore in a low temperature and low pressure environment. Flight durations of several weeks or more rule out detectors utilizing substances or devices requiring additional logistical support. Thus otherwise suitable detectors requiring sustained low temperatures - for example liquid Xenon time projection chambers such as LXeGRIT [9] or super-conducting magnetic tracking detectors such as HEAT [32] - would either greatly increase construction and operational costs or limit the maximum flight duration.

Tracking chamber detectors utilizing conventional, rather than superconducting, magnets are suitable for longer flights, but by definition have detector area equal to some fraction of the magnet's cross-sectional area, and with a limited geometric acceptance. Spaceborne experiments such as the AMS-02 experiment [2] aboard the ISS would provide sufficient flight duration (on the order of years) to make up for the smaller area and acceptance. However, magnetic tracking chamber detectors fail to resolve the charge sign of leptons above some critical energy. Above this energy, the transverse deflection of the electron's trajectory due to the magnetic field is comparable to the deflection due to multiple scattering within the tracking media itself. Given current tracker technology, the energy region of interest (several TeV and above) remains inaccessible to these detectors.

Since CREST's detection technique (detailed in Section 2) does not require that the primary electron traverse the detector itself - only a subset of its secondary synchrotron photons must do so - the effective area increases with energy. Because the synchrotron power increases with the square of the primary electron energy, the fraction of total primary electron trajectories which result in a sufficient number of detected synchrotron photons grows roughly linearly. The combination of long duration flight, large effective area, high acceptance and increasing sensitivity with primary energy make CREST's detection scheme

ideal for detecting such a low expected signal flux without the expense of a space-borne detector.

To satisfy the timing and photon containment constraints with a temperature- and pressure-tolerant material scintillating crystals made of Barium Fluoride (BaF_2) were chosen as the primary photon-sensitive detector. To provide a means to distinguish between crystal signals generated by photons and the charged particle background an organic plastic charged-particle veto was chosen to surround the array of crystal scintillators, covering 99% of trajectories by solid angle.

To convert scintillation light from the crystals and veto paddles to electrical signals, fast, low-power, high-gain, vacuum-potted photomultiplier tubes were chosen. Equipped with Cockroft-Walton voltage amplification generators, the high voltages (thousands of volts) necessary to power the dynode chains are up-converted from lower voltage power supplies inside each PMT base, minimizing the increased risk of electrical discharge inherent to high-pressure, low-pressure environments and obviating the need for heavy, bulky and expensive high-voltage rated cabling and connectors.¹

An electronics system capable of operating and supporting these sensitive detector components would have to meet the following constraints:

1. Sub-nanosecond pulse height discrimination of the prompt component of the BaF_2 scintillation light
2. Collection, amplification and digital conversion of integrated analog pulse heights for energy estimation from the delayed component of the BaF_2 scintillation light
3. Long-distance (on the order of meters) error-free digital communication pathways for readout, control and monitoring of over 1000 detector modules and other sensors
4. Uninterrupted quiet provision of operator-tunable control voltage to over one thousand PMT assemblies
5. Fast (on the order of 100 ns) instrument-wide trigger logic and signaling
6. Solar power generation and on-board energy storage sufficient to meet average of 1kW of power consumption
7. All equipment able to withstand strong accelerations (up to 10g) associated with transport, launch, parachute deployment, landing and recovery, low temperatures (down to approximately -40°C) and low pressures (below 50 torr)

¹On the other hand, the voltage down-conversion (from the 35-40V input provided by the solar power system to the $\sim 5\text{V}$ and below powering the majority of CREST's electronics components) incurs significant power losses, resulting in significant heat generation.

8. Communication to and from NASA satellite communications equipment for ground control and data download
9. Sufficient stability and data storage for weeks of independent operation at balloon altitude
10. Thermal stability without access to conductive or convective cooling, constant thermal monitoring to take steps to avoid potential overheating

CREST's electronic support system meets these constraints with a system of custom-built electronics boards equipped with Field Programmable Gate Array (FPGA) technology. FPGAs are programmable parallel processing state machines capable of accepting multiple inputs, changing the state of internal parameters and simulated memory sites and providing multiple outputs on each clock cycle. FPGA clock speeds can vary greatly based on cost and application demands; though these clock speeds are generally slower than typical CPU clock speeds, FPGAs' parallel nature makes them more suitable for fast distributed hardware applications such as CREST. The clock frequencies used in CREST's boards ranged from the 40MHz 'heartbeat' provided to each STAC board by the Crystal Overlord (CROL) to the 125MHz clock provided internally on each STAC board to the 200MHz clock generated on the Digitize FPGA (DIGOL) which powered the instrument-wide digitize decision logic. The parallel nature of the FPGAs enabled CREST to simultaneously operate, read data from and send commands to hundreds of electronics components with data transfer occurring at these clock rates. For example, the FPGA acting as the instrument trigger processes 64 binary input signals every clock cycle and outputs an instrument-wide trigger decision a few 5ns clock cycles later to 72 separate read-out boards. It would require a battery of CPUs or serial micro-controllers with clock cycles on the order of a GHz to achieve this, likely at much higher expense, higher power consumption and reduced (or no) customizability.

CREST boards utilize two distinct FPGA solutions provided by integrator Opal Kelly, Inc. Opal Kelly boards used by CREST embed Xilinx Spartan S-3 FPGAs along with features such as RAM banks, USB (for interaction with a PC via a customizable Application Program Interface (API)) and JTAG (for firmware programming) interfaces, Programmable Read-Only Memory (PROM) chips for automatic re-programming on FPGA reset and a host of easily accessible test probe points, among other features. CREST electronics boards embed a total of over 150 XEM-3050/3010 boards and their diagnostic features greatly accelerated CREST design and testing.

High-speed error free digital communications between central FPGAs and distributed read-out electronics boards takes place along the CRESTBus, VOLTBus and Digitize path-




Figure 3.2: An Opal Kelly XEM-3050 board, integrating a Xilinx Spartan-3 4,000,000 logical gate FPGA with various additional features. CREST utilized four such boards, along with 72 Opal Kelly XEM-3010 boards (lower capacity chips with 1,500,000 gates each). From www.opalkelly.com.

ways, described in detail in Section 3.3.1.

The heat management concerns peculiar to balloon flight prompted the outsourcing of the thermal plan to experienced Scott Cannon of NMSU and analyst Doug Bell. The plan they recommended was implemented with few modifications. It entailed conducting heat to the surface of the aluminum chassis where a dual-layer material (aluminized mylar or silverized teflon) would radiate excess heat to the environment while reflecting as much incoming solar radiation as possible. I review this scheme (as well as the procedure and results of a thermal/vacuum test performed at NASA’s Plum Brook facility) in Section 5.1.2.

The equipment, support and expertise (including the balloon itself) provided by the Columbia Scientific Balloon Facility (CSBF) to facilitate communications between CREST and satellite communications pathways was crucial to the success of the CREST mission. I detail CSBF’s many contributions in Chapter 5 along with the rest of CREST’s Antarctic campaign and the various stages leading up to our collaboration’s anti-polar journey.

3.2 Instrument Design

CREST has two sets of sensitive detectors: the crystal array and the charged particle veto system. These detectors and their associated readout electronics sit atop a $3.22\text{ m} \times 2.66\text{ m} \times 1.62\text{ m}$ rectangular aluminum support chassis. Figure 3.1 shows the veto system in transparent green surrounding the crystal detector plane (white/grey tubes). Below this the frame interior provides mounting points and structural support for the various systems described herein. This space enjoys significant structural protection during periods of intense mechanical stress, as evidenced by the near-intact state of the instrument (including the SFC and its data disks) observed during post-flight instrument recovery.



Figure 3.3: Lead mechanical engineer Rich Northrop and CREST's aluminum gondola which he designed and built.

3.2.1 Hierarchical Design

CREST's structural, electronic and operational design share a common hierarchical theme as I detail here.

The crystal system consists of 1,024 crystal/PMT assemblies capable of detecting the impact of X- and γ -rays with energies between between 40 keV and 40 MeV and charged particles with a wide range of energies.

Sixteen such couplings arranged in a 4x4 grid form the basic building block of the CREST crystal plane known as a Sedecim module. Each Sedecim module is equipped with three electronics boards for pulse height discrimination (AFE board), instrument-wide trigger input/output, signal time and energy estimation (STAC board) and PMT high/control voltage control and readback (SVI) board. I describe the functions of these electronics components in more detail in Section 3.3.

Chains of eight Sedecim modules integrated electronically and structurally comprise sub-units of the detector referred to as a crystal bus or C-channel. The name derives from the C-shaped aluminum shell cradling each unit of eight Sedecims; the open face of the C is oriented downwards. Figure 3.12 shows a schematic of a cross-sectional view of a C-channel with the crystal/PMT assemblies on top of the aluminum support structure and the readout electronics suspended below, inside the curve of the C. Figure 3.4 reproduces a schematic drawing of one such C-channel, showing crystal/PMT assemblies in light brown, their lead jackets in dark grey and the aluminum C-channel in light grey. Eight such C-




Figure 3.4: Mechanical design of a crystal C-channel (also referred to as a bus) containing one row of eight Sedecim modules. Units are in inches. Courtesy of R. Northrop, CREST lead mechanical engineer, University of Chicago.

channels in a row, forming a square, comprise the crystal plane. Figure 3.5 shows such a channel being assembled at Indiana University. Each bus has 16 PMT/crystal pairs per Sedecim \times 8 Sedecim modules per bus yielding 128 PMT/crystal pairs per bus. Eight buses with 128 PMT/crystal pairs each makes for 1,024 PMT/crystal pairs total. Figure 5.53 shows a fully assembled C-channel after being lifted out of the instrument for repairs in the LDB Hangar 096 at the LDB facility near McMurdo, Antarctica.

The eight crystal buses are surrounded by the veto system: a 99.12%-hermetic box composed of thin plastic scintillating paddles coupled via embedded fibers and fiber optic light guide arrays to pairs of PMTs, one on either side of the paddles.

3.2.2 BaF₂ Scintillating Crystals

CREST's $2.4 \times 2.4\text{m}^2$ detector plane is partially covered by an evenly spaced square 32×32 grid of crystal/PMT assemblies with lattice separation parameter in both directions of 7.5



Figure 3.5: Chuck Boyer and Alex Shroyer assembling a crystal C-channel at the Indiana University high bay in Fall 2010. On the blackboard to Alex's right the hexadecimal numbering scheme for tubes within a Sedecim has been drawn.



Figure 3.6: One fully assembled crystal C-channel rests on top of CREST's aluminum chassis. Picture courtesy of R. Northrop. Each silver tube is one crystal/PMT assembly.



Figure 3.7: Picture of two 1 cm thick and one 2 cm thick BaF₂ crystals supplied by Proteus. Picture by A. Yagi [66].

cm. The crystals have the shape of cylindrical disks, 5 cm in diameter and 2 cm in height. Combined, the 1,024 disks comprise a sensitive area (when viewed from directly above) of approximately $1024 \times 0.00196 \text{ m}^2 = 2.01 \text{ m}^2$, covering approximately 35% of the detector plane. Spatial resolution is on the order of the size of each disk.

BaF₂ is a high-density, high-light yield scintillating crystal with fast and slow scintillation components. The fast component has a rise time on the order of tens of picoseconds, a decay time of 0.8 ns, and produces approximately 18% of the total scintillation photon energy (yielding approximately 1,800 photons/MeV). The discriminators on the AFE boards are designed to latch on this prompt component, thereby enabling ns-level timing. The slow component has a 630 ns decay time and produces the remaining 82% of the scintillation photon energy (yielding approximately 10,000 photons/MeV). Table 3.1 (partially from [8]) lists the relevant properties of the crystals.

3.2.3 Photomultiplier Tubes

Hamamatsu-manufactured R7724CW-ASSY photomultiplier tubes with two-inch diameter borosilicate windows were used to convert the scintillation light emitted by the crystals to electrical current pulses. These tubes have a gain of 5×10^6 and a chain of ten dynodes. High voltage is provided individually for each tube by means of a low-power, vacuum-potted Cockroft-Walton generator. CREST used one anode tap in a configuration which resulted in a timing resolution of about 0.75 ns (as demonstrated by my predecessor, Atsushi Yagi, using a test setup he constructed at University of Michigan [66]).

The borosilicate PMT windows are opaque to the ultraviolet light (at 195, 220 and 310

BaF ₂ Crystal Property	Value
Diameter	5 cm
Thickness	2 cm
Mass	192 ± 1 g
Density	4.89 g/cm ³
Rise time of fast component	~ 10ps
Decay time of fast component	0.8 ns
Decay time of slow component	630 ns
Photons/MeV, fast component	1800
Photons/MeV, slow component	10000
Scintillation photon wavelength	195-220 nm
Average excitation and ionization potential (<i>I</i>)	375.9 eV
Minimum ionization energy loss	1.326 MeV / (g/cm ²)
Radiation length	9.91 g/cm ²
Critical energy (e ⁻ , e ⁺)	13.78, 13.34 MeV

Table 3.1: Properties of the BaF₂ scintillator crystals used in CREST’s main detection grid. 1,024 such crystals arranged in a square 32x32 grid and coupled to custom-built Hamamatsu PMTs comprise CREST’s main detector plane. Compilation of information from CREST design and testing documents, [39] and [8]. For a calculation of photon interaction probabilities over a range of energies and for various geometries see Section 4.9.

Hamamatsu PMT Property	Value
Model	R7724CW-ASSY
Weight	499.8g
Full Length	170.0mm
Window diameter	2.0"
Supply voltage (to CW generator)	5V
Max operating voltage	2,000V
Control:Output Voltage Ratio	1:1,000
μ -metal thickness	0.032"
Dynode stages	10
FWHM Resolution @ 661.7 keV	10.13%
Nominal gain	5 × 10 ⁶
Maximum gain	8 × 10 ⁶

Table 3.2: Properties of Hamamatsu photomultiplier tubes used to collect scintillation light from the BaF₂ crystals. From [8] and [66].




Figure 3.8: A Hamamatsu R7724CW photomultiplier tube, like those reading out CREST's 1,024 crystal/PMT assemblies (and its top veto paddles).

nm) emitted by the BaF₂. To avoid using expensive quartz windows, the waveshifting compound Tetraphenyl butadiene (TPB) was deposited onto the crystal surfaces. This waveshifts the incoming ultraviolet light into 430 nm visible light which could then traverse the PMT windows without attenuation.

3.2.3.1 Lead Shielding

Each PMT/crystal assembly is wrapped in a 0.4 cm-thick lead shield designed to minimize multiple secondary scattering between PMTs resulting from a single incident noise or background photon. This adds approximately 500 g to the weight of each assembly, or over 500 kg to total instrument weight and resulted in a significantly lower altitude (and greater atmospheric overburden) than would have occurred without this extra weight. The purpose of the lead shield was to reduce the rate of secondary scattering between crystals.

3.2.3.2 Sedecim Construction

The PMT/crystal assemblies are supported by 0.14 g/cm³ polyethylene foam blocks grouping the assemblies in a 4x4 square grid with side length of 30 cm. A common set of electronics boards (STAC, AFE and SVI boards) perform readout, slow control and monitoring for all 16 assemblies. This unit is referred to as a Sedecim module. The dark grey




Figure 3.9: Schematic of the crystal/PMT assembly. Not shown is the lead shielding subsequently installed to reduce secondary hits migrating from crystal to crystal. Reproduced from [66].

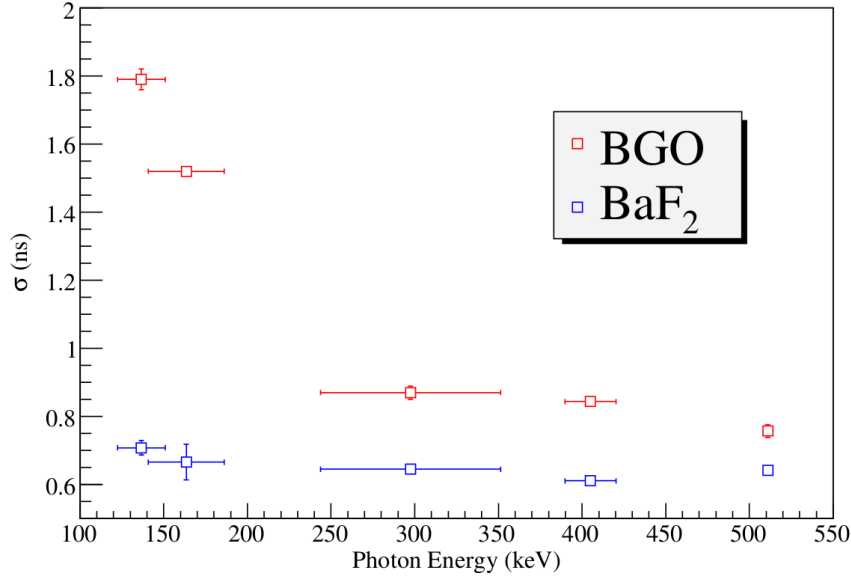


Figure 3.10: Initial tests of the Hamamatsu PMT’s timing capability when used to detect scintillation photons from BGO and BaF₂ crystals performed by Atsushi Yagi. Picture reproduced from [66].

rings near the top of each tube depicts the lead shielding added to each tube to minimize inter-tube scattering. See Figure 3.11 for a schematic depiction.

3.2.4 Veto System

The veto paddles are constructed from EJ-200 plastic scintillator produced by Eljen Technology. EJ-200 consists of Anthracene-doped Polyvinyltoluene (chemical composition $C_{27}H_{30}$).

The PMTs used to read out either end of the veto paddles were recruited from a mixture of the custom Hamamatsu tubes described above and Burle PMTs which had also been evaluated for use in the main crystal array. Two such Hamamatsu tubes reading out bottom veto paddles are shown in Figure 3.15. The plastic material and thickness was chosen to maximize the light response of the paddles to charged particles while minimizing the light response to X - and γ - rays. In this manner a veto signal preferentially indicates the passage of a charged particle such as a muon, electron or proton (or their anti-matter partner particles) over the passage of a photon. To achieve hermetic coverage while minimizing weight and cost, the paddles were made as large as practical. This reduces the number of PMT assemblies and readout/control electronics boards required to provide hermetic coverage of the entire crystal array. As an unintended consequence, the veto system has a greater spatial resolution (compared to the crystal system) when measuring the location




Figure 3.11: Mechanical design of a 4x4 Sedecim module. Units are in inches. Courtesy of R. Northrop, CREST lead mechanical engineer, University of Chicago.

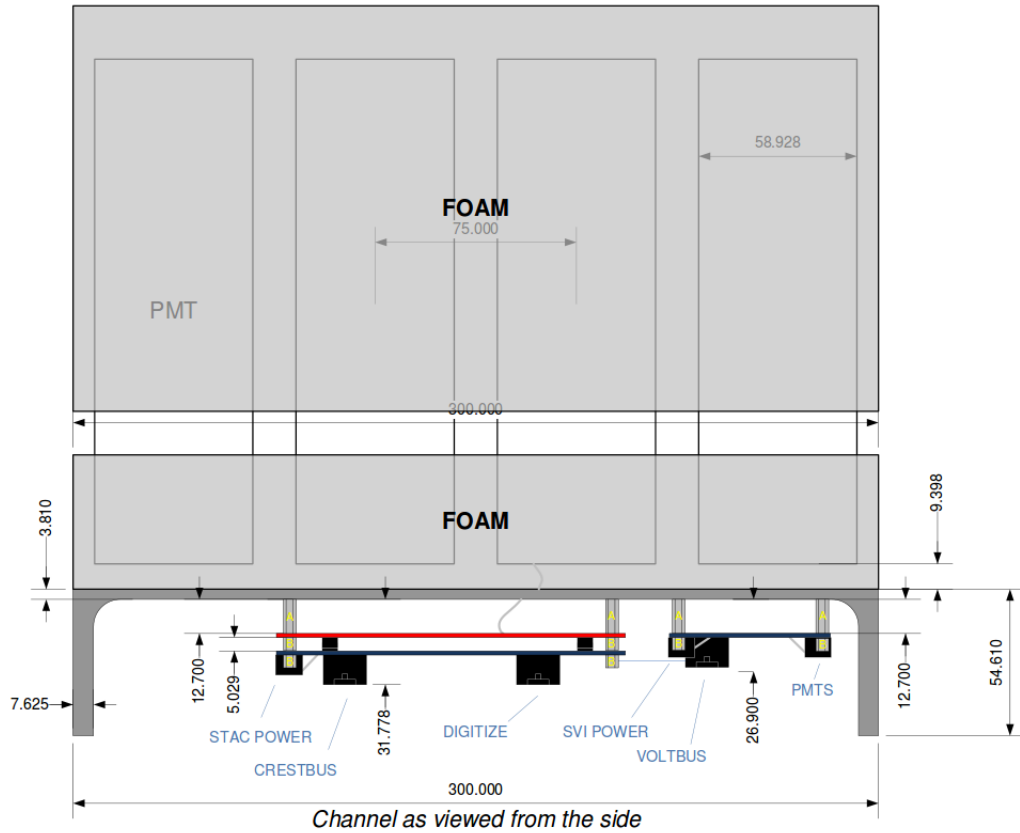


Figure 3.12: Detailed side elevation of a 4x4 sedecim module. Units are in millimeters. The grey outlines of the crystal/PMT assemblies show their foam-ensconced positions relative to the readout electronics. The red rectangle is an edge on view of a AFE board, which houses the discriminator, charge sampling and charge amplification circuitry. Just below that is the blue edge-on view of a STAC board, which interfaces with upstream electronics, generates a clock capable of time-stamping hits with 1ns precision and digitizes analog voltage levels on the AFE corresponding to charge signals. Just to the right, the shorter blue rectangle is an edge-on view of an SVI board, which sets PMT control voltages and provides the input to the Cockcroft-Walton high voltage generators located in the base of each PMT. Schematic courtesy of J. Ameel, CREST lead electrical design engineer, University of Michigan.

EJ-200 Scintillator Property	Value
Density	1.032 g/cm ³
Refractive index	1.58
Light output, % Anthracene	64
Photons/MeV (e ⁻ incident)	10,000
Max. emission wavelength	425 nm
Rise time	0.9 ns
Decay time	2.1 ns
Pulse width FWHM	2.5 ns
H atoms/cm ³	5.17 × 10 ²²
C atoms/cm ³	4.69 × 10 ²²
e ⁻ /cm ³	3.33 × 10 ²³
Average excitation and ionization potential (<i>I</i>)	64.7 eV
Minimum ionization energy loss	2.022 MeV / (g/cm ²)
Radiation length	43.90 g/cm ²
Critical energy (e ⁻ , e ⁺)	94.11, 91.62 MeV

Table 3.3: Properties of the EJ-200 plastic scintillator used to construct CREST's charged particle veto system. The material consists of a polymer base of Polyvinyltoluene (C₂H₃O) doped with the fluor Anthracene. From Eljen Technology EJ-200 data sheet, http://www.eljentechnology.com/images/stories/Data_Sheets/Plastic_Scintillators/EJ200%20data%20sheet.pdf, accessed on 4/5/15. The average excitation and ionization potential *I* is discussed and estimated in Section 4.1.

of hits in the veto paddles. Veto hit location estimation is still possible, but is based on comparing the arrival times of pulses in paired PMTs. This results in spatial resolution parallel to the paddles' long axes on the order of the time resolution of the veto PMTs multiplied by the speed of light in the fiber optic light guides. This distance uncertainty is approximately 30-40 cm - imprecise compared to the resolution of the crystal detector elements which have spatial resolution comparable to the radius of each crystal (2.5cm). In the direction perpendicular to the paddles' long axes, the spatial resolution is approximately half the width of the paddle. Most paddles are between 20 and 30 cm wide.

3.2.5 Pulsar System

To assist in timing calibration, light from an Agilent HLMP-CB15-R000 blue light LED was piped to every Sedecim module via a system of fiber optic fan outs. The SFC controlled the LED's on/off state and pulse duration and intensity. Periodically, CREST's flight software would cease normal data taking and activate the pulser briefly, providing a pulse of light in each PMT offset by a known value proportional to the length of fiber between the pulser LED and each PMT. During testing of this system (by myself and J. Musser) a noticeable amount of deadtime (which effectively reduces the instrument signal rate) was incurred when flashing the entire crystal array simultaneously. The pulser




Figure 3.13: Veto paddle showing paddle base with grooved channels containing fiber light guides. Courtesy of M. Geske, PSU.




Figure 3.14: Hamamatsu PMT and attached FOLGA, designed to mate with the intra-paddle fibers and provide an efficient pathway for photons trapped in those fibers to the PMT. Each paddle had one of these assemblies on each end. Courtesy of M. Geske, PSU.




Figure 3.15: Hamamatsu PMTs wrapped in black Tedlar tape reading out bottom veto paddles.

run was broken into three segments in an effort to circumvent this behavior.² In segment one, only C-channels 1-4 were active (not masked by the CROL). Masked channels operate normally (i.e. high voltage is provided and STAC modules still record timestamps for hits that latch discriminator circuits), but that the CROL will not address them when performing data readout. In segment two, only C-channels 3-6 were active, and in segment three, only C-channels 5-8 were active. By combining the data sets in analysis, a time calibration map (consisting of offsets between pairs of tubes) of the entire detector was generated (see Section 6.4.2).

3.3 Readout Electronics

The CREST electronics components were primarily designed, routed and commissioned for fabrication by Jon Ameel, lead electronics engineer at the Electronics Shop at the University of Michigan. Collaborating engineers based at multiple institutions, including Casey Smith (lead development and testing of the initial versions of the CROL, VTOL and STAC firmware) at the University of Chicago until 2010, Mark Gephard (who redesigned the discriminator circuits on the AFE boards) at Indiana University and Michael Lang (who performed much of the wiring and testing during assembly) at Indiana University. J. Musser and B. Kunkler at Indiana University also contributed to the final versions of the STAC, CROL and VTOL firmware, particularly by rectifying long-standing problems in hit timestamping and readout algorithms.

Given the high rate of signals and the extremely short duration of signal events (on order nanoseconds) the readout electronics needed to measure time signals from individual PMTs before deciding, instrument-wide, whether or not to record an entire event. This led to the design (primarily by my advisor, Professor Greg Tarlé) of an asynchronous, trigger-less minimum bias system utilizing both the fast and slow components of the BaF₂ scintillation light (with time constants of 0.8 ns and 630 ns - see Table 3.1).

Here I explain in what manner the data acquisition system is said to be triggerless. Any PMT signal sufficient to latch a discriminator circuit generates a hit timestamp (accurate to within 1 ns); such signals derive from the fast component of the BaF₂ light response. The charge generated by the slow component was collected and stored in a bank of multiplexed capacitor channels. STACs report the presence of these hit timestamps to the digitize system (described in Section 3.3.7, which then compared the timestamps so generated by all

²It was decided at the time that the instrument would never be asked to process freeze cycles with most or all of the crystals activated, thus the exact reason why it could not do so was never fully investigated. I ask that the reader keep this in mind when I discuss the instrument performance degradation during the Antarctic flight in Section 5.4.6.

1,024 crystal channels with each other on a timescale shorter than the decay time of BaF₂'s slow component. Based on settable coincidence parameters, the digitize system would then instruct the instrument to do nothing (allowing the collected charge to dissipate and the discriminator circuits to reset), or to initiate instrument-wide readout. This command (called a "digitize" command) instructs all instrument STAC/AFE pairs to generate hit energy measurements by sampling the stored charge deriving from the slow component of the BaF₂ light response. Once complete hits ("complete" meaning having both time *and* energy measurements) have been generated at the STAC level, a central data collection module (the CROL) reads out all instrument hits, collects them into a list and reports them to the SFC. In this manner, each individual detector channel is said to be "triggerless," in that the fast component of every detected signal results in a hit timestamp. Thus while individual channels are triggerless, data from the instrument as a whole (consisting of all the complete hits observed by any PMT channel from a certain block of time) reported by the detector modules to the CROL is gated by an instrument-wide trigger mechanism (the digitize system).

3.3.1 Hardware Components and Communication Scheme

Figure 3.18 provides an overview of the hierarchy of electronic components. At the top level is the Science Flight Computer, assembled by Michael Schubnell from extended-temperature and -pressure rated components and using the Scientifix Linux operating system (version 4.5). It houses four solid state disks each with 0.5 TB of storage and linked to the motherboard via Serial ATA.

The SFC communicates with the XEM Carrier Board via USB connection. On board the XEM Carrier are four Opal Kelly XEM-3050 4,000,000 gate FPGAs. Each FPGA runs a firmware state machine serving a specific purpose described below. Each Endcap board (12 total, one for each of 8 crystal buses and 4 veto buses) terminates a CRESTBus and VOLTBus connection originating from the XEM Carrier board. Also terminating on each Endcap is a Digitize signaling cable, meaning that there are a total of 36 cables for CRESTBus, VOLTBus and Digitize between the XEM Carrier and the 12 buses.

Applications (i.e., the flight software process) running on the SFC interact with the FPGAs onboard the XEM-Carrier by means of the Opal Kelly-provided API in order to carry out various commands intended to set instrumentation parameters such as PMT control voltage, discriminator thresholds and many others.

At the Sedecim level there are three separate types of boards. I describe the hit time-stamping and charge estimation process (performed by the STAC and AFE boards) in detail

in Section 3.3.5.

1. STAC boards house an XEM-3010 board, embedding a Xilinx Spartan-3 1,500,000 gate FPGA. They communicate with the CROL via CRESTBus to receive commands and upload hit time and energy data.
2. AFE boards house the pulse height discrimination circuits and the charge collection chips. These boards have no FPGA and operate at the command of the STAC to which they are mated.
3. SVI boards feed voltage inputs and read back voltage outputs in order to maintain each PMT assembly at nominal operating gain. They communicate with the Master Voltage Controller (MVC) via the VOLTBus communications channel.

The following list details the data and communications signals handled by each pathway.

- CRESTBus
 - Hit readout (STAC → CROL)
 - Housekeeping data (temperature, voltage levels, etc.) (STAC → CROL)
 - Control commands (CROL → STAC)
- VOLTBus
 - PMT Control voltage (CV) setting (VTOL → SVI)
 - PMT High voltage (HV) (Cockroft-Walton generator input) setting (VTOL → SVI)
 - PMT HV and CV readback (SVI → VTOL)
- Digitize
 - Individual STAC Hit buffer status (Single-/Multi-hit) (STAC → DIGOL)
 - Instrument-wide Digitize enable/disable (DIGOL → STAC)
 - DIGOL clock provision (STAC → DIGOL)

3.3.2 CRESTBus and VOLTBus

Communications between Sedecim component boards and central FPGAs aboard the XEM Carrier take place over two Low Voltage Digital Signal (LVDS) pathways, CRESTBus (between CROL and STAC boards) and VOLTBus (between MVC and SVI boards).³ CRESTBus and VOLTBus are LVDS communications pathways originating on the XEM Carrier board, routing through the Endcap board at one end of each crystal bus and each veto box and then daisy-chaining through Sedecim component boards. The cables transmitting these signals are 8-strand individually shielded copper braid, cut to length and terminated by hand using a custom made jig and externally sourced Molex connectors at the UM E-Shop under J. Ameel's supervision by myself and support staff in the E-Shop during Summer and Fall 2010. CRESTBus and VOLTBus constitute tree-like structures. For example, to communicate with the Mth STAC in the Nth bus, the CROL signals the first STAC in the Nth bus via CRESTBus channel N. Bus N STAC 1 receives and re-transmits the signal (at full power) to the bus N STAC 2 and so on to bus N STAC M. These bi-directional communications paths carry all signals between Sedecim modules and the central detector control FPGAs, including the instrument-wide synchronized 40MHz clock, data and diagnostic mode read-outs, power-on and reset signals, slow control commands (such as PMT control voltage or discriminator threshold changes) and temperature, voltage and current monitoring. The sole exception to this is digitize signaling.

3.3.3 Channel Endcap Boards

Figure 3.19 depicts one of CREST's Endcap boards. An Endcap resides on either end of the crystal C-channels (veto boxes have only 1 endcap board). These boards act as power supplies and cabling junction points. The sun-side Endcap (that side closer to the XEM Carrier) powers four of the Sedecim modules and feeds through CRESTBus, VOLTBus and Digitize signaling to all eight Sedecim modules. The far-side Endcap merely powers the other four modules. Since high voltage differences across short distances in a low-pressure environment could easily lead to electronic breakdown, Endcaps convert power from 28V instrument power to two power levels: 5V (for the SVI boards) and 3.3V (for the STAC and AFE boards). To achieve this they each have three DC-DC converter chips. A similar chip used on the solar power conversion module is shown in Figure 3.20. CRESTBus and VOLTBus cables from the XEM Carrier terminate on the sun-side Endcap; their signals are passed through the Endcap to the first STAC/SVI in the chain of 8 comprising the

³VETOBUS is essentially the same thing as CRESTBus, except that it provides a communications path between the CROL and STACs reading out veto, rather than crystal, PMTs.

C-channel. From there the CRESTBus and VOLTBus signals are propagated down the C-channel in a daisy chain. Figures 3.25 and 3.26 display such a configuration used to test the STAC, AFE and SVI boards at the UM E-Shop in Winter 2010. Not displayed there are the digitize cables. Digitize signals travel between the sun-side Endcap and XEM Carrier, then between the Endcap and each STAC individually with matched-length cables. The Digitize signal propagation time between STACs on a given C-channel are nominally identical. The digitize cable lengths between the XEM Carrier and each of the 8 sun-side Endcaps were different, with channels 3 and 4 having the shortest length, 2 and 4 having a median length, and 1 and 8 having the longest.

3.3.4 Digitize Signaling

CREST's individual channels, being triggerless, record a timestamp for every signal detected by the PMTs. As a consequence of the high rate of incidence of background X - and γ -rays, most of these signals will not form part of a signal event. For example, it is impossible for an event with only one signal within 100 ns to form a detectable signal event. The detector module that made the determination of whether a time-pattern of hits qualifies as potential signal is called the Digitize System. The Digitize System decides, based on two sets of criteria, whether a time-pattern of hits warrants hit energy measurement and subsequent presentation to the CROL for hit histogramming. The inputs and outputs to this system are referred to here as the digitize signaling pathway. This pathway provides for communication directly between the STAC boards and an XEM-3050 FPGA, referred to here as the DIGOL, mounted in the XEM Carrier. Unlike the CRESTBus and VOLTBus daisy-chain cabling scheme, the digitize cables connect directly to each STAC after passing through the Endcap boards. The cables have matched lengths between Endcaps and STACs (but not between the DIGOL and Endcaps).

There are three types of digitize signals. Each STAC board is assigned one signal of each type. Two of these signals represent input to the DIGOL from each STAC about the hits it has recently observed: Single Hit (henceforth 'SH') and Multi Hit. The other type of signal is an output from the DIGOL to each STAC called 'Digitize.' I will resume discussing the details of the digitize system in Section 3.3.7 after reviewing the details of CREST's readout scheme.

3.3.5 Front-End Electronics Readout

The original readout design called for the PMT anode (that is, after the 10th dynode stage) to be routed to the discriminator circuitry and taps at the 5th and 8th dynode stage to be

routed to the charge collection chips. This would provide a convenient means of differential amplification input. Testing of the Hamamatsu tubes revealed that the dynode taps had an irreducible 6mV ripple with a frequency of 220 kHz.⁴ It was thus decided to forego the planned utilization the dynode taps and split the anode signal, sending one copy to the discriminator circuitry and another copy to the charge collection chips. The differential amplification for the separate high and low gain charge collection channels was instead performed on a re-designed AFE board.

Each Sedecim module's AFE board has 16 independent discriminator circuits, one for each PMT. These discriminators compare the voltage arising from the anode current to a programmable voltage threshold. These thresholds were tuned in McMurdo prior to launch to be equal to half the voltage arising from one photoelectron initiating an electron cascade in the PMT dynode chain. When the voltage threshold is crossed the discriminator latches and sends notification of such to the STAC board. Each of the 16 PMTs has its own dedicated discriminator and there are therefore 16 separate edge inputs to the STAC timestamping module. In data analysis, this timestamp value attached to each hit is referred to as **TDC**, for Time-to-Digital Conversion.

The STAC board compares the time it receives this signal to the clock state of four on-board Clocked Programmable Logic Device (CPLD)s, each handling the rising and falling edge of four 8 ns-period clock signals offset from each other by 1 ns. The state of the eight clock edges (four rising and four falling) encodes each 8 ns period into 8 separate logic states, thereby enabling 1 ns resolution time stamping for each discriminator latching event. It was identified as early as March 2010 that the firmware implementation controlling this behavior faced many challenges in maintaining the 1 ns offset between the clock circuits and maintaining the phase relationship between STACs. For example, adjacent STACs in a chain were frequently as far out of phase as 180 degrees. J. Musser re-investigated the issue in Spring 2011 and learned that the timing constraints used by the auto-routing algorithm in the Xilinx firmware compiler were insufficient to guarantee offsets on nanosecond time scales. He was able to solve the problem through a combination of hand-routing FPGA signal paths (an intricate task normally handled by the Xilinx firmware compiler) and a feedback loop between the STAC FPGA and the CPLDs. Figure 3.16 shows the final firmware architecture (combining the original design by C. Smith and the feedback loop J. Musser introduced).

The AFE boards are mated to the corresponding STAC boards and house two IDEAS

⁴Though the source of this ripple was never conclusively determined, based on the increased impact of the noise on the less-amplified taps it is likely that the source of the ripple was noise from the oscillation circuits in the Cockcroft-Walton generators.

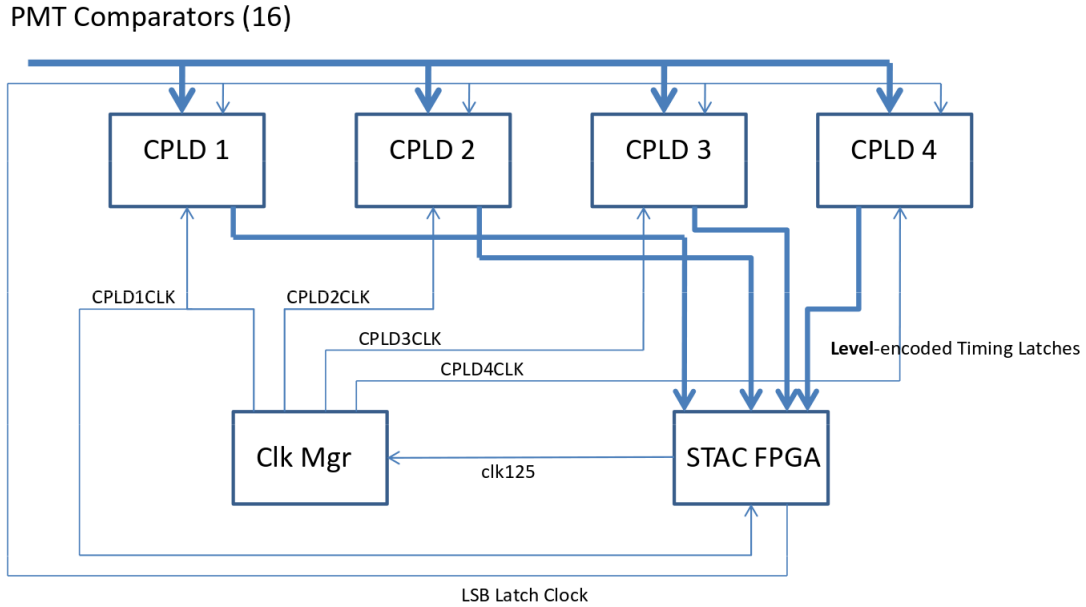


Figure 3.16: Firmware architecture for encoding the 8ns STAC FPGA clock signal into 8 1ns blocks, thereby enabling signal readout with ~ 1 ns time resolution. Architecture by J. Musser and C. Smith. Diagram by J. Musser.

VA32_HDR11 ASICs which are operated by signals from the XEM 3010 FPGA on the STAC board. Figure 3.23 shows one such pair. Each of the two Viking ASICs onboard the AFE (one for the low gain channel and one for the high gain channel) integrates and stores the slow component of each PMT's anode output in 32 independent charge collecting channels. The chips feature a multiplexing addressing system which enables sampling one channel while the others continue to collect charge. This continuous operation feature reduces deadtime. The charge thus collected on each chip is used for signal energy estimation in one of two energy ranges. The VA chip preceded by the low-gain amplification channel estimates energy over a sensitive range from 40 keV to 2 MeV and the other with the high-gain amplification circuitry covers a sensitive range of up to 40 MeV. This provides for an overall dynamic range after combining both signals of approximately 1:1000 (see Section 4.3.1 for a summary of the energy calibration procedure using both signals).

The collected charge samples decay over time with a programmable time constant on the order of hundreds of μs . The time since charge collection is known and used to estimate the initial charge collected from the decayed charge (a process referred to as 'charge shaping' performed onboard the VA ASICs using shaping parameters settable at run time via STAC commanding). This helps prevent the charge decay from influencing the estimate of the size of the slow-component signal (and therefore the estimate of the energy of the photon or particle that originally struck the crystal scintillator). The low-gain channel

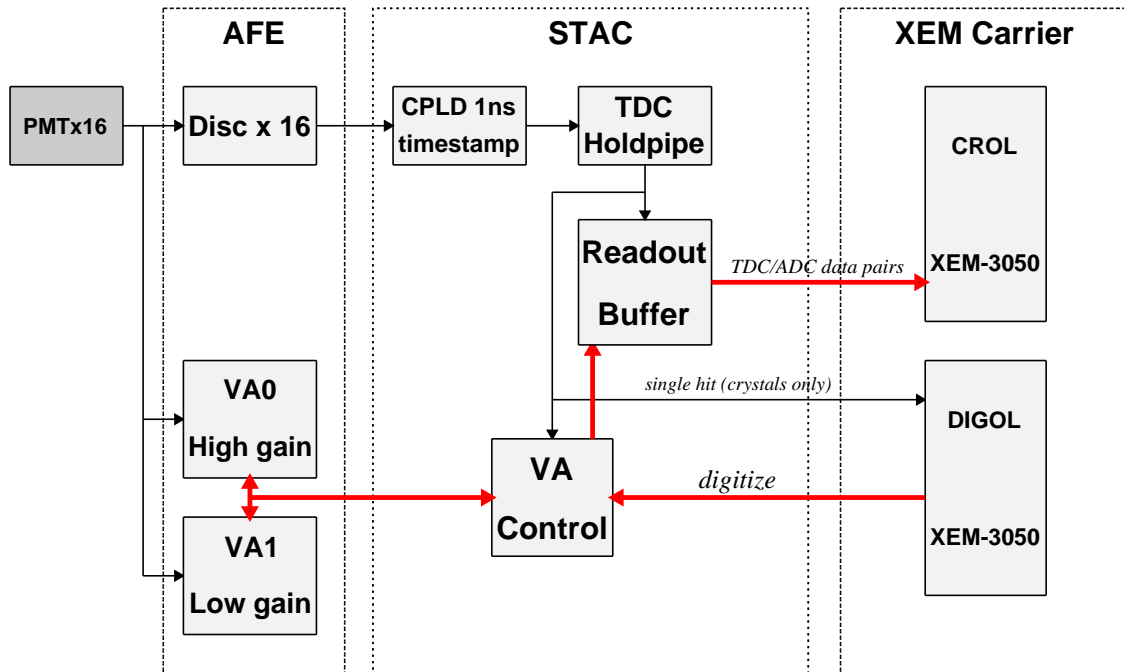


Figure 3.17: Simplified schematic of the readout architecture integrated with the Digitize decision. Each Sedecim module's 16 PMT anode outputs feed into the 16 independent discriminators on the AFE. Small boxes represent firmware modules or physical board components. Large dashed boxes indicate the carrier board in which the boxed modules reside. **Heavy lines (red online)** indicate signals that only proceed if the DIGOL has initiated an instrument-wide Digitize command (see Section 3.3.7 for details). Latched discriminator signals are timestamped (with a TDC word) by the STAC's CPLD array with 1ns time resolution. These TDC words are sent to a delay line (the TDC Holdpipe) to await their fate. The STAC also reports to the DIGOL that it has received a hit. The DIGOL integrates this SH with all other SH signals from all STACs and over the last 100 ns. If this integrated signal meets the coincidence trigger, the DIGOL issues a Digitize command. If the VA Control module receives this command *and* a TDC word from the hold pipe, it will then sample the charge accumulated on the VA0 (high gain) and VA1 (low gain) charge collection ASICs on the AFE. It digitizes these analog charge estimates into the ADC0 and ADC1 words, pairs them with the TDC word on the readout buffer. The next time the CROL polls this STAC for data the readout buffer will transmit all the TDC/ADC data words built in this manner via CRESTBus. Veto Sedecims read out in an identical manner, except that the *single hit* signal is not sent to the DIGOL and therefore veto hits do not contribute to the DIGOL's coincidence trigger.

and high-gain channel Analog-to-Digital Conversion (ADC) measurements are referred to during data analysis and energy calibration to as ADC1 and ADC0, respectively.

Depending on the mode of operation chosen by the instrument operator, the STAC will either “digitize” collected charge a set time after generation of the hit timestamp or await an instruction from the Digitize Decision logic to do so (these two states are referred to as the “Command” mode of the STAC boards - either they command themselves or are commanded by the Digitize FPGA). “Digitizing” collected charge means halting the charge collection for a particular VA chip capacitor channel and instead discharging that capacitor to measure the amount of collected charge.

Thus in one sense the readout system is triggerless: STAC/AFE board pairs register and generate a timestamp for every PMT anode signal that latches a discriminator. But only when those timestamps collectively meet the trigger decision (as determined by the digitize system settings) are charge measurements for those hits then digitized and pairs of timestamps and charge amounts (TDC and ADC words) reported up the readout electronics chain. During flight the crystal and veto STACs were set to external commanding mode. That is, they only read out collected charge when prompted by the DIGOL.⁵

Thus are the STAC boards in each Sedecim module tasked with building lists of pairs of digital words, each pair consisting of a digital time word (TDC value) and the digital charge amplitude words (the ADC0 and ADC1 values). It will continue doing so until either the FIFO buffers become half full or the CROL instructs it to send the list of time and energy words as part of its Accumulate cycle, at which time it will relay the times and charge amplitudes of hits that have occurred and been recorded since the last Accumulate instruction.

3.3.5.1 Integration of Commanding and Readout: Crystal and Veto Overlords

The CROL is a firmware process operating in an XEM-3050 Opal Kelly FPGA housed on the **XEM Carrier Board** (which also houses the Veto Overlord (VTOL), MVC and DIGOL). The primary interface between the software layer (SFC) and hardware layer (STACs), the CROL facilitates input and output between the crystal STACs and the SFC. STAC-level control commands are issued by the flight software running on the SFC to the CROL, which then relays them to the crystal STACs via CRESTBus. The CROL module also issues many STAC commands as part of its normal operation.

The VTOL is a clone of the CROL. Just as the CROL commands and reads data from the crystal Sedecim modules via the CRESTBus communications channel, so does the

⁵For more details on the interaction between crystal and veto STACs and the DIGOL see Section 3.3.7.

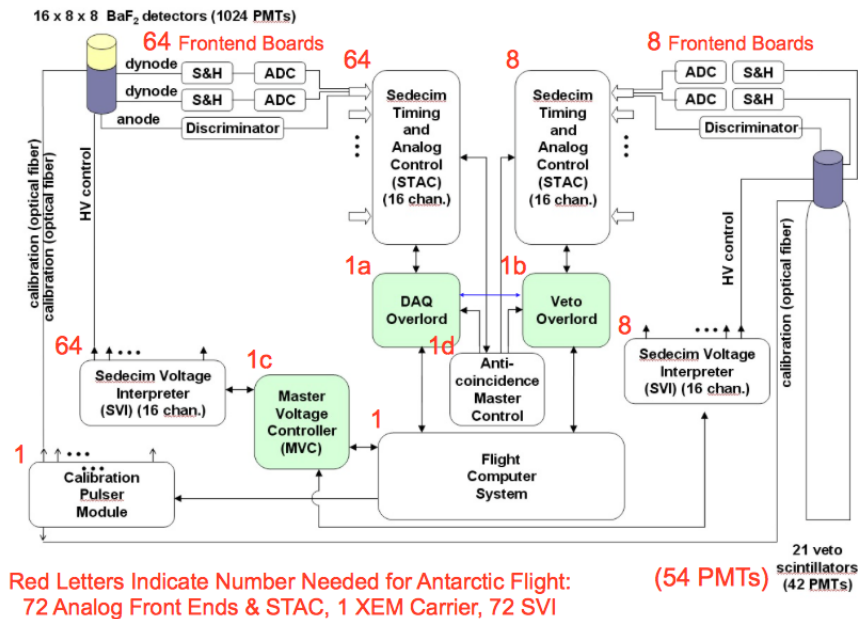


Figure 3.18: Diagram by M. Schubnell showing the schematic relationship between hardware components of the CREST detector. Note that the dynode signals from the PMTs were not used as they did not meet the specified noise requirements and that the number of veto scintillator panels is 27 (for 54 veto PMT channels) and not 21 as shown. Also not shown is the Digitize sub-system (at the time the diagram was made the hardware pathway existed but was not being utilized) which represents an additional independent two-way signaling pathway between the DIGOL (mounted on the XEM Carrier Board) and the 72 Sedecim modules.

VTOL do so for the veto Sedecim modules via VETOBUS. The key difference between the two is that the VTOL does not report data directly to the SFC. Instead it provides the data to the CROL which then combines it with the crystal data before relaying it to the SFC. The volume of data words the CROL handles contributes to deadtime; in our experience freeze cycles with below approximately 500 or so hits were handled well, and freeze cycles with more than this were not. A rapid source of many simultaneous veto hits (e.g. a light leak or higher-than-expected flux of low energy charged particles) could greatly increase the number of hits the CROL would have to process in each freeze cycle. This could overwhelm instrument-wide data readout and contribute to the degradation in detector performance during flight (see Section 5.4.6 for further discussion).

The CROL maintains two separate buffers for the combined TDC and ADC words for all 64 crystal sedecim modules. While one buffer stores collected data words (during an Accumulate process) the other would be frozen (a Freeze process) against data input and instead is used as data output to the histogramming module of the CROL firmware. This data is histogrammed in 180 ns blocks. Empty (those with fewer than 2 TDC words in the 180ns window) histograms are discarded and the surviving histograms are built into 'Freeze Cycles', a binary data block which the SFC could then record to a file in binary format. For more information on the data pathway after this point refer to Section 6.3.

3.3.6 PMT Voltage Control: SVI Boards

In addition to a STAC and AFE board pair, each Sedecim module also has an SVI board responsible for setting, maintaining and reading back the High Voltage input to each PMT's integrated Cockroft-Walton generator and the Control Voltage used to fine-tune the gain of each PMT. The associated currents are also monitored and can be read out via the MVC for housekeeping purposes. Each of these voltage control circuits is repeated 16 times on each SVI, one for each PMT channel. The control voltage output is highly accurate, within approximately 0.5 mV of the desired setting.

3.3.7 Digitize Decision: Development, Testing and Flight Operation

I describe the detailed operation of CREST's trigger logic, internally referred to as the Digitize System, and how to use it as a first layer of data reduction in my Digitize Decision Manual [43], provided to the CREST collaboration in early 2011. Here I summarize my involvement in implementing and testing CREST's digitize decision logic with the assistance of J. Ameel and J. Musser.

As of January 2011, a first attempt at building the Digitize firmware state machine that




Figure 3.19: A channel Endcap board with metal support plate custom designed by the author. 16-pin CRESTBus and VOLTBus connections to the XEM Carrier are on the far left, along with CRESTBus and VOLTBus chain connectors. Power is input via the left facing connector, converted to the 5V and 3.3V required by the SVI and STAC/AFE boards, respectively, then output along the Sedecim chain. The power converter chips are hidden beneath the board; one such chip is shown in Figure 3.20. The 64-pin facing connector interfaces with the DIGOL FPGA on-board the CROL, and the 8 rightmost connectors send Digitize commands and receive Single/Multi-Hit signals to/from each STAC board. These cables were of matched length to minimize variations in trigger behavior between STACs.




Figure 3.20: A DC-DC converter chip identical to those used on the channel Endcap boards.




Figure 3.21: The crimping jig used to attach the drop connectors mid-strand of the CREST-Bus and VOLTBus cables. All such cables were purchased as a single strand, cut to length, then fitted with two endpoint connectors (one for each Endcap) and 8 drop connectors (one for each Sedecim module in each C-channel) at the Michigan E-Shop.




Figure 3.22: The UM Test Stand used to test a single C-channel's worth of electronics. Essentially a 1-channel version of the XEM Carrier it houses an Opal Kelly XEM-3050 FPGA (lower left) with USB connection to a controlling PC. In middle front the CRESTBus and VOLTBus cables for the single attached C-channel depart the test stand.




Figure 3.23: Lower front: one AFE(red)/STAC (blue, underneath the AFE) mated combination. The 16 thin grey cables provide the PMT anode signal to the 16 discriminator channels on board the AFE. Upper rear: an SVI board with 16 voltage control and read-back channels. This is the test mode in which the UM E-Shop confirmed the read-out functionality of each STAC, AFE and SVI board sent to Indiana University for instrument integration. In total UM provided 100 functioning copies of each board.




Figure 3.24: Another view of the test stand pictured in Figure 3.23 also showing the terminating channel Endcap board (far right) and the connection to the neighboring STAC/AFE/SVI Sedecim unit (far left).




Figure 3.25: Another view of the test setup in Figures 3.23 and 3.24 showing a series of STAC/AFE/SVI Sedecim modules connected in daisy chain fashion.




Figure 3.26: Another view of the test setup in Figures 3.23-3.25 showing 7 Sedecim modules chained together. The angular red device in the left foreground is the Xilinx tool used to program every FPGA on board each STAC in the C-channel via the JTAG chaining scheme devised by J. Ameel. This enabled us to rectify STAC firmware issues after instrument integration by interfacing only with the first STAC in each C-channel, saving us from weeks of delay in re-assembling the instrument that would have otherwise been necessary.




Figure 3.27: CREST's readout electronics mounted in the instrument below the crystal C-channels. Each C-channel's boards are linked in a fashion identical to those in the test stand pictured above. Visible are the green stickers on each board indicating that they had been tested and received Michigan E-Shop Seal of Approval. There were a total of 64 crystal PMT readout assemblies (one for each Sedecim) and an additional 8 dedicated to veto PMT readout. Picture by R. Northrop.



Figure 3.28: The veto “honeycomb” before mounting to the instrument, upside-down, supporting four of nine bottom veto paddles with FOLGA’s mated to their PMTs. The white veto boxes are the veto system’s equivalent of the crystal C-channels; each would later hold four STAC/AFE/SVI electronics modules. Picture by R. Northrop.

would operate in the the DIGOL had been made but never completed or tested in concert with the STAC firmware. The collaboration tasked myself and Jon Ameel with development of a new version of the “Digitize Decision,” firmware for operating the DIGOL in concert with the STAC boards that would implement a real-time instrument-wide trigger decision.

Every internal clock cycle, each STAC’s XEM-3010 firmware state machine sends updated single-hit (SH) and multi-hit (MH) signals along the Digitize signal cable to the DIGOL based on the hits recently⁶ recorded by that STAC. If the STAC had no recent hits to report, both SH and MH would be false. If the STAC had recent hits from only one PMT channel to report, SH would be true but MH would be false. If the STAC had recent hits from more than one PMT to report, both SH and MH would be true.⁷

When STACs are set to internal commanding mode, PMT discriminator latches generate a timestamped data word which the STAC FPGA holds in a pipeline, referred to in the STAC firmware architecture as a “holdpipe.” The latch event allows the PMT anode current from the slow component of the BaF₂ scintillation light to charge up capacitors on the two

⁶‘Recent’ is defined by a parameter in the STAC firmware that determines the length of the delay loop (‘hold pipe’) that contains TDC timestamps as they await pairing with ADC energy estimates.

⁷MH inputs to the DIGOL were ignored in the flight digitize decision implementation.

VA chips in the AFE. This charge accumulates as it trickles in but is sampled and digitized a set time later. In other words each STAC decides on the basis of its own recent pattern of discriminator latches when it should sample and digitize accumulated PMT anode current. This mode of operation would likely induce deadtime at two stages.

The first potential source of deadtime could be incurred at the individual PMT channel level. Reading out VA chips is a lengthy (of order microseconds) process and prevents additional charge accumulation on the PMT channels being sampled. Given an expected combined background and signal event rate of tens per freeze cycle (on order ten events per hundred microseconds) this would mean that individually triggered sampling activities taking microseconds would occasionally (if rarely) overlap, resulting in deadtime on individual PMT channels.

The second potential source of deadtime would be more severe as it would be instrument-wide. The CROL has two buffers to store hits read out from STACs. When accumulating data to one buffer, it freezes the other as it histograms its data and prunes empty freeze cycles. If it is still histogramming one buffer when the current freeze/accumulate cycle ends, deadtime is introduced while it finishes processing the current buffer before unfreezing it for data read out. Given sufficient hit rates on any individual PMT, this behavior could overwhelm the data histogramming capabilities of the CROL. Thus the digitize system acts as the first line of defense against blindingly high rates (either from *bona fide* signals or those arising from malfunctions such as light leaks or noisy tubes).

These two sources of deadtime were defended against by configuring the STACs to operate in external commanding mode. In this mode STACs accrue hit timestamps (TDC words) as they do in internal commanding mode, but do not sample and digitize accumulated charge unless they receive a timely digitize command from the DIGOL. Receipt of this command readies the STAC for the exit of the TDC dataword from the hold pipe delay loop upon which time the STAC will initiate two ‘sample and hold’ procedures (one for each VA chip) to measure the accumulated charge on the matching PMT channel. Digitizing these charge measurements generates corresponding ADC0/ADC1 data words for the TDC dataword. The STAC places these combined datawords in a buffer which, when requested by the CROL, would be relayed to the CROL via CRESTBus. If on the other hand a TDC word exits the hold pipe without the Digitize command having been received, the STAC does not sample the accumulated charge and the TDC word vanishes without further consequence. For a schematic summary of this operating mode as it affects the readout procedure see Figure 3.17.

In what follows I assume that STACs are set to external commanding mode, which was the case in the flight configuration. The only exception is during pedestal housekeeping

runs - STACs are set to internal commanding for the duration of such runs. Thus our algorithm used $64 \times 2 = 128$ potential input signals to make an instrument-wide “yes” or “no” decision to read out the TDC words in STAC holdpipe limbo.


The main constraint this decision algorithm had to meet was time. STACs held data words in their internal hold pipe for a certain amount of clock cycles. Initially this was set to approximately 23.8 ns clock cycles, meaning 184 ns after the hit was assigned a TDC by the STAC CPLD clock division module, the hit would either be fully read out or would be discarded. The STAC would need to receive the Digitize command from the DIGOL some time before the 23rd clock cycle in order to rescue the hit.

Part of the difficulty in implementing the previous digitize algorithm turned out to be that this duration was insufficiently greater than the combined signal travel and processing time of the SH and MH signals from STACs to the DIGOL and vice versa. The signals not only had to be received and processed by the DIGOL, but also had to travel a round-trip distance of ~ 4 meters of cabling to reach the most distant STAC boards. This timing window could be increased by modifying the STAC firmware, but the longer the STAC waits to sample the charge stored by the capacitor banks inside the dual VA chips, the more severe a correction would have to be made by the charge shaping modules attached to those chips. I measured the dispatch and arrival of the relevant signals on the DIGOL and test points reflecting the internal state of the XEM-3010 board embedded in a test STAC board with the goal of finding a holdpipe length that would satisfy both conditions. That is, the optimal holdpipe delay would be long enough such that STACs would hold on to timestamps before receiving a digitize command, and short enough so that the collected charge would not have dissipated by a significant amount. Table 3.4 details the signal chronology for the test stand setup I used.


Figure 3.29 shows a typical signal timing diagram generated by probing the electronics test points with a 1 GHz oscilloscope. (Some signals have normal logic where high voltage indicates true, others have inverted logic where low voltage indicates true). For demonstrative purposes the coincidence requirement for initiating instrument-wide charge readout was set to two unique STACs.

(a) shows the DIGOL receiving a Single Hit signal from one STAC (negative pulse in the top trace, blue online), causing the ‘stachit’ signal to become true approximately 20 ns later (positive pulse in the top-middle trace, light blue online). This indicates that one (and not more than one) STAC was hit 5 ns later (positive pulse in the bottom-middle trace, green online). Since the coincidence threshold was two, the DIGOL correctly refused to initiate a digitize command (lack of pulse in the bottom trace, magenta online).


(b) shows a similar situation except that a second unique STAC reports a hit approx-



(a)



(b)



(c)

Figure 3.29: Oscilloscope-generated timing diagrams of digitize signaling within and between the DIGOL on board the XEM Carrier and a STAC board’s XEM-3010 FPGA. Horizontal divisions represent 40ns. Each digital signal is represented by an analog voltage level between 0 and 3.3V. See discussion in Section 3.3.7 for details.

Signal	Location	Incremental delay (ns)	Cumulative delay (ns)
Hit timestamped	STAC	0	0
SH signal dispatched	STAC	20	20
SH signal received	DIGOL	40	60
Digitize command sent	DIGOL	60	120
Digitize command received	STAC	50	170
Hit charge sample initiated	STAC	8	178

Table 3.4: Chronology of signal dispatch and arrival times for a single STAC connected to the DIGOL aboard the XEM Carrier used to optimize firmware timing parameters on STAC boards. Signal transmission along cables up to two meters long would only increase certain delays (between SH signal dispatched/received and between Digitize command dispatched/received) compared to these measurements (generated using a digitize cable of $\sim 12''$ in length). Therefore timing parameters were chosen to ensure STACs would hold on to hits longer than the minimum 180ns implied by these measurements.

imately 20 ns after the initial STAC hit report (the bottom-middle trace, `solohit`, is false briefly while the lookback time spans both STAC hits). This satisfies the unique STAC coincidence requirement of two and the bottom trace, `digitize_2H` pulses positive, indicating that the DIGOL has initiated an instrument wide charge readout to all STACs with hits residing in their hold pipes.

(c) shows the series of events inside the STAC XEM-3010 FPGA following receipt of an instruction from the DIGOL (negative pulse in the top-middle trace, [light blue](#) online) to collect the charges for all hits stored in its hold pipe. Eighty ns after receipt of this signal the ‘accept Hit’ window opens (bottom middle trace, [blue](#) online) until the hit pops off the end of the STAC’s hold pipe (positive pulse in the bottom trace, [green](#) online). These two signals trigger the STAC XEM-3010 FPGA to initiate a charge collection action upon both VA chips which will generate ADC0 and ADC1 digital charge estimates to pair with the TDC (time estimate). Together the charge and time estimates comprise a data word which will later be read out to the CROL via CRESTBus.

This iteration of the Digitize firmware used the framework of the previous iteration (called ‘Decretum’ and jointly developed by J. Ameel and other members of the UM E-Shop). J. Ameel and I jointly developed the single-hit counting lookback table method. The lookback table consists of a binary-valued two-dimensional array in memory space. The first dimension (visualized here and in Figure 3.30 as rows of the memory space) is time (measured in clock cycles) and the second (columns in memory space) is the STAC address number. A value of true (false) in row, column i, j indicates that i clock cycles ago STAC j had (did not have) a hit. Each location in this memory block is initialized to a false state in every memory address. The number of clock cycles kept in memory (that is the number of rows) determines the look-back period over which single hit signals

will be integrated. A look-back period of 20 clock cycles was chosen for flight operation, amounting to 100 ns of time integration. The coincidence threshold for the number of unique STACs reporting a single hit is settable via switches on the XEM carrier board. The possible coincidence thresholds ranged from 1 (digitization of every hit) to 7 (digitization only when 7 unique STACs had hits within the lookup period). These switches were set to a coincidence threshold of 3 for the flight. Thus during the flight whenever three or more unique crystal STACs reported hits to the DIGOL within 100 ns, every STAC (crystal and veto) would digitize hits.

The algorithm integrates the single hit signals from all crystal STACs in a manner detailed as follows. Each action is performed on each clock cycle of the DIGOL (once every 5 ns). In a parallel framework, none of these actions occur before or after any of the other actions on the same clock cycle. They are only guaranteed to take place after every action on the previous clock cycle and before every action on the next clock cycle. Thus actions that depend on the results of other actions (e.g. checking to see if we've met the coincidence threshold) must explicitly depend on the outputs of actions on previous clock cycles and not current clock cycles. Otherwise, unpredictable behavior such as race conditions which depend on the relative timing of signals within clock cycles can occur.

Here follows an overview of the digitize decision algorithm:

- **Advance:** Copy each row's logic state to the row above, thus advancing the state of each address forward by one clock cycle. Values in the row representing the most distant times are discarded.
- **Observe:** Record the logic state of the Single Hit signal for each STAC in the current clock cycle's row. The memory state (true or false) at each column in this row then represents the Single Hit state of each STAC board (present or absent) at some time in the past after correction for differential cable length delays between STACs.
- **Time integrate:** Calculate the logical OR of each memory column. The resulting row of truth values counts the number of unique STAC boards recording a TDC data word in a time corresponding to the lookback time.
- **Update:** Count the number of unique columns in the memory space containing a true value.
- **Prepare:** Compare this number to the Coincidence Requirement. If greater than or equal to it, raise the Digitize_2H logical value for this clock cycle to true. Otherwise, lower the Digitize_2H logical value to false for this clock cycle. In both cases retain the value for the previous clock cycle.

- **Broadcast:** The Digitize_2H flag now has four possible states corresponding to its truth value this clock cycle and last. When a rising edge is detected (previously false, currently true) an instrument wide Digitize flag is set (and held) to true. Otherwise nothing is done.
- **Stop broadcasting:** The Digitize flag is held true for 25 5ns clock cycles, or 125ns. Any STACs with hits during the period they see a true Digitize flag from the DIGOL will digitize them and later report them to the CROL. This includes STACs that have hits occur after the coincidence threshold is met and veto STACs (which never contribute to the coincidence threshold).

Other classes of events besides signal events are of interest to CREST. For example energy calibration depends on detection of various known spectral lines and it is possible that many events containing such useful information would not meet the coincidence threshold outlined above. Therefore separately from the coincidence requirement, another sufficient condition for instrument-wide readout (aimed at capturing such single hits at a reduced rate) is prescribed as follows. For every DIGOL clock cycle in which the unique STAC count reporting a single hit was different from zero, a counter (referred to as the pre-scale counter) was incremented. When this counter exceeded a certain value (the pre-scale value), a separate Digitize_1H signal was set to true. Rising edges in Digitize_1H would result in an instrument-wide readout command, just as for the coincidence requirement above. This parameter is settable via switches on the XEM Carrier board and sparsely covers the range of values from 1 to a bit over 32,000. For flight, this parameter (the pre-scale counter's maximum value) was set to 451.

Though all of the above put the Digitize system, and the STAC readout logic, in a functional state, one final roadblock remained between us and a flight-ready Digitize system: the 200 MHz clock provision for the DIGOL. When designing the digitize hardware the 200MHz DIGOL clock was prescribed to be routed from any STAC board to the DIGOL. Initial attempts to implement this clock path were unsuccessful. As a temporary workaround, an external interface, (i.e., a wire) was routed between the CROL and the DIGOL and an unused clock manager onboard the CROL was used as an input to the DIGOL internal clock multiplier.

In the interest of avoiding this irregularity, I attempted to utilize the clock from a STAC module as designed. When doing so I discovered the reason earlier efforts to use this clock input had been unsuccessful: a parameter on the STAC firmware which set the frequency of the clock input to the DIGOL had been set using units of picoseconds rather than nanoseconds. After changing the relevant parameter from 83,000 ns to 33.333 ns and set-

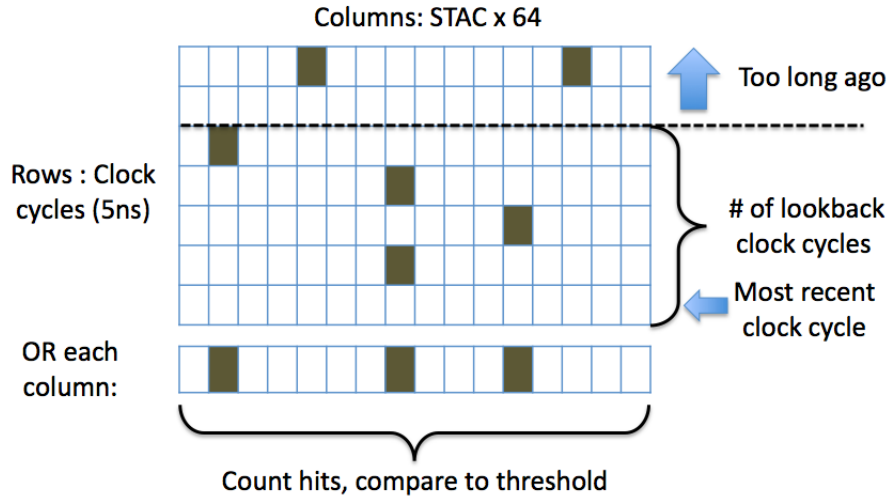


Figure 3.30: How coincidence is defined in the context of the Digitize system. The Digitize FPGA receives 64 singlehit signals (one from each crystal STAC). It also remembers for 20 5ns clock cycles the Single Hit history for each crystal STAC. To count the number of coincident single hits from unique STACs, it logical OR's together each truth value in each STAC's signal history, then counts the number of 1's in the resulting 64-bit array. If this number meets or exceeds the coincidence threshold (set to 3 for flight) then the Digitize FPGA will instruct every instrument STAC (crystal and veto) to digitize (read the charge resulting from) any hits it receives for a brief period.




Figure 3.31: Effect of modifying the digitize pre-scale value. On the x-axis is STAC ID number (for these runs crystal C-channels 2 and 8 were not operational). On the y-axis is the number of hits reported by each STAC in a data run taken over several minutes. The legend shows the pre-scale value used for each line; from highest to lowest rate the values were 1, 3, 65 and 257.




Figure 3.32: Effect of modifying the digitize coincidence threshold. On the x-axis is STAC ID number (for these runs crystal C-channels 2 and 8 were not operational). On the y-axis is the number of hits reported by each STAC in a data run taken over several minutes. The legend shows the coincidence threshold used for each line; from highest to lowest rate the values were 1, 2, 3 and 4.

ting the DIGOL's clock multiplier to 6.66, the desired stable 200 MHz local clock aboard the DIGOL was attained.

3.4 The SFC and the Flight Control Software

The Science Flight Computer (SFC) ran *crest_fs*, the flight software which acted as the brain of the CREST instrument. *crest_fs* was originally developed by Jim Musser. As he and I were the primary instrument operators during integration, we both constantly revised it as needed.

All commands received from the ground were routed through *crest_fs* and passed on to the destination hardware component along with the required parameters. *crest_fs* was also responsible for reading data from the CROL and MVC, converting it to a binary format, then writing it to one of four 500 GB solid state disks. This data included the TDC/ADC words that comprise signal data as well as housekeeping data polled from the CROL, VTOL and MVC FPGAs. The configuration of the instrument when it powered on or after a full reset was set by *crest_fs*, which read in configuration files for discriminator thresholds, STAC configuration states, PMT control voltages and innumerable other parameters.

crest_fs transmitted a selected subset of the signal data and housekeeping data to the

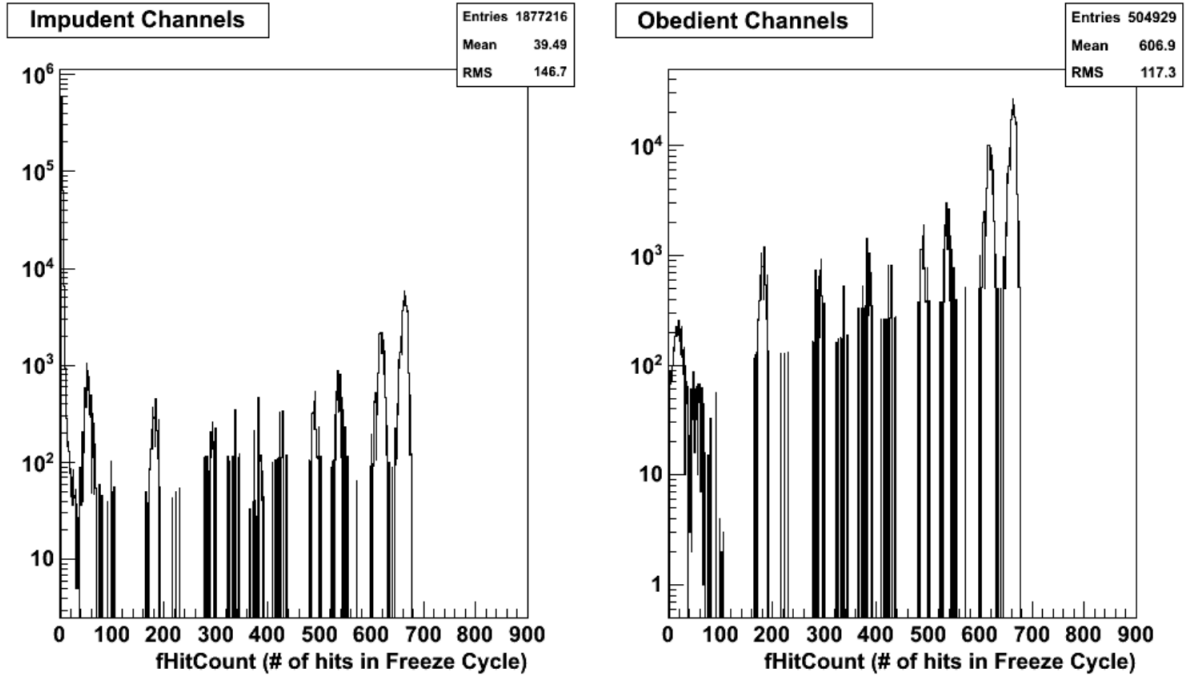


Figure 3.33: Histogram of the number of hits in individual Freeze/Accumulate cycles (periods of $131.074 \mu\text{s}$) from a pulser calibration data run with STACs operating in two distinct modes. In the left plot STACs digitize and report every hit to the CROL (‘internal commanding’). On the right STACs are set to the mode used in flight where they only digitize and report hits when instructed to by the Digitize FPGA (‘external commanding’). The DIGOL was set to a coincidence threshold of 7 unique STACs. Doing so eliminates events with fewer than 7 participating STAC boards.

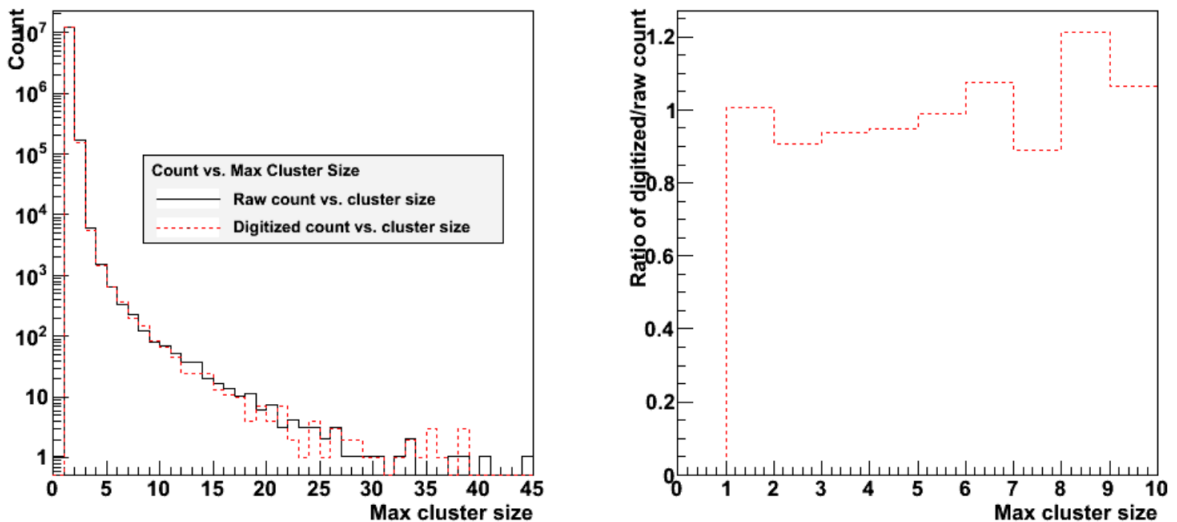


Figure 3.34: *Left:* Comparison of data taken with STACs in internal commanding (black line) and in external commanding with a coincidence threshold of 1 (red dashed line). ‘Cluster size’ is the number of STACs with signals in an individual freeze cycle. *Right:* Ratio of plots in left pane. An ideal distribution would be identically 1.

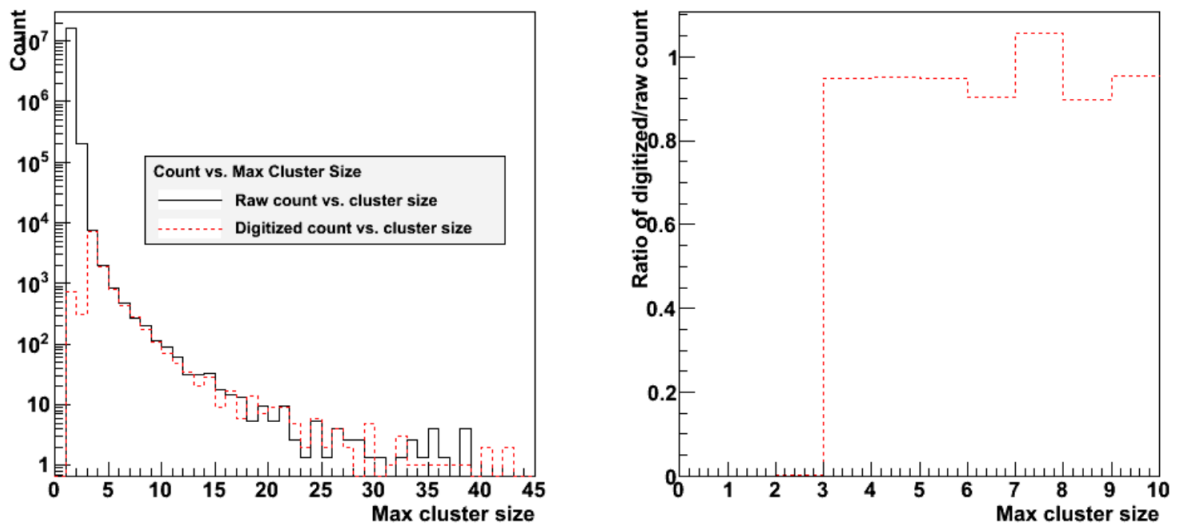


Figure 3.35: Identical plots to those in Figure 3.34 but with a coincidence threshold of three STACs instead of one STAC reporting a single hit. Freeze cycles with one or two STACs are greatly reduced in number. The remaining cases are likely straddling consecutive freeze cycles (so that hits on STACs contributing to the digitize decision were read out as part of the previous or next freeze cycle). The clusters with size one result from the single hit pre-scale set to its maximum value of 32000 (so that one out of every 32000 clock cycles that have one or more STACs active are read out as well).

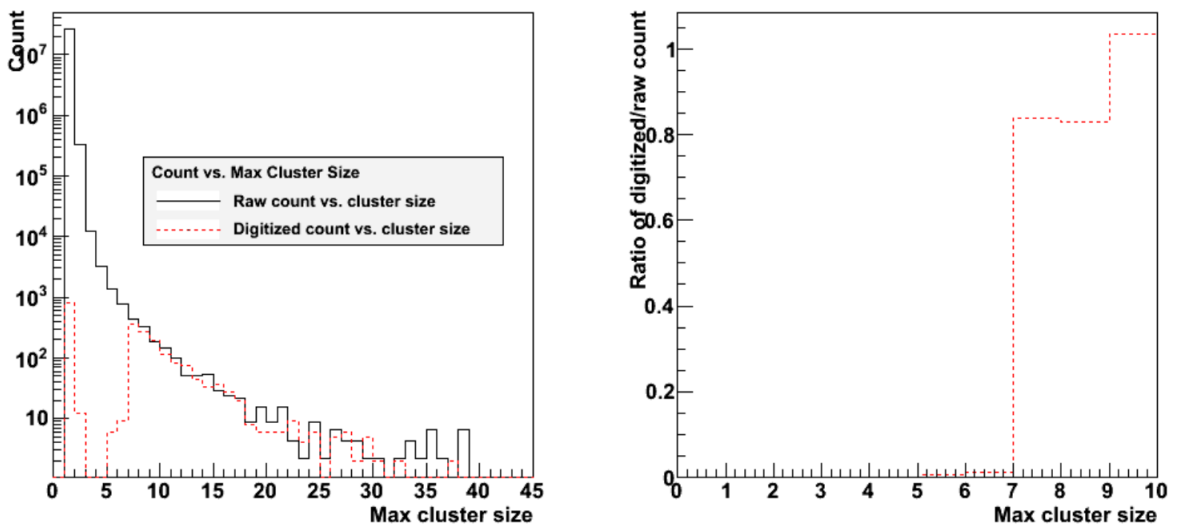


Figure 3.36: Identical plots to those in Figures 3.34 and 3.35 but with a coincidence threshold of seven STACs reporting a single hit, the strictest setting possible. Freeze cycles with fewer than 7 STACs participating are nearly completely eliminated. The remaining cases are likely straddling consecutive freeze cycles (so that hits on STACs contributing to the digitize decision were read out as part of the previous or next freeze cycle).




Figure 3.37: The Science Flight Computer assembled by M. Schubnell. On the left are visible the four 0.5 TB solid state disk drives using to store data during flight. Picture by J. Gennaro, McMurdo, Antarctica.

ground via CSBF's communications equipment. These pathways are described in Section 5.2.3. This data, especially the housekeeping data, informed shift operators' decisions with respect to resetting all or some of the instrument and its FPGA's, changing destination hard disk, changing data run numbers, etc. Figure 3.38 provides a very high level overview of *crest_fs*' control flow.

3.4.1 Pedestal Data Taking

One area of the flight software that I was tasked with developing was pedestal data taking. This amounts to measuring the ADC channel corresponding to a PMT's response to no incident signal. This channel would be subtracted off of ADC readings during energy calibration. To read out PMT's with no charge injection from the AFE-borne charge collection channels, STACs had to be put in a particular configuration state. However, putting all instrument STACs in this state in internal commanding mode quickly overwhelmed the CROL with a high volume of densely-packed freeze cycles, resulting in pedestal data for most of the instrument being lost.

I modified *crest_fs* to mask off the entire detector except for one STAC, beginning with

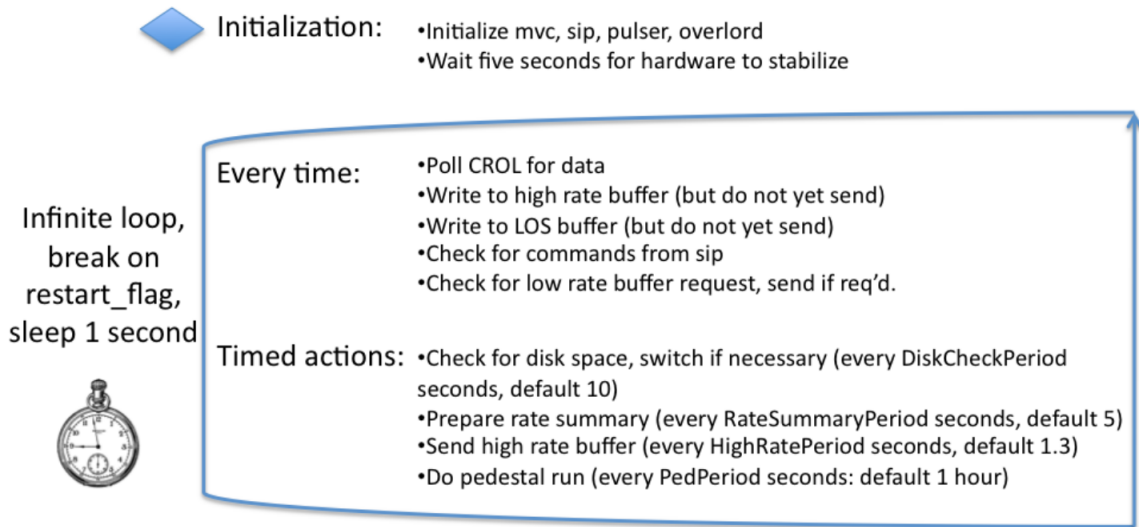


Figure 3.38: Control flow of the flight software.

the first STAC and iterating over all 72 crystal and veto STACs. In this manner *crest_fs* would build up a pedestal run with data from all STACs, just not all at the same time. During flight, *crest_fs* would execute a pedestal run once per hour. Each run took several minutes to complete, mainly owing to pausing long enough on each STAC to generate sufficiently high statistics for accurate pedestal measurement. After completing a pedestal run *crest_fs* would then execute a pulser run, reset the instrument, then resume normal data taking.

3.5 Magnetic Field Sensor

During the Antarctic flight, the magnetic field sensor pictured in Figure 3.39 provided by Nilson Remo of the University of Michigan Atmospheric Oceanic and Space Sciences department was mounted on CREST’s exterior. The power control unit was modified to provide power for the unit but otherwise the field sensor operated independently of CREST, autonomously recording data to its own internal hard drive. I present the results of these measurements in Section 5.4.8.

3.6 Summary, and a Brief Look Ahead

The CREST detector was designed to be sensitive to the anticipated signature of signal events produced by signal high-energy electrons. The collaboration achieved great suc-



Figure 3.39: The magnetic field sensor provided by the University of Michigan Atmospheric Oceanic and Space Sciences department. It flew on board CREST and recorded the three local components of magnetic field strength.

cess in this regard, by providing a finely-segmented, large detector plane, which efficiently detects photons in the signal energy range, with sub-nanosecond timing resolution. The triggerless portion of its readout system ensured unbiased sampling of incident photons and charged particles, and the triggered portion (driven by the digitize system) protected against the blindingly high rate of events below the coincidence threshold.

I describe in Chapter 5 the degradation in detector performance during flight I refer to as “The Incident.” While the source of this disruption is unknown, it is probable that it has its root in a re-emergence of a failure mode which presented itself at least twice during ground testing; first, while attempting to simultaneously flash the entire crystal array with the LED-driven pulser system, and second, in attempting to simultaneously sample the pedestal mode of the entire crystal array. Both these diagnostic modes failed in the same way: the read-out system was not capable of reading out freeze cycles which contained more than approximately 500 hits. In the frenzy leading up to the Plum Brook test, we worked around these problems by simply *avoiding* such situations. For example, pulser runs were formed by combining three separate instrument configurations, each comprising one third of the crystal STACs. Pedestal runs were implemented by sampling only one crystal STAC at a time.

With the benefit of in hindsight, it is clear that rather than finding workarounds to avoid this issue, we should have diagnosed it more carefully. Other than a major revision of the CROL/VTOL firmware or some other major fix, we likely could have taken some measures to protect the read-out system from such freeze cycles. For example, this could have been accomplished with an extension of the digitize system, by providing a *maximum* number of activated crystal STACs, rather than only a minimum coincidence threshold. I was intimately involved in both workarounds; if anyone could have recognized this error, it should have been me.

The other major shortcoming of the detector was the flatness of its crystal array. Again, with the benefit of hindsight, it is easy to see that a flat crystal array presents a very large mass thickness to events with predominantly horizontal momentum direction. This resulted in a small fraction (but very large absolute number) of side-going protons evading the veto system, dominating our background and plaguing our analysis efforts. I suggest an improvement to this configuration in Section 7.3.

Another area ripe for re-configuration is the veto system. Elegantly designed and expertly constructed, each individual paddle performed better than expected. The veto system as a whole, however, was simply asked to do too much with too little. If protons were charging cavalry, and the paddles were our pike-wielding lines of infantry defenders, the paddles were seriously mis-allocated, both in number and equipment. The top paddles were given

the best-performing PMTs (Hamamatsu tubes, identical to those reading out the crystal assemblies), while the side and bottom paddles were given slower, noisier Burle tubes. The top paddles had the least important role to play in rejecting side-going protons, since the tiny solid angle from which such protons originated was covered by the side, not top, paddles. These protons had only to avoid detection by two veto paddles out of fifty-four to bypass the veto system entirely.

Meanwhile, the bottom paddles form the crux of my “second layer” of spatially segmented detectors, enabling my momentum-pointing analysis techniques described in Chapter 6. Had they been equipped with Hamamatsu tubes, or been more finely spatially-segmented, it is possible that CREST’s pointing resolution would be significantly higher than the 30-50 degrees I have demonstrated thus far.

Of course, such was not the intended use of the paddles. The veto system was designed to do nothing more than note the passage of charged particles *somewhere in the detector*. The emergent utility that enabled my work could only arise from the fertile ground of insightful design.

CHAPTER 4

Photon and Charged Particle Energy Loss in CREST

CREST's analysis efforts predominantly involve working with the background and signal Monte Carlo data sets. These software packages rely on Geant4 to handle the underlying physics simulation. So as not to cede to Geant4 all understanding of the physics interactions required to have an intuition for CREST's response to simulated background and signal events, I reviewed these processes in detail. I relate the salient points here, both as a convenient reference for designers of future similar experiments and so that careful readers of Chapter 6 might be convinced that my various hypotheses about the behavior of primary and secondary particles and photons in the detector have a firm basis in the framework of the relevant physics interactions.

As related in Chapter 3, CREST has two main sensitive detector materials, Barium Fluoride crystals and EJ-200 plastic scintillating veto paddles. As lead makes up a significant fraction of CREST's crystal plane by mass thickness, I also include this material as a relevant target for energy deposition by traversing particles and photons. An event, as observed by the CREST detector, consists of the read-out system's response to the energy deposited in these materials by an ensemble of photons and charged particles.

In signal events, this ensemble includes the aforementioned Bremsstrahlung and synchrotron photons, secondary charged particles and photons, and possibly the primary electron itself. Background events bombard the detector with protons, pions, muons and photons of widely varying energies. I review here the response of CREST's sensitive detectors to incident photons and charged particles, focusing on photons in the relevant energy range (above approximately 20 keV) and relativistic protons (the dominant source of background to our signal events). I also review the energy loss mechanism of high-energy electrons in the veto paddles and crystals, as this mechanism underlies a substantial component of the traversal-type signal events I identified as an important subset, statistically speaking, of CREST's signal population.

I organize my review of these mechanisms for energy deposition according to particle species (charged particles first, then photons), and within those Sections, by detector material. Within the section on photon energy deposition, I review the roles of the three main interactions by which photons do so. Given the crucial part played by Compton scattering, in both light collection by the PMT/crystal assemblies and my primary electron momentum direction-estimating algorithm, I place particular emphasis on understanding the Klein-Nishina formulae for the energy and angular differential cross sections, for both the recoil photon and electron.

I also review my efforts in testing and calibrating CREST's light collection and readout system, focusing on the response of our photomultiplier tubes to incident radiation. I relate my method for, and the results of, the first comprehensive test of CREST's crystal array, PMT slow control and read-out architecture, which I undertook at Indiana in early 2011.

Photon interactions of interest result in the ejection of an atomic electron or creation of electron/positron pairs. Since the energy deposited in CREST's sensitive detectors depends on its response to these charged particles, I will first review charged particle energy loss, then draw upon these results in order to characterize the energy collected by the detector in response to incident photons.

4.1 Charged Particle Energy Loss

Charged particles interacting with absorbing material interact mainly with the coulomb fields of the material's atomic electrons. The mass and energy of the traversing particle compared to the electron rest mass $m_e c^2$ determines the appropriate model used to describe the particle's energy loss. If the traversing particle has mass m and kinetic energy E , then the maximum energy it can transfer to an electron in a single collision is approximately $4Em_e/m$ [52]. Thus heavy or slow particles lose energy via a very large number of weak interactions while light or fast particles can do so in fewer, stronger interactions.

4.1.1 Stopping Power for Protons

Below the Bragg peak, elastic nuclear recoil and electron capture dominate energy loss processes. In the intermediate range above the Bragg peak, the average rate at which a heavy or slow particle loses energy with distance is treated as a specific energy loss: $\frac{dE}{dX}$. Knoll [52] writes the estimate for this quantity as

$$-\frac{dE}{dx} = \frac{4\pi e^4 z^2}{m_e c^2 \beta^2} NB, \quad (4.1)$$

where e is the charge of the electron, $\beta \equiv v/c$ and ze are the velocity and charge of the primary particle (given our focus on muons, protons and electrons in what follows I assume $z \equiv 1$), N and Z are the number density and atomic number of the absorber and:

$$B \equiv Z \left[\ln \frac{2m_e c^2 \beta^2}{I} - \ln(1 - \beta^2) - \beta^2 \right]. \quad (4.2)$$

Here, I is the mean ionization and excitation value, which I discuss in the following Section. The multiplicative component proportional to $\frac{NZ}{\beta^2} \ln \frac{\beta^2}{I}$ dominates until the particle becomes highly relativistic (so that the $\frac{v^2}{c^2}$ term becomes significant); thus at low energies the specific energy loss falls rapidly with increasing energy. In the highly relativistic regime (at a few GeV for protons or about 1 MeV for electrons, depending on the properties of the absorber) the energy loss becomes fairly constant until the collisional losses become unimportant compared to the radiative losses. Particles in this regime are referred to as “minimum ionizing particles” [52]. Thus without much loss of fidelity we simply want to know above what minimum energy protons and electrons can be considered to be minimum ionizing particles and what kind of signal such minimum ionizing particles produce.

This form of the equation has some severe limitations in terms of the applicable energy regime. At very low energies the particle track may no longer be assumed to be straight and this stopping power model becomes inaccurate. For example, as a proton’s velocity approaches that of the absorber’s shell electrons, the probability of electron capture increases, meaning the proton would no longer lose energy via coulomb forces (as it would then be a neutral Hydrogen atom). Slow muons would undergo multiple large-angle scatterings, then eventually decay. At the other extreme, for high values of $\beta\gamma$ (above 100 or so), the most important correction term omitted here is the density effect. This effect takes into account the fact that at very high kinetic energies, the traversing particle’s electric field interacts with the combined electric field of a coherent, polarized region of the absorber rather than with individual electrons. This blunts the energy loss at higher values of $\beta\gamma$ leading to a plateau in the most probable energy deposition above $\beta\gamma \approx 100$.

I include electrons in this discussion with the caveat that radiative losses become important at far lower energies than for protons. The electrons whose signals we are interested in (with kinetic energies of a TeV or so) will induce electromagnetic cascades in the crystals. I discuss the properties of these signals in Section 4.1.2.

4.1.1.1 Mean Ionization and Excitation Value

I is the (often empirically determined) average ionization and excitation potential of the absorber. We are interested in the ionization and excitation potential of CREST’s sensitive

detectors (BaF₂ and EJ-200) and the lead shielding to determine the energies above which we can treat protons and muons as minimum ionizing particles.

Given the average ionization potential of the constituent atoms in a molecule, one can use Bragg's Rule to estimate the average ionization potential of the molecule as a whole (see e.g. Ahlen's review article [3], pg. 138):

$$N \ln I \approx \sum_i N_i \ln I_i \quad (4.3)$$

According to [62] Barium, with $Z = 56$ has $I \approx 491.0$ eV and Fluorine with $Z = 9$ has $I \approx 115.0$ eV. This works out to BaF₂ having $I \approx 344.8$ eV. The Particle Data Group (PDG) reports a value of $I = 375.9$ eV [39] showing that in this case the Bragg rule is accurate to about 10%.¹ In what follows I use the PDG's value.

BaF₂'s $\langle Z/A \rangle$ (also weighted by number of nuclei) is 0.42207. I calculate the stopping power for BaF₂ using Equation 4.1 with this value of I and show the results in Figure 4.1. Minimum ionization of approximately 1.326 MeV/(g/cm²) occurs at around 1 MeV for electrons, at around 200 MeV for muons and at around 2 GeV for protons.

4.1.1.2 Most Probable Energy Loss

This stopping power is an average total energy loss comprised of a large number of small interactions resulting in the loss of a random amount of energy by the traversing particle. When attempting to predict the energy deposited in a material of finite thickness (e.g. a CREST BaF₂ crystal) by any one particle, this number may not be very useful as the average can be dominated by infrequent interactions with very high energy deposition. When attempting to eliminate the proton background from our population of candidate signal events it will be useful to know the typical energies deposited in the crystals by these protons. This distribution is characterized chiefly by the most probable energy loss in an absorber of a given mass thickness. This can be estimated as in [15]:

$$\Delta_p = \xi \left[\ln \frac{2m_e c^2 \beta^2 \gamma^2}{I} + \ln \frac{\xi}{I} + j - \beta^2 - \delta(\beta\gamma) \right] \text{ MeV}, \quad (4.4)$$

where $\xi = (0.307075/A)(x/\beta^2)$ MeV / (g/cm²), x is the mass thickness of the absorbing pathlength in g/cm², $j = 0.200$ and $\delta(\beta\gamma)$ is the density effect correction. For simplicity I avoid including the full density effect and merely note that it would result in a flattening of the most probable energy deposition above $\beta\gamma \approx 100$. I calculate that above approximately

¹Computing the specific energy loss for minimum ionizing particles at these two values of I results in a variance of only a few keV/(g/cm²).




Figure 4.1: Stopping power as a function of energy for protons, muons and electrons in BaF₂ calculated from Equation 4.1 and $I_{BaF_2} = 375.9$ eV. Note that this model is inapplicable at the low end of the plotted energy range for muons and protons because the energy is below the location of the Bragg peak, below which particle velocities become comparable to that of shell electrons and protons can capture electrons. At the high end of the range electron interaction becomes dominated by radiative interactions (i.e. Bremsstrahlung) and the model is also invalid. Particles in the intermediate range near the minimum of these curves (at around 1.326 MeV / (g/cm³)) are said to be minimum ionizing particles, or ‘mips’.




Figure 4.2: Stopping power as a function of energy for protons, muons and electrons in polyvinyltoluene (the polymer base for the EJ-200 veto scintillator) calculated from Equation 4.1 and $I_{EJ-200} = 64.7$ eV. The average energy loss of minimum ionizing particles is approximately $2.022 \text{ MeV} / (\text{g/cm}^3)$.

Parameter	Value
Δ	1.326 MeV
ξ/x	0.065 MeV
$(\Delta_p/x)/\Delta$	0.80
Horizontal incidence	
ρt	24.45 g/cm ²
Δ_p	26 MeV
$w \approx 4\xi$	6.3 MeV
Vertical incidence	
ρt	9.78 g/cm ²
Δ_p	10 MeV
$w \approx 4\xi$	2.5 MeV


Table 4.1: Parameters of the Landau distributions for proton energy deposition in the BaF₂ crystals for horizontal and vertical trajectories, for which the crystals present 24.45 and 9.78 g/cm² of thickness, respectively. Δ is the Bethe average energy loss $\langle \frac{dE}{dX} \rangle$ for minimum ionizing protons in BaF₂. Δ_p is the most probable energy loss per g/cm² of thickness of BaF₂. Δ_p/Δ and ξ are the plateau values for $\beta\gamma \gtrsim 10$; at lower energies protons will lose more energy. The widths do not include the widening effects of convolution with finite energy resolution of the crystals and variations in pathlengths through crystals.

a few GeV, protons have a most probable energy loss of approximately 80% of the average energy deposited by a minimum ionizing particle, or approximately $0.8 \times 1.326 \approx 1.06$ MeV per g/cm² mass thickness of BaF₂. The Landau distributions of energy deposited in crystals by protons will be characterized by most probable energies as high as 26 MeV (for side-going protons traversing a full crystal diameter), median energies around 10 MeV (for vertically incident protons traversing one crystal height) or less for paths which intersect fractions of these pathlengths. The width of the Landau distribution is given approximately by $w \approx 4\xi$ [39]. I calculate $\xi \approx 65 \text{ keV}/(\text{g}/\text{cm}^2)$. I summarize these results in Table 4.1 for reference when discussing proton events as CREST’s main source of background events.


While the same parameters for the veto paddles are of some interest, a critical parameter needed for estimating the background contamination of our signal event set is the fraction of incident protons which go undetected by a veto paddle. Recall that the veto PMT assemblies and discriminator thresholds were tuned to the photoelectron level, and that each MeV of energy deposition produces 10,000 photons (see Table 3.3). Given that the most probable energy deposition is 0.85 MeV with a FWHM of approximately 0.16 MeV, protons traversing the paddles will produce at least on order 5,000 photons. To convert this into a detection probability we would need additional information concerning what fraction of these photons are collected by the fiber wave guides and actually reach the PMTs at either end.

Parameter	Value
Δ	2.022 MeV
ξ/x	0.083 MeV
$(\Delta_p/x)/\Delta$	0.825
Perpendicular incidence	
ρt	0.51 g/cm ²
Δ_p	0.85 MeV
$w \approx 4\xi$	0.16 MeV

Table 4.2: Parameters of the Landau distributions for proton energy deposition in the EJ-200 veto paddles for normal incidence (0.51 g/cm² of thickness). Parameter definitions are identical to those in Table 4.1.



(a)



(b)

Figure 4.3: Distribution of ADC1 counts for signals in the PMTs at the -X (top) and +X (bottom) ends of top veto paddle 4 for ground events that met a muonic cut set. According to the estimates in Table 4.2 the peaks correspond to approximately 0.85 MeV deposited energy. The veto PMT gains were adjusted only coarsely to ensure the Landau most probable energies were visible in such data runs.




Figure 4.4: The probability that both ends of a top (a) or non-top (that is, bottom, side or inside slanted) (b) veto paddle detect a signal given that either end of the paddle detects a signal as a function of combined ADC channel for background Monte Carlo (red online) and flight data (black). For most of the interesting ADC range the flight data displays a two-sided inefficiency of approximately 86%.

Figure 4.4 shows the results of one method for estimating the probability that a charged particle transits a veto paddle evades detection as a function of the ADC value recorded by the veto PMTs. Devised by M. Geske and J. Musser, it shows that for most of the interesting ADC range the probability of both PMTs recording a signal when either side records a signal is approximately 86%. Analogous estimates give approximately 92-93% for the one-sided inefficiencies. Thus the probability that the particle escapes detection completely is approximately equal to

$$(1 - 0.93)^2 \approx 0.49\%. \quad (4.5)$$

It follows that the probability of foiling the veto system entirely (that is, of passing through two paddles undetected) is of order 10^{-5} . Given the extremely high flux of protons at float altitude this means that a significant number of charged particles likely escaped detection by the veto system. The vast majority of these particles would transit the detector on steeply inclined trajectories and therefore would not leave long, linear, signal-like tracks in the crystal system. However the fraction of incident particles with trajectories parallel (or nearly parallel) to the crystal plane could leave long, linear, signal-like tracks; such particles comprise the majority of our background events.

4.1.2 Fast Electron Energy Loss

In Section 2.1.4, I propose accepting as signal a class of events I dub “traversal events” wherein the primary electron traverses the detector plane, passing through at least two veto paddles and potentially a crystal as well. To aid in characterizing these events I estimate the detector response to the energy deposited by very high energy electrons. In the energy regime of interest (electron energies above ~ 1 TeV) we must take into account radiative effects and the likelihood that an electromagnetic cascade will ensue when the electron traverses lead shielding or a crystal.

In the preceding discussion it was reasonable to treat the traversing particle’s momentum as being deflected only slightly in each Coulomb interaction since the particles absorbing the momentum were electrons with much lower mass than the traversing particle. This is false for electrons (and positrons). Instead of the Bethe energy loss equation and its attendant corrections, a correct analysis would calculate the Møller cross section for energy transfer to atomic electrons and the Bhaba cross section for electron-positron scattering (see [39], Equations 32.24 and 32.25, and references therein). In spite of these complications, the stopping power for minimum ionizing electrons is typically not very different from that of more massive particles. With the simplifying assumption that electron collisional loss will proceed as above, we can approximate electrons’ total energy loss as the sum of collisional and radiative losses:

$$\frac{dE}{dx} \approx \left(\frac{dE}{dx} \right)_c + \left(\frac{dE}{dx} \right)_r \quad (4.6)$$

As electron energy increases, radiative energy loss (via Bremsstrahlung) becomes more significant compared to collisional losses. I discussed this energy loss rate in Section 2.2 to describe electron energy loss in the atmosphere; briefly, there are two parameters of interest. First, the radiation lengths of BaF_2 , EJ-200 and elemental lead, which are 9.91, 43.90 and 6.37 g/cm^2 , respectively. Second, we wish to know how much of the energy loss is due to radiation vs collisions. Knoll estimates [52]:

$$\frac{\left(\frac{dE}{dx} \right)_r}{\left(\frac{dE}{dx} \right)_c} \approx \frac{EZ}{700} \quad (4.7)$$

with E measured in MeV. Collisional losses make up 1% of the total loss in BaF_2 with mass-weighted Z of 44.25 when $E \approx 70,000/44.25 \approx 1.6$ GeV. This occurs in lead at approximately 850 MeV and in EJ-200 at approximately 12.2 GeV. Therefore, at TeV energies it is safe to ignore collisional energy loss and focus our attention on the electromagnetic showers produced by the radiation lengths of material presented for different

geometrical paths through the detector materials. I summarize these parameters in Table 4.3. The Geometric Cases (B5, B2, etc) are defined in Section 4.2.3 and refer to differently oriented paths through the crystals traversing differing mass thicknesses of material.

For example, in the B2 case the electron traverses 2 cm of BaF₂ crystal, which presents a mass thickness of 9.78 g/cm². Thus $x = 9.78$ and $E(9.78)/E_0 = \exp \frac{-9.78}{9.91} = 0.37$, meaning the electron will lose 63% of its energy to Bremsstrahlung photon emission while traversing the crystal vertically. While this means the *original* electron will have a vastly reduced energy after traversing the crystal, the energy deposited in the crystal will likely be some small fraction of this. This is because the resulting Bremsstrahlung photons will produce electron/positron pairs, which will generate their own Bremsstrahlung, etc. Thus the energy *deposited* depends on shower containment, that is the size of the shower formation as a whole relative to the individual crystal size.

The crystals are radiatively thin: about 1 radiation length for vertical incidence and less than 3 radiation lengths for horizontal incidence. Since the great majority of electromagnetic shower energy deposition occurs beyond 3-5 radiation lengths, most of the shower energy will escape the crystal plane for vertically-oriented events. On the other hand if the traversing particle has momentum mainly parallel to the crystal plane, a series of crystals and their lead shields could present many radiation lengths to the shower. In such an event the shower containment would be driven by the height of the crystals (2 cm) compared to the Molière radius of BaF₂, which is approximately 15.25 g/cm² and of lead, which is approximately 18.18 g/cm². Thus the shower would strike approximately two rows of crystals (the reader will recall the crystals have 7.5cm center-to-center separation), but vertically only about 2cm out of a shower diameter of approximately 30 cm would be occupied by crystals. Very roughly this would lead to approximately 1/6th or 1/7th total shower energy containment for horizontally aligned particles.

From the above it follows that TeV electrons will emit on average approximately 70 GeV-worth of photons while traversing a veto paddle. It is not immediately obvious that this will produce a detectable signal in the veto paddles, as the interaction probability for photons with energy above 10 MeV or so is effectively zero (see Figure 4.9). In other words the paddles do not have sufficient radiation lengths to even initiate an electromagnetic cascade. The initial round of Bremsstrahlung photons would largely escape before producing electron/positron pairs. Bremsstrahlung photons with energy below 1 MeV or so have non-zero probability for interaction in the veto paddle; therefore, we would like to estimate the number of Bremsstrahlung photons emitted in the veto paddles with energy below 1 MeV, and see if this number is different from zero.

Ref. [39] provide the following means of estimating this value when the thickness of

Parameter	Material	g/cm ²	Value
X_0	BaF ₂	n/a	9.91 g/cm ²
	EJ-200	n/a	43.90 g/cm ²
	Pb	n/a	6.37 g/cm ²
$E_{1\%}$	BaF ₂	n/a	1.6 GeV
	EJ-200	n/a	12.2 GeV
	Pb	n/a	850 MeV
E_{loss}/E_0	BaF ₂	24.45	0.91
	BaF ₂	9.78	0.63
	BaF ₂	4.89	0.38
	EJ-200	0.51	0.07
	Pb	4.5	0.51
	Pb	9.0	0.75

Table 4.3: Radiative energy loss parameters for electrons in CREST detector materials. $E_{1\%}$ is analogous to the electron critical energy and refers to the electron energy at which the collisional losses make up 1% of the total energy loss. E_{loss}/E_0 refers to the fraction of incident electron energy radiated as Bremsstrahlung photons while it traverses the given mass thickness of medium. The geometry cases listed here are identical to those detailed in Section 4.2.3. The lead geometry cases assume traversing 0.4 cm-thick layer of lead either once or twice, for 4.5 or 9.0 g/cm² mass thickness.

the absorber is much less than one radiation length:

$$N_\gamma = \frac{d}{X_0} \left[\frac{4}{3} \ln \left(\frac{k_{max}}{k_{min}} \right) - \frac{4(k_{max} - k_{min})}{3E} + \frac{k_{max}^2 - k_{min}^2}{2E^2} \right]. \quad (4.8)$$

In our case $\frac{d}{X_0} = 0.0116$ so if the above applied, we would have on average 0.07 photons between 10 keV and 1 MeV emitted for every electron that traverses a paddle. If we enforce $k_{min} = 100$ keV instead of 10 this figure is 0.035 photons emitted per electron. This figure is essentially constant for electron energy above 10 GeV.

Since these fractional photon energies (denoted by $y = k/E$) are so small an additional suppression factor arising from the Landau-Pomeranchuk-Migdal (LPM) effect should be taken into account. For amorphous (non-crystalline) media (such as EJ-200), photons with energies $k < \frac{E^2}{E + E_{LPM}}$ are suppressed, with $E_{LPM} = (7.7 \text{ TeV/cm}) \times \frac{X_0}{\rho}$. For electron energies between 100 GeV and 10 TeV, this limit ranges from 30 MeV to 300 GeV for 0.5mm thick EJ-200. In other words the $\sim 7\%$ probability that a TeV electron would produce even one photon detectable by the veto system is a gross over-estimate because photons with such small fractional energy are heavily suppressed by the LPM effect. In traversal-type events we should not expect the electron to directly produce a measurable signal in the veto system.

The Bremsstrahlung photons produced by electron interaction in the veto paddles would interact in the crystals, as photons with significant fractions of the electron energy (above a GeV or so) have a very good chance of interacting with the crystal (see Figure 4.9).

4.2 Photon Energy Loss

4.2.1 Interaction Processes

Photons interact with the material they traverse primarily via three mechanisms: the photoelectric effect, Compton scattering and pair production (see e.g. [52]). In contrast to charged particle energy loss in which the total energy loss results from the accumulation of many soft interactions, these interactions are catastrophic in the sense that the photon is destroyed or its energy and momentum direction are significantly altered as a result of the interaction. The exception to this rule is coherent (Rayleigh) scattering where the photon is deflected by a small angle in the Coulomb field of an atom but retains its energy. The Rayleigh scattering angle vanishes to zero as photon energy increases; thus this interaction can normally be ignored in the photon energy regime relevant to CREST (above some tens of keV).

4.2.2 Mass Attenuation Coefficient

Knoll [52] (pg. 53) defines the linear attenuation coefficient as the probability per unit path length that a photon would interact with a given material

$$\mu \equiv (\tau + \sigma + \kappa) \text{ cm}^{-1}, \quad (4.9)$$

where μ is the total linear attenuation coefficient and τ, σ, κ are the contributions from photoelectric, Compton and pair-production interactions, respectively. The mean free path, $\lambda = \frac{1}{\mu}$. To eliminate dependence on the density of the absorbing material it is convenient to express this quantity as the mass attenuation coefficient $\equiv \frac{\mu}{\rho} \text{ cm}^2/\text{g}$ where ρ is the density of the medium.

The individual attenuation coefficients for each process depend on the photon energy and the absorber's atomic number. The photoelectric coefficient $\tau \approx C \times \frac{Z^n}{E_\gamma^{3.5}}$ with $n \approx 4 - 5$ ([52], pg. 49). It is evident that the photoelectric attenuation coefficient scales very well with atomic number but also decreases rapidly with increasing photon energy.

The Compton mass attenuation coefficient increases only linearly with Z , the number of electrons available to scatter off of and falls off with energy. It can be expressed as the

Klein-Nishina cross section per electron times the number of electrons per gram of material [11]:

$$\frac{\sigma}{\rho} = \frac{N_A Z}{A} \sigma_{KN}, \quad (4.10)$$

where N_A is Avogadro's number, A is the grams per mole of molecules or atoms in the absorber and σ_{KN} is the total Klein-Nishina cross section in units of $\text{cm}^2/\text{electron}$.

The pair production mass attenuation coefficient varies with Z^2 , is zero for photon energies below $2m_e$ and rapidly attains a constant value for a wide range of energies above the threshold $2m_e$. At low energies and high Z the photoelectric effect dominates, at medium energies the Compton effect dominates (somewhat independently of Z) and at higher energies pair production dominates.

With this definition, the probability that a photon will penetrate a medium to a given thickness t , denoted by P_T , is given by²

$$P_T = \exp \left[-\frac{\mu}{\rho} \rho t \right], \quad (4.11)$$

where t is the photon's pathlength through the medium. The interaction probability is then $P_I = 1 - P_T$. Given the known dimensions and densities of the materials relevant for CREST (the atmosphere, Barium Fluoride and EJ-200 plastic scintillator) and their respective energy-dependent mass attenuation coefficient, we can estimate interaction probabilities as a function of energy. Rather than directly estimating interaction probabilities for each of the processes in each of the materials at all energies (which are highly E - and Z -dependent) I extract mass attenuation coefficient data for each of the materials (air, BaF_2 and EJ-200 scintillator) from the online XCOM database, maintained by National Institute of Standards and Technology (NIST) [33].

Figures 4.5, 4.6 and 4.7 show plots of the mass attenuation coefficient (for the three dominant interaction processes, and for the total effect) for photons between 1 keV and 100GeV in energy in BaF_2 , EJ-200 scintillator and air, respectively. Evident is the transition from the low energy interactions dominated by the photoelectric effect to a median energy range dominated by Compton scattering finally giving way to the high energy regime dominated by nuclear and electronic pair production.

²Knoll envisions this as the fraction of surviving γ -rays in a beam and writes this as $\frac{I}{I_0}$, the fraction of photons in a beam penetrating to a given mass thickness. P_T could be interpreted identically, assuming that the "beam" of photons had a total of one photon in it.




Figure 4.5: Mass attenuation coefficients for the three major interaction processes (and their summed effect) for photons in Barium Fluoride (BaF_2) (plot produced from data provided by the XCOM database, courtesy of NIST [33]). Also shown by the vertical line at $E_\gamma = 0.511$ MeV is the electron rest mass. Visible at low energies are the discontinuities in the photoelectric cross section arising from the discrete energies associated with atomic transition edges. These edges are listed in Table 4.4. The highest energy edge (the 56K edge at around 37 keV) is quite close to the potential critical synchrotron energies for electrons at around 1 TeV, as shown in Figure 2.1. Signal synchrotron photons between ~ 100 keV and \sim a few MeV are very likely to interact via Compton scattering.

Edge	E_γ (keV)
56 M ₃	1.062
56 M ₂	1.137
56 M ₁	1.293
56 L ₃	5.247
56 L ₂	5.624
56 L ₁	5.989
56 K	37.44

Table 4.4: Photoelectric absorption edges in Barium Fluoride. From [33].




Figure 4.6: Cross section against the three major interaction processes for photons in EJ-200 plastic scintillator (plot produced from data provided by the XCOM database, courtesy of NIST [33]). Also shown by the vertical line at $E_{\gamma} = 0.511$ MeV is the electron rest mass. Note that the absorption edges relevant for this polymer (those of Carbon and Hydrogen) occur at energies which are below the minimum photon energy tabulated here.




Figure 4.7: Cross section against the three major interaction processes for photons in the atmosphere (plot produced from data provided by the XCOM database, courtesy of NIST [33]). Also shown by the vertical line at $E_{\gamma} = 0.511$ MeV is the electron rest mass.

4.2.3 Estimated Interaction Probabilities

Armed with the mass attenuation coefficients, I will now calculate the interaction probability for photons with air, BaF₂ and EJ-200 as a function of energy using the following geometric cases. Elsewhere in this work I refer to these geometric cases by the provided reference code (e.g., “A3” denotes traversing 3 g/cm² of air). In the flight data, background and signal Monte Carlo simulations, actual or calculated energy deposition will lie somewhere between these extreme cases.

1. Air with residual mass thickness traversed by the photon between emission by the primary electron and intersection with the detector plane
 - A1: 1 g/cm²
 - A3: 3 g/cm²
 - A5: 5 g/cm²
2. BaF₂, with mass thickness corresponding to three paths through the crystal discs
 - B5: Sideways through a full diameter ($d = 5$ cm, 24.45 g/cm²)
 - B2: Vertically through a full height ($h = 2$ cm, 9.78 g/cm²)
 - B1: Clipping the edge or corner for 1 cm of pathlength (4.89 g/cm²)
3. EJ-200, shortest path through a veto paddle
 - V50: Thickness of a veto paddle ($t = 50$ mm, 0.51 g/cm²)

For a given geometry case, we can employ Equation 4.11 with μ/ρ calculated at each energy and ρt given by the geometric cases listed above to estimate the energy- and geometry-dependent photon interaction probability. I show the results of these calculations in Figure 4.9.

Photons have a high probability for interacting with the BaF₂ crystals up to ~ 100 keV beyond which the pathlength strongly affects the probability; incident angles between horizontal and vertical will have probability between the solid blue and the first blue dotted lines (Geometry Cases B5 and B3). The veto paddles have unit probability of interaction up to ~ 5 keV, $\sim 10\%$ at ~ 40 keV, and effectively no interaction at higher energies. The probability for atmospheric interaction lies between that for BaF₂ and EJ-200, strongly depending on the atmospheric overburden the photon must traverse before reaching the detector plane.




Figure 4.8: Mass attenuation coefficients in Air, BaF_2 , EJ-200 and Lead (data from [33]). Between $E_\gamma \sim 600 \text{ keV}$ and $\sim 5 \text{ MeV}$ BaF_2 actually has a smaller mass attenuation coefficient than both air and EJ-200. This is due to the Compton mass attenuation length scaling like Z/A ; EJ-200 is composed mainly of Carbon and Hydrogen and therefore has a higher Z/A ratio than heavier elements such as Barium. In terms of interaction probability this effect is more than compensated for by the increased mass thickness of the BaF_2 crystals compared to the veto paddles.

This leaves a range of photon energies between approximately 40 keV and several MeV where the relative probability of interaction in the crystals compared to air or the veto paddles is relatively high. As shown in Figure 2.1 for primary electrons with energy between 1-10 TeV, the critical synchrotron frequency ranges from approximately 50 keV to a few MeV depending on the strength of the perpendicular component of the geo-magnetic field. This justifies the claim that the relative interaction probability (of the crystals compared to the veto paddles) for a photon in the relevant energy range is quite high.

4.2.4 Effect of the Lead Shield

Each PMT assembly is wrapped with a 0.4 cm thick lead shield around its top edge, designed to reduce the effects of multiple Compton scatterings occurring between assemblies from one incident signal photon. This has the effect of adding a mass thickness of $0.4 \times \rho_{Pb} = 0.4 \times 11.34 = 4.5 \text{ g/cm}^2$ to Geometry Case B5 (a photon going sideways through the crystal) and any other path along which the photon intersects the lead shield. Photons which pierce through both crystal side edges experience twice this mass thickness, or an additional 9.0 g/cm^2 . Based on this figure alone it is clear that this additional material




Figure 4.9: Interaction probabilities for photons in air, BaF₂ crystals and EJ-200 veto paddles computed using Equation 4.11, the total mass attenuation coefficients in Figure 4.8 and the various Geometric Cases listed above.

will strongly affect the shape of signal events with significant momentum parallel to the crystal plane.

Figure 4.10 shows lead’s mass attenuation coefficient as a function of energy and for the three dominant interaction processes. Because of its high atomic number ($Z = 82$), lead is a very efficient absorber of photons up to high energies, with the photoelectric effect remaining important up to CREST’s signal energy range. For instance, note that the crossover point between photoelectric and Compton dominance happens right at $m_e c^2$ (whereas this happens well below $m_e c^2$ in BaF₂). Since photoelectric interactions typically contain a large fraction of the incident photon’s energy, this means that many of the secondary photons which might escape one crystal interaction (e.g. an outgoing Compton photon) with energies below $m_e c^2$ would come to rest within the confines of the lead shield without inciting further light production in a crystal (as intended).

An unintended effect of the lead shield is to induce electromagnetic showering within the crystal plane. Given the large mass thickness added by the lead and its high atomic number, any photon entering the detector from an inclined angle or scattering between crystals with energy above a few MeV is liable to undergo pair production in the lead shield. This produces characteristic annihilation photons which can then interact in the crystal (see for example Knoll, pgs. 321-322 [52]) resulting in a large, lead-induced secondary signal at 0.511 MeV as shown in Figure 4.11.

In the absence of an extremely high energy source of cosmic ray electrons (above 50




Figure 4.10: Cross section against the three major interaction processes for photons in lead (plot produced from data provided by the XCOM database, courtesy of NIST [33]). Also shown by the vertical line at $E_\gamma = 0.511$ MeV is the electron rest mass. Note the absorption edge at just over 100 keV, in the center of the signal synchrotron photon typical energy range.

TeV or so) it is very unlikely that synchrotron photons would reach such high energies. However, it is likely that Bremsstrahlung photons from primary electrons with \sim TeV energies will reach these energies (see Figure 2.5). For such photons the lead shield acts, for good or ill, as a thin calorimeter surrounding each crystal.

Another unintended effect of the lead was a decreased flight altitude profile. The lead shielding added approximately 0.5 kg of mass to each assembly, thereby increasing the total instrument weight by over 500kg. This appreciably lowered CREST's flight altitude which also increases the flux of Bremsstrahlung photons (by increasing the atmospheric overburden and thus the amount of air traversed by primary electrons while illuminating the detector).

Bremsstrahlung Photon Energy Deposition




Figure 4.11: Energy deposited by simulated Bremsstrahlung photons and their daughter products in the BaF_2 crystals. The dominant feature at the electron rest mass energy is primarily caused by pair conversion in the lead shield surrounding each crystal/PMT assembly. The blue plot shows the energy deposition resulting from Bremsstrahlung photons emitted by 10 TeV primary electrons traversing Earth's atmosphere. The red plot shows the same for a electron energies ranging from 500 GeV to 1.25 TeV. Plot generated from simulated CREST signal events using CRESTMag8, written by Scott Nutter.

Given the steep penalties arising from the lead shield, an alternative method for eliminating or diagnosing signals from multiple Compton scattering is desirable. Compton spectrometers which rely upon knowing the direction of incidence of the primary γ -ray and measuring the angle of its deflection are well known (see e.g. Knoll pg. 324 [52]). Since CREST has multiple detectors in close proximity with well-measured spatial separation vectors, if the incoming γ -ray direction were known, such a scheme could be replicated with CREST's crystals.

Since it was assumed by the CREST group that the photon incidence directions were unknown, such a procedure would be impossible. As I shall show later, CREST can point at the incident direction with a resolution of approximately 40 degrees. While this is probably not sufficiently precise to perform Compton spectroscopy, a future detector with enhanced pointing capability and CREST's excellent spatial and temporal resolution could make use of such a scheme to avoid using such a costly lead shield.

4.2.5 Energy Deposition and Collected Light for Photon Interactions

The photoelectric and Compton interactions both produce one outgoing electron and photon while pair production produces at least one electron/positron pair (potentially initiating an electromagnetic cascade). The resulting electrons and positrons will then lose energy in the absorbing medium before either stopping or escaping in the manner described in Section 4.1. The resulting photons are subject to the same attenuation coefficients as the original photon, albeit at a lower energy. This chain of events results in incoming photons of a particular energy generating a distribution of light collected by the crystal/PMT assemblies. The shape of the light spectra and the fraction of incident photon energy thus measured depends on which of the three processes occurred and random variances within those processes (e.g. the Compton scattering angle).

After BaF_2 atoms absorb the incident photon via the photoelectric effect, they de-excite via x-ray emission, with the resulting x-ray photon generally being fully absorbed in the crystal (within a millimeter or so due to the very high mass attenuation coefficient at low photon energies). The ejected photoelectron receives an energy $E_{e^-} = h\nu - E_b$ where E_b is the binding energy between nucleus and electron. The amount of energy the photoelectron deposits in the crystal is proportional to the rate at which the electron loses energy and its pathlength in the crystal (see Section 4.1). The total deposited energy is therefore a fraction of the incident photon energy which is typically close to one. This gives rise to the photopeak, a band of collected charge whose mean channel number is proportional to the total incident photon energy and whose width can be used to characterize the energy

resolution of the detection apparatus.

Compton scattering events result in an outgoing electron and photon pair, with the original photon's energy distributed between them in an angle-dependent manner [52]. When departing with a scattering angle of θ (with $\theta = 0$ representing no interaction) the outgoing photon's energy is given by³

$$h\nu' = \frac{h\nu}{1 + (h\nu/m_e c^2)(1 - \cos \theta)}. \quad (4.12)$$

The recoil electron therefore has energy

$$E_{e^-} = h\nu - h\nu' = h\nu \left(\frac{(h\nu/m_e c^2)(1 - \cos \theta)}{1 + (h\nu/m_e c^2)(1 - \cos \theta)} \right). \quad (4.13)$$

The energy deposited in the crystal due to the incident photon then depends on how much of the energy imparted to the outgoing photon and electron is subsequently absorbed. With the assumption that the outgoing photon escapes the crystal and the relatively low-energy electron deposits all its energy in the crystal (the 'small detector' approximation) the result is a Compton continuum covering the possible energies imparted to the recoil electron. These energies range from 0 (at $\theta = 0$, a grazing interaction) to the maximum $E_{e^-, \theta=\pi} = h\nu \left(\frac{2h\nu/m_e c^2}{1+2h\nu/m_e c^2} \right)$. All energies between these two extrema should be visible in the light collection spectrum, representing intermediate angles between these extremes. Since the outgoing photon must depart with a fraction of the incident photon energy, there exists a gap between the upper edge of the Compton continuum and the onset of the photopeak of approximately ([52], Equation 10.5):

$$E_C = h\nu - E_{e^-, \theta=\pi} = \frac{h\nu}{1 + 2h\nu/m_e c^2}. \quad (4.14)$$

For example, Cesium-137 produces (via Barium-137) 662 keV γ -rays. With $h\nu = 662$ keV, $E_C \approx 184$ keV. To illustrate these distributions for 662 keV photons, I show in Figure 4.12 an example multi-channel analyzer spectrum resulting from γ -ray interactions generated as part of my energy calibration procedure described in Section 4.3. The Compton continuum spans the range from channel 100 to approximately channel 350. Since the gap between the continuum and photopeak is approximately 150 channels and $E_C \approx 184$ keV, this indicates this assembly when operated at the high voltage used to produce this spectrum has approximately 1.2 keV/channel.

³I discuss the Compton angular spectrum in Section 4.2.5.1.




Figure 4.12: Spectrum from crystal/PMT assembly number 56 while exposed to 662 keV photons from a Cesium-137 source during a ground calibration run. After subtracting pedestal noise (the green dotted line) from the raw spectrum (black line) two features are clearly evident: the Compton continuum (from approximately channels 100-350) and the 662 keV photopeak at approximately ADC1 (low gain) channel 500.

4.2.5.1 Angular Spectrum for Compton Scattering

Though CREST does not make direct use of the Compton angular spectrum, I describe it briefly here to make two points. First, the average recoil electron momentum is in the direction of the incident photon. This feature is critical to my analysis detailed in Section 6.7 which attempts to extract the primary electron momentum direction. Second, there is a polarization effect which breaks the azimuthal symmetry which obtains when averaging over all incoming and outgoing photon polarization states which may be of use in future synchrotron detection schemes.

The differential angular spectrum for the outgoing photon (relative to the incident photon direction) for Compton scattering is described by the Klein-Nishina formula [52]:

$$\frac{d\sigma}{d\Omega} = Zr_0^2 \left(\frac{1}{1 + \alpha(1 - \cos\theta)} \right)^2 \left(\frac{1 + \cos^2\theta}{2} \right) \left(1 + \frac{\alpha^2(1 - \cos\theta)^2}{(1 + \cos^2\theta)[1 + \alpha(1 - \cos\theta)]} \right), \quad (4.15)$$

where $\alpha \equiv h\nu/m_e c^2$ and $r_0 = \frac{e^2}{m_e c^2}$ is the classical electron radius. I plot this cross section in Figure 4.13(a) for various incident photon energies. CREST's minimum detection threshold of a few tens of keV has an angular spectrum somewhere between the outer 1 keV line and the second-most-outer 100 keV line, which already shows significant bias in the forward direction.

The recoil electron's outgoing angle ξ is related to the outgoing photon's angle by

$$\tan \theta_e = \frac{1}{(1 + \alpha) \tan \theta_\gamma / 2}, \quad (4.16)$$

and can be inverted to

$$\tan \theta_\gamma = 2 \tan^{-1} \left[\frac{\cot \theta_e}{1 + \alpha} \right]. \quad (4.17)$$

The electron angular spectrum can be expressed as a function of the photon angular spectrum [11]

$$\frac{d\sigma}{d\Omega_e} = \frac{d\sigma}{d\Omega_\gamma} \times \frac{(1 + \alpha^2) (1 - \cos \theta_\gamma)^2}{\cos^3 \theta_e}, \quad (4.18)$$

where $\frac{d\sigma}{d\Omega_\gamma}$ is given by Equation 4.15. Care must be taken with the differential elements. $d\sigma/d\Omega_e$ is the differential cross section as a function of *electron* angle, while $d\sigma/d\Omega_\gamma$ is the differential cross section as a function of *photon* angle.

I plot this recoil electron angular spectrum in Figure 4.13(b). The most likely angle for the recoil of the electron is directly forward at all energies. As energy increases, the probability that the electron will be scattered forward increases rapidly. At all energies the *average* recoil direction for electrons scattered by photons with the same incident direction will be overwhelmingly forward. This means any measurement of the average recoil electron momentum in an event with multiple signal photons with the same momentum direction provides an estimate of that momentum direction. Coupled with the fact that signal photons emitted by a primary electron have momentum parallel to that of the electron, this means that the average momentum direction of the recoil electrons can be used as an estimate of the primary electron momentum direction. In Section 6.7 I describe my technique for estimating the primary electron momentum direction which makes use of this property.

This form of the Klein-Nishina angular spectrum displays azimuthal symmetry. This is only the case when averaging over polarization states of the incoming and outgoing photons; in fact the azimuthal spectrum is modified when the incoming photon and outgoing photons have known polarization states. See [65] for a full treatment. Synchrotron photons emitted by the same primary electron are very strongly polarized in the plane of motion of the primary electron (see [51], pg. 678, Equation 14.80 and following discussion). Therefore it is conceivable that some use could be made of the recoil direction of Compton electrons if the polarization state of the outgoing Compton photons could also be measured.





Figure 4.13: The Klein-Nishina differential cross section per unit solid angle for outgoing photons and recoil electrons following a Compton interaction. *Top*: Cross section per unit solid angle as a function of *outgoing photon polar angle* (Equation 4.15) normalized to unit probability in the forward direction. The polar angle is the angle off horizontal at which the outgoing photon departs relative to the direction of the incident photon. *Bottom*: Cross section per unit solid angle as a function of *recoil electron polar angle* (Equation 4.18). The angle on the x-axis, in degrees, is the angle relative to the incident photon direction with which the recoil electron departs. As E_γ increases, the probability of a forward scattering of the electron increases drastically, making the recoil direction of the Compton electron a good estimator for the momentum direction of the incident photon. Polarization effects are neglected in both plots, resulting in an azimuthally symmetric distribution.

4.3 Energy Calibration

To characterize the relationship between ADC channel and the energy deposited in a crystal by an incident particle or photon, CREST can utilize signals of known energy to provide an energy calibration. Two methods were carried out using the full instrument electronics readout system. I performed the first method at Indiana University in May 2011 utilizing a 662 keV Cesium-137 source. The second method utilizing in-flight data was performed by Nahee Park, of the University of Chicago during post-flight analysis.

4.3.1 Ground Calibration and Gain Matching

In May 2011 during my stay at Indiana University, I performed a source calibration and PMT gain matching study utilizing a Cesium-137 radioactive source of 662 keV photons. Though energy calibration studies utilizing individual crystal/PMT assemblies, custom PMT high voltage supply and current readout tools had previously been performed by A. Yagi at UM, this represented the first attempt to use CREST as an integrated detector. By this I mean that the entire crystal system was read out by the readout electronics under command of the digitize system with the PMT's powered by VOLTBus, all operating in a flight-like mode, albeit with an artificially low Digitize coincidence threshold of 1 STAC (see Section 3.3.7 for details). This process provided not only a first attempt at energy calibration but also provided the first validation of CREST's integrated operation prior to our first significant mission milestone (the vacuum test in Plum Brook, Ohio described in Section 5.1.2).

Cesium-137 beta decays to Barium-137m, which decays with a half-life of 153 seconds and characteristic γ -ray emission at 662 keV. By means of a small DC motor and a rope and pulley system devised by A. Shroyer (an IU CREST electrical and mechanical engineer), the source was locomoted across the crystal plane, illuminating with good geometry each crystal PMT. Figure 4.14 shows the time dependence of hit rates in each of the eight Sedecim modules in crystal c-channel 4.

A summary of my procedure follows. To prepare, I produced data sets with Digitize coincidence threshold = 1 and STACs set to external commanding with the Cesium-137 source above every crystal/PMT assembly for some reasonable duration (on the order of seconds). This resulted in 35 runs, each transiting one row of 32 tubes. The first 32 runs transited 29 of 32 tubes in each row in the detector plane. The remaining 3 tubes at the start of each row were transited in 3 additional runs (after realizing that the source locomotion apparatus had an erroneous starting point past the third tube in each row). I produced an additional data set with no source to measure the background spectra for each tube.




Figure 4.14: Time progression of the hit rate resulting from a Cesium-137 source transiting across the center (between the 2nd and 3rd rows of assemblies) crystal c-channel 4 during energy calibration on the ground at Indiana University (May 2011). The x-axis is in units of Freeze Cycles (each lasting $131.074 \mu\text{s}$). Each row of the y-axis corresponds to an individual crystal/PMT assembly. Increased box areas correspond to more hits per Freeze Cycle. As the source passes above the crystal the hit rate reaches a maximum. Each run along a c-channel lasted approximately two minutes. For this mode of data taking the Digitize system was set to instruct the instrument to collect charge with a coincidence threshold of 1 STAC (in other words, STACs read out every hit).




Figure 4.15: Low-gain (ADC1) channel spectra from crystal/PMT assembly number 638 (which happens to be the 15th PMT on C-channel 5, STAC 8) without (*left*) and with (*right*) exposure to a Cesium-137 662 keV γ -ray photon source.

To further test the accuracy of our gain calculations for the phototubes and our voltage control system (via VOLTBus, the MVC and each Sedecim’s SVI boards), I also attempted to move each PMT’s peak channel number to channel 500. To estimate the high voltage setting required to achieve this given the peak channel number and previous high voltage setting, I used:

$$\left(\frac{V_{500}}{V}\right)^{N\alpha} = \frac{500}{\bar{C}}, \quad (4.19)$$

where \bar{C} is the measured peak channel number with high voltage V , N is the number of dynode stages (10 for the Hamamatsu PMTs) and α is an empirical constant related to the secondary emission yield between the photocathode and the dynode stages. I estimated a value of $\alpha = 7$ in an initial sample of trial runs and used that value throughout the procedure.

The calibration procedure was carried out as follows for each of 1,024 tubes:

1. Identify the run and time within the run corresponding to the source transiting the tube
2. Construct the single-hit spectrum for this tube corresponding to the transit period
3. Construct the single-hit spectrum for this tube from a source-less control run




Figure 4.16: Example of the Gaussian fit performed for each 1,024 crystal/PMT assemblies using a Cesium-137 662 keV γ -ray source. The location of the peak was determined as in Figure 4.12. Its width is found here as the σ of the best-fit Gaussian function. Using this procedure, I measured a FWHM peak channel resolution of approximately 20-25% for all 1,024 independent instrument crystal channels. In this case, the FWHM resolution = $2.355 \times \sigma / \bar{X} \times 100 = 23.5\%$.

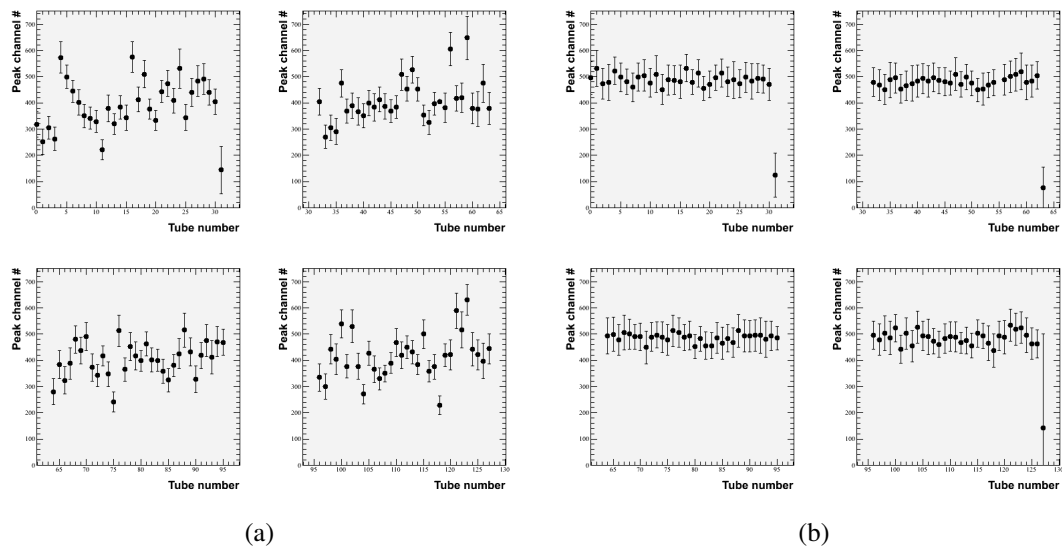


Figure 4.17: Gaussian mean (filled circle) and sigma values (width of error bars) for each of the 128 tubes in crystal c-channel 1 using the procedure outlined in Section 4.3.1. The four left plots used the old high voltage values, and the four right plots use the new high voltage values.

4. Subtract the source-less spectrum from source spectrum
5. Locate the photopeak with a peak-finding algorithm (ROOT's TSpectrum class)
6. Find the best fit Gaussian to the peak, extract the mean and standard deviation
7. Record the peak mean channel number, peak standard deviation (in channel number units) and the high voltage setting
8. Tune the PMT high voltage to attempt to achieve a peak mean channel number of approximately 500
9. Re-run the source transit with the modified high voltage values and repeat steps 1-7

Successful completion of this procedure validated multiple critical functions of the CREST detector, including charge collection by the AFE boards, Digitize decision logic and signaling, STAC hit timestamping and hit reporting, hit collection by the CROL and PMT high voltage supply, readback, and setting. I reproduce the results from this procedure in Figures 4.17 - 4.24.

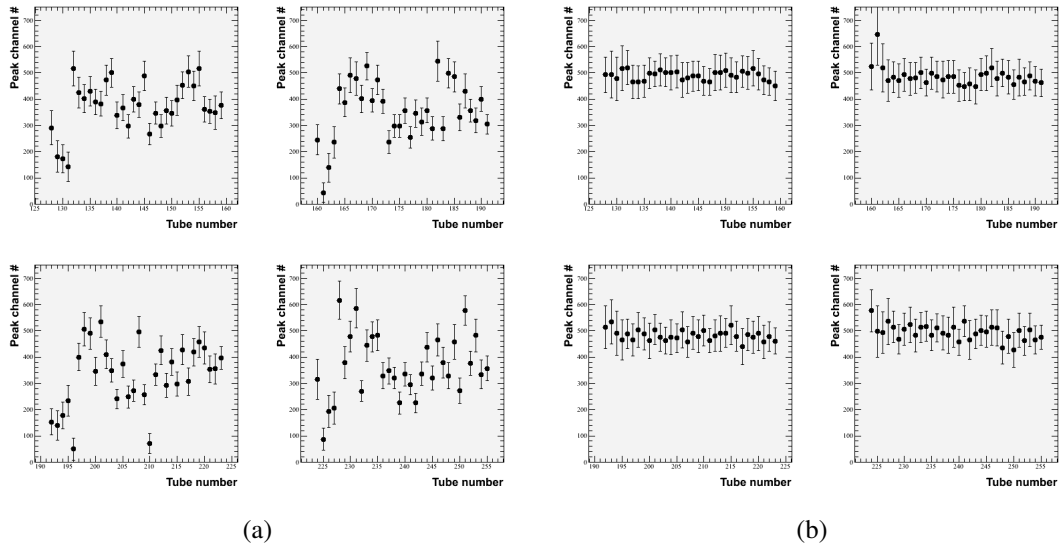


Figure 4.18: Gaussian mean (filled circle) and sigma values (width of error bars) for each of the 128 tubes in crystal c-channel 2 using the procedure outlined in Section 4.3.1. The four left plots used the old high voltage values, and the four right plots use the new high voltage values.

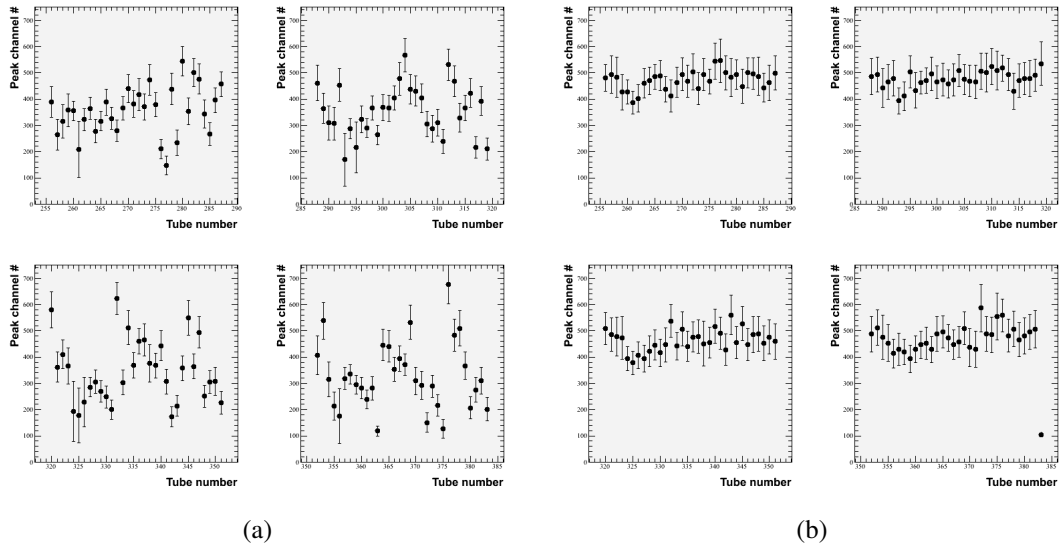


Figure 4.19: Gaussian mean (filled circle) and sigma values (width of error bars) for each of the 128 tubes in crystal c-channel 3 using the procedure outlined in Section 4.3.1. The four left plots used the old high voltage values, and the four right plots use the new high voltage values.

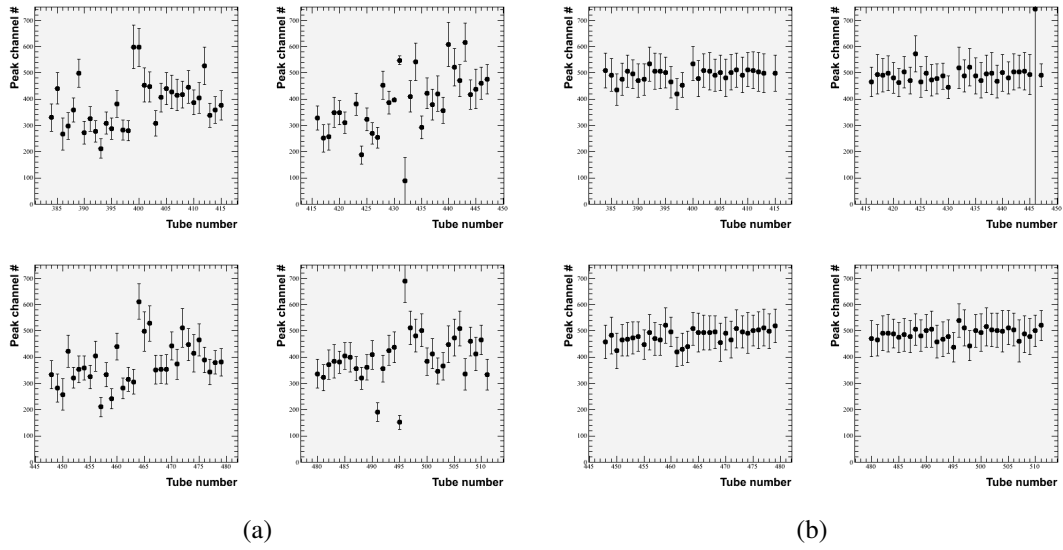


Figure 4.20: Gaussian mean (filled circle) and sigma values (width of error bars) for each of the 128 tubes in crystal c-channel 4 using the procedure outlined in Section 4.3.1. The four left plots used the old high voltage values, and the four right plots use the new high voltage values.

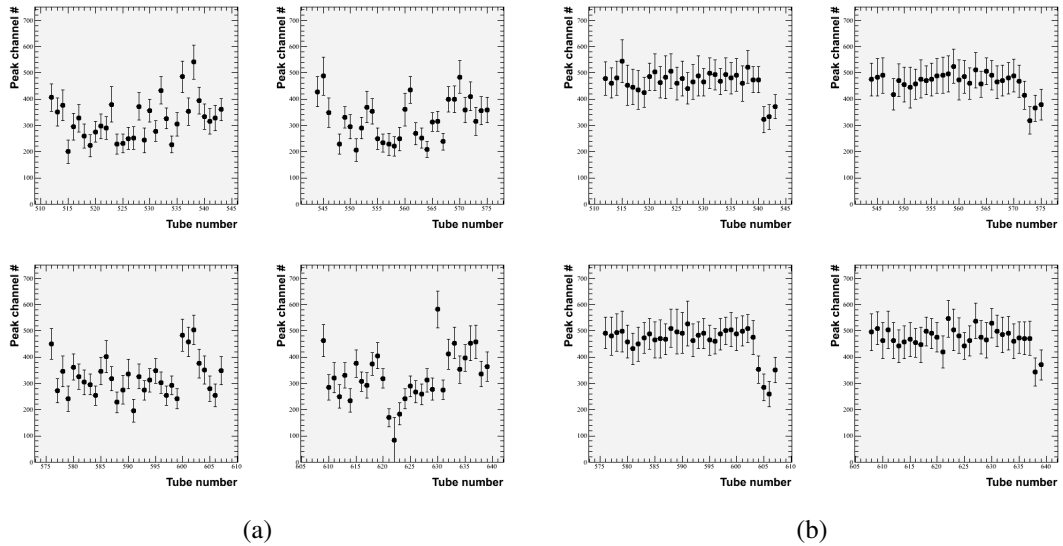


Figure 4.21: Gaussian mean (filled circle) and sigma values (width of error bars) for each of the 128 tubes in crystal c-channel 5 using the procedure outlined in Section 4.3.1. The four left plots used the old high voltage values, and the four right plots use the new high voltage values.

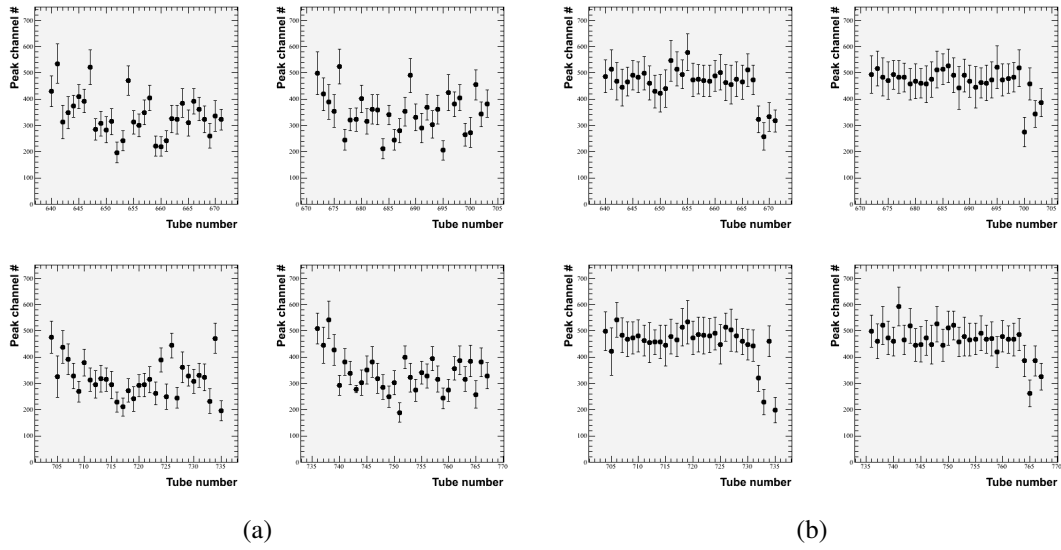


Figure 4.22: Gaussian mean (filled circle) and sigma values (width of error bars) for each of the 128 tubes in crystal c-channel 6 using the procedure outlined in Section 4.3.1. The four left plots used the old high voltage values, and the four right plots use the new high voltage values.

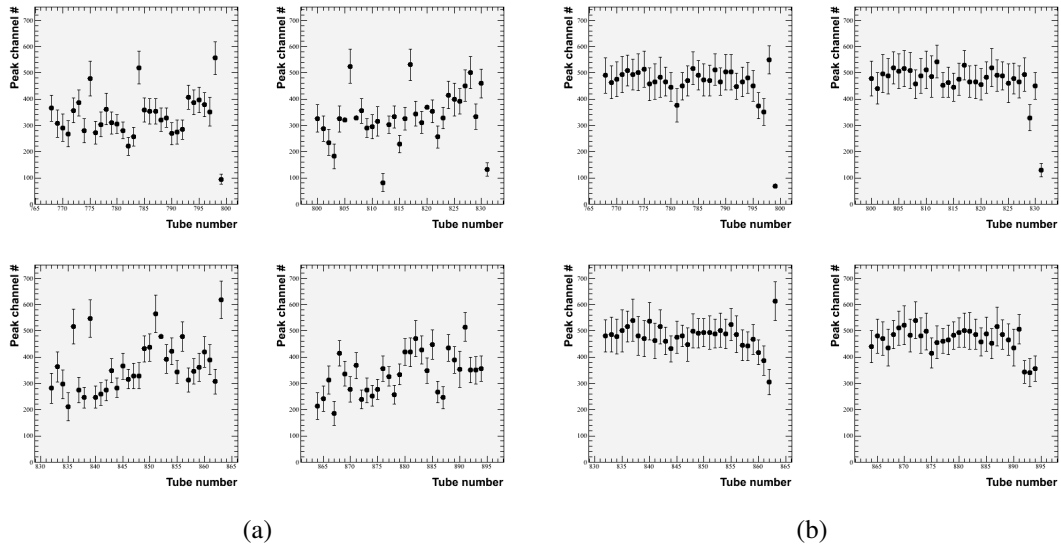


Figure 4.23: Gaussian mean (filled circle) and sigma values (width of error bars) for each of the 128 tubes in crystal c-channel 7 using the procedure outlined in Section 4.3.1. The four left plots used the old high voltage values, and the four right plots use the new high voltage values.

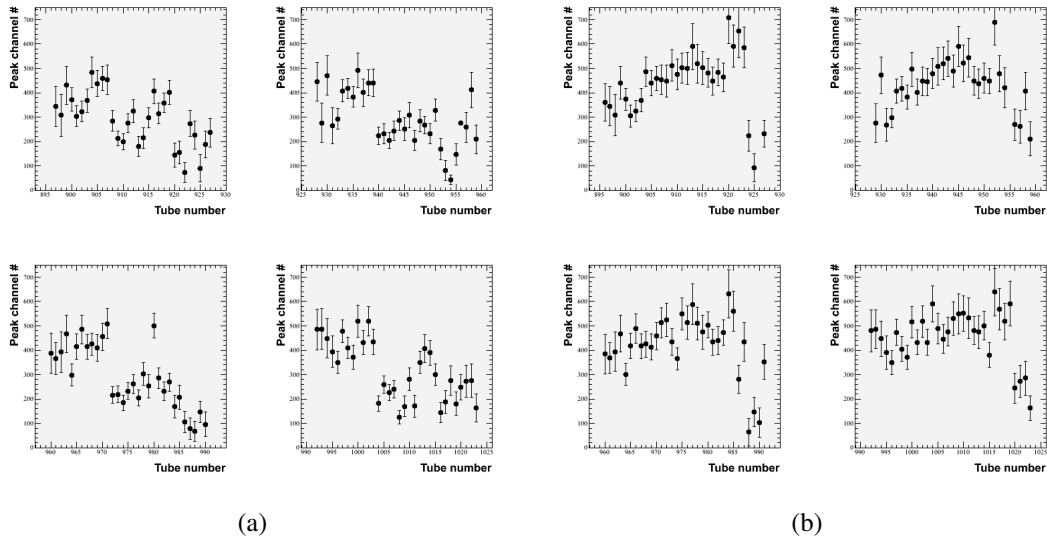


Figure 4.24: Gaussian mean (filled circle) and sigma values (width of error bars) for each of the 128 tubes in crystal c-channel 8 using the procedure outlined in Section 4.3.1. The four left plots used the old high voltage values, and the four right plots use the new high voltage values.

Peak	Energy (MeV)
Pair production line	0.511
BaF ₂ Background 1	1.352
BaF ₂ Background 2	1.634
BaF ₂ Background 3	1.848
BaF ₂ Background 5	2.590

Table 4.5: Mean energies of background peaks used by Nahee Park to perform the in-flight gain matching and energy calibration of the ADC0 and ADC1 channels.

4.3.2 In-Flight Energy Calibration and Linearization

Our in-flight calibration procedure was developed by Nahee Park and utilizes the 511 keV line and various known radioactive impurities inherent to BaF₂ to characterize the linearity and scale of the ADC0 (high gain) and ADC1 (low gain) channels for every crystal tube independently. These signals are summarized in Table 4.3.2.

Nahee found that ADC1 readings were linear up to a tube-dependent channel (she refers to this channel as the “break” channel) while ADC0 readings were linear (at least, to an energy beyond our highest signal of known energy). She constructed a fine-gain reading which was a linear function of ADC1 below the break and a second-order polynomial above the break. The location of the break and the parameters of this polynomial were determined by comparison with ADC0. Above ADC1’s saturation point, Nahee generated

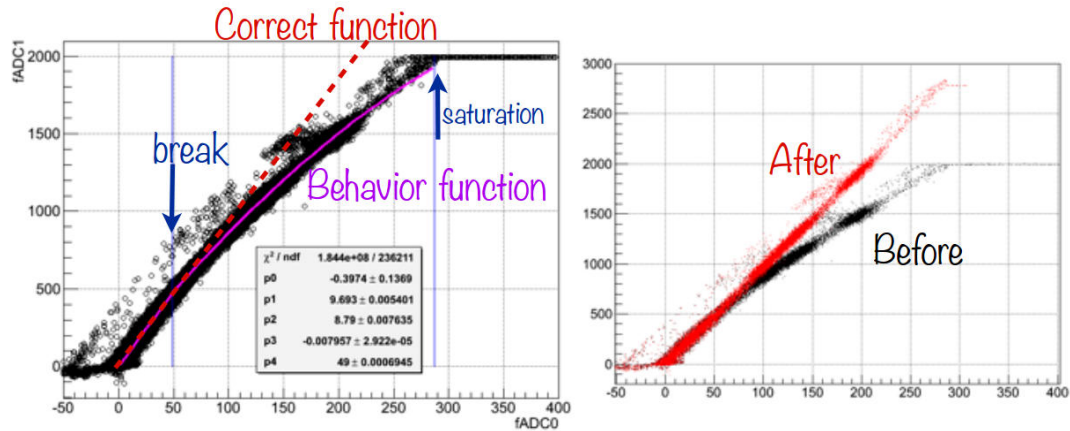


Figure 4.25: Example of parameters used in Nahee Park’s gain-stitching procedure. On the left she plots ADC1 (low-gain) channel on the vertical axis against ADC0 (high-gain) channel on the x-axis for crystal/PMT assembly 268. Below the break channel around ADC1 channel 500 ADC1 displays linear behavior with respect to energy. Above the break the relationship between ADC1 channel and energy is fitted by a second-order polynomial. ADC1 saturates at channel 2048, at which point ADC0 is relied upon to provide a coarse energy reading. Plots courtesy of Nahee Park.

a coarse-gain reading which was simply a linear function of ADC0.

Nahee generated a software module which utilized the pre-scale events (discussed above in Section 3.3.7) to estimate the break channel, saturation channel, and the parameters of the linear and second-order polynomial functions relating ADC1 to energy for each tube and for each hour of flight operation which I incorporated into my software analysis chain. The outputs of this module, “Energy Combined - Coarse” and “Energy Combined - Fine” replace ADC1 and ADC0 as the key variables representing incident particle or photon energy in signal hits. In the analysis software, these variables are referred to as “ECC” and “ECF.”

4.4 Summary

In the above, I attempted to highlight the most relevant aspects of charged particle and photon energy loss in CREST’s sensitive materials, and in lead, and the read-out system’s response to these energy loss modes. In doing so, I showed why it is correct to view the crystal array as a reliable x- and γ -ray photon detector, and as a radiatively-thin calorimeter for high-energy electrons and γ -rays. The veto paddles do indeed admit the passage of photons above a minimum energy threshold, but would appear unable to directly detect the traversal of very high-energy (above a Tev or so) electrons. This unexpected behavior

is primarily a consequence of the suppression of low-energy Bremsstrahlung photons by the LPM effect and deserves special care when modeling and analyzing the reponse of the detector to traversal-type events.

Side-going protons are CREST's most pernicious, and most numerous, background species. As shown here, high energy protons deposit energy as minimum ionizing particles in the veto and crystal systems. They lose energy to a multitude of low-energy coulomb interactions, which are unable to significantly deflect the protons' trajectories. This means their detector signatures will be straight lines, enabling them to mimic part of the signature of signal events. The highly stochastic nature of their energy deposition precludes using energy-based cuts to eliminate them in an efficient manner. Our first line of defense against protons capable of making long, straight crystal tracks are the four side veto paddles. While these slabs of EJ-200 scintillator are thick and efficient enough to detect the vast majority of these protons, a fraction of the small minority that sneak through these guardians will pose serious analysis hurdles. As I will relate in Chapter 6, a substantial portion of CREST's analysis efforts have been (and continue to be) dedicated to combating their pollution of the candidate signal event population.

In the energy regime particular to CREST's anticipated signal electrons, the Compton process dominates the angular evolution of secondary hits within the detector which form an important subset of the crystal and veto response to signal events. On average, these products scatter predominantly in the forward direction. This behavior enables (and greatly simplifies) my electron momentum-based analysis, as detailed in Chapter 6.

CHAPTER 5

Antarctic Expedition: Preparation and Flight

Through creative collaboration, CREST overcame many unforeseeable difficulties; I feel that a chronological summary of these efforts, besides being of historical interest, might prove useful to those undertaking design, construction, testing and launching of an instrument (or other endeavor) of similar scale. In that spirit, I here recount the CREST expedition, beginning with preparation for the Plum Brook, Ohio Thermal/Vacuum test in June 2011, continuing with the integration process with CSBF communications equipment and satellite networks in Palestine, Texas, culminating with the Antarctic flight and recovery originating at McMurdo Station in Winter 2011/2012.

In the final Section of this Chapter, I have included a number of photographs chiefly taken by me, but also some by other members of CREST, including of course all the recovery pictures. Taken together, they amount to a visual chronological summary of the Antarctic portion of our expedition, as witnessed from my point of view. The sublime beauty of the Southern-most continent and its frozen surroundings inspire awe in all of those fortunate enough to see it in person; my feeble photographic efforts do it little justice. I hope through these photographs to transmit at least a sliver of that majesty. I further hope to call to mind the sheer scale of investment - enormous amounts of logistical, financial and human resources - required to mount scientific expeditions such as this one. I reiterate my gratitude to all those who made CREST possible.

5.1 Thermal Shield Assembly and Instrument Testing

By May 2011, the instrument's sensitive detectors (crystal and veto system), read-out and control electronics and flight computer had been fully integrated. In order to successfully complete the planned thermal/vacuum test at NASA's Plum Brook facility, we then constructed a fully functional version of the thermal shield. Following the validation provided by this test an improved version was constructed in Palestine and McMurdo.

5.1.1 The Thermal Plan

At float altitude, the low air density, instrument heat generation and intense sunlight makes a heat management plan based around radiative cooling a critical part of any balloon flight without active cooling elements. Mimicking the designs of previous balloon and space flights, CREST's thermal shield was designed to achieve stable operating temperatures by reflecting as much solar radiation as possible, while simultaneously radiating as much heat as possible. To achieve this the entire outer layer of CREST would need to be covered with a dual-layer coating. The outer layer would consist of a transparent, highly emissive material (either mylar or teflon) and the inner layer, in thermal contact with the instrument chassis, would consist of a highly conductive and reflective metal, either aluminum or silver. Heat generated by the instrument would then conduct to the instrument chassis, then to the inner layer of the thermal skin, then to the outer layer, then be radiated away as infra-red light.

The upper part of the shield forms a box ensconcing the detector plane, veto system and front-end electronics. The box is formed from thermal panels consisting of aluminized mylar sheeting glued to standard construction foam. The panels were mated to each other by conventional packing tape at first, though this was upgraded to Nashua low-temperature duct tape in McMurdo. The foam was hand cut and assembled into a hermetic box by myself, J. Musser and Indiana University engineer Alex Shroyer. Since the aluminized mylar sheeting was expensive and in short supply, great care was taken during the two-man process of applying the sheeting to the foam. Trial and error resulted in several failures; clumsy lamination resulted in wrinkles which reduced heat conduction efficiency. Figure 5.3 shows the results of an application attempt wherein the glue was not allowed to dry before mating the aluminized mylar to the foam. The interiors of the foam panels were bare, with the exception of those covering the dark side of the instrument. We laminated both sides of those panels with aluminized mylar in hopes of retaining some heat on the coldest side of the instrument.

Since CREST would be rotated underneath its host balloon during flight so that the solar panels would be pointed at the sun, the +X face of the instrument (the direction normal to the solar panels) was often referred to as "sun-side." This side was protected by 2" thick foam panels in hopes of providing as much passive protection as possible in case of a mechanical breach of the thermal skin. The -X face (dark side) of the instrument as well as the +Z, +Y and -Y faces received 1" thick foam panels.

The thermal plan called for the sheet metal panels covering the lower portion of CREST's four sides and its entire bottom surface to be laminated with silverized teflon. Silver being a more efficient conductor than aluminum, and teflon being a more efficient emitter than

mylar, this material would provide better thermal shedding near the parts of the detector which would generate the most heat (the power distribution system and the battery charge controllers). Since silverized teflon was significantly more expensive than aluminized mylar, it was elected to laminate the panels with aluminized mylar and regard a successful thermal test with this test configuration as validation of the plan to use the better, expensive material in the flight configuration. The panels were painted with appliance white paint, then laminated with the aluminized mylar sheeting and glue in a process similar to that for the foam panels. The sheeting was not subsequently perforated. Air trapped between the sheeting and the panels formed expanding bubbles during the Plum Brook test, separating the sheeting from the panels (see Figure 5.33 for visual evidence of this). For this reason we chose to use perforated silverized teflon tape for the flight configuration rather than unperforated sheeting.

5.1.2 Plum Brook Vacuum Test

NASA's Glenn Research Center at the Plum Brook Facility in Sandusky, OH houses the B-2 Spacecraft Propulsion Research Facility. Historically used to test validate rocket engines (such as those in the Centaur and Delta-3 series) with hundreds of thousands of pounds of thrust capability in near Earth orbit-like conditions, NASA made it available to CREST at a subsidized rate in order to verify CREST's operation at a simulated high-altitude environment. The B-2 facility is centered around a 62-foot tall, 38-foot diameter stainless steel cylindrical vacuum chamber capable of reaching pressures as low as 10^{-7} torr and temperatures as low as -320°C for extended periods. Though CREST utilized only a small fraction of the capabilities of the B-2 facility, it was readily apparent to all involved that the entire Plum Brook staff were dedicated to a safe and successful test procedure.

5.1.3 Arrival and Preparation

CREST was delivered via truck to the B-2 test facility on 6/14/11 along with a full contingent of the collaboration team. Significant loosening of many bolts was observed to have occurred during transit; we added lock washers and re-tightened them all. The thermal skin's outer surface was cleaned and augmented with aluminized mylar tape, covering every seam between thermal panels and each exposed piece of metal.

After lowering the instrument into the chamber, we verified its normal operation, formed a standardized housekeeping output format and finalized details of the test procedure with Plum Brook personnel. The first procedure was planned for 6/21/2011 at approximately 2:00 PM. We divided our personnel into three, eight-hour shift groups of three people each.




Figure 5.1: The CREST instrument on wheels in the Indiana University high bay. All C-channels have been installed, but the thermal shield has not. The black tedlar-wrapped veto paddles surround the crystal plane with triangular FOLGA's connecting the paddles to their readout PMTs. The edges of the bottom veto paddles peek out from beneath the crystal array. The coiled fiber optic cables are pulser system fan-outs waiting to be connected to the pulser box. In the lower-near corner the power panel accepts the bench power supply. In the lower-far right corner the SFC is visible, and to the left of that (blue board) is the XEM Carrier. Yet to be installed in the lower portion are the batteries, the rest of the power system and the CSBF SIP module (later integrated at Palestine).




Figure 5.2: One of the 2”-thick foam panels (of the second batch, laminated in McMurdo) covering the dark side of the top portion of the instrument.



Figure 5.3: A ruined panel not used in the Plum Brook test constructed by laminating aluminized mylar sheeting onto foam backing. When applied prematurely (before allowing the glue to dry) a reaction between the aluminum, glue and/or foam occurs resulting in the formation of an intriguing green pattern. J. Musser dubbed this beautiful mistake “RTFM.”




Figure 5.4: Various tapes used as part of CREST's thermal shield. From left to right: aluminumized mylar, perforated silverized teflon, conventional packing tape, Nashua duct tape. None of these tapes' adhesive functioned at all during our first launch attempt (scrubbed to due surface wind conditions), which took place during the coldest weather we experienced in McMurdo.

Component	Zone	Surface material	Flux, hot	Flux, cold
Sun @ 22°	Sun-side upper	Al Mylar	685 W/m ²	591.2
	Sun-side lower	Ag Teflon (see note)	343.2	296.2
	Top	Al Mylar	187	161.4
Albedo	Bottom (95% case)	Ag Teflon (see note)	101.3	52.0
Dark side	Rear	Ag Teflon (see note)	116.9	0

Table 5.1: Heat plan for the test procedure in Plum Brook’s B-2 test chamber. Infra-red lamps would simulate the indicated component of solar radiation for an assumed sun elevation of 22° above the horizon. The listed flux values are a maximum; at any given time the actual flux values would be a fraction of the maximum but in the listed proportion. Note: Aluminized mylar was used to cover these surfaces for this test. The flight configuration utilized silverized teflon, as planned.

Each group was responsible for ensuring the safe operation of the instrument, especially during cool down and lamp-on periods, during which the thermal state of the test environment (and therefore potentially the instrument) would be changing rapidly.

A critical part of the test procedure was the simulated heat flux. In consultation with Plum Brook staff, the two flux plans comprising hot and cold cases detailed in Table 5.1.3 was settled upon. In this manner we could verify that the sun side and bottom of the instrument would not overheat in the hot case while the anti-sun side of the instrument would not freeze in the cold case.

5.1.4 First Test Procedure: Two Critical Failures

Approximately 3 hours after initial pump down the temperature reported by the sensor we believed to be in thermal contact with the SFC’s CPU heat sink quickly rose to approximately 60°C. The SFC became unresponsive soon thereafter. Note that no heat lamps had been engaged yet; the only environmental change was the pressure reduction to approximately 4 torr. After allowing the SFC temperature sensor to cool to below 40°C we powered back on, only to have the same failure occur at approximately the same temperature. Under the pressure of producing some result while the chamber was being evacuated we decided on a safe CPU temperature range of 35-55 °C and powered the SFC on and off repeatedly.

After handing control to the next shift at approximately 10:00 PM on 6/21/2011 the LN₂ cold fill began, for which the team was required to exit the control area and retreat to the ground level. The process was interrupted minutes later by a rupture in an LN₂ line which necessitated evacuation of the entire B-2 facility for approximately an hour. After being readmitted to the control area we learned that the line rupture spilled liquid N₂ approximately 10 feet from our control equipment, meaning our team would have been at

risk of lethal oxygen deprivation had they remained at their posts and been unable to escape the area quickly enough.

Upon resuming the cold fill, when environmental temperature reached approximately 2.6°C, instrument power, provided by a feed through to a power supply in the control area, was lost. After many futile attempts to cycle power we realized that no data could be gained until we diagnosed this new issue. We requested that Plum Brook staff terminate the test procedure. They evacuated the LN₂ from the copper tubing and slowly allowed the chamber to return to normal temperature and pressure. As the environmental temperature exceeded 19°C, instrument power returned. Fearing some catastrophic occurrence within the detector electronics, to our relief it seemed to be operating normally. It was speculated at this point that the cooling was warping some hardware component, probably a connector, causing it to lose contact at low temperatures.

5.1.5 Diagnosis and Adaptation

After the chamber returned to normal pressure and temperature we were permitted to enter the chamber and diagnose the two failure modes mentioned above: SFC overheating and loss of instrument power at high temperatures.

The power cable from supply to instrument power interface was inspected and discovered to be improperly secured in such a way that it made normal contact at room temperature, but would likely warp out of contact at very low temperatures. This hypothesis was confirmed by observing instrument power loss during extensive application of liquid coolant sprayed on the attachment point. To remedy this, we assigned two people to be responsible for power cable attachment thereafter.

The SFC was removed from the instrument and examined. The temperature sensor thought to have been placed on the CPU was in fact on the SATA hard drive interface chip. Thus what we believed were safe operating temperatures of approximately 60°C were probably closer to 100°C, the temperature at which most Intel chips automatically shut down to prevent damage. The cause of this overheating was identified as improper thermal contact between CPU and heat sink, caused by mechanical stress on the SFC during instrument transport. Rather than attempt repairs, an identical replacement SFC (hand delivered by M. Schubnell, rather than having traveled inside the instrument) was swapped in after receiving augmented thermal contact between CPU and heat sink.



Figure 5.5: CREST is delivered via flatbed truck to the Plum Brook testing site on 6/14/2011. The white box is the aluminum shipping container constructed and painted at Indiana University. Visible in the background is the upper section of the vacuum chamber with the mushroom hatch in the open state.




Figure 5.6: Nahee Park, Stéphane Coutu, Scott Nutter and Matt Geske examine the detector's condition following delivery by flatbed truck to Plum Brook in the B-2 facility's loading dock.



Figure 5.7: CREST is lowered via crane (a Critical Lift Procedure) into the testing chamber at Plum Brook on 6/15/2011.




Figure 5.8: A view from the ground level of the test chamber. Michael Lang, IU engineer oversees the installation of skin temperature sensors. The cylindrical wall of the chamber consists of densely packed fine copper tubing through which would later flow liquid N₂ gas to cool the chamber to a minimum environmental temperature of about -40 °C.




Figure 5.9: Below the instrument inside the Plum Brook B-2 facility test chamber. The hexagonal rods with orange tape on the end are heat lamps. These lamps would simulate the albedo heat flux resulting from sunlight reflecting upwards from the ocean and ice shelf during the actual flight.



Figure 5.10: Deep in the bowels of Plum Brook's B-2 facility, CREST members watch over the performance of the instrument. The LN₂ line which would later rupture during the first cold fill bides its time in the background. From left: Stéphane Coutu, Matt Geske, Michael Schubnell and Nahee Park.



Figure 5.11: The LN₂ pipe which ruptured during the first cold fill. Due to Plum Brook's strict safety protocol (and the grace of God) no one was injured.

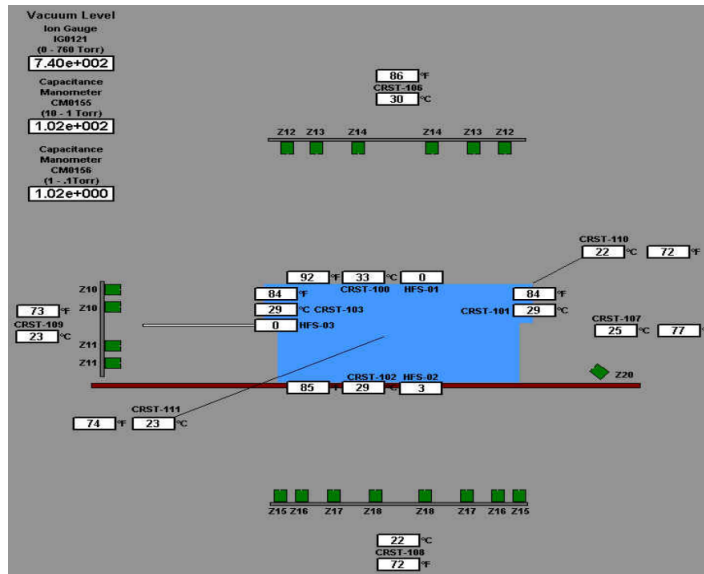


Figure 5.12: Schematic showing the locations of external temperature and heat flux sensors during the test initiated on 6/27/2011, courtesy of NASA Plum Brook staff. The sun-side (+X direction) is to the left and the dark-side (-X) is to the right. The green tubes indicate the location and orientation of the heat lamps used to simulate the heat produced by the solar and albedo radiation flux that would be experienced during flight.

5.1.6 T/V Test: Take Two

Following two days of monitoring the instrument to make sure the new flight computer was running smoothly and that no other problems would threaten the second test procedure, the Plum Brook staff allowed us to repeat the test beginning on 6/27/2011 at 9:00 AM. The planned test conditions were unchanged and this time the flight computer temperature stayed within the normal range and instrument power was not lost upon reaching cold temperatures. Figure 5.12 shows a schematic of the sensor and heat lamp locations. Figures 5.13, 5.14 and 5.15 show the thermal performance of the instrument during the entire test procedure. Based on these results the test was regarded as a validation of CREST's operational stability in a flight-like environment, including the thermal plan and the incarnation of the thermal system protecting CREST during the test, with one exception.

Inspection of CREST following the second test procedure revealed evidence of air bubble formation between the aluminized mylar sheeting and the aluminum thermal panels covering the lower portion of the instrument. Apparently air pockets that had gone unnoticed during the lamination process expanded, detaching sizable circular regions of the laminate from the sheet metal surface. In Figure 5.33 it is evident that up to approximately 1/5th of the surface area of the thermal skin could have been separated from the panel by bubbling. This would present a major impediment to heat-shedding via conduction, which




Figure 5.13: Plot showing the pressure over time measured by a vacuum ion gauge (sensor IG0121) during the second Thermal/Vacuum test at Plum Brook. The minimum pressure achieved was approximately 0.035 torr. The time axis shows seconds since 6/27/2011 at 9:00 AM.

relies on close thermal contact. In particular, the amount of conduction varies with the surface area of the conductive region. This may explain why the bottom skin temperature did not fall as fast as the sun-side or dark-side skin temperature during the cold fill in test procedure 2, approximately between seconds 25,000 and 45,000 in Figure 5.15. To avoid this type of thermal skin degradation during flight we resolved to use strips of perforated silverized teflon tape to cover all of the lower sheet metal panels, so that any similar air bubbles would be limited to less than the size of the gap between perforations (and hopefully be self-resolving).

Following this successful second test, CREST was removed from the test chamber, packed back up into its shipping container and shipped back to the high bay at Indiana University where the final instrument component to be integrated awaited our return: the power system.

5.2 Integration Test at CSBF: Palestine

Following the successful Plum Brook test, we had just a few weeks to prepare for the integration test at NASA's Columbia Scientific Balloon Facility in Palestine, TX. Our most important task prior to this step was to integrate the power system, of which Scott Wakely, Nahee Park and others had recently finished producing at the University of Chicago.

Previous to this, we looked for any transportation-related degradation similar to that experienced when delivering the instrument to Sandusky. Our use of lock washers on chas-




Figure 5.14: Plot showing the measurements over time of two temperature sensors (CRST-103, the sun-side skin temperature and CRST-109, the sun-side environment temperature) (left axis) and of one heat flux sensor (HFS-03, the sun-side heat flux sensor, right axis) during the second test procedure. The time axis shows seconds since 9:00 AM on 6/27/2011. While the skin temperature was sensitive to both the incoming flux and environment temperature, at the simulated flight conditions (-30 °C and the hot flux case detailed in Table 5.1.3) the sun-side skin temperature was stable at approximately 40 °C, well within stable operating range.




Figure 5.15: Plot showing the measurements over time of three temperature sensors (CRST-103, the sun-side skin temperature, CRST-101, the dark side sensor, and CRST-102, the bottom sensor) during the second test procedure. Temperatures were well within the nominal range for the hot and cold cases in all regions of the detector.

sis bolts seemed to prevent the bolt loosening issue. The upgraded SFC heat sink support structure seemed to have prevented any damage along those lines. One new issue did crop up: the thermal panels appeared to be rubbing up against the black tedlar wrap around two separate bottom veto paddle FOLGAs, causing minor light leaks. These were easily remedied with black tedlar tape, and when reconstructing the thermal shield in McMurdo we took care to avoid this issue.

5.2.1 Power System Integration

The power system consists of solar panels, Nickel Metal Hydride (NiMH) battery banks and charge controllers, a power control unit, a relay box and a power panel.

The two battery banks each consist of two 12V batteries in series, each providing approximately 80 Ah of power at 24V [64]. Each bank is protected by a MorningStar ProStar-15 Charge Controller unit designed to protect against overcharging.

The CREST solar panel array consists of two wings containing a total of 540 Sunpower A300 cells, with a total estimated power output of 1124W, with an output voltage capped at 33V and surface temperature of 90°C. In flight the surface temperature was significantly lower; thus the power output was much higher than anticipated and had to be moderated by rotation of the instrument 10-30° off-axis from pointing directly at the sun. One of these wings is shown after being mounted to the instrument in Palestine in Figure 5.26.

After installation of the PCU, relay box and power box in the instrument's underbelly, we rolled CREST outside of the high bay, wired the solar panels into the power system and received approximately 400W of power from both, sufficient to operate the SFC and peripheral equipment. Though an encouraging result, running the entire instrument would require over twice this output (see Figure 5.21). After charging the batteries from ground power we attempted running the instrument from one battery, which we were successful in doing for approximately 2.2 hours, with the discharge curve and power consumption plot shown in Figure 5.21.

In addition to the power system testing, operating CREST in direct sunlight revealed a few light leaks inducing very high rates in veto tubes 17, 20, 33, 53 and 54. These tubes' paddles and folgas were treated to an extra heavy layer of tedlar wrap, which seemed to largely normalize their rates. A plot of the relative veto hit rates is shown in Figure 5.17.

5.2.2 Delivery to Palestine

With the power system functionally integrated, we transported CREST by truck once more, this time to Palestine, TX. IU engineer Alex Shroyer and myself were tasked with trans-



Figure 5.16: Scott Wakely, Nahee Park and Michael Schubnell install the power panel inside the CREST instrument. Photo courtesy of R. Northrop.

porting the bulk of our supplies, tools and ground station computer to Palestine via 14' Ryder truck. This trip took approximately 18 hours which we split into two days driving, stopping in Mississippi. We arrived to find the instrument and collaboration waiting for us and in good order.

5.2.3 CSBF Integration

Our main goal in Palestine was to verify that CREST was ready for its Antarctic mission: that it was functioning as a detector and could communicate with the ground via satellite. Of the former we were confident following the Plum Brook test and follow-up testing at IU. To accomplish that latter we worked directly with CSBF personnel, especially our main contact Joseph Smith, to establish a connection between the SFC and the Support Instrumentation Package (SIP). The SIP was provided entirely by CSBF. It had its own power system fed by the solar panels attached below CREST's chassis. Communication between the SIP and the SFC occurred on two serial channels referred to as the Low-rate and High-rate channels (with 1200 and 19,200 baud bandwidths, respectively). The SIP has two on-board computers referred to as COMM1 and COMM2, either of which can relay packets of data between the SFC and one of three communications pathways




Figure 5.17: Full-sunlight veto tube hit rates (in hz) during the power system integration test at Indiana University on 7/13/2011. Tubes 17, 20 and 33 displayed uncharacteristically high rates on only one side of a veto paddle, indicative of light leaks in the FOLGA mating the PMT to the paddle. Tubes 53 and 54 comprise two ends of the same (small) paddle, indicating either a light leak on the paddle itself or light propagating out of the folga and back into the paddle, thus lighting up both ends. Plots of the veto tube rates such as this one formed an important part of shift operators' sense of whether the instrument was operating normally or not during the Antarctic flight. Ignoring the abnormal rates of the outlier tubes, the normal rates are largely determined by paddle size (top, bottom and side paddles are larger than the inside slanted tubes), proximity to the crystals (top paddles are closest) and geometry (side paddles have their surface normal facing sideways and thus have less flux from down-going particles).

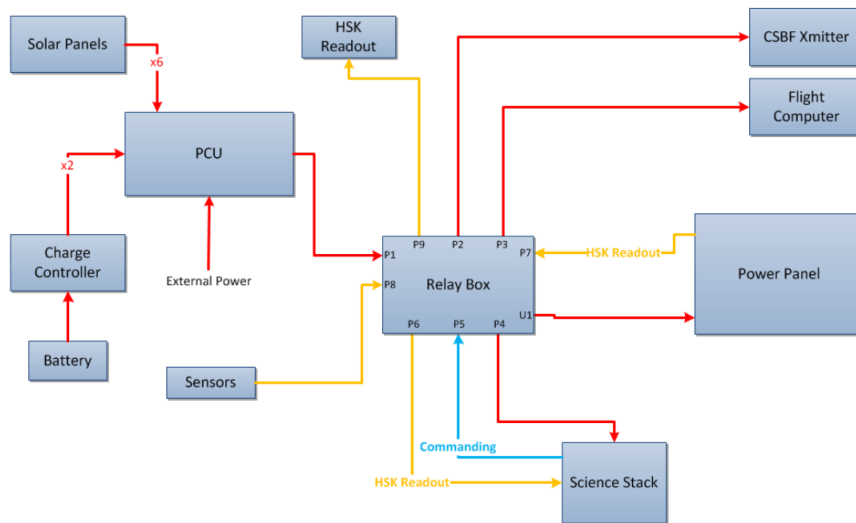


Figure 5.18: Schematic overview of CREST's power system. From [64], courtesy of S.P. Wakely (University of Chicago).






Figure 5.19: Assembly, with handles, of CREST's two Nickel Metal Hydride battery banks. The grey modules on top are the charge controllers.



Figure 5.20: From left to right, Jim Musser, Scott Nutter, Alex Shroyer, Scott Wakely and a helpful IU faculty member whose name escapes me help perform initial integration of the power system with the CREST instrument outside the IU high bay prior to mounting the solar panels. Supporting the panels here are four wooden 'veto boxes' which were used to transport the veto paddles to Plum Brook, Palestine and McMurdo. They were also used to retrieve the veto paddles during recovery.



(a)



(b)

Figure 5.21: (a) Plot of the voltage and current provided by one nickel metal hydride battery bank to the instrument running in normal operation from an unknown charge level to discharge. During flight two such batteries would be active at once. (b) Plot of the power supplied to the instrument over the same time period showing that the total instrument power consumption is between 800 and 850 watts.



Figure 5.22: The Ryder truck full of tools and supplies Alex Shroyer and I drove from Bloomington, IN to Palestine, TX.



Figure 5.23: CREST's shipping container, prior to being unpacked in Hangar 2 at CSBF: Palestine.




Figure 5.24: CREST in CSBF Hangar 2 prior to SIP integration. The power panel (middle left), power control unit (middle right) and relay box (slightly further right, just to the left of the SFC) are visible. Photo by M. Geske.



Figure 5.25: One of CREST's solar panel wings awaits attachment to the instrument in its shipping container. Photo by M. Geske.




Figure 5.26: One of CREST's solar panel wings (on the +Y side) mounted on the instrument in Hangar 2 at CSBF: Palestine. The vector normal to the solar panel face points in the +X direction. Photo by M. Geske.




Figure 5.27: Alex Shroyer examines CSBF's SIP (Systems Integration Package) prior to assisting CSBF install it onboard CREST. Photo by M. Geske.



Figure 5.28: CSBF personnel fit CREST with crush pads (stacks of cardboard tubes), the ‘skirt’ of solar panels that will power the SIP and the ballast container (center). The SIP is shown here mounted in the lower right portion of the instrument. The white-bearded fellow at far left facing the camera is CSBF’s Dave Sullivan, who ran LDB during most of our time in McMurdo and was of immense help to CREST in many ways during our time in Palestine and in Antarctica. Photo by M. Geske.




Figure 5.29: CREST crystal rates during the final hang test at Palestine. All tubes had acceptable rates except for Bus 5, STAC 1, Tube 13. The problem was later diagnosed as a control voltage provision issue on the associated SVI card, which was replaced in McMurdo.



Figure 5.30: Alex Shroyer (left) and yours truly with CREST during its final hang test aboard 'Tiny Tim' at CSBF: Palestine. Although Alex was unable to travel to Antarctica, his efforts were instrumental to CREST's successful completion of the Plum Brook and Palestine tests and the Antarctic expedition. Photo by M. Geske.




Figure 5.31: Preparing CREST for a hang test in Hangar 2 at CSBF: Palestine. CSBF had just finished attaching this metal superstructure, to which all antenna equipment would be mounted in McMurdo. Photo by M. Geske.

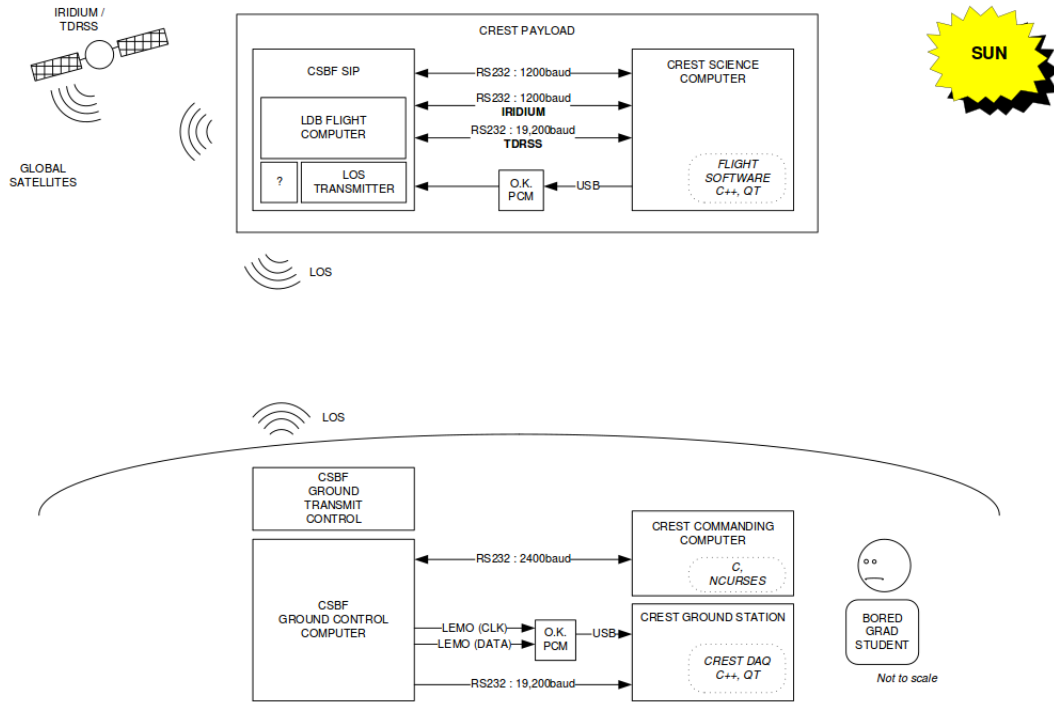


Figure 5.32: Schematic of information flow from CSBF ground station in Palestine, TX, to the airborne CREST instrument via satellite array. Courtesy of staff at CSBF: Palestine.

between the SIP and ground control: the Iridium satellite network (always available), the TDRSS satellite network (available during scheduled pointing events) and direct Line-of-Sight (LOS) transmission, available within a few hundred kilometers of the ground station at McMurdo.

5.2.3.1 Commanding, Instrument Control and *CRESTM*

Our first job was to transition the SFC from receiving commands via the simulated SIP (provided by CSBF to ease such transitions) we'd been using for months to the actual SIP. This entailed several flight and ground command software debugging projects of varying complexity, which I will not detail here. During the flight, commands would originate from a commanding ground station (at McMurdo during line of sight, or at Palestine during other times), through one of the above communications channels to the SIP, then from the SIP to the SFC. During the test at Palestine, commands originated from a ground station computer there, to the SIP and then to the SFC. Commands could be, and frequently were, lost at various stages of this journey. Compounded with the fact that the flight software had to be running and processing these commands properly in order to generate an identifiable instrument response, three important modifications to the flight software were made as a

result of the difficulties encountered during this integration process.

1. Due to the frequent incidence of ‘missed’ commands somewhere in this chain, we modified the monitoring software to display the last command received and which subsystem it targeted. This proved exceedingly useful during testing and especially during flight, as commands which took less than a second to register on the ground often took minutes to register when issued via the Iridium satellite network.
2. Commands that had the potential to disrupt communications with the SFC were gated by a software toggle switch. A command to toggle this switch “on” was required prior to certain commands taking effect. These commands were those used to power down or reboot the SFC, or interrupt the flight software process.
3. Since the flight software process *was* the command processor, once interrupted it could not be restarted, except by rebooting the SFC or means of a monitoring daemon. These interruptions could result from operator error or an unexpected software crash. Further, if such an interruption occurred while the instrument was out of operator contact, either because of a satellite outage or some other unforeseen circumstance, the instrument could feasibly waste a portion of the flight simply by not taking data. To attempt to maintain data collection in such a situation, I developed *CRESTMOn*, a software monitoring daemon, which I describe below.

The *CRESTMOn* daemon defined and periodically monitored two simple criteria to attempt to determine whether flight software was operating normally. When *CRESTMOn* detected that these criteria were out-of-bounds, it would attempt to re-run the flight software. Both the definition of “out-of-bounds” and the corrective actions it would take depended on settable parameters. For example it could be set to run in a monitor-only mode, never taking corrective action; at the other extreme, it could be set to actively terminate an extant flight software process that was not recording sufficient data to the hard drives.

The first criterion was the existence of a flight software process using Linux’s process ID database. The second was decreasing hard disk space over time as would be consistent with normal data collection. Depending on *CRESTMOn* parameter settings, it would wait for one or both conditions to be bad over a specified amount of time (ranging from several minutes to several hours) then attempt to restart the flight software. To ensure that *CRESTMOn* would be running even after an SFC reboot it would be invoked by the operating system’s startup script. I extensively verified *CRESTMOn*’s functionality and lack of harmful action. It was elected by the collaboration not to be used during flight; it was not

invoked by any of the SFC's startup scripts, and could not be initiated by operator command during flight either.

In my manual for operating the CREST instrument [44], I provide a detailed review of the commanding system including a listing of commands, their expected effect on instrument operation and how the command system interacts with the flight software process (including *CRESTMom*), which I omit here as they are unlikely to be of further interest. Figure 5.32 provides an overview of the communications scheme between ground, satellite network and instrument.

Hardware- and electronics-wise, CREST was already in exceptional shape and we spent most of the time at Palestine waiting for CSBF to perform SIP integration. During this time, the power system's power control unit underwent some major upgrades, mainly to prevent overheating of any of the six main MOSFETS which acted as resistive voltage dividers. We also configured two PC's, dubbed Ground Station Equipment 2 and 3 (GSE2 and GSE3), to act as command uplink and data downlink relays. By remotely connecting to the active ground station, CREST collaboration members could communicate with CREST during flight via POCC's satellite connections. We thoroughly tested the command link operating in this fashion and configured GSE2 to be the main link, with GSE3 as a backup.

5.3 Pre-Antarctic Preparation

5.3.1 Instrument Operation Training

After completing integration with CSBF and buttoning up the instrument and all of our gear for shipment by boat to McMurdo we left Palestine. It became apparent that most collaboration members (besides myself and Jim Musser) had almost no experience actually operating CREST. With the instrument already on a McMurdo-bound container ship, there would be little chance of gaining that experience first hand, as I had done while working directly with the instrument in the months leading up to the Plum Brook test. Since CREST would be operated during the flight by teams of any two or three collaboration members, I resolved to remedy this situation by writing the "CREST Manual" (distributed within the collaboration as [44]). This document summarized the information needed to keep data taking going in the face of various obstacles, such as the flight software crashing, instrument rates vanishing or various other conditions, most of which could (typically) be remedied with a few simple commands. I also detailed other more involved actions such as masking off detector components, changing discriminator thresholds or requesting that certain housekeeping data be sent, among many others. I included minutiae such as sensi-



Figure 5.33: CREST on the flight pad at CSBF following our successful hang test. Clearly visible on the lower lefthand thermal panel is a circular region of the aluminized mylar sheeting which bubbled off of the panel's surface during the low pressure portion of the Plum Brook thermal/vacuum test. These panels were replaced in McMurdo, along with most of the rest of CREST's thermal shield. The new panels were laminated with perforated silverized teflon tape rather than aluminized mylar sheeting, both to avoid this failure mode and for improved thermal performance.

tive detector numbering schemes, CSBF lingo and other information useful for instrument operators. I also summarized the state of our software repository at the time, focusing on the software needed to display the stream of housekeeping data used to make real-time instrument operating decisions.

To facilitate simulation of the flight environment Jim Musser, developed a program to transmit housekeeping data from GSE3 at the POCC to member's individual workstations. We organized two tutorial conference calls assisting collaboration members in viewing this data stream with the software monitoring tools he and Scott Wakely had developed. This experience proved to be critically important when "The Incident" occurred during the second day of our flight, when both Jim and I were out of commission (i.e. exhausted) following the flurry surrounding the launch.

5.3.2 Ground Station Maintenance at the POCC

On or about the evening of Monday October 8, 2011 the POCC (Palestine Operational Command Center) experienced a power surge followed by a power outage. This damaged the motherboard of GSE2 which until then was to have served as the command uplink between collaboration members' computers and the satellite communications equipment CSBF would use to relay our commands to the SIP aboard CREST via the TDRSS and Iridium satellite networks.

During the week of Monday, October 10th, 2011 as I prepared to conduct commanding and instrument monitoring tutorials with the collaboration we noticed that we had lost the ability to communicate via Internet with GSE2 and GSE3.

After consultation with CSBF engineer Joseph Jones (CREST's lead integration engineer at the POCC) it was determined that GSE2 had suffered hardware damage as a result of the power surge/failure and needed to be repaired. It was shipped back to Indiana University, had its motherboard replaced by Dell (the computer supplier), then shipped back to POCC. GSE3 on the other hand simply lost power and seemed to be operating normally after restart. However GSE2, not GSE3, had been configured with the serial connections to POCC's equipment, and Joseph Jones was unable to get commanding working merely by swapping cables. As we had planned on having two functional ground stations at POCC in case of just such an incident, and commanding was currently non-functional, it was deemed mission-critical for a CREST team member to re-visit POCC and remedy the situation.

I flew to Palestine on 10/19/2011 to investigate. Once there, I reworked the serial port configuration of the commanding software to enable GSE3 to communicate with POCC's communications equipment. While we were unable to fully test commanding the instru-

ment through this new interface (as CREST was currently on its way to a boat to be shipped to McMurdo), Joseph Smith and I were able to verify receipt of commands at POCC from collaboration members accessing GSE3 remotely. GSE2 was now nominally functional (having been repaired and shipped back to POCC); we left it powered down and unplugged to isolate it from any further damage. We protected GSE3, now the main uplink and downlink between the collaboration and CREST during flight, with an uninterruptible power supply and surge protector.

5.4 Antarctic Expedition

Many of the details about daily life and work on the instrument recalled here were recorded by myself and reported to the collaboration in my daily serial “Updates from McMurdo” [45]. Each evening I reported the events of the previous day via email, ranging from detector-centric issues, to interesting events at LDB or McMurdo or anything else that collaboration members not able to be in the hangar with us would enjoy knowing. In all I recorded the events of approximately 45 such days, between 11/6/2011 (our first full day in Hangar 096 at LDB) and 12/25/2011 (launch). On the two occasions I was unavailable the duty was taken up by a comrade (once by Jon Ameal, and and once by Stéphane Coutu).

All of our collaboration members took copious amounts of pictures and video. I took approximately 3,500 photos and dozens of hours of video. Many of the photos reproduced here were taken by myself; I have attempted to credit collaboration members for their photographic efforts but in some cases I simply don’t recall who took the picture, in which case credit should be assumed to hew to the CREST collaboration.

5.4.1 Living in McMurdo

As mentioned above, following the successful hang test at Palestine, CREST and all of our gear was sent to McMurdo via shipping container aboard a boat. Our trips to McMurdo routed us through Christchurch, NZ, the site of a series of catastrophic earthquakes in previous years which had reduced much of the center of the city to rubble. Despite the extensive physical destruction, residents of the city welcomed us with open arms. The natural beauty of the area and the optimism with which they undertook their vigorous rebuilding efforts uplifted us during our three-day stay, as we awaited a US Air Force flight to McMurdo Station.

Much of our early time in McMurdo centered on getting to know the town and our worksite at LDB, and navigating the plethora of training classes required to do pretty much

everything. Dave Sullivan of CSFB kept us on a strict schedule of making it onto the bus in McMurdo at 7:30AM each morning, with latecomers being left behind for the day. Though it was more difficult for some than for others, no CRESTacean missed a single trip without intending to take the day off. The hour-long bus ride aboard the slow-but-steady Ivan the Terra Bus made for excellent morning napping, reading and conversational opportunities.

Living in McMurdo reminded me of being an undergraduate. Our living quarters were rather like college dormitories, complete with roommates, shared bathrooms and common areas with couches, board games and televisions. The perpetual sunlight threw many of our diurnal cycles off, both because of the obvious light issues and due to the constant noise from fellow residents coming and going at all hours of the day/night.

Meals at the main McMurdo mess hall were satisfying, especially the bread, which was baked fresh daily and very delicious. The freshness of fruits and vegetables was excellent, whenever a delivery had just come in (which happens only a few times until the ice breaks later in the summer). The most tedious mealtime rituals were handwashing, and finding your 'Big Red' jacket amongst the hundreds of identical jackets on the coat racks, differentiated only by nametag. After a week or so of such hunting, one became quaintly familiar with the stains, discolorations and bits of fuzz that decorated one's own Big Red, enabling recognition of one's own jacket without aid of the nametag.

Handwashing, on the other hand, never became any easier. It was a critical step to attempt to avoid catching (or for the less fortunate, spreading) "The Crud," a highly infectious super-cold that developed at McMurdo presumably as a consequence of the mixing of viral populations from various countries. Those afflicted suffered mild-to-serious respiratory distress, especially congestion, coughing and soreness. This often led to loss of voice. All of this was made less bearable by the very dry atmosphere. It reminded us of a combination of bronchitis and respiratory flu, of viral rather than bacterial origin (antibiotics did nothing for it). Almost all of us caught it eventually (including yours truly), though some held out longer than others. At any given time, I estimate approximately 15-20% of McMurdo's population either had the crud or was recently recovering from it; the impossibility of taking off work or avoiding the mess hall for the sick made its isolation impossible. Periods of rapid turnover surely exacerbate this fraction, as the fraction of the population lacking immunity would then be greater than normal.

5.4.2 Working at LDB

Hangar 096 at LDB was an ideal working setting for our collaboration. The lower floor housed the instrument, storage space and the makeshift electrical engineering stations set

up by Jon Ameel and Michael Lang. We initially utilized the mezzanine balcony above for reconstruction of the foam thermal panels, then used it as working and writing space, meeting space and jamming space (for our ragtag rock band, called The Longdrops). The more adventurous nappers among us formed nests of blankets and bags beneath the folding tables, taking on the risk of a somnolent prank or two in exchange for a few minutes of rest.

Every few days, a tour group from McMurdo would arrive just before lunchtime, putting whatever unfortunate collaboration member happened to be in the hangar when they arrived in the unenviable position of justifying CREST's practical and scientific value. The groups were recruited from those employed in McMurdo and served mainly as a diversion and reward, a recess of sorts. CREST's imposing stature, nests of cables and banks of blinking LEDs aided in the speaker's attempts to awe the unfamiliar with the majestic universal truths accessible only through cosmic ray detection. But the fastest way to terminate the question and answer sessions tended to be, "Have you had lunch yet?"

5.4.3 First Launch Attempt: Bitter Cold

On 12/18/2011 we left McMurdo for LDB at approximately 1:00 AM to begin preparing for our first launch window. The last time we had truly thermally secured the instrument was prior to the 2nd test procedure at Plum Brook. That experience, gained in the comfort of the early Ohio summer, did little to prepare us for the difficulty we would face this bitterly cold morning.

As shown in multiple pictures herein, during various rollout procedures, CSBF attached cardboard crush pads around the bottom of the instrument in addition to a ballast module directly below the instrument center. The attachment schemes for all of this material called for mechanics to screw bolts through CREST's lower chassis struts, then tighten them from within the instrument's interior. This meant that CREST had to be thermally secured *outside* the hangar. The vertical clearance, from crane to hook to floor, was insufficient to crane CREST out with the ballast container attached below.

This would have been only a minor issue during most of our time at LDB. Though the sun and wind were harsh, demanding sunglasses and sunscreen even on temperate days, our thermal tape was easy to apply and screws were relatively easy to attach at 15-25° F. This day proved to be the coldest we would experience on our expedition, with temperatures well below 0° F, exacerbated by a biting wind. Avoiding pain and eventual frostbite required multiple layers of gloves and headgear, making screwing in the eight small screws securing the eight lower thermal panels an arduous task. Jon Ameel and Mike Lang bore the greatest brunt of this hardship and soon became too cold to continue. We went outside

in groups of two and three, clumsily holding screws in place with mittened hands. Our halting progress came to a complete stop when it came time to secure the diagonal cross struts that went over each panel. Tension from the crane lift had displaced the screw holes, meaning mechanically the struts would not fit. They needed to be in place *before* the crane lift, not after.

Abandoning attaching these struts, we attempted to finalize sealing the instrument's thermal skin with silverized teflon tape. However the air was so cold that the tape's adhesive froze before application. Trusty packing tape and supposedly cold-temperature-proof Nashua duct tape performed little better. We attempted keeping strips of tape warm inside our coats until immediately before application, but this did no good. As a final attempt, I suggested heating the tape with a heat gun after application. This worked, but was incredibly laborious. It was also somewhat dangerous, as by this point CSBF had lifted CREST onto The Boss and rolled away from the patio. To reach the instrument we had to clamber up a stepladder, and hold on to the instrument as it swayed in the stiff wind. Realizing the danger, and the shortness of time, Jim Musser enlisted CSBF's help in lashing the foam panels in place with ripcord. Significant gaps remained between the foam panels and the instrument surface, and at the edges of the lower sheet metal panels. CREST was not thermally secure.

This soon became a moot point as our weather forecasters scrubbed the launch attempt due to the extremely windy surface conditions. Our entire contingent was frozen, tired and disappointed. Several of us made our displeasure at CSBF's unhelpful attachment scheme known. To their credit, they revised the attachment plan to enable us to thermally secure CREST inside the hangar for future launch attempts.

5.4.4 Second Launch Attempt: Battery and Power Issues

After remedying the errors that threatened to scrub our first attempt even before the winds failed to cooperate, we awaited word from CSBF for another launch opportunity. As we were aware of no detector problems that we intended to fix¹, we sought diversion in many areas as we bided our time. Our chance would come on Wednesday 12/20/2011, so we prepared the night before to catch a 4AM shuttle to LDB.

At approximately 9:30 PM that evening CSBF engineer Joseph Smith contacted Jim Musser to inform him that a battery temperature sensor had exceeded 80° C. As we had taken to leaving the instrument running overnight, we were able to check the instrument's

¹Bus 3 STAC 2 was non-operative, but as experience had taught us that removing an entire C-channel tended to break on average one instrument module, we chose to simply leave it inactive, rather than attempt to debug or replace it.




Figure 5.34: Surface temperatures of CREST’s two NiMH battery banks shortly before and after bank 1’s apparent charge controller failure. Bank 1 and its charge controller were replaced with spares hours before the second launch attempt. Bank 2’s temperature and operating characteristics remained stable throughout.

power system status remotely, and noticed nothing amiss other than the high temperature reading. Confident that we were merely going to replace a faulty temperature sensor, but wanting to be absolutely sure, Jim, Nahee and myself procured transportation to LDB to investigate. Upon our arrival, Nahee observed that the charge controller had malfunctioned, allowing the bench power supply to overcharge it to abnormal levels. Figure 5.34 shows the temperature of both batteries before and during the malfunction.

Under Nahee’s direction, we removed battery bank 1 from the instrument and replaced it with the spare battery bank and charge controller. After doing so and attempting to power up, CREST’s behavior was extremely erratic. We had lost commanding ability, and some instrument buses refused to power up entirely. We informed CSBF that we did not believe we would be ready to fly that morning; they responded that they would show up and be ready to launch anyway.

After continued iteration, and with valuable assistance from Joseph Smith, we restored commanding ability, and were able to return all instrument buses to nominal operation shortly thereafter. Hours later, we were convinced that the instrument was in fact flight ready, with the exception that one battery bank had almost no charge. This could have resulted in a power down during ascent, when solar power would be intermittent at best. Thus we maintained connection to the bench power supply in order to charge the battery as much as possible.

Thanks to CSBF's revised skirt and ballast attachment procedures, we were able to thermally secure the instrument within the hangar's cozy confines, and were ready for rollout well ahead of time. Once we got CREST rolled out and hung from The Boss, we made sure to continue charging the new battery via a generator ('Maverick') onboard The Boss. Mistakenly, this led to an unfortunate confluence of battery charging conditions. The orientation CREST was hanging from The Boss presented the panels with maximum solar flux. The batteries were receiving full charge from the generator and the sun, and the PCU began to heat up dramatically. Scott Wakely, the power system guru, had traveled back to Chicago the previous week but was closely monitoring the power system remotely. He was the first to notice the problem, and frantically attempted to alert us via Skype. Amidst the buzz of the rollout his warnings went unheeded for several minutes. Once again Nahee came to the rescue, independently noticing the heating condition. We promptly disconnected the generator power supply, and ordered a re-orientation of the payload away from the sun.

Although after several minutes the thermal state of the power system had returned to nominal, we would not know until testing the MOSFETs in the PCU later that day that no permanent damage had been done. While this exact situation could not occur during flight and therefore did not call for any immediate modifications to the power system, it did teach us to be far more vigilant with respect to the power system's status. This served us well during flight operating shifts. Requests for instrument rotation were often met minutes or hours later, meaning we had to predict when the batteries were nearing full, and take pre-emptive rotational action to reduce the solar power generation.

After an hour or so, the launch attempt was scrubbed. Despite the complete lack of clouds and surface winds, pilot balloon launches revealed that wind patterns at higher altitudes were too fierce to risk a launch. Disappointed, but glad to have avoided permanent damage to our power system, we lived to launch another day.

5.4.5 Third Launch Attempt: Success

Other than a launch window which was open for precisely one minute, then closed due to intense surface winds on 12/22/2011, until Christmas Day we had no other launch opportunities.

We spent part of our copious free time streamlining the thermal sealing process. A major improvement was the installation of hinged doors on the lower thermal panels, enabling them to remain on the instrument permanently while still allowing last-minute access to power cables and other sundries. Jim Musser (at LDB) and Scott Wakely (remotely, at

Chicago) continued to refine and augment our housekeeping displays with additional information and stability. Nahee continued testing the PCU for signs of damage, including a test of the solar panels during a rollout which indicated that the power system was fully operational. Whenever we received word of a TDRSS satellite communications window opening, our State-side collaboration members were able to connect to GSE3 at the Palestine POCC, and try out instrument commanding and monitoring, providing valuable additional practice opportunities.

We received word of our next launch window on Christmas Day, opening at approximately 5:00 PM. Though we had grown skeptical given the multiple scrubs, we could have gotten CREST ready in our sleep at this point. We had the instrument all buttoned up and ready to go when the clouds parted, revealing a majestic Christmas day on the Ross Ice Shelf. Though there was a slight breeze, we shared the feeling that this attempt would prove decisive.

Word came a few hours later that pin release was planned for 5:30 PM. The activity at LDB, on the road to the launch pad and on the launch pad itself, grew to a fevered pitch. Road smoothing trucks swept up the tracks left by heavy vehicles delivering liquid helium, the balloon and its unbelievably long train, and a plethora of support equipment, not to mention The Boss, CREST swaying placidly from its crane.

Our team kept watch over the instrument via the housekeeping data provided by the LOS transmitter, specifically the performance of the power and readout systems. Should anything go haywire we would have to inform CSBF immediately - once the pin was pulled there was no way of getting CREST back without aborting the entire flight. No such reason presented itself, and we gave CSBF the green light at all stages. Overcome with happiness, Jon Ameel took to flying a kite he had shipped from Palestine for exactly this purpose. This prompted a low-level alert and a radio warning from the launchmaster about “a rogue kite or something on the launch pad roadway.” We handed Jon Stéphane’s high-definition video camera and assigned him cameraman duties instead, which he conducted with aplomb.

Besides monitoring the instrument, there was one last task for us to perform. The power cable connecting the Maverick power supply aboard The Boss was still connected to CREST’s power box, trickle charging the battery banks in preparation for the worst case of no solar power generation during ascent. Mike Lang and I waited by the instrument for word from Hangar 096 to cut the cord. I took this opportunity to take photographs of CREST from an opposite viewpoint of every other photo- and videographer on hand. When Nahee confirmed via radio that the batteries had adequate charge and launch was imminent, Mike cut the cords then hopped onboard my snowmobile. I expertly but cautiously drove us back to the safety of the patio of Hangar 096 to join the rest of the collaboration.

Finally at just about 5:30 the launch controller's loudspeaker-augmented voice initiated a countdown. Jim frantically reminded us to keep an eye on our instrument readouts, but to little avail - we were transfixed by the sight of the massive balloon being freed from its stay. With an impossible lightness it took to the sky, the tension in the line steadily increasing until it seemed almost as if the balloon would pick up The Boss itself. Releasing CREST with the nimble heft of an elephant's trunk, The Boss's driver spewed a cloud of thick soot as it lumbered safely clear from CREST, now entirely supported by the balloon's tether. CREST took to the air amid cheers, whoops and screams. A team of approximately 100 people had just successfully launched a 7,000 pound hunk of aluminum, foam and physics into thin air, and they let the entire ice shelf know about it.

All cameras and eyes pointed at the cube with wings, which in a few minutes was reduced to a lighthouse-like flash of light once every few seconds. CREST spun as it ascended, the sun's reflection signaling to us in Morse code: "Get back to the ground station!" Or was that Jim shouting? It's very hard to remember at this point.

In any case, we had an instrument to monitor. We filed back into the hangar, clapped each other on the back, poured some champagne and watched the data begin trickling in.

5.4.6 The Incident

At around 6:00 UT on 12/26/2011, approximately 38 hours into the flight, incoming shift operator Scott Nutter and outgoing operator Nahee Park noticed that some housekeeping data showed abnormal results during Disk 1 Run 3509. Specifically, the "trigger %" (the percentage of recorded non-empty freeze cycles that met various criteria, for example more than three crystals activated) was returning "nan" (C++ parlance for Not A Number). The denominator, the total number of non-empty freeze cycles, was 0, making the trigger ratio incalculably large. While such abnormalities were not uncommon, they were also normally short lived, or easily rectifiable via a CROL reset. When it persisted through several such resets, Scott performed a full instrument reset, a much more drastic action which resets every FPGA in the entire readout system, which fixed the issue temporarily. However, the "nan" behavior soon returned.

At the same time, "the CSBF rotator reset itself" (this is a direct quote from NASA's rotator engineer), and started rotating by arbitrary amounts centered on the requested angle off-axis. Control was regained minutes later, but this too was very unusual. Several hours later Scott handed control off to Michael Schubnell, who experienced similar trouble with the rotator and the "nan" housekeeping values. After attempting multiple instrument and SFC resets without improvement, Michael noticed that although the flight software

correctly reported “SFC Reboot” as the last command received, the run numbers never deviated from zero. The SFC was overwriting Disk 2 Run 0 every time we rebooted.

Realizing something was seriously amiss, Jim Musser took over and begin investigating in earnest. I was scheduled to fly out to Christchurch later that day, and rushed to LDB via shuttle and ski when awakened and alerted to the problem by Stéphane. After many hours of pounding away, Jim found little success in operating the full instrument for more than a few seconds before instrument rates crashed and data taking ceased.

He was able to simultaneously operate approximately half the instrument (i.e. crystal buses 1-4, but not 5-8). He wrote these settings to the default instrument configuration, so that future resets would cause the instrument boot up in this state. Raw event rates in this configuration fell by approximately 40% (from approximately 5000 hz to 3000 hz - see Figure 5.41).

Contemporaneously with the rates issue, the SFC had issues of its own. Under normal conditions when creating a new data file, the flight software would write the current run number to a text file in its home directory. After the incident it was no longer capable of reading this file and always resorted to the default run number of zero after reboot. This could be fixed manually by issuing a command to change the current run number. However, this meant that runs 0-1 on disks 2, 3 and 4 were constantly being overwritten every time the SFC was rebooted, and, as a corollary, that there was probably some errors related to disk input or output on the SFC. Figure 5.35 shows that the free disk space on disk 2 stopped decreasing at precisely the same time as the SFC current draw from the power box increased. Disks 1, 3 and 4 appeared unaffected. During copying of the flight data post-recovery, disk 2 displayed no errors; thus the problem may have been between disk 2 and the SATA controller, which connected all four hard disks to the SFC motherboard.

These two observations show that something funky happened with the SFC right around the time the instrument performance degraded. To my knowledge no one (including me) has come up with a good suggestion as to what could have caused damage to both the SFC and the crystal overlord at the same time.

Early on in this process, I had suggested that we prepare a checkerboard pattern with 4 STACs in each C-channel on in an alternating pattern, hoping to lose slightly less than half of CREST’s effective area by covering a larger area less compactly. Greg Tarlé enacted this mode with Jim’s blessing soon after I arrived in Christchurch. This resulted in slightly higher raw event rates (approximately 4000 hz compared to 3000hz), though the post cut rates were not significantly altered (and the altitude was different, perhaps making the comparison not entirely appropriate).

CREST limped along in this half-dead configuration for the remainder of the flight,




Figure 5.35: Plot showing the disk space remaining on SFC disk 2 (in GB) and the current drawn by the SFC from the power system. At the same moment disk space stopped decreasing and current draw increased by a measurable amount. It is believed that at this moment SFC's connection to disk 2 was lost.

except for brief periods of tinkering with various instrument parameters, mainly crystal discriminator thresholds, in hopes of find some working state that had a larger percentage of the instrument active.

Meanwhile the balloon was steadily losing altitude, as can be seen in Figure 5.39. Since the anti-polar cyclone center's deviation from the geographic south pole increases as altitude decreases, this meant CREST was at greater risk of heading out towards sea, were it allowed to attempt another circumnavigation of the pole. Since this would entail either an incredibly difficult recovery (if CREST had to be cut down prematurely over a relatively inaccessible area) or no recovery at all (if the landing site were too remote, or if CREST fell into the sea).

Put on notice by CSBF that CREST's hours were numbered, Jim gave Nahee and I free rein to attempt various ideas to try to get CREST working fully again, or at least get a better idea of what the root cause might be. To us the most striking difference between ground operation and flight operation were the veto hit rates: each veto tube was producing as many signal hits as entire Sedecims of crystal tubes. On the ground each veto tube typically produced a third to a half of that of an entire Sedecim. After futilely attempting to reduce veto hit rates by increasing their discriminator thresholds (it seemed that no matter the threshold the rates remained abnormally high), we hit upon the idea of reducing the high voltage provision to all veto tubes. After much experimentation we determined that the full crystal array could be successfully read out with the veto PMT high voltages at or below 30% of their nominal values. Figure 5.36 shows a housekeeping plot of STAC and

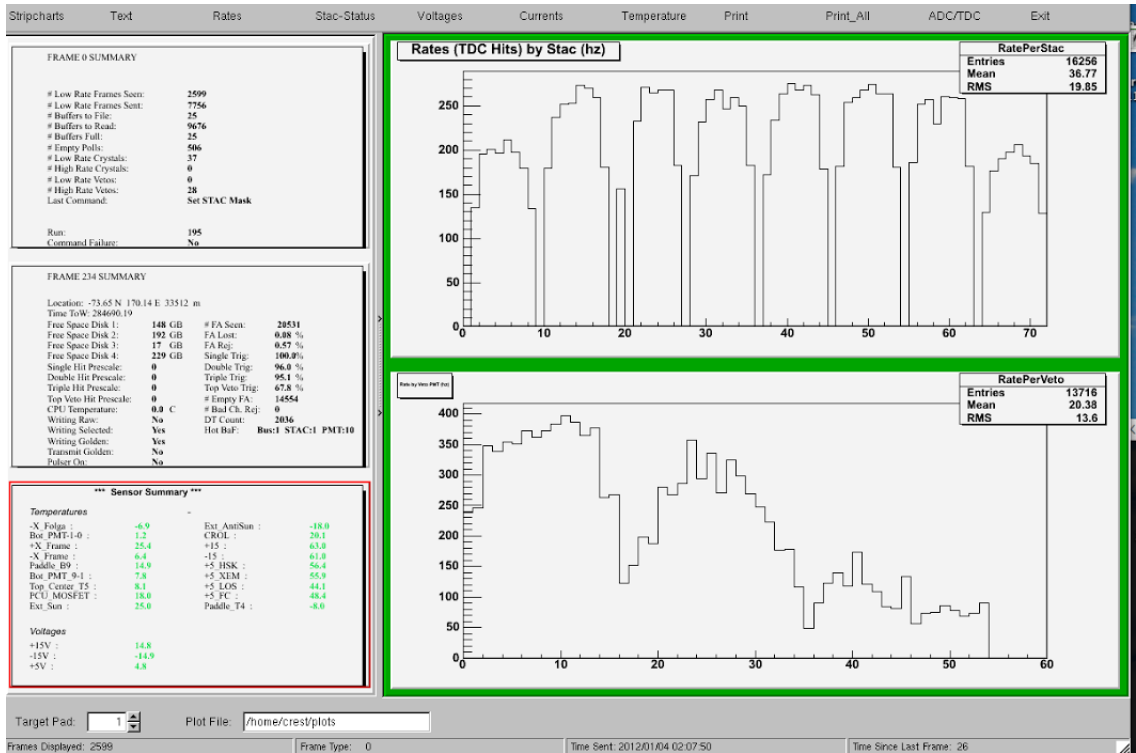


Figure 5.36: A housekeeping plot generated during the last few hours of flight, showing crystal STAC (top right) and veto tube (bottom right) rates in hz. With the veto high voltage provision at 30% of normal the full instrument could be operated in a stable manner. With the veto high voltage at normal values, the CROL would crash within minutes, reducing crystal STAC rates to zero.

veto rates at that setting, showing that all instrument STACs reported typical rates.

While a full instrument autopsy was never and will never be performed any theories as to what really went wrong would need to be consistent with the following critical in-flight observations:

- Before “The Incident,” the instrument could take data from all 63 crystal STACs with nominal threshold and high voltage settings
- After “The Incident,” activating more than 32 crystal STACs with nominal settings broke data collection
- After “The Incident,” the instrument could take data from all 63 crystal STACs with veto high voltages below approximately 30% of nominal values

While I can not explain the simultaneous disabling of SFC’s disk 2, CSBF’s rotator behaving anomalously and the degradation of readout performance, I believe that the ex-

planation of the latter lies in the subtle, non-linear interaction between the digitize system, the active crystal area and the veto hit rate.

As discussed in Section 3.3.7, veto hits do not contribute to the digitize decision. Thus even if a veto tube was signaling hits to its host STAC at the maximum rate allowable by the discriminator logic, this on its own could never influence the rate at which the DIGOL initiated an instrument-wide readout. Therefore an increase in veto hit rates *on its own* cannot cause an increase in the rate at which the CROL initiates a readout procedure.

Nevertheless our experience with the pulser system could explain how very high veto rates, either from a light leak or from a very high flux of low energy charged particles from solar storms, could choke the CROL. Initial attempts to use the pulser system failed when flashing the entire crystal plane at once. This taught us that freeze cycles with more than a certain number of hits in them result in a huge amount of deadtime induced at the CROL, as it attempts to histogram all the hits. It only has approximately 131 μs to finish processing all hits, and when a freeze cycle contains more than approximately half the detector's worth of crystals, this deadtime increases exponentially.

The reason CROL deadtime would be so sensitive to the number of active crystal STACs in such a scenario is because veto hits are only read out when they are coincident in time with a trigger-satisfying pattern of hits in the crystal system. The digitize FPGA only cared about input "Single Hit" signals from the crystal system. Below a critical number of unmasked crystal STACs, the digitize decision trigger rate convolved with the number of hits in each freeze cycle was below the maximum rate digestible by the CROL. Above that number of unmasked crystal STACs, the trigger rate convolved with the number of hits in each freeze cycle was above that maximum rate. In other words, a hyperactive veto system couldn't increase the rate at which CREST triggered, but it could increase the average number of hits in each freeze cycle, and the impact the veto system had on that statistic was itself dependent on the number of active crystal STACs. While the analysis required to verify such a claim is probably impossible to complete given the unknown state of the instrument (and the concomitant low degree of confidence in any rate measurements made with the data from around this time), this explanation is consistent with the CROL being able to read out all 63 crystal STACs when the veto high voltages were reduced to 30% of their nominal values.

If this were the case, a simple way to avoid this issue would have been to have the coincidence trigger settable in-flight. The first layer of data reduction was the digitize system. Though it succeeded in its role of guarding against blindingly high rates, it failed to protect against the un-anticipated case of blindingly large freeze cycles. While the capability to set digitize parameters in-flight was suggested by several of us during instrument integra-

tion, it was decided that it was more important to keep the CROL and DIGOL FPGA's out of direct communication. Since commanding is routed through the SFC, and the SFC has a USB connection to the CROL but not the DIGOL, such a command would have to be routed through the XEM Carrier circuitry somehow, which could have had unforeseen negative consequences.

Another way of guarding against these “blindingly large freeze cycles” would have been to add another criteria to the digitize decision. The two existing criteria, the single-hit prescale and the coincidence threshold of three STACs, both guarded against high rates of small freeze cycles. A third criteria, a *maximum* number of activated STACs, including veto STACs, could have reduced the number of very large freeze cycles being read out to the CROL. Events consisting of the entire detector being activated would certainly be culled in most analysis strategies, and could have safely been eliminated at the digitize level. Future flights with similar read-out architectures may benefit from this type of safety valve.

Even if CREST had been completely stable in this “veto-less” mode of operation, it is likely that data taken after this point would be of limited incremental scientific utility, since the veto system provided valuable information used to characterize events in the analysis stage. Based on this consideration, and on CSBF's urgings that we avoid delaying or endangering instrument recovery, Jim Musser gave the go-ahead for flight termination at around 11:30 EDT on 1/4/2011. As collaboration members comforted each other in the Skype chat with dry humor, CREST's telemetry readings provided a macabre countdown to the demise our mutual creation.

Personally, the decision to terminate the flight so prematurely felt like a rational, but unadventurous, decision. Had I been in charge, CREST would probably be at the bottom of the ocean right now, or slowly being consumed by the shifting snows of the far side of the continent, recovery impossible, on-board data lost forever.

5.4.7 Recovery

Though CREST had landed in a relatively convenient location (only 250 miles from McMurdo on flat terrain), recovery was delayed repeatedly due to a combination of weather and the lack of available aircraft. Logistics planners had not planned on CREST having such a short flight, and our requests for air support competed with missions to recover *people*, not machines, from far-flung bases. For our remaining personnel, Jon Ameel, Matt Geske and IU undergraduate student Thomas Bishay, this meant nearly a month of false alarms, endless recovery planning sessions and tedious waiting. Finally, these three and




Figure 5.37: A composite graphic overlaying electron and proton flux data from NOAA’s Polar-orbiting Operational Environment Satellite (POES) (yellow shading) with CREST’s flight path up to that point (red squiggly line). This was subsequently referred to as the “Yellow Ring of Doom” as a possible explanation for CREST’s degraded performance beginning the second day of flight. Areas of intense yellow indicate high power of protons and electrons of solar origin during an M2-class solar flare which occurred on 12/26/2011 at 8:00 PM UT. The onset of instrument performance degradation (on 12/26/2011 at 6:00 UT) coincides with CREST leaving the safety of the eye in the center of the solar storm. Whether or not this abnormally high solar activity in fact had anything to do with CREST’s degraded performance is unknown. The instrument’s behavior is consistent, however, with the veto system detecting an abnormally high flux of low energy charged particles as would be expected from this kind of situation.

various CSBF/Raytheon personnel recovered CREST on February 2nd-3rd, 2012, traveling to the site via Bassler aircraft. A team of landing-strip groomers preceded them in a Twin Otter aircraft on 1/30/2011.

The team took pictures from air as they approached and on the ground. The instrument was in remarkably good condition, the major structural damage occurring due to the balloon rotator falling through the +Y-side solar panel. Their first task (after ensuring the batteries were safe to approach) was to secure the flight data disks.

Having done this, with a total of 12 men on-site, recovery went relatively quickly. Jon, Matt and Thomas separately guided teams of non-science personnel in disassembling the sensitive detector components (veto paddles and FOLGAs, veto PMTs, and crystal C-channels). The aluminum chassis was sawed into manageable pieces and loaded onto the aircraft as scrap. A chief concern for CSBF is to avoid polluting the pristine landscape with any remnants of the instrument.²

They initially used a snowmobile to move the material back to their aircraft (which had landed 2 miles from the recovery site). After surveying the scene the Bassler's pilots decided to taxi to the recovery site for more rapid loading. After loading up the Bassler it had an estimated 8000 pounds of cargo. While not too heavy for the craft to take off on the runway prepared by the groomers, on the loose snow it caused the wheels to get stuck (much like other vehicles on the ice road). They had to dig the wheels out twice to be able to taxi back to the runway.

Once on the runway they were informed that McMurdo was experiencing 40+ knot winds with low visibility, meaning that the team would have to camp out overnight. They elected not to build an ice wall due to fatigue and slept in sleeping bags inside the plane. They shivered through a frigid night, with sub-zero temperatures and a biting wind.

Come morning they left the snowmobile on site to save weight, attempted take-off and succeeded, finally returning to McMurdo for a hero's welcome from the rest of the collaboration via email. They left McMurdo the following Monday, on one of the last flights back to Christchurch of the summer.

²This led to an amusing and oft-broken rule at LDB of no unsealed beverage containers (especially not coffee mugs) allowed outside the hangars or kitchen. Supposedly coffee spills mandated an incident report and collection of the soiled ice. Luckily, none of us ever spilled even a drop of any beverages.



Figure 5.38: CREST’s flight trajectory, starting from McMurdo on 12/25/2011 NZDT (12/14/11 EDT) and ending in nearby Victoria Land ten days later.

5.4.8 Flight Profile

After launch from McMurdo the balloon was at the mercy of the polar anti-cyclone and began heading roughly towards the South Pole. Since the cyclone did not set up directly over the pole, CSBF deemed that its trajectory risked spiraling into the sea.

Stéphane Coutu utilized Antarctic atmospheric models to convert local pressure (as measured by CSBF’s onboard sensors) into an estimated atmospheric overburden. The local pressure and calculated overburden is plotted in Figure 5.39. For the first day or so of flight CREST’s altitude was approximately 36 km (118,000 ft) and overburden was between 5 and 6 g/cm². Approximately 1.5 days later (around the time of “The Incident”) CREST’s altitude profile changed, dropping to an average altitude of between 32 and 33 km (modulo daily oscillation related to the zenith angle of the Sun). This led to an increased overburden between 7-10 g/cm² for the remainder of the flight.

The magnetic field sensor (pictured in Figure 3.39) provided by Nilson Remo of the University of Michigan Atmospheric Oceanic and Space Sciences department took data for approximately 2 days. This provided the measurements shown in Figure 5.40 for the magnetic field at CREST’s location along the sensor’s X, Y and Z axes.






Figure 5.39: The environmental pressure and atmospheric overburden at CREST’s flight altitude. The pressure was measured by onboard instruments while the overburden was calculated from Antarctic atmospheric models. CREST’s overburden was below 6 g/cm² for only a short duration which roughly matches up with data taken when the full (rather than half) detector plane was operational. Plot by S. Coutu.

Event ID	Disk	Run Number	Description	Date	Time (EDT)
1	1	3477	Launch	12/24/11	23:20
2	1	3482	Reach float altitude (117k ft)	12/25/11	03:00
3	1	3513	Onset of instrument instability	12/26/11	13:00
4	2	21	Running with first four buses	12/27/11	17:00
5	3	48	Switch to disk 3	12/28/11	22:30
6	3	109	Switch to checkerboard pattern	12/31/11	12:15
7	3	157	Adding additional STACs	1/2/12	08:00
8	3	186	Last run before veto-less mode	1/3/12	17:08

Table 5.2: Chronology of significant detector status changes during CREST’s Antarctic flight. Runs between Disk 1, 3482 and Disk 1, 3513 were generated with the full detector plane active at float altitude. Various periods after that consist of runs with only half the detector plane active, either in a checkerboard or solid-half pattern, and at lower altitudes.



(a)



(b)

Figure 5.40: The magnetic field sensor’s measurements during the Antarctic flight. *Top, a*: The Z-component of the field was approximately constant at 0.6 G (downwards), with the X and Y components oscillating daily with a maximum transverse magnitude of approximately 0.12 G. Error estimates for each data point were not available. *Bottom, b*: The verticality of the magnetic field at CREST’s location ranges between approximately 75% and 90%.




Figure 5.41: CREST’s raw event rate as a function of time (blue, left axis) along with its flight altitude (red, right axis). Major events at times shown by the large red numbers are detailed in Table 5.4.8.




Figure 5.42: CREST’s ‘No Top Or Side’ event rate as a function of time (blue, left axis) along with its flight altitude (red, right axis). Major events at times shown by the large red numbers are detailed in the text.

CREST Flight Chronology: Two Side




Figure 5.43: CREST’s ‘Two Side’ event rate as a function of time (blue, left axis) along with its flight altitude (red, right axis). Major events at times shown by the large red numbers are detailed in the text.

5.5 Summary, and A Look Back

While CREST was successfully launched and recorded nearly two days worth of full-instrument data, a variety of factors prevented us from achieving the ~ 40 days of flight as originally planned. A combination of factors, including “The Incident,” the sub-optimal location of the anti-polar cyclone and CREST’s heavier-than-was-perhaps-optimal total weight contributed to this result. None of these factors can be attributed to any one person, or even any one group of people, and must be evaluated in the context of a risky expedition having gone fairly well, all things considered.

Other than these critical factors, a series of less-critical, but still mission-threatening, factors could have been avoided, and should be avoided in future missions. For convenience and with the benefit of hindsight, I list them here, in no particular order:

- Practice last-minute preparations in the worst possible simulated conditions (e.g., while freezing and being lifted by a crane) to reveal necessary revisions and improvements before they are needed.
- Spread accumulated “instrument know-how” among those who will be responsible for operation during the mission. Our shift operation training sessions prepared but

barely those present at the time of “The Incident.”

- Defend the data acquisition system against the worst-case of the anticipated scenarios, but provide flexibility for un-anticipated failure modes. While no one can say with certainty what it would have taken to ensure CREST’s data-taking integrity while flying through a solar storm of some consequence, it is likely that some heightened degree of resilience could have been achieved.
- Design instrument components modular enough to make repairing even minor defects worthwhile, even in the days leading up to launch. Our plan of removing 1/8th of the instrument (one entire C-channel), at a minimum, resulted in lots of wasted effort, and flying with 1/64th of the instrument known to be inoperable, but not worth fixing due to the risk of damaging other components.
- Take pictures. Lots and lots of pictures.

The list above notwithstanding, I believe the CREST mission to be one marked by inspired and effective leadership. In spite of never achieving his goals of reaching the Pole or crossing the continent, Shackleton was revered as a hero, both by his comrades and by his fellows at home, “merely” for providing decisive, wise and humane leadership in the face of overwhelming adversity. While (thankfully) none of us endured hardships of comparable magnitude, at all times those in leadership positions, of CREST, of LDB and of CSBF, provided rational, effective and friendly guidance, while doing their best to effect a safe and successful mission. Therefore, they have all earned my gratitude and respect.

5.5.1 Pictures from McMurdo and LDB



Figure 5.44: We received our cold weather gear at the CSBF departure facility in Christchurch. This was my first meeting with my ‘Big Red’, the standard issue coat issued to every McMurdo resident, which would serve me well over the next several months.



Figure 5.45: The USAF C-17 which transported the initial CREST ice team (and several hundred other passengers) from Christchurch to McMurdo.



Figure 5.46: A view of McMurdo with Scott's Hut in the foreground. Between the Hut and McMurdo is Winter Quarters Bay, where when the ice is broken in summer resupply ships will anchor.



Figure 5.47: Our first glimpse of Ivan the Terra Bus, our principal mode of transport between McMurdo proper and the LDB facility.



Figure 5.48: The first elements of the CREST team arrived in McMurdo on 11/2/2011.



Figure 5.49: The exterior of 'HQ' at LDB, housing the conference room, mess hall and kitchen. LDB's social hub also includes a foosball table.



Figure 5.50: CREST PI Jim Musser stands in front of CREST as it arrives at LDB inside of its shipping container.



Figure 5.51: Lifting CREST into Hangar 096 via sliding ceiling crane.




Figure 5.52: Home, sweet home. When CREST next leaves this hangar it will be by balloon.



Figure 5.53: Professor Scott Wakely of University of Chicago pauses while examining a CREST C-channel on the ground level of Hangar 096 at the LDB facility near McMurdo Station, Antarctica. The aluminum jig used to lift the channel was designed by Alex Shroyer. This allowed us to safely and quickly remove the C-channels from the instrument for maintenance.



Figure 5.54: Nahee Park inspects CREST's thermal outer covering during a rollout test at LDB. Visible in the foreground is CREST's flight computer's umbilical (ethernet) cord, which we had to subsequently wean ourselves off of prior to flight.



Figure 5.55: A view of Mt. Erebus from the patio of Hangar 096. Mt. Erebus, approximately 30 miles away in the photo, is the southernmost known active volcano in the world. Clouds such as this unfortunate one often cling to its slopes, sometimes in stacks of three or four at a time. The puff of smoke at the summit is volcanic in origin and varies significantly in size and direction, at times a useful indicator of the wind pattern at higher altitudes.



Figure 5.56: New Zealand's Scott Base is far smaller than McMurdo population-wise but more than makes it up for it in their zealous hospitality (and with the quality of their gift shop). It lies directly between McMurdo and LDB and we passed it twice each day on our way to and from Hangar 096.



Figure 5.57: Michael Schubnell poses in front of McMurdo on our way up Observation Hill (referred to locally as Ob Hill). Fuel tanks dominate the foreground, with McMurdo's organic architectural style evident behind them.



Figure 5.58: Michael Schubnell, myself and Scott Nutter in front of Mt. Erebus from the slope of Observation Hill.



Figure 5.59: A flock of four antarctic terns near the summit of Castle Rock. Picture taken by Stéphane Coutu during our duo climb of said rock while awaiting our second launch attempt.



Figure 5.60: The trail back to McMurdo from Castle Rock. Orange flags mark the safe path; by deviating from the flags one risks falling into an unmarked cravasse, likely never to escape. Photo by S. Coutu.



Figure 5.61: Cutting ice blocks at Happy Camper training session. To survive in the Antarctic waste, construct a wall out of such blocks to block the incessantly heat-thieving wind, pitch a tent, then radio for help.



Figure 5.62: These handy ice blocks can also serve as solid foundations for kitchen appliances.



Figure 5.63: A ‘Foremost Delta Two’ transport vehicle. Supposedly the safest transportation vehicle for ice roads, its light weight and immense tires make it susceptible to getting stuck in ice potholes in warmer temperatures, or whenever the driver spins the tires too quickly (see the next Figure). Riding in the undulating toolshed hitched back of one is a rapid way to lose your lunch. If one had to sprint 100 meters to steal your buddy’s seat aboard Ivan the Terra Bus thereby dooming them to a ride aboard the Delta, one did so without regret. Overall, McMurdoans least-favorite mode of transportation vehicle, barely beating out “on foot.”



Figure 5.64: A heavily laden Delta stuck in the snow, awaiting rescuing.




Figure 5.65: CSBF's rotator, responsible for pointing CREST at the sun during flight to heroic effect. CSBF claimed it had pointing accuracy of a few degrees and precision of less than one degree. Without this device maintaining a 10-30° aspect towards the sun our batteries would have overcharged mid-flight due to higher than expected solar power production. Micro-managing this device to avoid over- or under-charging our battery banks proved to be the chief occupation of shift operators during otherwise quiet periods.



Figure 5.66: The crate used to transport the solar panels from Chicago to Indiana, from Indiana to Palestine and from Palestine to McMurdo. Once there, its foam blocks and cozy frame served mainly as a convenient resting place. Note: physicist not included.



Figure 5.67: The magnetic field sensor supplied by the University of Michigan Department of Space Science. Data taken by this device is presented in Section 5.4.8.



Figure 5.68: Before using the snowmobiles we were given snowmobile training. Snowmobile training started off slowly for some of us.



Figure 5.69: One of the two snowmobiles assigned to CREST for use during launch procedures and for 'training'. This particular vehicle left me stranded on the launch pad during a recreational jaunt after its carburetor detached during a routine jump.



Figure 5.70: Thanksgiving dinner at the main McMurdo mess hall (and probable breeding ground for the Crud super-virus). From left: Jim Musser, Stéphane Coutu, Michael Schubnell, Joseph Gennaro, Scott Wakely, Nahee Park and Scott Nutter.



Figure 5.71: Thanksgiving dessert at LDB mess hall. All the lunches there were truly extraordinary. Completely unrelated to this is the fact that all McMurdo personnel tours of LDB occur just before lunch time.




Figure 5.72: CREST's first of many roll-outs onto the balcony of Hangar 096. Here we were testing re-integration of the solar panels.



Figure 5.73: Our neighbor, STO. Due to damage suffered by their cryogenic system during shipment, their launch was delayed until after CREST's. During our launch they served nobly as filmtakers.



Figure 5.74: CREST PI Jim Musser skis out towards the launch pad at LDB. Skiing was an essential recreational component of our time at LDB, providing much-needed respite from the stress of hangar life.



Figure 5.75: It turns out the best way to smooth an ice road is to roll a very, very heavy wagon on top of it.



Figure 5.76: 'The Boss', CREST's launch vehicle and for many years the largest vehicle in Antarctica by mass.




Figure 5.77: The sun-side of the CREST instrument (facing the +X direction). The top portion of the detector is shielded by 2"-thick foam panels assembled into a hermetic box surrounding the veto system. The lower portion shows the layers of perforated silverized teflon tape covering sheet metal panels. The same tape was also used to cover all metal surfaces of and attached to the instrument chassis (e.g. the bar installed by CSBF supporting their communications array) with the exception of the solar panel mounting brackets. Both solar wings mounted on the -Y and +Y instrument sides are visible.




Figure 5.78: CSBF personnel testing attachment of the skirt solar panels (that power the SIP) and the high-tech crush pads (made of layered cardboard). The white oblong shape at center bottom is the ballast container, which can release up to several hundred pounds of glass beads during flight. Notice that the side panels are unattached; CSBF required access to the interior of the instrument to attach the skirt panels. This required the instrument to be lifted out of the hangar doors before closing it up (as the doors were not tall enough to extract CREST with the ballast container attached). This caused an enormous amount of trouble and nearly scrubbed our first launch attempt before it began. This was amended for the next two launch attempts, enabling us to button up the instrument inside rather than outside the hangar.




Figure 5.79: The Boss taking CREST out for its first hang test at LDB.




Figure 5.80: CREST and our two assigned snowmobiles during its first hang test aboard The Boss at LDB.




Figure 5.81: CREST's bottom surface. The sheet metal panels are laminated with perforated silverized teflon tape. Exposed steel boltheads were covered with aluminized mylar tape to prevent them from compromising the thermal scheme. Visible near top left is some tape that has been blown loose by the wind. The cold temperatures and persistently vigorous breeze vastly reduced the effectiveness of all of our tapes' adhesives.




Figure 5.82: Close-up of a solar panel showing the amount of damage sustained in and between Palestine and McMurdo. Despite the scars the panels performed above specifications during flight.



Figure 5.83: In order to maintain our position in the launch queue, CREST formed a vicious bicycle gang known as the CRESTaceans. From left to right: Scott Wakely, Jon Ameel, Michael Schubnell and Scott Nutter.



Figure 5.84: In this photo I am precisely at the minimum legal distance from the seal specified in the Antarctic Treaty.



Figure 5.85: CREST declared flight readiness on 12/6/2011. From left to right, top: Michael Lang, Jon Ameel, Scott Nutter, Scott Wakely, Jim Musser, Michael Schubnell. Bottom: myself, Nahee Park, Stéphane Coutu.



Figure 5.86: The CREST group has generously supported me in many ways and I am truly grateful for the opportunity to have worked with and gotten to know each of them.



Figure 5.87: A Kress tractor hauling a massive load, stuck in a pothole on the ice road to the Pegasus airfield. Yours truly formed part of the rescue team. For those of you asking which vehicle displaced The Boss as the heaviest on the continent, this is it.



Figure 5.88: The rescue vehicle for the heaviest vehicle on the continent: another Kress tractor. Jules, fearless leader of the ice road maintenance team, coordinates a rescue effort via radio. This was the first ever rescue of a fully laden Kress truck on the continent.



Figure 5.89: Cool physicists don't look at explosions...



Figure 5.90: This seal had nothing whatsoever to do with the scientific aims of the CREST expedition. Or did it?




Figure 5.91: Celebrating declaration of flight readiness with an improvised musical instrument. From left: Jon Ameel, Stéphane Coutu, myself and Nahee Park.




Figure 5.92: The CREST Ground Station inside Hangar 096. Front left: GSE3. Facing away: GSE4. Jon Ameel is visible working at his engineer's station near the rear of the hangar.




Figure 5.93: Sheet metal panel showing cutout region for power cables that will remain connected to CREST until moments before launch to ensure full battery charge during ascent. As the instrument will likely be spinning, pointing (and thus reliable solar power) will be impossible until it stabilizes. The region was to be blocked by a foam block covered in aluminized mylar after disconnecting those cables, secured by tape. It is doubtful that said block remained in place given that we were typically unable to adhere tape securely outside the hangar (not to mention at flight altitude).



Figure 5.94: The foam block used to cover the power cable access did in fact shrink during flight. Photo by J. Ameel (during recovery).




Figure 5.95: CRESTaceans plot their next move. From top left, clockwise: Scott Nutter, Nahee Park, Jim Musser, Stéphane Coutu.



Figure 5.96: ‘*Standing on a hill in my mountain of dreams, telling myself it’s not as hard, hard, hard as it seems*’ (Led Zeppelin, *Going to California*). Michael Schubnell during our approach on foot to Castle Rock. The foggy weather precluded an ascent on that day.



Figure 5.97: Ivan the Terra Bus waiting to take us back to McMurdo from LDB. Though comfortable the ride was also very slow, taking approximately 50 minutes to travel between LDB and McMurdo (vans did the same trip in approximately 10 minutes). While climbing the switchback up the hill past Scott Base the bus travels more slowly than a jogging CSBF engineer.




Figure 5.98: View of LDB from the launch pad. All non-tent buildings are deployed on skis and moved to higher ground every winter. Otherwise, non-mobile buildings eventually sink below the ice surface.



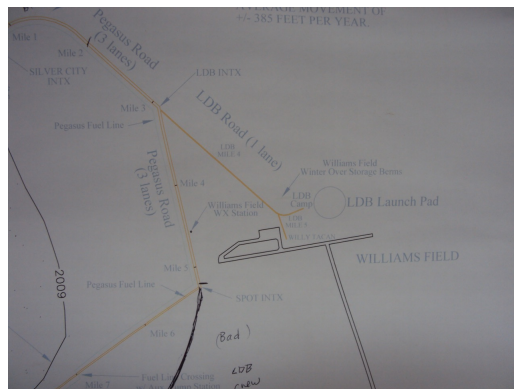
Figure 5.99: Inside a Delta transport desperately trying to avoid losing today's lunch. It is customary for experiments, unions and military outfits to paste their sticker all over the inside of the vehicle, lending it an decor eerily reminiscent to that of a men's bathroom. CREST had no sticker.




Figure 5.100: View of the dorm buildings in McMurdo where we spent the ‘nights’. We were two to a room. Bathrooms were shared and we were encouraged not to shower every day to save clean water. Each building also had laundry facilities (each room was assigned one laundry day per week) and common living areas with seating, board games and televisions.




(a)



(b)



(c)



(d)

Figure 5.101: (a): Detailed map of the road system for McMurdo and the surrounding ice shelf. The only solid land here is colored in green. The distance from McMurdo to LDB is approximately three miles following the road. We arrived on the Ice Runway (top left) and departed months later from the Pegasus Runway (bottom left). An arrow near the top middle states that the ice shelf flows to the left at a rate of approximately 385 feet per year. Closer to Pegasus Runway the rate is only 95 feet/year. (b): Closer view of the LDB surroundings, which abuts against Williams Field (called Willy Field). Clearly visible is the circular launch pad, around which we skied many times. (c): Closer view of McMurdo, Hut Point and the Ice Runway. (d): The edge of the permanent ice shelf lies just beyond the main Ice Road between LDB and Pegasus Runway. It was on this stretch that we rescued the stranded Kress truck.



Figure 5.102: IU Electrical Engineer Michael Lang. All of the tools transported by Alex Shroyer and myself to Palestine were then shipped to McMurdo, enabling significant work to be done on CREST in Hangar 096 prior to flight readiness.



Figure 5.103: An early prototype for the GoPro mobile camera, used by myself to make videos of launch procedures, ski trips, snowmobile rides and kite flying.




Figure 5.104: During our first launch attempt on 12/18/2011, CSBF's skirt and crush pad attachment procedure called for us to button up the instruments bottom thermal panels outside the hangar. Temps of below -10°F and a bitterly frigid wind made this an incredibly difficult task. Without mittens in addition to thinner gloves my hands froze; with mittens I was unable to screw anything in. Furthermore, the tension induced in the chassis by the crane lift misaligned some of the screw holes, making the task even more difficult. Consequently a procedure which normally took a few minutes inside took nearly two hours.



Figure 5.105: CREST, finally buttoned up for its first launch attempt.




Figure 5.106: CREST, finally buttoned up for its first launch attempt.



Figure 5.107: One example of the *fata morgana* (Morgana's Curse) mirage, caused by layers of air with differing temperature near the ice surface.



Figure 5.108: A slightly more striking example of the *fata morgana* mirage.



Figure 5.109: The foam block covering the power cable pass-through hole during our first launch attempt.




Figure 5.110: Stéphane Coutu attempts to activate Nashua duct tape's adhesive with a heat gun while I stand ready to catch him should the ladder tip over. The low temperatures during our first launch attempt prevented our silverized teflon tape's adhesive from adhering. At first we attempted to tape over it with packing tape, but it also would not stick. Neither would Nashua duct tape (which CSBF staff claimed had never happened to them before). I hit upon the idea of blasting the tape with a heat gun in order to activate the adhesive. This worked but was very time consuming. Pressed for time CSBF assisted us in securing the foam panels by tying ripcord around the instrument.




Figure 5.111: M. Lang, N. Park, J. Musser and J. Ameel (ladder) attaching the -X, +Y side panel in the bitter cold.



Figure 5.112: Our bench power supply, “Maverick,” braved the elements along with us on the failed first launch attempt. Not pictured is the previous power supply, “Viper,” which powered CREST at Indiana and in Plum Brook. No “Top Gun” references were harmed during the creation of this instrument.




Figure 5.113: One of CREST's two Nickel Metal Hydride battery banks. The wires marked 'CUT' are for the benefit of recovery personnel.



Figure 5.114: Jim Musser, Nahee Park and Myself after replacing one of CREST's two battery banks, which had been overcharged by a faulty charge controller.




Figure 5.115: The power cables which would remain attached to CREST until moments before launch. Michael Lang bravely performed this task, thereafter transported away from the launch pad by a snowmobile driven by CREST's most elite snowmobile pilot.



Figure 5.116: To assuage our pain at the failure of our first launch attempt we fled to the comfort of the neighborhood wine house.



Figure 5.117: View of Mt. Erebus from atop Castle Rock.



Figure 5.118: View northwest from midway up Castle Rock.



Figure 5.119: The same view northwest from atop Castle Rock.



Figure 5.120: View south from Castle Rock looking over the ice shelf. The tiny, barely visible buildings are LDB.



Figure 5.121: A Skua. These hardy residents of the icy coast feed on anything to survive, including each other's young and the contents of cafeteria trays of unwitting McMurdoans. Protected by the Antarctic Treaty and free from retribution, they grow bolder by the day.



Figure 5.122: Life in McMurdo was not without its comforts. Jon Ameel, Jim Musser, myself and Nahee Park at the Café, one of three establishments at McMurdo serving alcoholic beverages. I cannot recall the photographer; judging by the rapidly emptied wine glass, neither can he/she.




Figure 5.123: CREST being picked up by The Boss the morning of our third and final launch attempt. It would not touch ground again until the end of its flight.



Figure 5.124: Mt. Erebus' plume the morning of our third and final launch attempt.



Figure 5.125: Overcome by happiness at the prospect of our imminent launch, Jon Ameel flies a kite.




Figure 5.126: A massive spool unwinds the tether that will connect CREST to its host balloon.



(a)



(b)



(c)



(d)

Figure 5.127: Prior to launch dozens of CSBF engineers and support staff expertly prepared the balloon and tether. (a): Helium tanks that will inflate CREST's host balloon. (b): A view down the flightline: preparing CREST's tether. Visible in the distance is The Boss with CREST hanging behind. (c): Moving down the flightline towards The Boss. (d): The rotator and ring harness.




Figure 5.128: CREST hanging from The Boss minutes before launch. The power cord linking the Maverick generator-fed power supply aboard The Boss snakes in from the lower left and would need to be detached shortly.




Figure 5.129: CREST hangs majestically from The Boss minutes before launch.



Figure 5.130: Erebus' plume shows evidence of calm winds minutes before launch. CSBF used an actual weather balloon to determine this, but I prefer to use the sulphurous emission of a volcano.



Figure 5.131: Initiating helium fill of the balloon.



Figure 5.132: The balloon begins to inflate.




Figure 5.133: The balloon continues to inflate.



Figure 5.134: The balloon continues to inflate via the fill tubes (white wisps coming out either side of the inflated portion of the balloon). The loud ‘whoosh’ of rushing helium dominates the scene. At left, CREST’s orange and white parachute lies in wait.



Figure 5.135: Shortly before launch.




Figure 5.136: A minute or so after launch one could barely make out the solar wings.



Figure 5.137: Several minutes after launch all that was visible of CREST to the naked eye was periodic flashes of reflected sunlight as CREST spun every few seconds.



Figure 5.138: Though it was hard to say goodbye, I would much prefer CREST continue to rise than the alternative.




Figure 5.139: The remaining ice team celebrates with champagne in LDB-issued spill-proof coffee cups. A well-deserved celebration.



Figure 5.140: A skua, likely fleeing the scene of an unimaginably horrible crime. Even they appear majestic from the right angle.



Figure 5.141: Bessler aircraft shrouded by fog at the Pegasus airfield. The fog cancelled inbound and outbound flights that day.



Figure 5.142: While awaiting the recovery mission, Jon Ameel encountered an amiable LDB visitor, separated from his tour group and asking for directions to the kitchen tent.



(a)



(b)



(c)



(d)

Figure 5.143: Pictures of and from aboard the LC-130 emergency medevac aircraft that also transported Stéphane Coutu and myself from McMurdo to Christchurch fuels up while awaiting takeoff at the Pegasus Runway on 12/28/2011. (b): Three phases of water: fragmented ice sheets float on a blue ocean beneath the clouds. (c): The melting, fragmented edge of the Ross Ice Shelf. (d): Rugged splinters of the Transantarctic Mountain Range.




Figure 5.144: The shifting snows digest CREST as it awaits recovery. The +Y side solar wing (far side) was martyred by the falling rotator attachment - ironic given that the rotator moderated the solar panels' power production all flight long by pointing them up to 30° off sun. Photo by J. Ameel.






Figure 5.145: The rotator crashed right through the +Y-side solar wing. We were fortunate it did not crush the crystal system, the batteries or the SFC. The +Y-side solar wing was not. Photo by J. Ameel.



(a)



(b)

Figure 5.146: The SFC during recovery. Disk 2 (the one that stopped working) is second from the front. Photos by J. Ameal.




Figure 5.147: The recovery team finishes picking CREST's bones clean. Only a battery bank, the SIP and some odds and ends remain. Photo by J. Ameal.



Figure 5.148: CREST awaits recovery in a desert of windblown snow, casting a shadow in the downwind ice pattern. Photo by J. Ameal.

CHAPTER 6

Data Analysis

In this chapter I describe my original analysis of CREST’s data, focusing on diversifying the classes of events contributing to our signal population and extracting important additional information about the primary electron’s trajectory as compared to the collaboration’s previous, and ongoing, analysis efforts. I first describe the anatomy of signal events, in order to familiarize the reader with the concepts utilized in my analysis techniques. I also describe the data pipeline, starting from the files generated aboard the Science Flight Computer (SFC) to event-based analysis and display.¹

In what follows, I make the case that confining one’s desired signal event population to those events with no veto activity and only the simplest crystal morphology results in a critically reduced potential signal set. While the main background to this reduced signal set, side-going protons, can be ameliorated through use of various analysis techniques, including a crystallographic cut of my own design, the surviving flight events appear to remain background dominated. In particular, demonstrating that CREST can replicate the results of multiple other experiments at lower energies would go a long way towards proving that CREST’s technique is practically realizable, even with a less-than-ideal flight profile. However, as the primary energy decreases, the fraction of “clean” events also falls, chiefly because the electron is far more likely to pass through the detector than at higher energies.

Accepting events with veto activity and more complicated crystal morphology requires additional analysis techniques. My techniques extend the mainstream analysis, which focus on identifying and characterizing the “main event line,” (defined below), primarily by making use of the other primary direction present in signal events: the momentum of the primary electron.

During CREST’s development it was assumed the primary electron’s momentum direction would be a hidden parameter during event analysis. As my analysis progressed, I real-

¹My efforts utilize the groundwork laid by Jim Musser, Scott Wakely and Nahee Park, who jointly developed CREST’s data file structure and guided my early efforts in learning to use ROOT in combination with C++ as an effective coding framework and analysis toolkit.

ized that not only would the electron’s momentum predictably alter the event signature in the bottom veto system, but it could also be estimated directly by comparing the combined crystal and veto responses. Events which were previously considered non-separable from background and ignored during simulation and early analysis efforts (“traversal” events) became the focus of my analysis efforts. The signature of the electron passing through (or very near) the detector could also unlock another event parameter thought to be hidden: the direction of the geo-magnetic force on the electron. This last piece of the signal event puzzle is that needed to determine the charge sign of the primary lepton, making possible a measurement of the positron fraction at arbitrarily high energies.

While much work remains to be done to positively demonstrate this technique’s viability using CREST’s flight data, such demonstration would provide the first (and as far as I am aware, the only) means of making such a measurement at energies of 1 TeV and above. As I reviewed in Chapter 1, such a measurement would have profound implications on our understanding of our region of the Milky Way, cosmic ray production and propagation and possibly the nature of the Dark Matter theorized to permeate our Galaxy, as well as all others. Such a development, combined with other detector improvements designed at greatly reducing the main background source, would provide a strong justification for additional missions to follow in CREST’s pathfinding footsteps.

I would like to make an important point with regards to the simulated background. Our Monte Carlo efforts told us exactly what we wanted to know: signal events have a characteristic temporal and spatial signature, which could occasionally be replicated by chance coincidence of the γ -ray background. The other main anticipated background source was the charged particle background, which is dominated by protons. The large flux of protons at the relevant energies, and CREST’s response to their passage, was exceptionally well-modeled. However, the simulations did not, and could not, call attention to how dominant side-going protons would be *in the final stage of analysis*. This was caused by CREST’s extremely high sensitivity to side-going protons in the tail of the angular distribution (i.e., those trajectories with very horizontal zenith angles). The events comprising the tail of this tail, that is, those protons which both evade detection of the veto system *and*, for whatever reason, have poorly estimated inverse track velocities, still dominate the low flux of signal events CREST hoped to identify.

6.1 Signal Event Morphology

Before describing analysis techniques or modules, I first wish to make the reader familiar with the basic elements of signal events as predicted by theoretical exercises such as that in

Section 2.1 and by Monte Carlo signal simulations.

6.1.1 Anatomy of a Signal Event

Figure 6.1 shows an idealized picture of a signal event. The electron enters the geo-magnetic field from above and begins executing its curved trajectory within its plane of motion. Before any photon emission or other interactions, two angles are already set. First, the polar angle θ measures how highly inclined the electron's plane of motion is with respect to the detector plane. When $\theta = \pi$ the planes are perpendicular. Second, the azimuthal angle ϕ measures the alignment of the intersection line within the detector plane. I measure this angle such that when $\phi = 0$ the intersection line is parallel with the detector's x axis.

The electron's extended trajectory illuminates the detector plane along the intersection line between the two planes; any synchrotron or Bremsstrahlung photons emitted by the electron will intersect the detector on, or very near, this line. In my analysis I refer to this line as the "main event line." Estimating the main event line is an integral step in dissecting signal events. Accurate reconstruction of largely all other event parameters relies on a proper estimate for ϕ , the azimuthal orientation of the main event line.

The main event line is a line, and not a main event point, because of the deflection in the electron's momentum resulting from interaction with the geo-magnetic field. I describe the vectors involved with the aid of the diagram in Figure 6.2. There, the electron begins illuminating the detector at point P_1 when it has momentum \vec{p}_1 , then stops illuminating the detector at point P_2 with momentum \vec{p}_2 . Between these extremes it experiences a total change in momentum $\vec{\Delta}_p \equiv \vec{p}_2 - \vec{p}_1$. This change in momentum comes from the Lorentz force $\vec{F}_B \propto q\vec{p} \times \vec{B}$. Therefore

$$\vec{\Delta}_p \parallel -\vec{p} \times \vec{B}. \quad (6.1)$$

(Note that in the case of a primary positron rather than electron, the "-" sign would be absent and the force direction would be reversed.) The orientation of the main event line in the detector plane (as measured by the azimuthal angle ϕ) is equal to the projection of $\vec{\Delta}_p$ on the detector plane. This equation is valid to the extent that $\vec{p} \times \vec{B}$ is constant over the path traversed by the electron between when it is directed at the points P_1 and P_2 (that is, between the times when the electron starts and stops illuminating the detector). Since the electron is so far away from CREST and the gyroradius is so large (of order one million km) it is reasonable to treat the change in Lorentz force direction due to the variation in electron momentum direction as indistinguishable from zero. The variation in \vec{B} , on the other hand, may not be small enough to ignore, since the illuminating path length ranges from tens to




Figure 6.1: Geometry archetype for signal events. A primary electron’s trajectory illuminates the CREST detector plane. The plane containing the electron’s motion intersects the detector plane on a line herein referred to as the “main event line.” The angle between these two planes θ is equal to the electron’s zenith angle against the horizon when CREST’s detector plane is parallel to the ground. I refer to this angle as the polar angle. To characterize the trajectory with respect to detector axes a second angle, ϕ is required, which characterizes rotation of the electron’s plane of motion around the detector’s z axis. I refer to this angle as the azimuthal angle. My analysis efforts revolve around reconstructing these two angles. Note also that a positron with the same initial momentum would behave identically to the electron with the exception that it would be pushed in the opposite direction by the geo-magnetic force.




Figure 6.2: A negative electron illuminates CREST while interacting with the perpendicular component of the geo-magnetic field, B_{geo} which points in the $-\hat{z}$ direction (this is close to the actual orientation for the initial segment of CREST's flight). The resulting Lorentz force on the negatively charged electron is directed as shown by \vec{F}_B . When illumination starts, the electron has initial momentum \vec{p}_1 and its extended trajectory intersects the CREST detector plane at point P_1 . When illumination stops, the electron has final momentum \vec{p}_2 and its extended trajectory intersects CREST at point P_2 . θ measures the polar angle between detector horizontal and \vec{p}_1 . Synchrotron and Bremsstrahlung photons emitted by the electron will have nearly-parallel trajectories such as those marked by gray arrows; the spacing between photon emissions is random. These photon trajectories will intersect CREST's detector plane somewhere along \vec{L} , the main event line, marked by the teal line between points P_1 and P_2 . Note that a positron would have an identical diagram, except that the directions of \vec{F}_B and \vec{L} would be reversed, and the roles of initial and final momentum and intersection points would be exchanged.

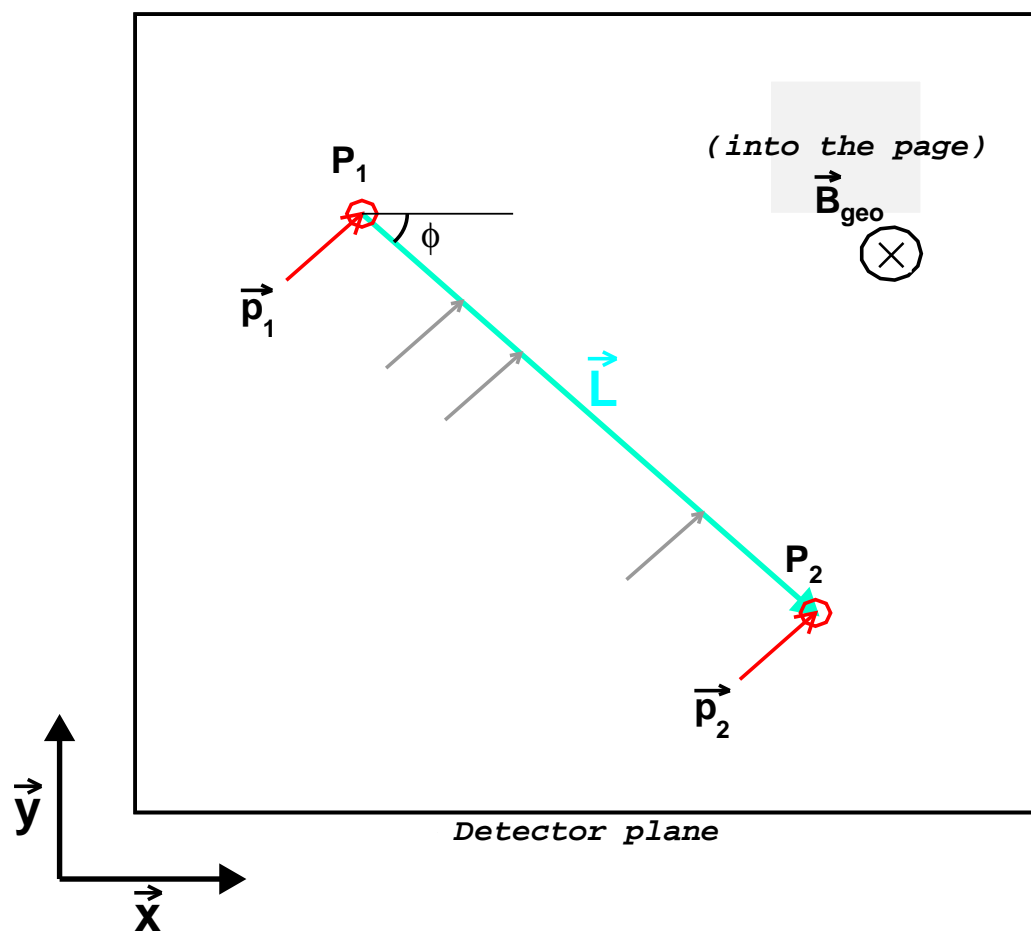


Figure 6.3: A top-down view (we are looking in the $-\hat{z}$ direction) of the same event as in Figure 6.2. The initial and final electron momentum vectors, \vec{p}_1 and \vec{p}_2 are foreshortened due to our perspective. Also foreshortened are the trajectories of photons emitted by the primary electron (gray arrows). In the projection of the detector (\vec{x}, \vec{y}) plane, the electron momentum is perpendicular to the main event line \vec{L} when the magnetic field \vec{B}_{geo} points in the $-\hat{z}$ direction.

Cut name	Min value	Max value
# Crystal hits	4	9
# Crystal clusters	5	–
# Top veto hits	–	0
Event length	75 cm	–
Largest gap	30 cm	–
2 nd Largest gap	20 cm	–
Linear fit reduced χ^2	–	6

Table 6.1: Cut set devised by Jim Musser that I refer to as the “mainstream analysis.” I discuss the definition of event length and a procedure for calculating the reduced χ^2 goodness-of-fit variable in Section 6.4.5.2. The “largest gap” and “second largest gap” are parameters generated via the crystallographic analysis technique I devised, which I present in Section 6.4.5.3.

hundreds of kilometers. Over such distances, the magnetic field may change appreciably in both magnitude and direction, especially in locations where the field is changing rapidly. In such a case, the main event line is close to, but not quite, a straight line.

6.1.2 Identifying Signal Events: Two Key Features

Here I focus on two features unique to electron signal events that enable analysis algorithms to distinguish them from charged particle background events. The first is the simultaneous arrival of the Bremsstrahlung and synchrotron photons (and the primary electron itself, if it traverses the detector). I refer to any cut making use of this fact as a “simultaneity cut.” A statistic (referred to as inverse track velocity) measuring this simultaneity features prominently in the mainstream analysis. However, accurate determination of hit simultaneity is difficult when presented with hits which can be either primary or secondary in nature. Secondary photons (from air showering and from secondary scattering within the detector plane) are numerous and not simultaneous. As shown in Figure 6.17 many signal events are not properly recognized as being simultaneous by my analysis algorithms. This is because determining exactly *which* hits are supposed to be simultaneous is very difficult. The mainstream analysis attempts to circumvent this issue by selecting only the simplest signal events with no shower products, few or no Compton scatters and no top veto activity. Unfortunately, this draconian approach results in a signal acceptance of approximately zero. Figure 6.4 shows one important reason why: the vast majority of signal events activate the veto system with high multiplicity. The “no top veto activity” cut alone eliminates over 90% of simulated 10 TeV signal events when including the interaction of the primary electron with the detector.

A large part of my analysis efforts have focused on augmenting the mainstream anal-




Figure 6.4: Normalized veto activity as measured by the number of paddles with pairs (hits in PMTs on both ends) in simulated 10 TeV signal events, by subsystem. Cutting out events with any top veto activity eliminates over 90% of these events. The bottom veto has high average multiplicity which suggests using it as part of the reconstruction process.




Figure 6.5: Plot by Jim Musser of the inverse track velocity distribution for flight data events. I discuss the meaning of and a procedure for calculating the inverse track velocity in Section 6.4.5.2. The cuts applied to generate this plot are listed in Table 6.1. Out of the entire flight, a few hundred events remain in the signal region of this plot (near zero inverse track velocity). Whether this represents an excess over the amount of background events expected to bleed into this region via random coincidence or event mis-characterization will depend on the results of additional background Monte Carlo event production (which, to my knowledge, is still in progress at the time of this writing).

ysis by also making use of a second feature unique to high energy electron signal events: *the electron's change in momentum is perpendicular to its motion*. This property is a consequence of the Lorentz force. Heavier particles such as protons will have no detectable sideways deflection from the magnetic field due to their lower charge-to-mass ratio.

In valid signal events the main event line estimates the direction of momentum change of the primary electron. If we were able to measure the electron's momentum direction, we could compare the two vectors and require

$$\hat{p} \times \hat{\Delta}_p \neq 0 \quad (6.2)$$

to further characterize signal events beyond the simultaneity cut.

A relativistic side-going proton does not share this feature. Its change in momentum is nearly zero, and what little change there is is anti-parallel to its momentum; that is, it simply slows down unappreciably as it passes through the crystal array. The main event line resulting from a side-going proton's transit through the crystal array is parallel to the momentum of its secondary products. There is only one vector at work here, as opposed to the two perpendicular vectors present in electron signal events.

6.1.3 Estimating Electron Momentum Direction

It is not immediately obvious how to estimate \hat{p} , or even whether it can be estimated in signal events. Here I outline the basic idea behind my process for doing so. I present the full technique in Section 6.7.

As discussed in Chapter 2, Bremsstrahlung and synchrotron photons emitted by the primary electron have average momentum direction parallel to the electron's instantaneous momentum. Although the electron's momentum direction is different at the end of the main event line compared to the beginning, the *angular* difference is tiny (\sim fraction of a degree) - far smaller than CREST's precision for determining momentum direction (\sim tens of degrees). Thus I treat all primary photons as having essentially parallel momentum. If I can estimate primary photons' momenta directions, I can estimate the primary electron's momentum direction.

The secondary products of Compton scattering events (both the recoil photon and electron) are emitted preferentially in the forward direction at high incident photon energies (as pointed out in Chapter 4). Since CREST's crystals are radiatively thin when traversed vertically, most of the electromagnetic cascade products will pass through the crystal system. The secondary products resulting from the primary photons' Compton interactions or pair production within the crystal plane will, on average, retain the momentum direction of

the primary photons. My technique for estimating the primary photons' momentum relies on this property and the high likelihood that these secondary products, especially recoil Compton electrons, will produce signals in the veto system as they make their way out of the lower portion of the detector. If I can estimate secondary photon and electron momentum direction, I can estimate primary photon momentum direction, which in turn allows me to estimate primary electron momentum direction. To summarize:

- An event produced by a side-going charged particle has only **one** vector: that of the proton momentum;
- an event produced by a curving electron has **two perpendicular** vectors: that of the electron momentum, and that of the change in electron momentum;
- the main event line provides an estimate for the direction of one of the vectors (Δ_p); and
- the activity in the bottom veto compared to the activity in the crystal system provides an estimate for the direction of the other vector (\hat{p}).

In what follows I present my work in estimating \hat{p} (Section 6.7) and exploring the use of Equation 6.2 in concert with the simultaneity requirement in an attempt to characterize signal events with greater acceptance than achieved through the mainstream analysis.

Before I continue discussing event analysis, I present my work in creating a tool to visualize the three-dimensional event and analysis morphology. By presenting pictures of events and their analysis using this tool, I strive to clarify the above and motivate my choice of analysis strategies, parameters and cuts.

6.2 *EventDisplay*: Events in Three Dimensions

Between the instrument operation training I ran following our integration test in Palestine and our departure for McMurdo (that is, in September and October 2011), I became intrigued by the idea of utilizing the veto system as a means to trace the paths of charged particles as they entered the instrument through a veto paddle, interacted with the crystal system in some way, then exited out another veto paddle. The collaboration to this point had generally considered any interaction with the veto system as grounds for immediate disqualification of the event as a potential signal event. As such, no tool had yet been created to visualize this kind of analysis. I sought a practical way to display the combined crystal and veto information in three dimensions and found a suitable solution in the ROOT libraries.

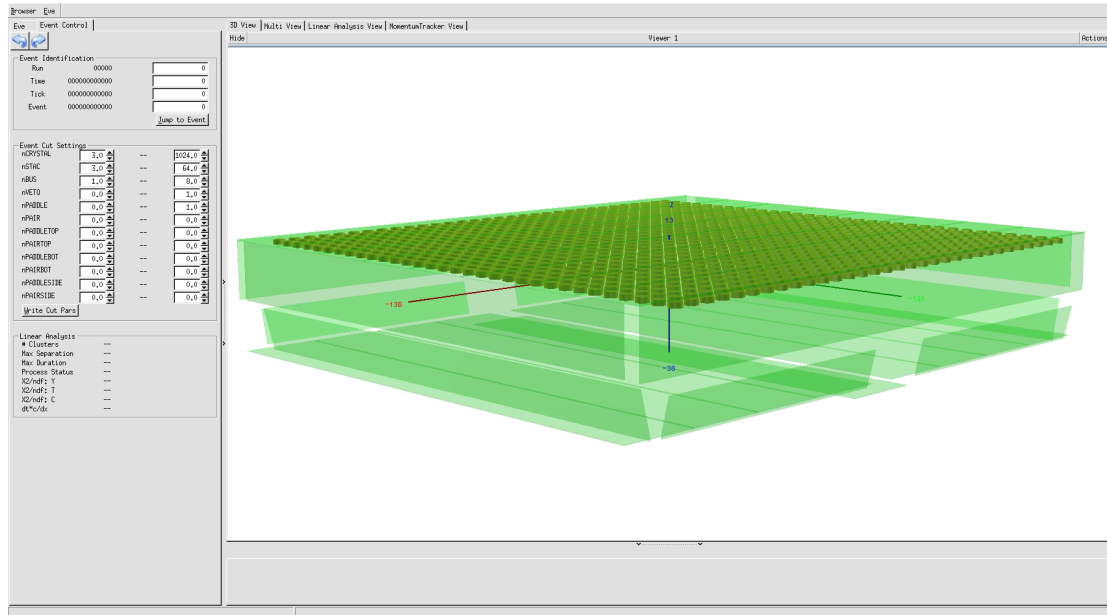


Figure 6.6: CREST as visualized in “EventDisplay,” the three-dimensional display tool I developed to help visualize simulated and flight events. The veto paddles are represented by transparent green rectangular solids. Beneath the top veto lies the crystal plane, each crystal depicted as a solid yellow disk. The user can interact with the display by rotating, zooming and panning, as well as inspecting instrument sub-systems (e.g., the bottom veto system) or individual detector components (e.g., bottom veto paddle 3). On the left is the control pane, allowing the user to navigate between events, view event metadata (upper-left) and view and/or set event parameters (middle left). Clicking on a detector module in the display window allows further interaction with it in the lower left (e.g. to view hit information or change its appearance).




Figure 6.7: Another view from EventDisplay making clear the location of paddles belonging to various veto subsystems, delineated by color. For visibility the facing two side panels and several top panels have been made invisible. The eight top paddles lying immediately above the crystals with long dimension along the X direction are colored blue. The four side paddles are green (the facing two are invisible here). The nine bottom paddles running along the Y dimension are colored red. One of the bottom paddles is slightly offset below the rest; this is to accommodate a structural beam. The six smaller, inside slanted vetos are colored magenta. The same beam that displaces the center bottom paddle causes a significant gap between four of the inside slanted paddles. Particles that sneak through these gaps and interact with the crystals have a relatively large zenith angle, and can not make tracks through the crystals longer than a few tens of cm (and can therefore be eliminated in detailed event analysis).

Originally developed for the ALICE collaboration, ROOT's TEve libraries allow one to input a three-dimensional detector model, define viewing ports with various projection axes and provide "on-click" information panels for each detector module specified in the three dimensional model. By modifying this code and inputting the spatial detector model developed by Scott Nutter for his Monte Carlo simulations, I was able to create *EventDisplay*. Figure 6.6 presents the CREST detector model as shown by *EventDisplay*. The geometrical placement of the veto paddles (transparent green rectangular solids) around the crystals (small yellow disks) is readily visible. Detector elements can be drawn with varying coloring or transparencies to indicate signal properties (e.g., TDC and ADC values). Estimated hit locations within veto paddles are endowed with special markers (see Section 6.6.1 for discussion of the technique I used to locate hits along veto paddles). Various options are given for visualizing crystal energy deposition (e.g. by disk coloring), main line estimation and charged particle trajectory reconstruction. The viewport provides an interactive three-dimensional model of the detector and the user is free to zoom, rotate and pan to gain a complete understanding of the geometrical arrangement of the detector and individual events as he/she sees fit.

This display provides an easy way visualize the various veto subsystems. For example, the placement of the inside slanted veto paddles deep within the structure of the instrument is readily apparent in Figure 6.7. The distortion of the geometric regularity of the veto system by the gondola's central support beam is also apparent; it necessitated the addition of a ninth bottom veto paddle offset vertically from the rest. This beam also cleaves two of the inside slanted paddles in half, thus requiring six instead of four total inside slanted paddles.

Figure 6.8 presents a flight event displayed in *EventDisplay*. I added the heavy teal curve after the fact to demonstrate an exaggerated potential reconstruction of the electron trajectory. The activated veto paddles are colored purple; this event features one top and one bottom veto paddle. The activated crystals are shaded from white to red with increasing energy. Small energies (less than one MeV) are colored white. The red crystals represent crystals with signal energies above 10 MeV. The purple crosses above the crystals represent the hit time relative to the earliest hit.

All event pictures in the following Sections were created with *EventDisplay*. As analysis progressed, I integrated additional functionality beyond what I've shown thus far, especially including displaying truth data from Monte Carlo simulation runs.




Figure 6.8: A template “traversal”-type event in Event Display. The primary electron follows the (very exaggerated) curved teal trajectory, crossing detector altitude inside the veto box and crystal plane (qualifying this as a traversal event). The thin teal line within the crystal plane is the intersection of the electron’s extended trajectory with the crystal plane. The magenta-colored veto paddles have been activated on both ends and thus have a pair. The relative time difference between these matched-end PMT signals provides an estimate of the location along the paddle where the hit occurred; these locations are marked by red elliptical cones.

File Type	Prefix	Typical Size (MB)
Main data	Run	2000
Pulser data	PulserRun	60
Pedestal data	PedRun	20
Golden data	GoldenRun	~ 1

Table 6.2: Types of binary run files written to the flight data disks by the Science Flight Computer, their filename prefix and their typical size. An underscore and a five-digit number were appended to each file to generate the full name. E.g. the pedestal file for run 03482 would be “PedRun_03482.dat” (if a pedestal run had been executed for run 03482). “Golden” data refers to data with more than 3 crystal STACs activated and no veto activity.

6.3 The Raw Data

During flight “runs” were written to one of four solid state disks by the Science Flight Computer. Each run was labeled with a five-digit number and consisted of up to four separate binary data files described in Table 6.2. Besides the main data run, each hour *crest.fs* automatically executed Pulser and Pedestal runs to enable run-dependent time and energy calibration. *crest.fs* also had the capability to crudely cull events before writing to disk - this was used to create “Golden” runs consisting of freeze cycles with more than 3 crystal STACs activated and no veto activity. Since the majority of my analysis hinges crucially on the signals in the veto system, I refrained from using these runs.

Following recovery, Jon Ameel and Matt Geske brought the flight data disks back to Hangar 096 at LDB where they made several copies. Matt Geske transported the original disks and copies to Penn State University, made additional copies and shipped disks containing copies of all flight data to each collaborating institution, including University of Michigan. Michael Schubnell stored copies of all flight data files on our local server; when I refer to flight data, I refer to the files in these copies.

6.3.1 Note on Monte Carlo Simulations

I draw from two data sets in addition to the flight data, namely signal and background Monte Carlo data sets. Both of these simulations were developed by Scott Nutter and I will not discuss their development in detail here, except for certain key aspects that greatly affect my analysis techniques.

The background and signal Monte Carlo simulations have two initial stages, both relying on Geant4 to model physics interactions. The first stage generates “primary” photons and charged particles, then propagates them through realistic models of the atmosphere and geo-magnetic field, generating secondary products via Geant4’s stochastic physics model. The second stage models the interaction of these primaries and secondaries with a three-

dimensional detector model. The output of this stage is a list of times and energy depositions in each activated detector element. A final stage then stochastically models the readout system's time and energy measurements in response to these energy depositions. Finally it generates a simulated list of hits within the detector, allowing the resulting files to enter the analysis pipeline, with the addition of the Monte Carlo truth data.

To generate signal Monte Carlo data sets one first uses Scott's CRESTMag tool to propagate electrons starting from a fixed point above the Earth. The electrons traverse models of the Earth's atmosphere and geo-magnetic field (the latter uses the 2010 version of the World Magnetic Model, WMM), and produce the attendant "primary" synchrotron and Bremsstrahlung photons. After setting the atmospheric overburden (thereby choosing detector altitude), CRESTMag scans a spherical shell in a grid pattern, searching for hypothetical detector locations that would intersect the trajectories of three or more primary photons. Thus, one firing of a simulated electron can result in any number of "events," including zero, one or more than one. Grid positions that generate events may or may not intersect with the trajectory of the electron itself. Those that do intersect with the electron's trajectory are, by definition, traversal events.

For the vast majority of the collaboration's work with these simulations, the interaction of the electron itself did not contribute to selecting grid points as valid events. For example, if an electron generated two synchrotron photons at a certain grid point, and the electron itself passed through that grid point, the algorithm would reject it as an event because only two photons intersect that grid point (while three would be required to activate the instrument trigger). While grid points with three photons in addition to the electron weren't *rejected* on the basis of electron traversal, ignoring the electron's potential interaction with the detector is, *a priori*, a serious error that could distort our signal event analysis and sensitivity calculations.

Following my work showing that a substantial fraction of signal photons come from events where the primary electron traverses the detector plane (see Section 2.1.4), Scott modified his second stage simulation module, CRESTDet, to take into account the passage of the electron itself through the detector. I used this modified code to generate the signal Monte Carlo files used in the analysis presented herein. Figure 6.9 presents the fraction of events in these runs in which the primary electron traversed the detector as a function of primary electron energy. At 10 TeV the fraction is 59%, very close to what I estimated using Professor Akerlof's model in Section 2.1.4. At lower energies the fraction increases rapidly, until at 500 GeV, over 80% of events are traversals. Figure 6.12 shows that a substantial fraction of detector activity is due to secondary hits from the primary electron interacting with the detector. This plot was generated using all signal events, not only




Figure 6.9: Fraction of simulated signal events in which the primary electron passes through CREST’s veto system. As I predicted in Section 2.1, the fraction is approximately 60% for 10 TeV primary electrons, and rapidly increases as primary energy decreases, until at 500 GeV the fraction is over 80%.

traversal events, showing that even in non-traversal events, the interaction of the electron with the detector should be taken into account.

6.4 Analysis Software Chain

I developed a C++ and ROOT based analysis software chain to perform calibration, data reduction, event formation, event display and event analysis on the root files produced by RawReader. This suite of tools consists of over 20,000 lines of C++ code and is stored in a Subversion repository called “Crestweware.” Certain components of Crestweware, especially the base data structures (representing freeze cycles and events), are closely modeled after components of the collaboration’s main software repository, “Crestware.” Table 6.4 summarizes the flow of signal data from inception in *crest_fs* (the flight software process running on the SFC) to databases of C++ objects of increasing levels of complexity. I describe the input, action and output of each of these modules here.

Tool	Description	Key Output
CREST	Balloon telescope	Binary freeze cycles
<i>RawReader</i> (by NP)	Binary converter	C++ freeze cycles
<i>Calibrator</i>	Time and energy calibration	Calibrated freeze cycles
<i>EventFinder</i>	Break FCs into events	Calibrated events
<i>EventCutter</i>	Data reduction	Cut events
<i>Lineal</i>	Linear crystal analysis	Main event line
<i>Geometizer</i>	Geometric veto analysis	Veto hit locations
<i>MomentumTracker</i>	3D trajectory analysis	Primary momentum direction
<i>EventDisplay</i>	3D visualization	Visual insight

Table 6.3: Summary of the software analysis chain I developed to analyze CREST’s signal events. The input to my first module (*Calibrator*) is the output of *RawReader* (Nahee Park’s binary file converter). The parameters generated by *EventCutter*, *Lineal*, *Geometizer* and *MomentumTracker* are the basic building blocks of cut strategies designed to discriminate between signal and background events. Moving down the pipeline reduces the number of events to process but also increases the computational intensiveness; the total computation time required for each run strongly depends on the number of events that survive *EventCutter*.

6.4.1 *Raw Reader: Freeze Cycles*

Adhering to the maxim that objects in software should mimic objects in reality, the basic building blocks for data files are the lists of TDC and ADC words (and the detector elements that generated them) contained in each iteration of the CROL’s Freeze/Accumulate process. Each process represents 131,074 nanoseconds of data taking. The aforementioned binary data files contain Freeze Cycle objects written in binary format with headers and footers separating each Freeze Cycle.

Freeze cycles also have a manifestation as C++ software objects, which form the starting point for all subsequent analysis. In this context, freeze cycles are lists of hits in the detector occurring during the same CROL readout cycle. Nahee Park developed the *RawReader* software which our collaboration uses to convert the binary data files into databases of C++ freeze cycles in the form of “.root” files. ROOT (<http://root.cern.ch>) is an open-source set of inter-dependent C++ libraries developed at CERN and is widely used in the high energy physics community. Besides the PMT channel ID number, TDC value and ADC values, *RawReader* also adds metadata including a sequential integer incremented with each freeze cycle (“FreezeTick”) and whether it encountered errors while processing the freeze cycle (“ErrCount”).

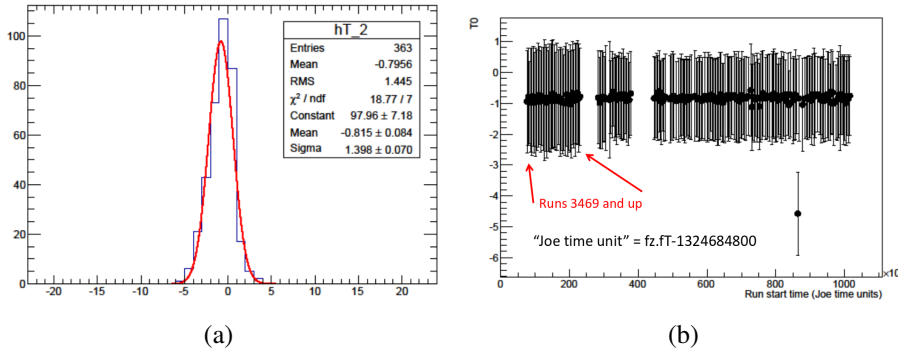


Figure 6.10: Scott Nutter’s plots which demonstrate utilizing the pulser information to determine timing offsets between crystal tubes. (a): Gaussian fit to the distribution of timing differences between tube 2 and an arbitrary reference tube in a single pulser run used to determine the mean t_0 . The x axis is measured in nanoseconds. (b): Plot of the time variation of tube 2’s mean timing offset over the course of the flight. The x axis is time relative to the what the collaboration referred to as “Joe Time,” which was merely a round number near midnight of launch day. Variation over the course of the flight is a fraction of a nanosecond.

6.4.2 Calibrator: Time and Energy Calibration

Since the scintillation properties of the BaF_2 crystals vary with temperature, in principle the calibration parameters could vary significantly over time. In practice this variation was rather small given that the period of data I work with here (from before “The Incident” discussed in Section 5.4.6) derives from a small portion of the flight and that STAC temperatures did not vary much over the flight.

The energy calibration technique (devised by Nahee Park) is discussed in Section 4.3.2. I integrated the code she provided into my software chain to generate a set of calibration parameters for every run I analyzed. Using these parameters, I converted the raw ADC0 and ADC1 values recorded in freeze cycles into ECF (Energy Combined, Fine) and ECC (Energy Combined, Coarse) values, with base units of keV.

Scott Nutter implemented the time calibration technique utilizing the Pulser system described in Section 3.2.5. The chief difficulty to overcome here was accurately correcting for the fiber lengths which differed from STAC to STAC (and from the documentation of the system provided to us). An example of the distribution of timing offsets (relative to an arbitrary reference tube) for crystal tube 2 is shown in Figure 6.10.

Jim Musser examined the “time-walk” correlation between hits’ TDC and ADC values and found it to be constant between tubes and over time. For each hit in a freeze cycle the TDC value was modified by a tube-dependent offset and an ADC-dependent time-walk correction.

To store this time-based data in a manner accessible to the entire collaboration, I created *cresthk*, a MySQL database. It stores the relevant time and energy calibration parameters as a function of run number (along with a host of other housekeeping related variables, such as pressure and temperature). *Calibrator* accessed *cresthk* at run-time (via the Internet using a MySQL hook I created called *HK_Connector*) and uses the calibration values there to perform time and energy calibration.

Since only relative veto TDC values were used in analysis we made, no attempt to calibrate the absolute TDC or ADC values for veto hits. In retrospect this would have aided some of my later analysis techniques, permitting a time-of-flight-like analysis variable to be estimated. However, many of the veto PMTs were lower-quality Burle tubes with worse time resolution than that of the Hamamatsu tubes used in the crystal system (~ 2.5 ns compared to ~ 0.75 ns), meaning that such analysis could be strained. This deficiency in the Burle tubes limits the effectiveness of the trajectory-pointing algorithm I detail in Section 6.7.


6.4.3 *EventFinder*: Splitting Freeze Cycles into Events

After time and energy calibration my next task was to break freeze cycles into analysis-sized chunks called Events. Events are subsets of freeze cycles in that they contain all the same data members as freeze cycles; the length of the list of hits is merely truncated. Each Freeze Cycle could contain an arbitrary number of events.


I wrote the tool “EventFinder” to break up freeze cycles into chunks called Events in a simple, unbiased manner. EventFinder sorts hits in freeze cycles by time, subtracts off the time of the first hit, then looks for gaps between hits longer than a settable parameter I call the gap width. Hits on either side of a gap width are assigned to separate events.

To avoid biasing event selection, the gap width should be longer than the time scale of a signal event, but not so long that one risks of combining events with other unrelated events or detector activity. As discussed in Chapter 2, the maximum time between arrival of signal photons from the same primary electron is less than one nanosecond. Convoluting this with CREST’s time resolution and adding in a fudge factor for potentially misaligned timing offsets between far-flung regions of the detector, units of analysis data should be of order tens of ns. This is far shorter than the duration of freeze cycles ($131 \mu\text{s}$). As shown in Figure 6.11(a), freeze cycles were sparse in that on average they contained little more than one event each; therefore any upper limit of reasonable size (smaller than $\sim 1 \mu\text{s}$) would be unlikely to run into another event.

One difficulty in using our signal simulations to guide this choice of time scale lies in



(a)



(b)

Figure 6.11: Two parameters used to guide my choice of 100 ns for the minimum time delay between events inside of freeze cycles. (a): Distribution of Freeze Ticks for events in a subset of flight run 03499. Bins on the x axis are Freeze Tick with bin widths of 12000. Each bin contains approximately 12000 events, showing that each freeze cycle in this particular run contained, on average, approximately one event. (b): Distributions of event durations in 10 TeV signal Monte Carlo simulation. Event duration is defined as the time in nanoseconds between the first and last TDC values of hits in the event.

the fact that they are not properly time-normalized. Our signal simulations generated a set number of events, not a set number of events per unit time; the time between simulated signal events is not a meaningful number. On the other hand, the simulated *duration* of signal events is a meaningful number. Figure 6.11(b) plots the distribution of the duration of Monte Carlo signal events, from first to last hit.

Instrument electronics also provide two relevant time scales. First the Digitize Decision used a look-back time of 100ns, meaning that STAC signals would typically be grouped by the read-out trigger in blocks of 100ns. Second the CROL histogramming function required two separate hits within a window of 180 ns in order to record freeze cycles.

To avoid biasing event selection, I used a minimum gap width between events of 100 ns. This allowed me to be as permissive as possible at this stage, leaving more sophisticated time-based rejection to tools higher up in the analysis chain.

6.4.4 *EventCutter*: Data Reduction Using Simple Parameters

One hour of flight data contains approximately 15 million freeze cycles, and therefore approximately 15 million events. After generating such an enormous number of events my next goal was to reduce this number by eliminating the obviously uninteresting events. Since different applications have different definitions of “interesting” (e.g. energy calibration requires single-hit events, while signal events have at least three crystal hits by virtue of the Digitize coincidence requirement) I developed *EventCutter* to provide a customizable data reduction tool.

I defined a set of surface-level event parameters listed in Table 6.4 . By “surface-level” I mean that these parameters require no analysis to generate. Rather, they reflect the multiplicity of active detector elements. This module processes events individually, compares the values of these parameters to user-provided limits, then discards or keeps events that fail or meet the selection criteria.

Recall that veto paddles are read out by tubes on both ends. A paddle having a pair in Table 6.4 refers to both its tubes having activity in the same event and therefore a high probability of a charged particle or low energy (below approximately 10 keV) photon traversing that paddle. At this point I do not require that the hits on either end be correlated in time, only that they be in the same event as determined by *EventFinder*.

I use *EventCutter* as a starting point for various analysis techniques by selecting events with the kind of detector morphology I desire. For example, if I wished to select single photon events to measure the 511 keV for energy calibration line I would set nCRYSTAL, nSTAC, and nBUS to a minimum and maximum value of one and all veto activity

ID	Name	Number of...	Min value	Max value
1	nCRYSTAL	Activated crystal PMTs	0	1024
2	nSTAC	Activated crystal STACs	0	64
3	nBUS	Activated crystal C-channels	0	8
4	nVETO	Activated veto PMTs	0	54
5	nPADDLE	Activated veto paddles	0	27
6	nPAIR	Veto paddles with hits on both ends	0	27
7	nPADDLETOP	Activated top veto paddles	0	8
8	nPAIRTOP	Top veto paddles with hits on both ends	0	8
9	nPADDLEBOT	Activated bottom veto paddles	0	9
10	nPAIRBOT	Bottom veto paddles with hits on both ends	0	9
11	nPADDLESIDE	Activated side veto paddles	0	4
12	nPAIRSIDE	Side veto paddles with hits on both ends	0	4
13	nPADDLESLANT	Activated inside slanted veto paddles	0	6
14	nPAIRSLANT	Inside slanted veto paddles with hits on both ends	0	6

Table 6.4: “Surface” event parameters defined in the ParFiller module describing the multiplicity of activated detector elements. The minimum and maximum values represent the sensible range for values of the cut parameter, not actual cut parameters which vary from analysis to analysis. By providing ParFiller with a configuration file consisting of lower and upper bounds for these parameters, the user can produce a culled file of Events meeting the desired cuts.

to maxima of zero. For the vast majority of signal-seeking techniques I set nCRYSTAL to a minimum of four or higher. The mainstream analysis cuts out all top veto activity; to replicate those cuts I would set maxima of 0 for the top veto parameters. My analysis incorporating signals from the veto system uses a variety of cut values, which I detail after presenting the elements of signal events.

6.4.5 *Lineal*: Clustering, Linear Analysis and Crystallography

The *Lineal* module takes as input (from *EventCutter*) events which have only the basic hit data (tubes, times and energies) and the surface level data outlined in Table 6.4. *Lineal* (short for “Linear Analysis”) then executes these key analysis procedures:

- Group hits into clusters (Section 6.4.5.1)
- Fit the main event line (Section 6.4.5.2)
- Calculate hit separations (Section 6.4.5.3)

I now explain how each of these procedures is accomplished and present the data that guided my choice of algorithmic parameters.

6.4.5.1 *Lineal*: Hit Clustering

In what follows I adhere to the convention of the collaboration in making the distinction between “primary” and “secondary” based on where the particle or photon in question is produced.² Primaries are generated above the detector and secondaries within a sphere around the detector with radius of approximately 2 meters. Thus the synchrotron and bresmsstrahlung photons directly produced by the primary electron are referred to herein as primary event members.

Primary particles and photons produce secondary particles and photons in the air immediately above CREST’s crystal plane, in the veto paddles surrounding the crystal plane, or after interacting with the crystal plane itself (including with the lead shield surrounding each crystals). In consultation with my advisor, I opted for a cluster-based approach, treating nearby (in space and time) crystal hits as if they were induced by the same parent particle or photon. To implement this I created a friends-of-friends (FoF) clustering algorithm as part of the *Lineal* module.

The FoF algorithm has two key parameters. The first parameter, the FoF window width Δ_t , determines which hits are valid cluster member candidates based on their time compared to the time of the earliest hit in the event. If t_0 is the time of the earliest hit and $t_i > (t_0 + \Delta_t)$ then the i -th hit is not included in any hit cluster. To guide my choice of this parameter I examined the relative times of hits in simulated signal events. Figure 6.12 plots the distribution of relative hit times in simulated 10 TeV signal events and reveals the existence of three time regimes:

1. **Regime I:** Primary photons arrive at the detector simultaneously. Hits with relative time of 0 belong to this category.
2. **Regime II:** Secondary scatters within the detector occur within the light crossing time of the detector after the earliest hit. Hits with times later than 0 ns but earlier than 6 ns or so belong to this category.
3. **Regime III:** Hits arriving more than approximately 8 ns after the earliest hit are products of electromagnetic cascades in the residual atmosphere above the detector. The primary electron and high energy Bremsstrahlung photons are likely to induce these air showers; low energy synchrotron photons are not. Hits resulting from long-range, high-angle deflections (e.g. backscatter reflections) also contribute to this portion of the time distribution.

²This convention arose because our simulation code was divided into a primary generator phase, where the electron interacted with the atmosphere and magnetic field to produce “primary photons,” and a secondary interaction phase, where the “primary photons” emitted by the electron interacted with the detector.

Thus the choice of Δ_t really depends on which types of clusters we want to capture. If we want to focus on primary photons only then we would set a very short time window of approximately 1-2 ns. If we also wish to capture secondary scatters then we should choose a window of approximately 6-7 ns. If we further wish to capture the air shower products then we should extend the window to 10-15 ns or perhaps even larger. In Figure 6.14, I plot the maximum separation (in units of the crystal lattice separation parameter, 7.5 cm) between primary photon and estimated cluster location as a function of Δ_t and cluster location method. As Δ_t decreases the frequency of catastrophically incorrect cluster locations decreases, and the incidence of slightly-less-wrong locations increases. Note that this is the maximum error in an event; the average error is much smaller (on the order of 1-3 crystal separations).

The air shower clusters can provide useful information, such as evidence of a traversing primary electron. To locate them would require a second clustering algorithm utilizing a time window that opens at approximately 8 ns and closes at approximately 20 ns after the earliest hit. The location of the cluster should be based on the average hit location, not the location of the hit with maximum energy or earliest time. This is because the primary electron or photon has already passed through the detector and the remaining shower detritus will have random arrival times and energies.

The second FoF parameter is the separation parameter, Δ_{xy} . Given a cluster with a set of hits as members, additional hits will be added to the cluster if the minimum separation in the x, y plane between the new hit and each of the existing hits in the cluster is less than or equal to Δ_{xy} . This parameter defines the maximum separation between a hit and the most distant hits in the same cluster. Values for Δ_{xy} that are too small causes the clustering algorithm to cleave hits into separate clusters when in fact they belong to the same parent particle. Values that are too large risk erroneously clustering hits together that in fact belong to different primaries. To guide my choice of this parameter I relied on the same studies the collaboration used to determine whether or not to add the lead shields to the crystals. These Monte Carlo studies indicated that over 90% of secondary signals from primaries with energies between 40 keV and 50 MeV occur in crystals within two crystal separations of the primary-activated crystal struck. I show a plot by Scott Nutter demonstrating this in Figure 6.15. Based on this analysis, I chose to use $\Delta_{xy} = 15$ cm (two crystal separations).

For events with long tracks of consecutive hits (such as the side-going proton shown in Figure 6.31) this results in very long clusters. Conversely, events with significant gaps along the main event line between hits result in shorter clusters. Since signal hits would be more likely to fall into the latter category, this behavior hints at a signal-seeking cut based on cluster shape. I return to this point in Section 6.4.5.3.




Figure 6.12: Hit times (compared to the time of the earliest hit) in simulated signal events for primary Bremsstrahlung photons and their secondary products (green online), primary synchrotron photons and their products (blue online) and the primary electron and its secondary products (red online). Three time scales are apparent. First, the simultaneous arrival (at 0 ns) of the primary electron and photons. Second, intra-detector scattering takes place within one light crossing time of the detector (between two and six ns). Third, showers formed in the air above the detector can result in the delayed arrival of shower products, ranging from eight ns all the way up to hundreds of ns. Synchrotron photons contribute relatively few hits to the late (air shower) regime because of their lower average energy; they more often interact via the photoelectric effect or Compton scattering than pair production (see Figure 4.7).

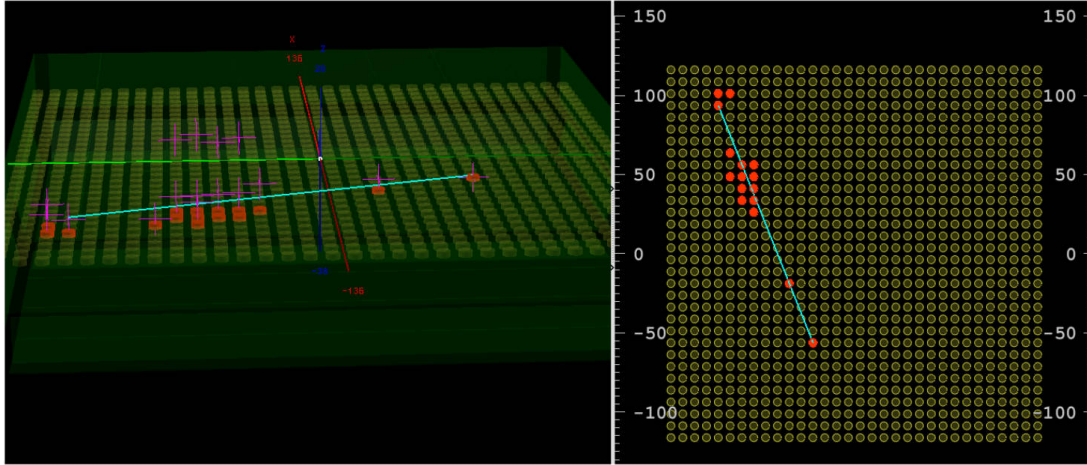


Figure 6.13: Electron signal candidate event from run 03488 showing cluster formation consistent with all three time regimes shown in Figure 6.12. The two clusters towards the right (lower right in the top-down view) happen at time 0, indicating they are probably primary photons. The left cluster (upper left in the top-down plot) has a primary photon with two slightly later hits that are most likely secondary Compton scatters. The center cluster is likely a Bremsstrahlung primary photon that air showered above the detector. These hits arrived 10-15 ns later and belong to the latest time regime.

6.4.5.2 *Lineal*: Fitting the Main Event Line

After separating an event's hits into clusters, the primary task of *Lineal* is to estimate the main event line via the spatial pattern of crystal hits. Since I want to know the locations where primary photons and particles interacted with the crystal system, I use the locations of clusters of hits, rather than the locations of individual hits, to fit the location of the main event line. I attempted many different methods of line fitting; all fail in various situations and for various reasons. In the end, I chose my best guess for the main event line as the line that connects the most distant crystal hits which are within Δ_t (the time window clustering parameter) of the earliest hit in the event. While this does not always find the right line, it does have the distinct benefit of being very simple. Figure 6.16 shows the distribution of errors in the direction of the main event line in simulated 10 TeV signal events. For this plot the estimate could never be more than 90° wrong; in other words I have estimated $\pm \hat{\Delta}_p$ without specifying the sign. Identifying the sign is not required for separating background from signal events, but *is* required for separating positron signal events from electron signal events. (I return to this point in Section 6.8.)

1. Group the crystal hits into clusters (as in 6.4.5.1)
2. Draw a line between the clusters with the largest separation

Max Cluster Sep, t08c0




Figure 6.14: Distance, in units of crystal separation parameter (1 separation = 7.5 cm), between the location where a primary photon pierced the crystal plane, and the location assigned to the cluster containing its secondary hits. These distributions resulted from applying my clustering algorithm to simulated 10 TeV signal events while varying the temporal FoF window parameter Δ_t , and the method used to determine the cluster center. I plot the largest error within each event. Using a shorter time for Δ_t reduces the frequency of catastrophically mistaken cluster locations. “Earliest” and “heaviest” refer to whether I take the earliest hit or most energetic hit as the cluster center; for most clusters, there is little practical difference between the two.




Figure 6.15: Plot of the distribution of distances between the crystal struck by a primary particle or photon and the crystals struck by its secondary products with a full lead shield configuration. The x axis is in units of crystal separations (one center-to-center crystal separation is equal to 7.5 cm). Plot produced by Scott Nutter using the signal Monte Carlo he developed. Over 90% of secondary hits occur in crystals within two crystal separations of the primary crystal leading to my choice for maximum separation between hits in the same cluster of 15 cm (2 crystal separations).

3. Call this line the main event line

The angle this line makes with the detector x axis is *Lineal's* estimate for the azimuthal angle ϕ shown in Figure 6.3. I will refer to this estimate as $\bar{\phi}$. Figure 6.16 plots the distribution in 10 TeV signal data of the angular error statistic $\bar{\phi} - \phi$, where ϕ is the true azimuthal angle of the main event line (calculated from the Monte Carlo truth data).

Figure 6.18 summarizes some of the analysis parameters generated by *Lineal* after attempting to find the main event line.

6.4.5.3 *Lineal*: Crystallographic Analysis

One important difference between linear events produced by by side-going protons and down-going signal electrons is just that: protons must traverse the crystal system horizontally while photons emitted by the electron illuminate it from above. Thus while the underlying skeleton of the crystal response to both of these types of events is linear, the distances between hits along the line should have different characteristics.

Protons have essentially unit probability of being detected in each crystal they traverse (see Section 4.1). Side-going protons can not “sneak” through the crystal system; the




Figure 6.16: Distribution of the statistic $\Delta_\phi \equiv \bar{\phi} - \phi$ in simulated 10 TeV signal events when using the main event line estimating algorithm outline in Section 6.4.5.2. Main event lines in events with very high crystal multiplicity are mis-characterized more often than those in events with low crystal multiplicity.




Figure 6.17: Distribution of the inverse track velocity in simulated 10 TeV signal events of varying maximum extent (event length is defined as the length of the fitted main event line). There are two main populations. First, the “simultaneous” population at zero inverse velocity (corresponding to simultaneous line endpoints). This value occurs when the main event line endpoints happen to be correctly chosen as primary photons. Second, the “scatter” population at a wide distribution centered approximately at $v/c = 1$. This value occurs when one endpoint of the line is chosen to be a cluster that results from intra-detector scattering. Since *Lineal* used a time window of 6 ns to generate this analysis, hits from the second regime in Figure 6.12 become valid cluster locations and times. Evidently *Lineal* is more likely to discover the correct line endpoints in longer events.

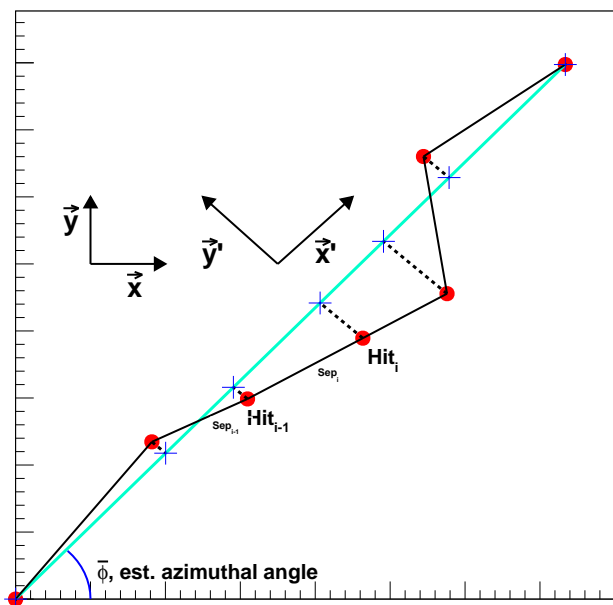


Figure 6.18: Schematic summarizing how Lineal dissects crystal hit morphology in an event. The two most distant clusters (at lower left and upper right) define the endpoints of Lineal’s estimate for the main event line (solid teal line). The length of this line defines the analysis parameter “event extent.” All of the other clusters in the event are sorted according to their projected coordinate along the main event line (that is, along the rotated axis \vec{x}'). This enables generation of a plethora of analysis parameters. The lengths of gaps between clusters (Sep_{i-1} , Sep_i , etc.) are measured as shown for use in the crystallography cut (discussed in Section 6.4.5.3). Second, the projected coordinates orthogonal to the main event line (that is, along the rotated axis \vec{y}') are measured by the lengths of the dotted lines. This provides an estimate of the “event width” and the starting point for calculating the goodness-of-fit of the line (as characterized by the reduced χ^2 variable). The last parameter extracted via this process is the ratio of the spatial distance and time difference between the line endpoints. This statistic (the “inverse track velocity”) measures the inverse of the implied velocity of the particle that made this track in units of $1/c$. If the line endpoints are simultaneous then the inverse track velocity is zero, implying that if a side-going particle created this track it must have had infinite velocity (in other words, a side-going particle likely did not make this track). If the line endpoints are on a null interval (as would be the case for hits produced by side-going charged particles) then the inverse track velocity is approximately 1. Figure 6.17 shows the distribution of this statistic in Monte Carlo signal events.




Figure 6.19: The two-point function for distances between crystals as a function of the azimuthal orientation of the vector between the two crystals. Orientations of 0 and 90 degrees both correspond to alignment with a principal axis of the crystal lattice. Gaps along the main event line in signal events should be characterized by this spectrum.




Figure 6.20: The un-occluded two-point function for distances between crystals. Generated by removing those distances from the plot in Figure 6.19 spanned by vectors which intersect a third crystal at some intermediate distance. Proton tracks at a given angle can leave gaps of these sizes between consecutive hits, but not larger.




Figure 6.21: Subset of the plot shown in Figure 6.20 focusing on tracks aligned with the crystal lattice's principal axes. Such tracks should be characterized by large distances between long runs consecutive hits. At these angles, occluded spectra (from proton tracks) and un-occluded spectra (from signal electron events) are not easily distinguishable.

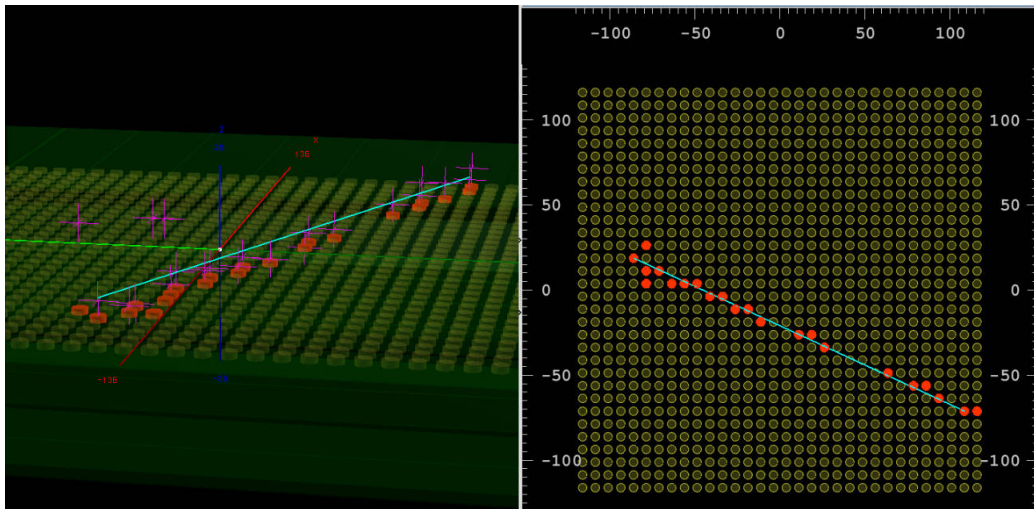


Figure 6.22: A linear event with no veto activity from run 03488. Since it is not aligned with the crystal lattice principal axes it has characteristically short gaps between hits.

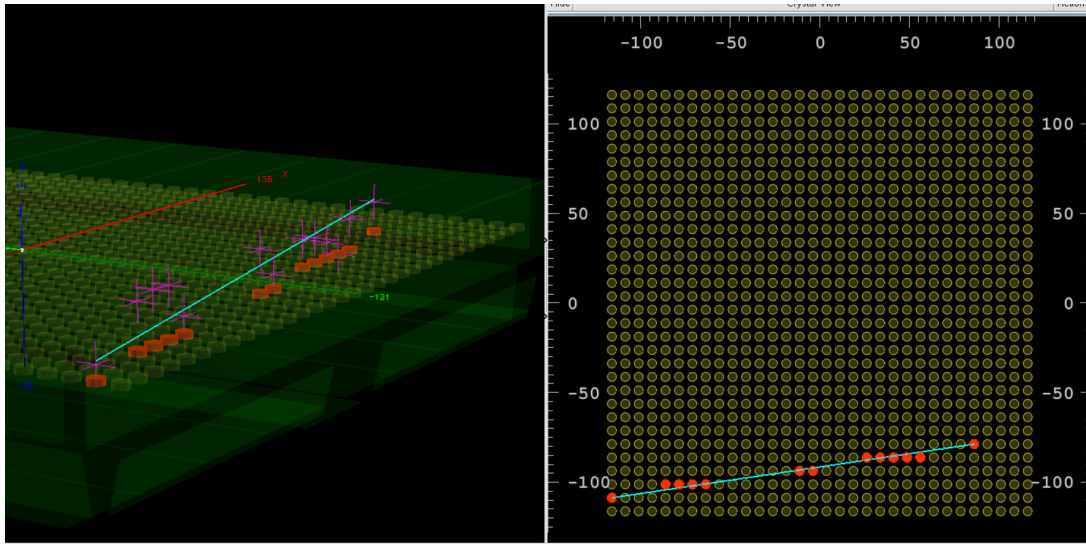


Figure 6.23: Another linear event with no veto activity from run 03488. Near alignment with a crystal lattice principal axis leads to gaps and runs of hits with a characteristic size as the proton enters and exits rows of crystals.

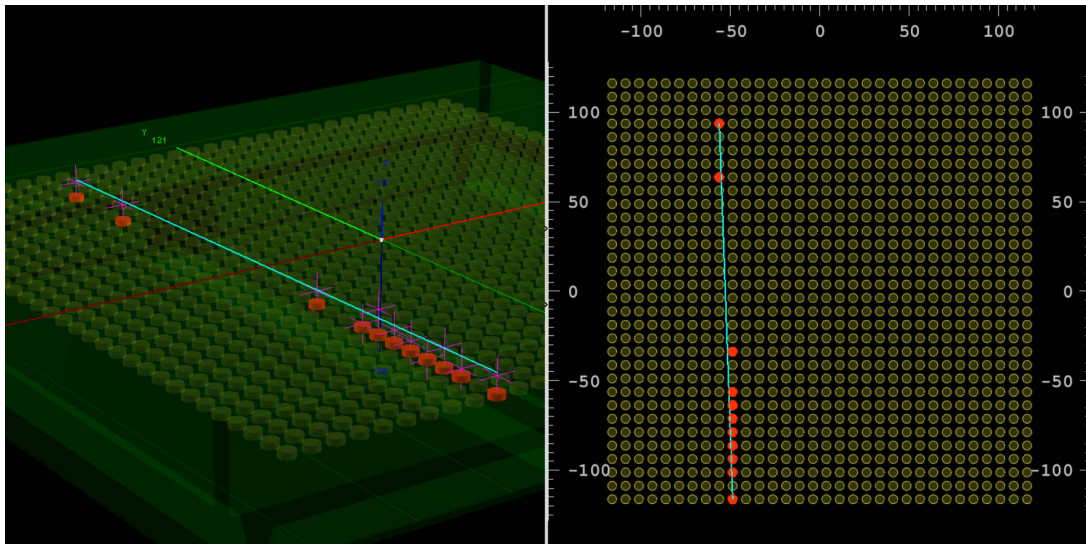


Figure 6.24: Another linear event with no veto activity from run 03488. Even closer alignment with a principal axis has enabled the proton to leave a gap (and a string of consecutive hits) nearly the size of the instrument.




Figure 6.25: Visual representation of the definition of an “illegal hit,” forming the basis of my crystallography cut. Gaps between hits that reside inside the red box in this plot are “illegal” for side-going protons: in order to create a gap between hits of those distances and at those track angles, they would first have to pass through another crystal. Events that create gaps inside the red box are likely not the result of side-going charged particles, independent of veto activity. Requiring multiple illegal separations in an event helps eliminate side-going charged particles that evaded detection by the veto system.

gaps between hits generated by such protons can only be as long as the un-occluded distance between consecutive crystals. Since the crystals are arranged in a regular lattice, the maximum un-occluded distance a proton can travel through the detector without striking a crystal varies with the direction of the proton track. For this reason I describe this line of analysis as crystallographic.

In Figure 6.19, I plot the two-point function for distances between crystals as a function of angle. This includes trajectories that are occluded by an intervening crystal. The x axis is measured in units of crystal separation (1 separation = 7.5 cm) and the y axis in degrees above horizontal (equivalent to the azimuthal angle ϕ shown in Figure 6.3). Figure 6.20 shows the same plot, except that trajectories occluded by a crystal are removed. The remaining data points represent clear paths between crystals: paths that could be traversed by a proton without striking a crystal somewhere along the way. Figure 6.21 zooms in on angles in close alignment with the principal axes of the crystal lattice; such angles permit long distance between activated crystals. In such events, protons are able to travel unimpeded for long distances between two rows of crystals, leaving unnaturally large gaps between hits.

Synchrotron and Bremsstrahlung photons emitted by the primary electron all lie along the main event line, but their location along the main event line depends on the random process governing the pathlength traversed by the electron between photon emissions. Thus the distance along the main event line between primary photon hits is independent of the properties of CREST's crystal lattice (other than the simple dependence on the crystal-to-crystal separation). This distance is represented in Figure 6.2 as the distance between the gray arrows' intersection point with the main event line \vec{L} . The distance between gaps between hits in signal events should obey the full spectrum shown in Figure 6.19, not the occluded spectrum in Figure 6.19.

6.5 Summary of Linear Analysis Parameters

Figure 6.26 shows various elements of *Lineal's* analysis structures overlaid on a top-down view of a flight event. As usual, the main event line is shown in teal and connects the two most distant clusters. The length of the main event line defines the event's length (requiring events longer than 80 cm is a typical signal cut). The maximum orthogonal distance of hits from the main event line defines the event's width; in this case no orthogonal distance was greater than 40 cm. The lengths of the gaps between hits in the lower right portion of the track, combined with the track's non-alignment with the crystal lattice's principal axes, constitutes three separate "Gaps of Unusual Size." This enables the event to pass the




Figure 6.26: Event display with overlay describing the meaning of several key linear analysis parameters. The event length is measured as the length of the teal main event line. The event width is measured as the maximum orthogonal distance between the main event line and hits in the event. The veto activity in this event has been suppressed in this display for demonstrative purposes. The number of illegal separations between hits is calculated as in Section 6.4.5.3.




Figure 6.27: The same event from flight data shown in Figure 6.26 with some veto information made visible. The red ellipses correspond to estimated locations of hits within veto paddles. The cones with the sharp ends upward represent hits in the top veto system; the cones with the sharp ends downward represent hits in the bottom veto system. Here I have highlighted those hits that are consistent with the proposed traversal location of the primary electron. The size of the cones reflects the uncertainty in the position estimate: approximately the width of the paddles in the short dimension, and approximately 30 cm (σ_{veto} in the long dimension).

crystallographic cut, as a side-going proton would be unable to achieve such large gaps at this azimuthal angle. The large cluster at the top left is likely the result of an air shower produced by a traversing electron or Bremsstrahlung photon; the crystal colored bright red has the highest energy deposition (about 10 MeV).

But all this is merely educated conjecture. How can we further convince ourselves that this cluster was due to a traversing electron, or the Bremsstrahlung shadow of a nearly-traversing electron, and not some background event? The next tool in my analysis chain, *Geometizer*, looks for evidence of the second, downward-oriented vector present in electron signal events (the electron momentum \vec{p}) - a vector which is absent from side-going charged particle events.

6.6 *Geometizer*: Geometric Analysis

If the cluster in the upper left portion of Figure 6.26 is due to a traversing electron, then we might examine the activity in the veto system for evidence of such. While an electron with energy above a TeV has a low probability of activating a veto paddle (since the paddles are too radiatively thin to contain the Bremsstrahlung photons produced by the electron's

traversal of the plastic scintillator), it is probable that any air shower products produced by the electron immediately above the detector will activate the top and bottom veto. The same is true for a Bremsstrahlung photon with sufficient energy to initiate electromagnetic showering in air or the veto paddle. Furthermore any secondary products (recoil Compton electron and photon or shower components) produced while interacting with the crystal system should pass through the detector and activate the bottom veto.

Figure 6.27 shows a three-dimensional view of this event. As shown by the large red ellipses approximating the location of hits in the veto paddles, there is in fact top and bottom veto activity in the place predicted by the large crystal cluster. In this Section I review how I arrive at those estimated veto hit locations, and how I utilize that information to select signal events. This work was the pre-cursor to my pointing technique, which builds upon these efforts to actually point at the electron momentum, rather than merely hint at its existence separate from the momentum implied by the track of crystal hits. I describe that pointing technique in Section 6.7.

6.6.1 Procedure for Veto Hit Location Estimation

To estimate the locations of hits within veto paddles, I utilize a technique outlined by those collaboration members primarily responsible for the veto system (Stéphane Coutu, Scott Nutter and Matt Geske). The technique simply consists of comparing the arrival times of signals at the PMT on either end of a veto paddle, then calculating the origin point of the signal using the empirically determined effective speed of light in the fiber optic wave guides embedded in each paddle. The inverse operation is performed in the Monte Carlo simulations to estimate TDC values in veto paddle PMT's given the arrival time of a hit somewhere along the veto paddle.

Let t_1 and t_2 be the times assigned to signals in each PMT, L be the length of the veto paddle and v_{eff} be the propagation velocity of light along the paddle. If a signal is generated within the veto paddle at location x as measured from side 1 at time t_0 , then

$$\begin{aligned} t_1 &= t_0 + x/v_{\text{eff}}, \text{ and} \\ t_2 &= t_0 + (L - x)/v_{\text{eff}}. \end{aligned}$$

Since we wish to estimate x in terms of $t_2 - t_1$, subtracting the two equations yields:

$$t_2 - t_1 = t_0 + (L - x)/v_{\text{eff}} - t_0 - x/v_{\text{eff}}. \quad (6.3)$$

Simplifying and solving for x yields:

$$x = \frac{L - v_{\text{eff}}(t_2 - t_1)}{2} + x_0, \quad (6.4)$$

where I have added the constant value x_0 to account for the fact that t_2 and t_1 will have some constant, unknown time offset relative to one another. In practice, I determined x_0 by requiring that veto paddle hit location estimates have average estimated location at the center of the paddle.

Jim Musser provided empirical measurements of v_{eff} , the effective speed of light in the fiber optic wave guides within the top paddle, by comparing the location of singleton crystal hits with the time differences $t_2 - t_1$ in the top paddle PMTs, and finding the best-fit v_{eff} values for each top paddle. They range from approximately 12 cm/ns to 18 cm/ns, reasonable values for visible light rays bouncing in a spiral along a wave guide, compared to the speed of light in vacuum of 30 cm/ns. This works well for the top paddles because the crystals are only millimeters away. Such a procedure is not applicable, however, to the non-top veto paddles, as they are further away from the crystals. Further, since the top paddles are read out by the higher-quality Hamamatsu PMTs with better time resolution than the Burle PMTs reading out most of the non-top veto paddles, it would be incorrect to simply assume that the rest of the veto system performed as well as the top paddles.

Before I could proceed with combining hit location information in the side and bottom veto paddles with the crystal hits I needed to determine how accurately and precisely these paddles could locate hits. For the side paddles I used the following procedure:

1. Identify long, linear crystal tracks in flight data
2. Linearly extend the crystal track outside of the crystal plane to its two intersection points with the veto system, generating estimated locations for where the putative side-going particle would have struck the side veto paddles
3. Compare the estimated hit locations within the veto paddles to the hit locations predicted by the extended crystal track

In the absence of measured v_{eff} values for the side paddles I used an average of the top paddle velocities. Figure 6.31 shows an example of the result of this procedure. The blue cone facing us represents the intersection of the extended crystal track with the side paddle facing us. The red ellipse on that paddle represents the estimated hit location as determined by the procedure above using the relative time of hits in the paddle's PMTs. In this event the agreement is excellent. To determine the precision of the hit location I repeated this

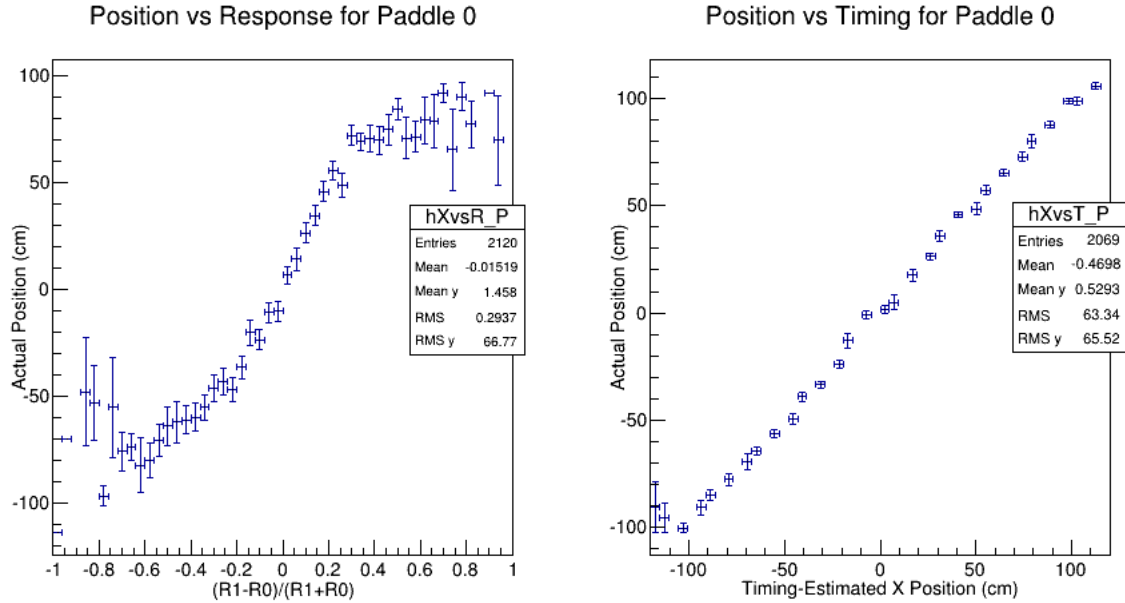


Figure 6.28: Comparing two methods for estimating hit locations along top veto paddle 0. On the left, I plot the average location estimate and error as a function of relative response in the PMTs on either end (using the formula shown). On the right, I plot the average location estimate and error as a function of relative timing in the PMTs on either end. The timing method provides smaller errors and more uniform results across the entire paddle.

procedure for an entire flight run's worth of data and produced the distributions shown in Figure 6.29. The RMS error in each distribution is slightly larger than 30 cm. I refer to this distance as σ_{veto} , since it is approximately the precision with which veto paddles locate hits. This is roughly consistent with timing errors of order ~ 2 ns, as expected from the Burle phototubes, convolved with errors in the alignment of the crystal track of a few degrees in events with long crystal tracks (see Figure 6.16). As a side benefit of this analysis, the accuracy of the results shows that the vast majority of long, linear crystal tracks are in fact produced in response to side-going charged particles.

Another method for estimating position along the paddle uses the relative response (that is, the relative ADC channel) of the hits in each PMT. This relies on the fact that as light signals propagate along the paddle, they experience exponential amplitude attenuation. Figure 6.28 compares hit location estimation for relative response (left) against relative timing (right). The non-linear results near either end of the paddle indicate that the attenuation was non-uniform in these regions; most likely because these signals originated in the FOLGA mating a PMT to the paddle, rather than inside the paddle itself. Based on this comparison, I chose to rely only on the relative timing method.

Figures 6.30 and 6.31 present two flight events displaying crystal and veto morphology




Figure 6.29: Hit location accuracy of the four side veto paddles. Plotted is the distribution of the distance, in cm, between the intersection of the extended track of crystal hits with a paddle and the location generated using the end-to-end relative timing method detailed in Section 6.6.1. The distributions have been corrected for the unknown timing offset between the PMTs on either end (x_{offset} in Equation 6.4. All four paddles achieve hit location estimation with error of order ~ 30 cm (a distance I call $1 \sigma_{veto}$).

typical of the main background for linear events: side-going protons. The proton in the first event nearly deviated detection where it exited the instrument: the magenta paddle without a red elliptical cone indicates that one, but not both, of the paddle's PMTs were activated. The proton in the second event had both its entrance and exit points properly flagged by the veto system. The blue arrow pointing out of the detector indicates the location, and direction, of *Geometizer's* best guess as to the location where particle made that track in the crystal system pierced the facing side veto paddle. The close agreement of this arrow and the red elliptical cone indicate that the crystal and veto hit patterns agree: here went a charged particle. Such an event would contribute to the central peak in that side paddle's error distribution plotted in Figure 6.29.

I am able to make accurate guesses as to where the hit locations in the side paddles should be mainly because main line fitting is relatively accurate for the long crystal tracks which extend into the side paddles. It is difficult to apply a similar technique to the inside slanted and bottom veto paddles, primarily because doing so would require estimating θ , the zenith angle for proton events. The steeper the incline of the proton trajectory relative to detector horizontal, the shorter the track it can leave in the crystal system, as shown in Figure 6.32. Deviation from horizontal of only a few degrees shortens the longest pos-




Figure 6.30: A flight event from run 03488, likely the event of a side-going proton which evaded detection in the far-side paddle. The large red elliptical cones indicate the estimate location within the veto paddle where the particle or photon traversed the paddle. This location is based on the relative timing between the PMTs on either end of each paddle (and as you can see this occasionally results in the ludicrous location estimate outside of the paddle itself, indicating a timing error). The gradual increase of the purple crosses indicating the hit time as the crystal track progresses is consistent with a particle moving across the detector with velocity c . The uniformly short distance between hits is consistent with a side-going particle as well. The lack of veto activity beneath the far side of the event provides strong evidence against the possibility that the particle inducing the event had momentum in the vertical direction.




Figure 6.31: Another flight event from run 03488, also likely the event of a side-going proton. This time however the veto system has properly flagged the event with entrance- and exit-locations. The blue cone in the facing side paddle indicates *Geometizer*'s best guess as to where it believes the crystal track indicates the particle exited the detector. The close agreement with the veto hit location estimate provides strong evidence, *independent of the timing of the crystal hits*, that this was indeed a side-going proton.

sible track by over 50%. Proton trajectories which point at the bottom veto paddles will inherently have stunted main event lines, which are estimated less accurately. Rather than directly testing the bottom and inside slanted paddles, I assumed they behaved similarly to the side paddles. While most likely a reasonable assumption given that the Burle tubes in these paddles are functionally identical, the fact that the pointing technique outlined in Section 6.7 works at all is convincing evidence that the paddles behave as expected.

6.6.2 *Geometizer*: Characterizing Veto Reponse to Signal Events

Events such as side-going protons where the direction of the crystal track is parallel with the momentum of the particle that made the track should display agreement between the blue cone and the red ellipse. Events where these two directions are *not* parallel should be characterized by disagreement between the crystal track and the veto hit locations. In particular, in signal events where the crystal track and the electron momentum are *orthogonal*, the disagreement should be significant.

I describe my first attempt to characterize this disagreement with the aid of Figure 6.34. I define “pierce error” as the distance between each non-top veto hit location and the particle’s exit location predicted by the crystal track. The minimum pierce error represents the level of agreement between the crystal track and the hits in the veto system as to the




Figure 6.32: Protons with more vertical trajectories leave shorter tracks in the crystal system. I plot here a geometrical estimate of the maximum length crystal track a proton with the given angle could make given that the crystals are only 2 cm tall.




Figure 6.33: A top-down view of the event in Figure 6.34 clearly showing how the location of the blue cone is established by simply extending the main event line until it intersects with a veto paddle. The way the bottom veto hits shadow the clusters in the main event line suggested to me that measuring the primary electron momentum utilizing the bottom veto signals might be possible.




Figure 6.34: Event with overlay showing how I define “pierce error,” the three-dimensional distance between the crystal-track generated particle exit location and the nearest estimated veto hit location. The minimum pierce error is the smallest of these distances. A small minimum pierce error indicates that the veto and crystal systems agree that on the trajectory of the particle that made the crystal track and veto hits. A large minimum pierce error indicates disagreement between the two systems and is a clue that the crystal track is the result of signal electron illumination, not a background event generated in response to a side-going proton.




Figure 6.35: Distribution of minimum pierce errors in background Monte Carlo, low- and high-energy signal Monte Carlo and flight data sets. Minimum pierce error is defined in Figure 6.34 and is measured in units of 30 cm, the RMS error for veto hit location estimates. The data sets used here are discussed in Section 6.6.2.

ID#	Data set	Details	Number of events
1	Background MC		19.8M
2	Signal MC	10 TeV e^- , $5g/cm^2$ o.b.	146k
3	Signal MC	0.5 - 1.25 TeV e^- , $5g/cm^2$ o.b.	134k
4	Flight	runs 03483-03488	90.5M

Table 6.5: Summary of the four data sets used when generating the plot shown in Figure 6.35. The flight runs form a five-hour subset of flight data from the period after reaching float altitude but before “The Incident.” The signal Monte Carlo data sets were generated by me using the modified simulation code that allows the primary electron to interact with the detector and therefore has limited statistics.

trajectory of the particle that created the crystal track.

In a background proton event, the end of the main event line further from the top veto activity should be accompanied by a hit in the veto system marking the location of the particle’s exit from the detector. Since I know that I can identify veto hit locations to within approximately 30 cm, if I plot the distance between the particle exit location (as predicted by the crystal track) and the nearest estimated veto hit location, I should thereby generate a distribution with a large peak at approximately 30 cm.

On the other hand, in electron signal events the exit location of the “particle” should not be covered by a veto hit, simply because no particle actually exited in that location. For these events there should be a large distance between the particle exit location predicted by the crystal system and the nearest estimate veto hit location. I plot the distributions of minimum pierce errors in Figure 6.35 for events in four separate data sets, as summarized in Table 6.5 and using the signal-seeking cuts listed in Table 6.6.

In Figure 6.35 it is clear that the simulated background and signal data sets have markedly different distributions for the minimum pierce error. The peak at $1 \sigma_{veto}$ for background events is expected by virtue of the definition of σ_{veto} as the width of the veto hit location estimate error distribution. The peak for signal events at just under $3 \sigma_{veto}$ represents a distance very close to the vertical gap between the crystal plane and the bottom veto system, indicating the the hit nearest the crystal-predicted exit location is most often directly underneath it. It was clear at this point that the bottom veto signals had important information about the event. My next analysis efforts were directed at unlocking more of this information.

6.7 *MomentumTracker*: Pointing at the Primary Electron

In this Section I describe my technique for utilizing the crystal and veto hit information in concert to estimate the momentum of the primary electron in simulated signal events.

Cut	Min	Max	<i>Fraction of prev. surviving</i>			
			1	2	3	4
# Crystal STACs	4	63				
# Top veto pairs	1	2				
# Non-top veto pairs	5	19				
	<i>Base cuts</i>		.240	.347	.452	.183
Event length	80 cm	–				
Event width	–	40 cm				
# Clusters	4	–				
# Illegal gaps	2	–				
	<i>Linearity cuts</i>		.091	.245	.221	.095
Minimum pierce error	$2\sigma_{veto}$	–				
	<i>Geometric cuts</i>		.112	.546	.571	.107
c / track velocity	–	0.1				
	<i>Simultaneity cut</i>		.038	.054	.052	.037
	<i>Total remaining events</i>		1856	358	399	6292
	<i>Total remaining (fraction)</i>		9.37×10^{-5}	2.45×10^{-3}	2.9×10^{-3}	6.95×10^{-5}

Table 6.6: Response of the four data sets listed in Table 6.5 to the listed cuts. The minimum pierce error cuts results in a multiplicative signal-to-noise enhancement of approximately 5 in simulated events. “Fraction of prev. surviving” indicates the fraction of events that reach a cut stage which survive that particular cut stage. Thus the “total remaining (fraction)” is equal to the product of all the individual survival fractions in each column. This particular set of cuts achieves background rejection of nearly 10^{-6} , but also suffers signal acceptance of less than 3×10^{-3} .

I again used the custom-produced high-energy and low-energy signal MC files utilizing simulation code modified to allow the electron to interact with the detector (data sets 2 and 3 in Table 6.5).

6.7.1 An Additional Signal-Seeking Cut?

Before I do so I, wish to direct the reader’s attention once more to Figures 6.2 and 6.3. In particular focus on the orientation of \vec{p} with $\vec{\Delta}_p$: the two are orthogonal, regardless of magnetic field orientation. $\vec{\Delta}_p$ is hidden from us. We only measure its projection into the detector plane \vec{L} , and we estimate this line’s orientation with $\vec{\phi}$, our best guess for the orientation of the main event line.

When the magnetic field points vertically downward $\vec{\Delta}_p$ is parallel to \vec{L} and we have $\vec{p} \times \vec{L} = 0$. When the magnetic field is close to vertically downward we instead of $\vec{p} \times \vec{L} \approx 0$. As shown in Figure 5.40 the magnetic field is not vertically downward at CREST’s location, but it is also not that far off from vertical.

I tested this hypothesis using Scott Nutter’s signal simulation code. CRESTMag fires the primary electron from the same location but along a randomly oriented trajectory for every event. Thus while it does sample different regions of the geo-magnetic field, it is a




Figure 6.36: Scatterplot of the projection of electron momentum \vec{p} into the detector plane (electronPhi, on the y axis) vs. ϕ , the azimuthal angle defining the orientation of the main event line \vec{L} on the crystal plane (labeled as beamPhi on the x axis). In simulated signal events the two vectors are near-orthogonal. This is because the magnetic field orientation sampled by the electron while it illuminates the detector is near-vertically downward.

biased sample in that the electron always starts somewhere near the magnetic south pole. With this limitation in mind, plotting \vec{p} against Δ_p produces the plot shown in Figure 6.36. There it is apparent that the two vectors are in fact orthogonal to each other for the vast majority of events. This suggests attempting to estimate these two vectors, and requiring that they are perpendicular, as an additional signal seeking cut.

6.7.2 Estimating \vec{p} : An Example Event

I have described in several places the close relationship between the primary electron momentum, the momenta of its daughter synchrotron and Bremsstrahlung photons, and the momenta of secondary products of the synchrotron and Bremsstrahlung photons. Here I show how these momenta align in an example event. To adequately display these events in *EventDisplay* I augmented the tool to also display Monte Carlo truth data. In what follows I adopt the following color code to distinguish photons and secondary products of the three species from each other.

- The primary electron and its secondary interaction products are colored orange.
- Primary synchrotron photons and their secondary interaction products are colored

blue.

- Primary Bremsstrahlung photons and their secondary interaction products are colored green.

Figures 6.37 - 6.42 show various views of the same simulated 10 TeV signal event. Four synchrotron primaries, one Bremsstrahlung primary and the primary electron itself illuminate the detector plane. The momenta for all of these primaries is nearly identical, as drawn in Figure 6.2. These event displays show how, at least in this case, the bottom veto activity is located roughly where one would predict by extending the primary momentum through the crystal activity and down to the bottom veto system.

Any technique combining the crystal and veto hit locations would convolve errors from at least two sources. First, as shown above the bottom veto hit locations are only accurate to within approximately 30 cm. Second, there are errors in the crystal cluster locations as well (see Figure 6.14). In order to eliminate the effect of the crystal clustering algorithm, I decided to “cheat” by using the *true* crystal cluster locations instead of the *estimated* ones. In other words, at first I’m assuming a perfect clustering algorithm. This allows me to estimate what order pointing resolution is achievable given the finite precision bottom veto system, in isolation from errors in the crystal cluster location estimates. After settling on an algorithm, I then stop cheating and use the estimated crystal cluster locations as found by *Lineal* (as in Section 6.4.5.1).

I attempted various methods of reconstructing the primary momentum using the combined crystal cluster locations and bottom veto hit locations. After many failed attempts of increasing complexity (culminating with a maximum likelihood model which attempted to pair individual veto hits with individual crystal clusters), I decided to try the simplest method I could think of, as much a test of my sanity as anything else:

- Calculate the average of the crystal cluster locations, then
- Calculate the average bottom veto hit locations, then
- Connect both average locations with a line.

In Figures 6.41 and 6.42 I draw, the reconstructed momentum as the dashed teal line. This is computed simply by connecting the average locations detailed above. In this event, the procedure happened to reconstruct θ fairly well, but missed a bit on ϕ . Statistically, this method works far better than all other algorithms I had tested previously. I now characterize its pointing resolution in terms of the polar angle θ and the azimuthal angle ϕ .




Figure 6.37: A simulated 10 TeV signal event with five primary photons and the primary electron. Four primary synchrotron photons (blue disks) were generated, as well as one primary Bremsstrahlung photon (green disk) and the primary electron (orange disk). The primaries are drawn at the top ends of the red dashed lines indicating their trajectories through the detector. The crystal activity produced by these primaries are colored identically, so that for example all the green crystals were activated by secondary products of the Bremsstrahlung primary (or the Bremsstrahlung primary itself). The veto activity has been removed to make the primary and crystal activity more clear.

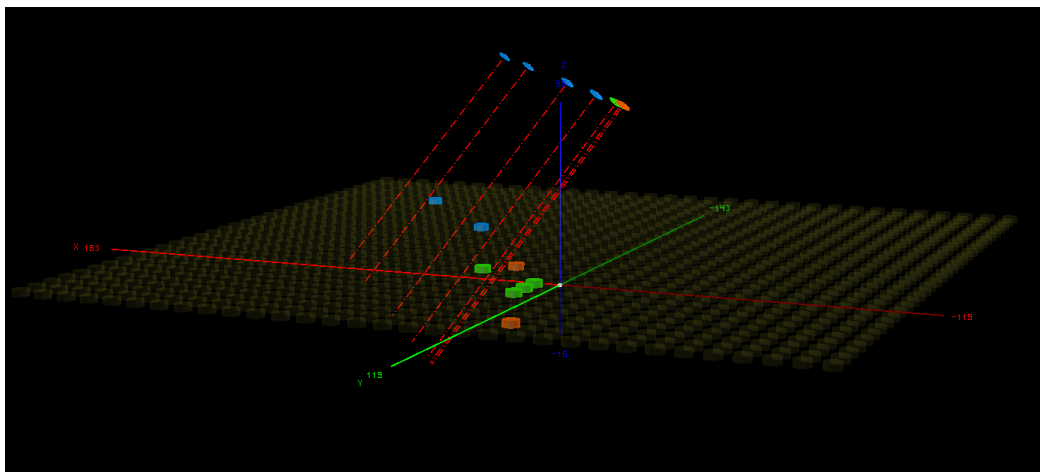


Figure 6.38: Another view of the same simulated signal event.




Figure 6.39: The event displays bottom veto activity offset from the crystal activity roughly in the direction of the primary (photon and electron) momentum.




Figure 6.40: Another view of the location of the bottom veto activity with respect to the crystal activity and the primary momentum.




Figure 6.41: To estimate the primary momentum I connect two points, the average crystal cluster location and the average non-top veto hit location, with a line. I draw that line here as the dashed teal line.






Figure 6.42: View showing that the estimate for the polar angle θ happened to be quite good in this event - the teal line (momentum reconstruction) and the red lines (true primary momentum) are nearly parallel in this projection.



(a)



(b)

Figure 6.43: Pointing resolution (when “cheating” by utilizing the truth-values for crystal cluster locations) for the polar (y-axis) and azimuthal (x-axis) angles of the primary electron momentum in the low-energy (a) and high-energy (b) simulated signal data sets. FWHM resolution is of order 20-30 degrees for θ and of order 40 degrees for ϕ .





Figure 6.44: Angular dependence of the pointing resolution when “cheating” by utilizing the truth-values for crystal cluster locations. (a): The distribution of errors in ϕ -pointing as a function of the true θ value. Very vertical primary trajectories ($\theta \approx \pi$ radians, right side of the plot) are likely to induce large errors in ϕ -pointing. (b): The distribution of errors in θ -pointing has no clear dependence on the true ϕ value.





Figure 6.45: ϕ -pointing resolution in low-energy (a) and high-energy (b) signal MC data sets as a function of true θ . When using estimated crystal cluster locations (red marks) rather than the true locations (blue marks), ϕ -pointing worsens for horizontal primary trajectories by only a slight amount. This indicates that while the crystal clustering algorithm is less accurate for horizontal primary trajectories (which tends to create larger clusters), the clusters tend to elongate in such a way that the ϕ estimate is undisturbed.

In Figure 6.43, I plot the distribution of angular pointing errors in the low-energy and high-energy data sets. Keep in mind that I'm using the true, rather than estimated, crystal cluster locations when making these plots. These distributions indicate that given perfect clustering, the bottom veto is accurate enough to point at θ with a resolution of approximately 20-30 degrees, and at ϕ with a resolution of approximately 40 degrees.

As might be guessed, the ϕ error distribution has a significant dependence on the true value of θ . Figure 6.44 shows that for vertical trajectories the ϕ estimate fails rather often. This makes sense, since for such trajectories small deflections to the direction of Compton recoil electrons and photons, for example, could outweigh the small transverse component of the primary momentum. Note that I still employ the true crystal cluster locations here.

In Figure 6.45 I explore what happens to the ϕ -pointing error distribution when I use estimated, rather than true, crystal cluster locations. While the errors are slightly increased overall there are no alarming regions of phase space (other than very vertical trajectories). Figure 6.46 shows that the same cannot be said of the θ -pointing error distribution. It has significant issues for vertical trajectories, but also experiences degraded resolution at horizontal trajectories when using estimated crystal cluster locations. This shows that the distortion in crystal cluster locations in events induced by horizontally-directed primaries affects the θ -estimate to a much greater degree than it does the ϕ -estimate.

In light of this limitation, and in the absence of a greatly improved clustering algorithm specializing in properly locating clusters in very horizontal events, there remains a "sweet-spot" in θ phase space, between approximately $\theta = 2.4$ and $\theta = 2.8$ radians off of vertical, where both θ and ϕ have their best pointing resolutions.

6.8 Outline of a Technique for Positron Detection

All of my event schematics thus far have assumed that the primary electron has negative charge. I have noted where applicable that identical schematics would hold for positrons, except that the direction of the magnetic force for a given trajectory would be in the opposite direction. Let us assume, for the moment, that we have perfect ϕ and θ pointing capabilities. Is charge sign discrimination possible in such an ideal case? With improved pointing, I argue that it is.

In Figure 6.48, I summarize my efforts to ensure proper orientation of the estimate for the main event line. This work is preliminary but is based around the idea that the terminal end of the line (the end the electron passes nearer to) is more likely to have air shower activity than the originating end of the line. Re-orienting the line so that it matches this template provides some improvement in the alignment of the main event line estimate. Note





Figure 6.46: θ -pointing resolution in low-energy (a) and high-energy (b) signal MC data sets as a function of the true θ . When using estimated crystal cluster locations (red marks) rather than the true locations (blue marks), θ -pointing worsens preferentially for very horizontal particle tracks. This contrasts with ϕ -pointing, which worsened only slightly for very horizontal tracks when not cheating. When clusters elongate due to horizontal primary trajectory, evidently they do so in a way that significantly disrupts θ -pointing.

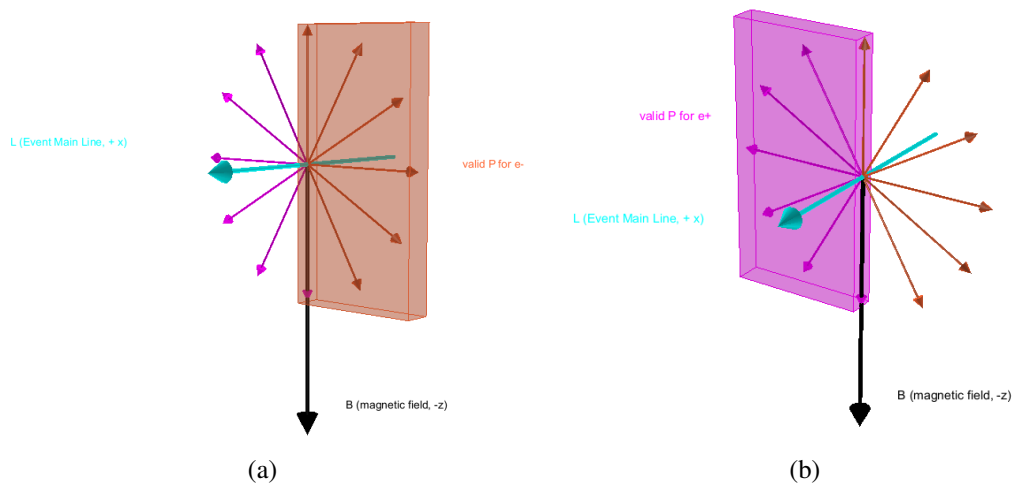


Figure 6.47: Two diagrams illustrating the regions of primary trajectory phase space that belong to signal electrons and positrons. With infinite-precision ϕ and θ pointing these diagrams show that charge sign discrimination would be possible with an improved CREST-like instrument by identifying which half of the potential signal phase space the primary lepton's momentum occupies. (a): Given a downward magnetic field and the teal main event line as shown here, any electron trajectory within the orange box would be a valid signal trajectory (that is, the magnetic force would point along the teal line). (b): Given a downward magnetic field and the teal main event line as shown here, any positron trajectory within the purple box would be a valid signal trajectory (that is, the magnetic force would point along the teal line).

that traversal-type events are easier to properly orient than non-traversal events, meaning that the class of events initially omitted from signal-seeking cuts may, in the end, be the only events for which this new measurement is viable.

Note that the analysis required to make use of the geometry shown in Figure 6.47 assumes that one can correctly identify the orientation of the main event line. In the Figure this corresponds to correctly establishing that the teal main event line points out of the page and not into the page. Preliminarily, I have achieved limited success in doing so by making use of both the location of top veto hit with respect to the two main line endpoints, and the “showeri-ness” of either end of the main event line. Since the electron passes nearer the detector on the terminal end of the main event line, that end is more likely to have air showers induced by Bremsstrahlung photons (which have a shorter mean free path against production in denser air) or by the electron itself. Thus one end of the line is more likely to be covered by a veto hit, and to have a large, long-duration cluster resulting from the by-products of the air shower.

6.9 Analysis Summary

In this Chapter, I have attempted to present CREST’s signal identification and background rejection strategies in a clear and complete manner. I have highlighted the pivotal role of the “Main Event Line.” The orientation of this line results from the geometric orientation of the primary signal electron, the geo-magnetic field and the CREST detector array. The role played by this analysis feature is analogous to that played by the shower axis in calorimetry-based detectors, or the particle track in tracking-based detectors. Due to relativistic beaming, “primary” signal photons are bound to simultaneously impact the crystal array along this line. The mainstream analysis utilizes this feature to attempt to single out clean signal events.

The dominant source of long, linear events in our flight was evidently side-going protons. While many of these protons are caught by the veto system, their great number allows some of the small fraction that evade veto detection to become the dominant event species in the population that meets most signal cuts. A key step forward in this line of attack was provided by my crystallographic cut, which attempts to identify these protons using the spectrum of gap between hits. This spectrum is purely determined by the angle-dependent lattice structure inherent to the crystal array. The stochastic nature of minimum ionizing particle energy deposition precludes energy-based techniques of eliminating this background. This line of analysis relies on one final cut statistic, inverse track velocity, to discriminate signal events from the remaining background events. I discuss this interplay

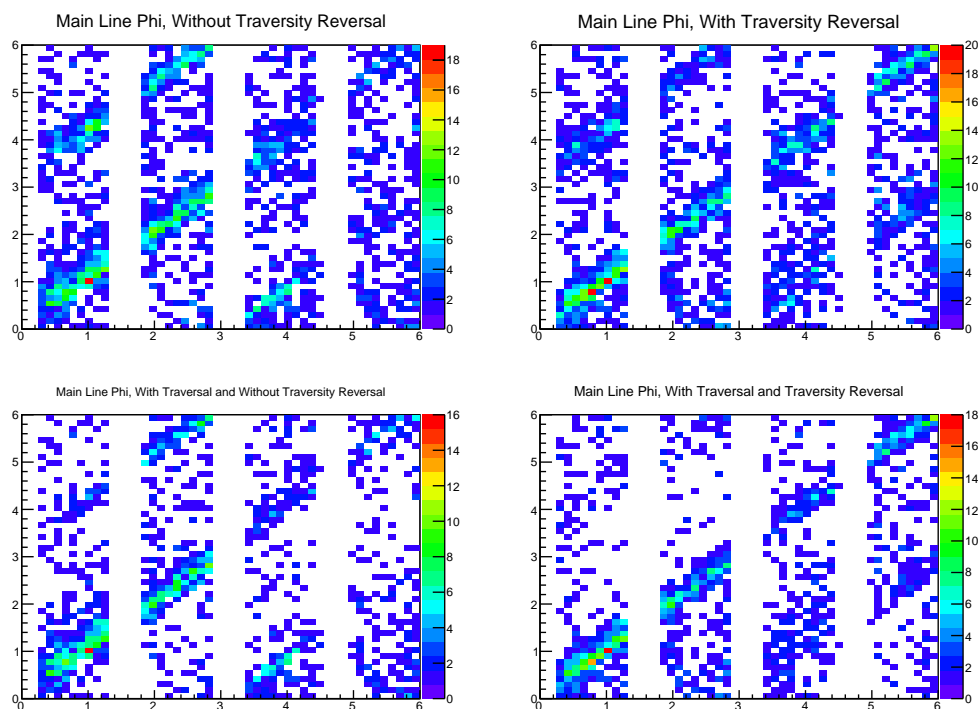


Figure 6.48: Four plots showing the results of my attempts to properly orient the main event line, \vec{L} , so that it points in the same direction as the magnetic force on the primary electron. Correctly oriented and aligned main event lines have data points on the main diagonal. Data points that are off the main diagonal represent line-misalignments (or line reversals). In the upper left I show the distribution produced by *Lineal*'s output, which uses the earlier line endpoint as the start of \vec{L} . In the upper right I apply a “traversy” reversal, which attempts to decide which end of the line looks like it more likely came from a shower-type event, such as from a nearby electron or high-energy Bremsstrahlung photon. The lower plots are identical, except that I restrict my attention to traversal-type events only. This suggests that traversal-type events are easier to properly orient, especially after applying the traversy reversal algorithm. This could mean that charge-sign discrimination will be viable for traversal-type events, and not for non-traversal type events.

further in Section 7.1.

I explored other means of signal/background separation, with the initial goal of achieving greatly increasing signal acceptance compared to that of the mainstream method. At first, I focused on relaxing the veto cut. When doing so, I observed that not only do signal events frequently activate the bottom veto sub-system, they do so in a predictable manner with respect to the momentum of the primary electron (i.e., on average, secondary scatters are directed parallel to the electron momentum). Due to the lack of a transverse force in side-going proton events, this vector is absent from those events. In proton background events, all secondary activity is oriented parallel to proton momentum, which is parallel to the main event line.

One of the reasons for the frequent, high-multiplicity veto activation in signal events is the close proximity of the primary electron to the detector plane during illumination. Using a simplified geometrical model, I estimated that most of the detectable events consist of photons produced while the electron is nearer than ~ 100 km to the detector plane. A large fraction (more than half, at primary energies below 10 TeV) of these detectable photons are generated in “traversal”-type events; that is, events where the electron itself passes through the detector plane and veto system. This fraction increases from approximately 59% at 10 TeV primary energy to over 80% at around 500 GeV. Many of the detectable signal events which are not traversal events are near-traversal events (i.e., the electron passes within a few meters of the detector). This realization led to a major improvement in our signal Monte Carlo simulations, and justified my focus on recovering events with morphology reflective of the electron’s traversal, or near-traversal.

This analysis arc led me to attempt to directly estimate the primary electron momentum direction. If successful, I can use its orientation with respect to the main event line (i.e., approximately perpendicular to it) as another means of signal/background separation. By combining the information available from both the crystal and veto systems, I devised a method to estimate the direction of primary electron momentum in Monte Carlo signal events. In those events, I find that CREST is capable of pointing at the momentum direction of the primary electron with a FWHM resolution of approximately 20-50 degrees, depending on various properties of the event. The pointing resolution depends most strongly on the zenith angle of the electron’s trajectory, with very horizontal and very vertical events resulting in larger azimuthal pointing errors.

I further observe that when both the main event line and electron momentum direction are known, it is possible (given sufficient information about the magnetic field orientation sampled by the electron) to discriminate between negatively-charged electrons and positively-charged positrons. Before verification of this technique is possible, incremental

algorithmic improvements in the determination of the polarity of the main event line are required; my early efforts along these lines show promise, but are as yet incomplete.

In the following (and final) Chapter, I discuss the potential implications of these findings, and present some potential detector configuration improvements that designers of future CREST-like detector may wish to incorporate.

CHAPTER 7

Conclusions

7.1 The Mainstream Analysis: Preliminary Results

Note: the following is preliminary analysis, and in no way do I claim this to be conclusively reflective of the CREST collaboration's as yet unpublished final results, which depend strongly on the production of additional Monte Carlo data.

While a final statement of the results of the mainstream analysis await completion of production of simulated background and signal events, I will outline the process here. This process consists of applying the following set of cuts to events from three data sources (flight data, the background Monte Carlo data, and the signal Monte Carlo data):

As explained in Chapter 6, the greatest proportion of signal events are eliminated by the veto activity prohibition (nearly 99%, with a greater proportion of signal events activating the veto system at lower energies). Essentially all “noisy” signal events (i.e., those with Bremsstrahlung or electron air shower clusters) which happen to avoid activating the veto system, are eliminated by the rather strict goodness-of-linear-fit cuts. A final signal acceptance value is still being estimated; it will likely be smaller than 10^{-5} , depending on the energy spectrum of the simulated signal electrons. This is because lower energy primary electrons tend to have a higher proportion of “shower-y” clusters.

Since this process only admits exceptionally linear signal events, it also preferentially selects exceptionally linear background events, i.e., side-going protons. Figure 7.1, pro-

Cut parameter	Accepted values
Active veto PMTs	0
Crystal clusters	> 3
Main event line length	> 75 cm
Largest gap between hits	> 40 cm
Spatial linear fit reduced χ_{xy}^2	< 5
Temporal linear fit reduced χ_{xt}^2	< 4

Table 7.1: “Mainstream” cuts used to produce distributions shown in Figure 7.1.

duced by J. Musser, shows *preliminary* distributions of the inverse track velocity in three data sets, all corresponding to the detector being operated in the “solid-half” mode. (I discuss the meaning of this statistic, and the procedure for generating it, in Section 6.1.2.) It is apparent from visual inspection that there is very little room, if any, for signal events. If present, they would show up near zero inverse track velocity in this plot. Given that other experiments have already detected electrons near 1 TeV of energy, this implies that this method has very little sensitivity to such electrons with such energies. The main reason for this is that most of these events are traversals, or near traversals, and the mainstream cuts preferentially extinguish those kinds of signal events (i.e., those with veto activity and large clusters).

By performing a best-fit regression using the two sets of simulated data, one can estimate the number of observed signal events, within a given confidence interval, integrated over all primary electron energies above a minimum detection threshold. The resulting flux limit (yet to be calculated) will depend on the estimated effective width of the instrument, the total live time, and the signal acceptance (i.e., the fraction of Monte Carlo signal events that survive the cuts). The low signal acceptance of this procedure will put an upward pressure on the flux limit.

7.2 Extending the Mainstream Analysis

I have argued in the above that it is possible to accept a greater fraction of the population of signal events against the random γ -ray and side-going charged particle background. In doing so, one introduces enormous complications related to the messy nature of secondary scattered hits and shower activity. I combat these complications by making use of signals in the veto system in concert with the main event line information provided by the response of the crystal array. In doing so, I realized that not only could one detect the presence of a primary signal electron’s downward momentum, one could directly estimate this vector. Doing so would provide an important next layer of signal-seeking cut, which could result in identification of a signal excess in CREST’s limited flight data. The degree to which this cut reduces the remaining background events depends strongly on the pointing resolution discussed in Section 6.7.

The key measure of success of this approach would be making contact with the previous measurements of the all-electron flux spectrum, especially those made by direct detection methods (i.e., AMS-02). Since such a hybrid direct-indirect detection approach is quite novel, showing that its results are consonant with proven methods would be an important anchor. Other experiments are underway to measure the all-electron spectrum beyond




Figure 7.1: Distributions of inverse track velocities for events from three data sets, after applying all “mainstream” cuts (listed in Table 7.1), with the instrument in the “solid-half” configuration: flight data (black), Monte Carlo background (red online) and simulated signal (blue online). The Monte Carlo data sets are still in production and have statistical bin errors approaching zero. The flight data set has statistical bin errors proportional to $1/\sqrt{N}$ (not drawn), and systematic bin errors related to the frequency of errors in estimating the inverse track velocity (not yet estimated as of the time of this writing). The exceptionally close match between the background Monte Carlo data and the flight data will likely place a strong upper limit on the estimated amount of *observed* signal events.

the current frontier. This includes the continued efforts of ground-based detectors such as HESS, and also CALET, a space-based direct-detection experiment [36]. CALET hopes to measure the all-electron spectrum up to approximately 20 TeV; should it detect a measurable flux in this energy regime, an indirect method such as that underlying CREST would be necessary to explore the spectrum's continuation at still higher energies.

Estimating the momentum of the primary electron in combination with the orientation of the main event line opens up an entirely new and exciting possibility: that of measuring not only the electron flux, but also the positron flux. Much work remains to positively demonstrate this capability. If successful, a result based on this technique this would be a major step forward. To my knowledge, no other method of performing this measurement at energies above a TeV has been conceived. CALET, the only proposed direct detection method that can probe energies above a TeV, will not be able to perform charge sign discrimination. Other direct-detection experiments (such as FERMI And AMS-02 can discriminate charge sign, but thus far only at energies below the TeV scale. If CREST has, or a future, enhanced CREST-like detector would have, this capability, it may well be the *only* practical way to make this measurement.

Such a measurement would be critical to the community's current efforts to characterize our Galactic neighborhood. Given the positron fraction data generated thus far, it is still unknown whether the electron flux at higher energies will be electron-dominated, or be approximately charge-neutral. As I related in Chapter 1, the former would indicate a nearby primary negative electron source, confirming much of the current theories about charged particle acceleration, escape and propagation from such Galactic neighbors as Vela. The latter possibility - a mixture of electrons and positrons - would favor the secondary production models, either from pulsars, WIMPs, secondary production by cosmic ray protons, or some other source. Even a crude measurement of the integral positron fraction, perhaps above some minimum detection energy (e.g. above approximately 5 TeV), would be of enormous utility in terms of helping to distinguish between these two classes of models. Probing the structure of the positron fraction as a function of energy would help distinguish among the particular models in each class. While it may yet be possible to achieve this with CREST's limited flight data, it is more likely that an improved instrument with a longer exposure time, and an improved detector configuration. In the following Section, I outline some key improvements that would benefit such a mission.

7.3 Ideas for a Future CREST-Like Detector: Defeating the Proton Background

Significant analysis challenges remain in implementing the above strategies for greater electron signal acceptance, electron momentum direction estimation and charge sign discrimination. CREST's main source of background is side-going protons, and the flux of these protons is so much higher than the predicted signal flux that even a 99%-efficient veto system, on its own, provides insufficient background rejection. Therefore, it is desirable for a future detector to have limited sensitivity to such events by instrument design, rather than relying on more analysis techniques of ever-increasing complexity.

One area of improvement lies in the geometrical arrangement of the veto system. While the top and bottom veto paddles ended up providing very useful information about event morphology, the bulk of the work in rejecting charged particles was performed by the side paddles. This was because protons capable of making long tracks in the crystal plane arrived from a very small section of the sky, within a degree or so of horizontal. Improving the efficiency of the side vetos, perhaps simply by adding a second layer (affectionately, yet aptly referred to by G. Tarlé as a "belly band"), would greatly aid positive detection-based proton rejection.

But short of simply throwing more material at the problem, what can be done to ameliorate the issue in other ways? What is it about CREST's crystal plane that makes it so susceptible to long, linear tracks produced by side-going protons? The answer is, its flatness. If CREST's crystals were not vertically aligned with one another, then the lengths of tracks made by horizontal protons would be inherently limited by the structure of the crystal array. One possibility is to stagger the crystals in height in a regular pattern; this would give side-going protons a signature "on-off" pattern as the proton entered and exited crystals of various heights. However, proton tracks *already* display a signature pattern of gaps and hits as discussed in Section 6.4.5.3, and the analysis required to make use of this fact is not trivial. Furthermore, the probability that signal events would randomly emulate this behavior is non-zero.

An alternative proposal, which I outline here, is to geometrically deform the crystal plane from a flat plane into a bowl shape, as shown in Figure 7.2. For a bowl-shaped crystal surface, the maximum length track a side-going particle could trace through consecutive crystals is determined by the height of the crystals and the curvature of the bowl. The tighter the curvature, the less distance a given proton can spend inside the crystals before exiting the curving crystal surface. Signal electron events illuminating the bowl from above would still enjoy long track lengths; the straight line would be warped into a curved




Figure 7.2: Four viewpoints of my proposed “Crystal Bowl” configuration for a future CREST-like detector. The bowl shape provides inherent insensitivity to sidegoing protons by shortening their track length by a factor proportional to the curvature of the bowl. Signal electrons illuminating the bowl from above would still be able to produce long “main event lines.” Due to the curvature of the bowl these would have a trajectory-dependent shape rather than simply being linear. To generate the displayed curvature I calculated the z-offset as $\Delta_z = \sqrt{R^2 - r^2}$, where $R = 350$ cm and r is the transverse distance, in cm, from the center of the flat crystal plane to any given crystal disk’s center (before being vertically offset).

trajectory in a predictable way, based on the orientation of the primary electron momentum and the azimuthal angle of the main event line projected onto the bowl. This warping could even provide another way to estimate the momentum direction of the primary electron and differentiate the signal hit pattern from noise patterns. The optimal curvature would be that which best balances the benefits of proton track shortening against the drawbacks of signal electron line warping and instrument width reduction, as determined by Monte Carlo simulation. The degree of curvature would also have to be balanced against the unique mechanical and electrical engineering demands such an unconventional arrangement would entail.

If this “Crystal Bowl” (or similar) configuration allows one to eliminate proton noise events by simply cutting out events with short track lengths, then there may be no need for the veto system to eliminate side-going protons. Indeed, the notion of a “side veto” loses meaning when the detector has no well-defined sides in the first place. The weight and complexity savings thereof could be repurposed to building an improved second-layer detector plane below the main crystal surface to enable enhanced pointing resolution.

One drawback of the bowl configuration is that the physical width of the detector when viewed from any angle would be slightly smaller than for a flat plane. This could be compensated for by a slightly larger inter-crystal spacing. If feasible to construct, the built-in rejection of protons offered by such a bowl configuration would greatly benefit a future CREST-like instrument or any other instrument wishing to decrease sensitivity to background events caused by charged particles moving on a straight line.

7.4 Final Thoughts: Looking Ahead, by Looking Nearby

The unknown nature of the cosmic ray electron and positron spectra at energies above a TeV represents one of the few unexplored frontiers in cosmic ray physics. As I have related above, the favored theories describing a multitude of Galactic phenomena, ranging from remnants of exploded stars, to pulsars, to the elusive Dark Matter theorized to permeate our and all galaxies, to something as mundane as secondary production from other, more well-measured cosmic ray species, all predict in common a detectable flux of electrons and positrons at the TeV scale and beyond. This makes measurement of this region of the flux spectrum an exciting prospect for proponents of each of these theories; no matter the result, the impacts will be profound.

Measuring this energy region of the electron and positron flux spectra is reminiscent to performing astronomy with a previously un-measured region of the electromagnetic radiation spectrum. Except that in this case, the origins of the signal-carrying particles are

confined to reside within our immediate galactic neighborhood. This also means more mundane, but no less important, theories about how charged particles propagate through the ISM will be put to the test. By contrast, the analogous situation for protons, as determined by the GZK limit, has a far larger envelope of influence, approximately the size of the Galaxy. For lack of a better analogy, today's cosmic ray detectors suffer from being farsighted.

Whatever the results of the mainstream analysis, and my extended analysis, I am confident that CREST will have served an important pathfinding role. It is not an iterative improvement over previous direct and indirect methods; indeed it appears as if these methods have reached their zenith, the former being limited by exposure time (the ISS' lifetime will not be measured in millenia), and the latter being limited by γ -ray contamination. Rather, CREST has forged a new path upward to higher energies by combining the best aspects of both approaches.

No one knew what wonders awaited the advancements of x-ray and radio astronomy. No one could have predicted the insights Victor Hess' pioneering balloon experiments, and the flood of fantastically bountiful cosmic ray-detection experiments that followed in their wake, would bestow on the physics community. It is tempting to give in to the common belief that most of the information available in the cosmic rays has already been discovered, and that we have only to measure the last few bits to gain a complete understanding thereof. From relativity, to the muon, to the still-mysterious Dark sector of cosmology, history shows that Nature invariably responds to such hubris with the presentation of unexpected phenomena. Whether it is CREST or another future experiment that measures the high energy electron and positron flux spectra, the results will surprise us.

BIBLIOGRAPHY

- [1] V. Ptuskin A. Strong, I. Moskalenko, *Cosmic-ray propagation and interactions in the galaxy*, *Ann.Rev.Nucl.Part.Sci.* **57** (2007), 285–327.
- [2] M. Aguilar, *AMS-02 provides a precise measure of cosmic rays*, *CERN Courier* **9** (2013).
- [3] S.P. Ahlen, *Energy loss of relativistic heavily ionizing particles*, *Rev. Modern Phys.* **52** (1980), 121–173.
- [4] S.P. Ahlen and G. Tarlé, *Secondary production as the origin for the cosmic ray positron excess*, In preparation **TBD** (2015), TBD.
- [5] C. Akerlof, *CREST Calculations*, personal communication, 2009.
- [6] D. Allard, E. Parizot, A.V. Olinto, E. Khan, and S. Goriely, *UHE nuclei propagation and the interpretation of the ankle in the cosmic-ray spectrum*, *Astron. Astrophys.* **443** (2005), L29–L32.
- [7] G.E. Allen, R. Petre, and E.V. Gotthelf, *X-ray synchrotron emission from 10-100 TeV cosmic-ray electrons in the supernova remnant sn 1006*, *Astrophysical Journal* **558** (2001), 739–752.
- [8] J. Ameel, C. Bower, S. Coutu, M. Gebhard, J. Gennaro, M. Geske, M. Lang, D. Muller, J. Musser, R. Northrop, S. Nutter, N.H. Park, M. Schubnell, A. Shroyer, C. Smith, G. Tarlé, S.P. Wakely, and A. Yagi, *The cosmic ray electron synchrotron telescope*, In preparation **TBD** (2015), TBD.
- [9] E. Aprile, A. Curioni, K.L. Giboni, M. Kobayashi, U.G. Oberlack, and S. Zhang, *Compton imaging of mev gamma-rays with the liquid xenon gamma-ray imaging telescope (LXeGRIT)*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **593** (2008), no. 3, 414 – 425.
- [10] A.M. Atoyan, F.A. Aharonian, and H.J. Völk, *Constraining dark matter models from a combined analysis of milky way satellites with the fermi large area telescope*, *PRL* **107** (2011), 241302.
- [11] F.H. Attix, *Introduction to radiological physics and dosimetry*, first ed., Wiley, 199.

- [12] A.R. Bell, *Turbulent amplification of magnetic field and diffusive shock acceleration of cosmic rays*, Mon. Not. R. Astron. Soc. **353** (2004), 550–558.
- [13] D.R. Bergman, *Observation of the GZK cutoff by the hires experiment*, ICRC **4** (2007), 451–454.
- [14] G. Bertone, D. Hooper, and J. Silk, *Particle dark matter: evidence, candidates and constraints*, Physics Reports **405** (2004), 279–390.
- [15] H. Bichsel, *Straggling in thin silicon detectors*, RMP **60** (1988), 663,699.
- [16] R. Blandford and D. Eichler, *Particle acceleration at astrophysical shocks: a theory of cosmic ray origin*, Physics Reports (Review Section of Physics Letters) **54** (1987), 1–75.
- [17] X. Chi, K.S. Cheng, and E.C.M. Young, *Pulsar-wind origin of cosmic-ray positrons*, Astrophys. J. **459** (1996), L83–L86.
- [18] M. Cirelli, M. Kadastik, M. Raidal, and A. Strumia, *Model-independent implications of the all-electron and antiproton cosmic ray spectra on properties of dark matter (updated including AMS 2013 data)*, Nuc. Phys. B. **813** (2009), 530–533.
- [19] L. O’C. Drury, *An introduction to the theory of diffusive shock acceleration of energetic particles in tenuous plasmas*, Rep. Prog. Phys. **46** (1983), 973–1027.
- [20] L. O’C. Drury, G.M. Webb, and P. Biermann, *Diffusive shock acceleration of energetic electrons subject to synchrotron losses*, Astron. Astrophys. **137** (1984), 185–201.
- [21] A.D. Panov et al, *Possible structure in the cosmic ray electron spectrum measured by the ATIC – 2 and ATIC – 4 experiments*, Astrophys. Space Sci. Trans. **7** (2011), 119–124.
- [22] Beatty et al, *New measurement of the cosmic-ray positron fraction from 5 to 15 gev*, PRL **92** (2004), 241102.
- [23] E. Komatsu et al, *Seven-year wilkinson microwave anisotropy probe (wmap) observations: Cosmological interpretation*, Astrophys.J.Suppl. **192** (2011), 18.
- [24] F. Aharonian et al, *First detection of a vhe gamma-ray spectral maximum from a cosmic source: Hess discovery of the vela x nebula*, Astron. Astrophys. **448** (2006), L43–L47.
- [25] ———, *Energy spectrum of cosmic-ray electrons at tev energies*, PRL **101** (2008).
- [26] L. Accardo et al, *High statistics measurement of the positron fraction in primary cosmic rays of 0.5500 gev with the alpha magnetic spectrometer on the international space station*, PRL **113** (2014), 121101.
- [27] M. Ackermann et al, *Fermi LAT observations of cosmic-ray electrons from 7gev to 1tev*, Phys. Rev. D **82** (2010).

- [28] ———, *Measurement of separate cosmic-ray electron and positron spectra with the fermi large area telescope*, PRL **108** (2010), 011103.
- [29] ———, *Constraining dark matter models from a combined analysis of milky way satellites with the fermi large area telescope*, PRL **107** (2011), 241302.
- [30] M. Boudad et al, *A new look at the cosmic ray positron fraction*, Astron. Astrophys. **In preparation**.
- [31] M. Vogelsberger et al, *Phase-space structure in the local dark matter distribution and its signature in direct detection experiments*, Mon. Not. R. Astron. Soc. **395** (2009), 797–811.
- [32] M.A. DuVernois et al, *Cosmic-ray electrons and positrons from 1 to 100 gev: Measurements with HEAT and their interpretation*, Astrophys. J. **559** (2001), 296–303.
- [33] M.J. Berger et al, *XCOM: Photon cross sections database*, www.nist.gov, 1998.
- [34] O. Adriani et al, *Cosmic-ray electron flux measured by the PAMELA experiment between 1 and 625 gev*, PRL **106** (2011).
- [35] ———, *Cosmic-ray positron energy spectrum measured by pamela*, PRL **111** (2013), 081102.
- [36] ———, *Status and performance of the CALorimetric Electron Telescope (CALET) on the International Space Station*, Nuclear Physics B - Proceedings Supplements **256-257** (2014), 225–232, Cosmic Ray Origin: Beyond the Standard Models.
- [37] S.W. Barwick et al, *Measurements of the cosmic-ray positron fraction from 1 to 50 gev*, Astrophys. J. **482** (1997), L191–L194.
- [38] T. Kobayashi et al, *The most likely sources of high-energy cosmic-ray electrons in supernova remnants*, Astroph. J. **601** (2004), 340–351.
- [39] K.A. Olive et al (Particle Data Group), *Review of particle physics*, Chinese Phys. C **38** (2014), no. 9.
- [40] J.L. Feng, M. Kaplinghat, and H.B. Yu, *Sommerfeld enhancements for thermal relic dark matter*, Phys. Rev. D **82** (2010), 083525.
- [41] E. Fermi, *On the origin of the cosmic radiation*, Physical Review **75** (1948), 1169–1174.
- [42] J.F. Beacom G. Steigman, B. Dasgupta, *Precise relic wimp abundance and its impact on searches for dark matter annihilation*, Phys. Rev. D **86** (2012), 023506.
- [43] J. Gennaro, *CREST Digitize Decision Manual*, technical manual, 2011.
- [44] ———, *CREST Manual*, technical manual, 2011.

- [45] ———, *Updates from McMurdo*, email, 2011.
- [46] E.V. Gotthelf and G. Vasisht, *Kes 73: A young supernova remnant with an x-ray-bright, radio-quiet central source*, *ApJ* **486** (1997), L133.
- [47] K. Greisen, *End to the cosmic-ray spectrum?*, *Physical Review Letters* **16** (1966), 748–750.
- [48] A.M. Hillas, *The origin of ultra-high-energy cosmic rays*, *Ann. Rev. Astron. Astrophys.* **22** (1984), 425–444.
- [49] ———, *Can diffusive shock acceleration in supernova remnants account for high-energy galactic cosmic rays?*, *J. Phys. G: Nucl. Part. Phys.* **31** (2005), R95–R131.
- [50] D. Hooper, P. Blasi, and P.D. Serpico, *Pulsars as the sources of high energy cosmic ray positrons*, *J. Cosmo. Astro. Phys.* **JCAP012009** (2009).
- [51] J.D. Jackson, *Classical electrodynamics*, third ed., Wiley, 1999.
- [52] G.F. Knoll, *Radiation detection and measurement*, third ed., John Wiley & Sons, 2000.
- [53] T. Linden and S. Profumo, *Probing the pulsar origin of the anomalous positron fraction with AMS-02 and atmospheric cherenkov telescopes*, *Astrophys. J.* **772** (2013), 18–26.
- [54] A. Lopez, C. Savage, D. Spolyar, and D.Q. Adams, *Fermi/lat observations of dwarf galaxies highly constrain a dark matter interpretation of excess positrons seen in AMS-02, HEAT, and PAMELA*.
- [55] I.V. Moskaleno and A.W. Strong, *Production and propagation of cosmic-ray positrons and electrons*, *Astrophys. J.* **493** (1998), 694–707.
- [56] D. Müller, *Direct observations of galactic cosmic rays*, *Eur. Phys. J. H* **37** (2012), 413–458.
- [57] Y. Ohira, R. Yamazaki, N. Kawanaka, and K. Ioka, *Escape of cosmic-ray electrons from supernova remnants*, *Mon. Not. R. Astron. Soc.* **427** (2012), 91–102.
- [58] S.P. Reynolds and J.W. Keohane, *Maximum energies of shock-accelerated electrons in young shell supernova remnants*, *Astrophys. J.* **525** (1999), 368–374.
- [59] Stephan Rosswog and Marcus Brüggen, *Introduction to high-energy astrophysics*, first ed., Cambridge University Press, 2007.
- [60] R. Rothenflug, J. Ballet, G. Dubner, E. Giacani, A. Decourchelle, and P. Ferrando, *Geometry of the non-thermal emission in sn 1006*, *Astron. Astrophys.* **425** (2004), 121–131.

- [61] M.A. Ruderman and P.G. Sutherland, *Theory of pulsars: Polar gaps, sparks and coherent microwave radiation*, *Astrophys. J.* **196** (1975), 51–72.
- [62] S.M. Seltzer and M.J. Berger, *Improved procedure for calculating the collision stopping power of elements and compounds for electrons and positrons*, *Int. J. Appl. Radiat. Isot.* **35** (1984), 665–676.
- [63] S.A. Stephens and V.K. Balasubrahmanyam, *Earth's magnetic field as a radiator to detect cosmic ray electrons of energy $> 10^{12}$ eV*, *J. Geophys. R.* **88** (1983), 7811–7822.
- [64] S.P. Wakely, *CREST Power System Manual*, technical manual, 2011.
- [65] A. Wightman, *Note on polarization effects in Compton scattering*, *Physical Review* **74** (1948), 1813–1817.
- [66] A. Yagi, *Prototype development and simulation of the Cosmic-Ray Electron Synchrotron Telescope (CREST)*, Ph.D. thesis, 2009.