A SPECTRAL STUDY OF THE SOFT X-RAY BACKGROUND WITH THE X-RAY QUANTUM CALORIMETER

by

Dallas W. Wulf

A dissertation submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

(Physics)

at the

University of Wisconsin–Madison

2018

Date of final oral examination: 12/20/2018

The dissertation is approved by the following members of the Final Oral Committee:

Dan McCammon, Professor, Physics
Peter Timbie, Professor, Physics
Sebastian Heinz, Professor, Astronomy
Daniel Den Hartog, Scientist, Physics

Abstract

High spectral resolution observations in the soft X-rays are necessary for understanding the hot-ionized component of the interstellar medium and its contribution to the Diffuse X-ray Background (DXRB). This extended source emission cannot be resolved with most wavelength dispersive spectrometers, making energy dispersive microcalorimeters (with 20–100× better spectral resolution than solid state detectors) the ideal choice for these observations. I present here a combined analysis of four sounding rocket observations of the DXRB, using the University of Wisconsin-Madison/Goddard Space Flight Center X-ray Quantum Calorimeter (XQC). XQC is a large area silicon thermistor microcalorimeter with energy resolution of 6–9 eV FWHM below 1 keV.

The spectra support the existence of a ∼0.095 keV Local Hot Bubble and a ∼0.2 keV Hot Halo. At low galactic latitudes, an additional component of 3/4 keV emission can be consistently explained by unresolved dM stars. Discrepancies between repeated observations are compatible with expected contributions of time-variable emission from Solar Wind Charge Exchange (SWCX).

Additionally, I present an algorithm for processing microcalorimeter data which was used for three of the sounding rocket flights. I demonstrate that this technique can significantly increase detector throughput without sacrificing energy resolution relative to conventional processing techniques. A considerable effort has been made to generalize the implementation of this algorithm in software, which makes it easily adaptable to other microcalorimter data sets.

Acknowledgments

This dissertation would not have been possible without the support and guidance of many people. To them I owe a debt of gratitude. To my adviser, Dan, who has taught me so much about everything from cryostats to cobras. To Kelsey, who showed me the ropes. I couldn't have asked for a better role model as a graduate student. To Felix, who seems to have solutions for everything. Thank you for your patience and help. To Ambarish and Avirup, for taking the helm of XQC. I wish you the best in the Australia expedition. To Rachel, Nick, and all the undergraduates that I have had the privilege to work with over the years. The lab would be lost without you. And to everyone in the SiCal collaboration, whose contributions over a quarter-century are responsible for XQC. I stand on your shoulders.

I have also been fortunate to have the love and support of friends and family. In ways big and small, but always much needed and appreciated, they have helped me along the way. In particular, I am grateful for my parents, who have encouraged me to follow my curiosity for as long as I can remember. And for my bothers, Justin and Travis, who have always been eager to help me in any way they can. Finally, I thank Jill, who helped me make Madison our home.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 The Interstellar Medium

The Interstellar Medium, or ISM, is an astronomer's shorthand for all the gas, dust, radiation, and magnetic fields that fill the space between the stars in our Galaxy. To put it another way, if we were divide everything we can see in our Galaxy into two groups, those groups would be "stars"—those brilliant, spheroidal furnaces speckling the visible night sky—and "not-stars," which is to say, everything else. It is this latter group, the "not-stars," that astronomers are referring to when they talk about the ISM. As with many astronomical classification systems, this came about for historical, rather than logical reasons, and as our understanding of our Galaxy has grown, so too has our knowledge of what constitutes the ISM.

The ISM plays a vital role in a range of Galactic processes, and is essential to our understanding of the evolution of our Galaxy. Perhaps most notable is the interaction of the ISM with stars—those favorite sons of astronomy. After all, it is the inhomogeneities in the density and temperature of the ISM that give rise to initial star formation. Stars

in turn heat and enrich the surrounding interstellar material throughout their life-cycles through radiation, stellar winds, and ultimately supernovae ejecta, fundamentally changing the composition and distribution of the very ISM that bore them and which will bear the following generation of stars. It is only through this cycle that we can understand the origin of present-day stars and, in a very real sense, our own origin. Carl Sagan may have famously said, "We are made of starstuff," but before all that carbon, nitrogen, oxygen, calcium, and iron could become part of our physical makeup, it had to get out of a star and take its turn as "not-starstuff" in the ISM.

To aid in the discussion of this important, but necessarily broad and nebulous topic, it is helpful to further divide the ISM into components. A fairly natural way to do this is based on temperature and the corresponding state of hydrogen, given that hydrogen accounts the majority of interstellar matter. At temperatures of 10–20 K, hydrogen is in its molecular form, H_2. This so-called "cold-molecular" component of the ISM can be observed directly through ultraviolet absorption, or indirectly through co-spatial tracers such as CO at 115 GHz (Figure 1.1a) or dust emission in the infrared. Between 100 K and 10^4 K, hydrogen can be observed in its atomic form, either through its 21 cm hyperfine transition (Figure 1.1b) or its ultraviolet Ly-α transition, constituting the "cold-atomic" and "warm-atomic" components. Above $\sim 10^4$ K, hydrogen becomes ionized. Near this temperature, recombining electrons in the "warm-ionized" component of the ISM can produce optical Balmer lines such as Hα (Figure 1.1c). At still hotter temperatures, hydrogen becomes fully-ionized and does not emit or absorb efficiently, so we must look for line emission from partially-ionized metals. Metals can be observed in the ultraviolet up to temperatures of $\sim 5 \times 10^5$ K, and in the X-rays up to $\sim 10^8$ K. It is this "hot-ionized" component that primarily concerns this work, and the associated X-ray emission that will be discussed in the following section.

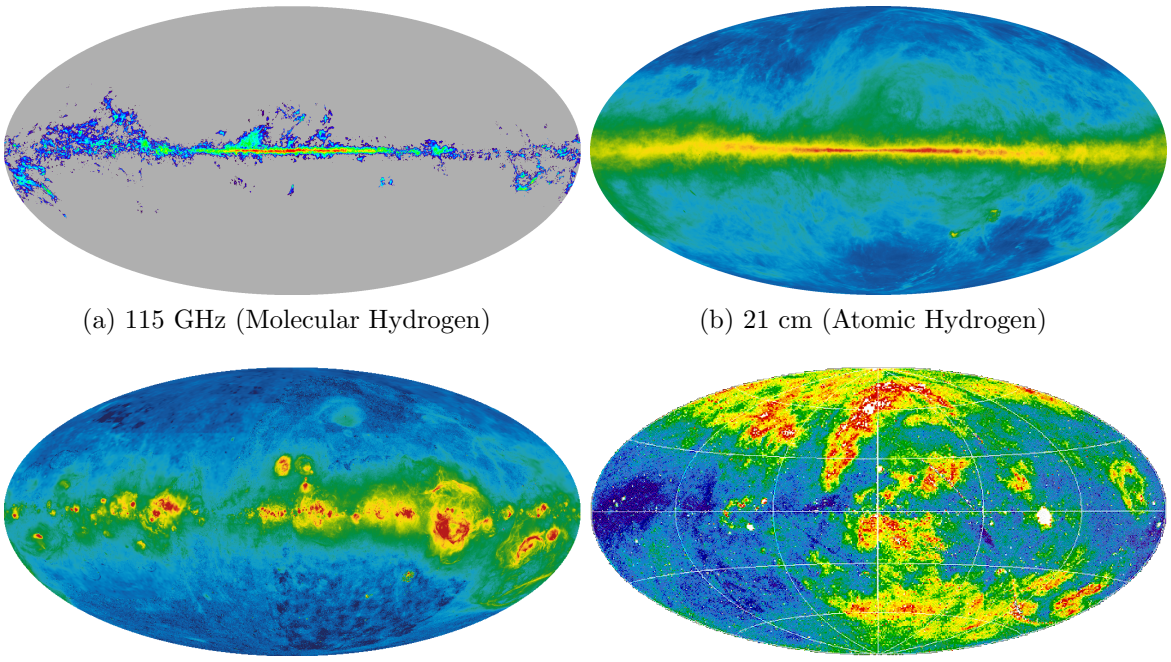

(a) 115 GHz (Molecular Hydrogen)

(b) 21 cm (Atomic Hydrogen)

(c) Hα (Ionized Hydrogen)

(d) < 2 keV ($> 10^6$ K Gas)

Figure 1.1: All-sky maps of diffuse interstellar emission. Maps are in Galactic coordinates with the plane of the Galaxy located along the equator and the center of the Galaxy at the center of the map. (a) 115 GHz emission from the CO J $1 \rightarrow 0$ rotational transition tracing the cold-molecular medium (Dame et al., 2001). (b) 21cm emission from the hyperfine transition in atomic hydrogen (Bekhti et al., 2016). (c) Hα emission produced by recombining electrons in the warm-ionized medium (Finkbeiner, 2003). (d) Composite soft X-ray emission below 2 keV, partly due to the hot-ionized medium (Snowden et al., 1997). Images a–c made available by the LAMBDA collaboration, d by the ROSAT collaboration and MPE.

1.2 The Diffuse X-ray Background

The Diffuse X-ray Backgound, or DXRB, was discovered in 1962 by a sounding rocket experiment launched from White Sands Missile Range in New Mexico. The experiment, sensitive to X-rays in the 2–6 keV bandpass, was intended to observe X-ray fluorescence of the lunar surface in anticipation of the Apollo missions. Instead, it observed a bright source in the constellation of Scorpius (what would later be identified as Sco X-1, the brightest X-ray source aside from the Sun), and a nearly isotropic difuse background (Giacconi et al., 1962). The isotropy of the background, combined with the transparency of the Galaxy at these energies, suggested an extragalactic source. A series of followup experiments at higher energies confirmed this (e.g. Bowyer et al. (1964) & Metzger et al. (1964)), making it the first observation of a cosmological diffuse flux—predating Penzias and Wilson's discovery of the Cosmic Microwave Background by two years.

This discovery motivated the development of instruments that could probe lower energies, where this emission could be used to "X-ray" the Galaxy. Due to larger absorption cross-sections at lower energies, the observed flux was expected to fall off at low Galactic latitudes with increased column depth. When the measurements were made however, the results were well above an extrapolation of the higher energy component, and it was later shown that the low energy flux did not go to zero at low Galactic latitudes where the ISM was entirely opaque. This meant that the low energy flux had to be originating locally within the ISM. The exact origin remained a mystery, however, since limits on interstellar magnetic fields, electron energy densities, and gamma ray fluxes ruled out known non-thermal mechanisms for X-ray production, and the ISM was thought to be much too cold to produce X-rays by thermal emission.

Everything changed in the mid-1970s, when the coincidence of theoretical and obser-

vational discoveries advanced the argument for an extensive hot component of the ISM. Cox & Smith (1974) reasoned that successive super nova remnants could overlap, forming large ($\gtrsim$ 100 pc) networks of clusters and tunnels throughout the disk of the Galaxy, displacing the cooler ISM with shock-heated gas. Such networks could account for a significant volume fraction of the ISM and potentially play an important role in Galactic dynamics. About the same time, Henry et al. (1976) discovered a void in the solar neighborhood that contained no neutral hydrogen and Jenkins & Meloy (1974) observed OVI absorption lines in stellar spectra characteristic of gas above 10^5 K. Assuming the same gas was responsible for the unexplained X-ray emission implied that the low-energy DXRB could be used as an additional observational tool to study this newly discovered component of the ISM.

1.2.1 An Observational Handle on the ISM

At the temperatures of the hot ISM ($\sim$ 1–4$\times 10^6$ K), most of the emission is in lines due to collisional excitation of highly-charged ions. This emission is also the primary cooling mechanism for the gas. Roughly speaking, emission around 0.25 keV is associated with $\sim 1 \times 10^6$ K plasma, while 2–4 $\times 10^6$ K plasmas emit most strongly from 0.5–1.0 keV. Which exact lines are emitted depends on a number of factors, including the electron temperature, ionization balance, and metallicity of the gas. Consequently, observations with sufficient spectral resolution to identify individual lines can provide a wealth of information about this component of the ISM.

The primary extinction mechanism for X-rays below 10 keV is photoelectric absorption. As the cross-section for this kind of interaction rises steeply with decreasing photon energy as $\sim E^{-3}$, photoelectric absorption causes hardening of observed X-ray spectra (in contrast to reddening caused by extinction in the optical). This steep energy dependence

also limits the average interstellar distances that can be probed in the the soft X-rays, with the mean free path for 0.25 keV X-rays being a mere 70 pc through an average interstellar density of $n_H = 1$ cm^{-3} (i.e. one optical depth at $N_H = 2 \times 10^{20}$ cm^{-2}). As a result, much of the volume of the Galactic disk is not directly observable in emission at these energies, though by 0.75 keV, the mean free path has grown to $\sim$1 kpc and one can easily see into the Galactic halo at latitudes above $\sim 8°$.

1.2.2 Origins of the DXRB

As already mentioned, the DXRB above 1 keV is isotropic and primarily of extragalactic origin. Early work found excellent agreement between the observed power-law spectrum and bremsstrahlung emission from a 40 keV plasma, supporting the theory of a hot intergalactic medium (e.g. Cowsik & Kobetich (1972) & Marshall et al. (1980)). However, further investigation revealed this to be entirely coincidental, and the true origin of this "diffuse" emission to be the superposition of unresolved distant accreting supermassive black holes, or Active Galactic Nuclei (AGN). In fact, this particular origin of the high-energy component of the DXRB had been theorized shortly after Giacconi's discovery (Gould & Burbidge, 1963), but the instrumental sensitivity required for observational verification was only achieved in the current century. To date, the Chandra deep fields have resolved at least 80% of the 2–8 keV intensity as discrete sources (Hickox & Markevitch, 2006).

At energies below 1 keV, interstellar absorption becomes important and local emission sources begin to dominate. Even at high Galactic latitudes where absorption is minimized, an extrapolation of the extragalactic power-law only accounts for 55% of the observed emission at 3/4 keV (McCammon et al., 1983). Moreover, high-resolution spectra from this part of the sky reveal 32% of the total flux to be in oxygen lines alone, and thus

distinct from the extragalactic continuum (McCammon et al., 2002). A prevailing theory attributes the line emission to a 2–3 MK hot Galactic halo,[1] though this can only be a partial explanation. High-latitude shadowing experiments have found no evidence for Galactic emission beyond 100 pc in some directions, suggesting the halo must be patchy in nature (e.g. Lei et al. (2009) & Liu et al. (2016)). Moreover, the 3/4 keV flux does not noticeably drop over much of the plane of the Galaxy, despite the fact that all extragalactic and circumgalactic emission is completely absorbed at these latitudes (Figure 1.2b). Other possible sources of 3/4 keV emission include dwarf M stars (e.g. Rosner et al. (1981) & Masui et al. (2009)) and superbubbles throughout the Galactic ISM, though the details of how these components combine to effectively cancel out the expected latitude dependence are poorly understood (McCammon et al., 1983).

In contrast to the 3/4 keV band, the sky at 1/4 keV exhibits a pronounced latitude dependence, with emission highest in patches at high latitudes that show a strong anti-correlation with the HI distribution (Figure 1.2a). Qualitatively, one would expect to see this pattern for an absorbed, distant source such as a Galactic halo (albeit likely a cooler component than is seen at 3/4 keV), and indeed shadows cast by the Draco Nebula imply 1/4 keV emission originating more than 300 pc above the plane (Snowden et al. (1991) & Burrows & Mendenhall (1991)). However, such a simple model fails to accurately describe the observations quantitatively across the entire sky (e.g. near the plane of the Galaxy, the ISM is over 100 optical depths at these energies, and yet the observed flux is still $\sim$1/3 of what is observed at high latitudes). Moreover, experiments using C and B absorption edges to divide the 1/4 keV band into higher- and lower-energy components do not observe the energy-dependence predicted by photoelectric absorption (Sanders et al., 1977). These

[1] In fact, two decades before the discovery of the hot component of the ISM, Spitzer (1956) hypothesized the existence of a Galactic halo as a means of providing pressure support to clouds observed at high Galactic latitudes.

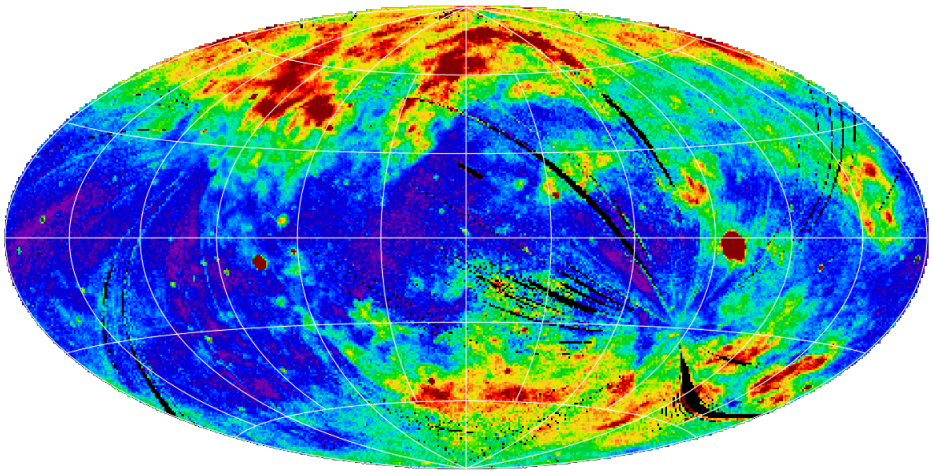

(a) 1/4 keV

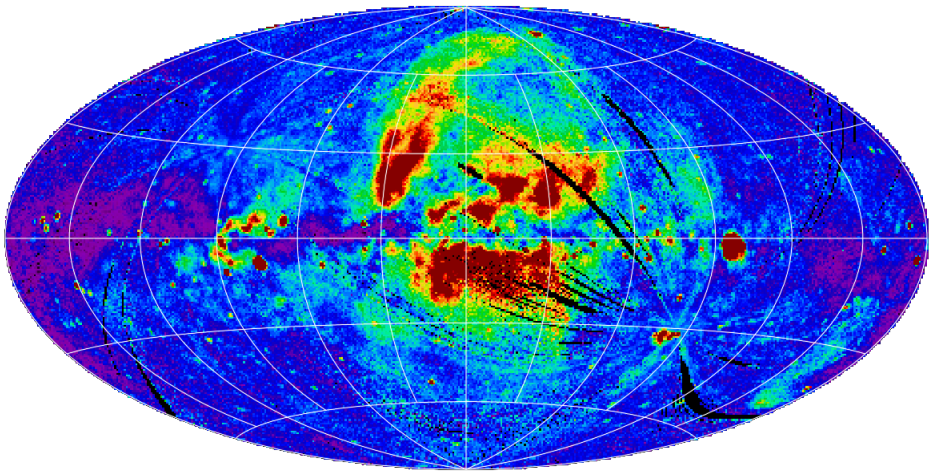

(b) 3/4 keV

Figure 1.2: Galactic coordinate all-sky maps of soft X-ray emission made by ROSAT (Snowden et al., 1997). (a) Although it exhibits a detailed anticorrelation with N_H, much of the 1/4 keV emission shows very little evidence of absorption and must originate locally. (b) Apart from few bright enhancements, the sky at 3/4 keV shows no pronounced latitude dependence, despite strong anisotropic absorption at these energies. A combination of sources covering a range of Galactic distance scales are believed to be responsible for this emission

observations, in addition to the discovery that the Sun is in a region of particularly low n_H (Henry et al., 1976), have led to the conclusion that a significant component of the 1/4 keV flux must originate locally. To explain the observed anti-correlation with hydrogen column depth without the color change expected for absorption, a displacement model was proposed, now known as the "Local Hot Bubble," or LHB (Sanders et al., 1977). According to this model, the LHB is a local superbubble $\sim$ 100 pc in extent, responsible for displacing the neutral ISM in the solar vicinity and replacing it with 10^6 K gas.

1.2.3 Solar Wind Charge Exchange

Another source of diffuse X-ray emission that we must consider is Solar Wind Charge eXchange (SWCX). SWCX is the process by which highly-charged ions in the solar wind capture a ground-state electron from an neutral atom or molecule into an excited state of the ion, resulting in the emission of one or more X-rays through radiative de-excitation. Although this process was first identified in an X-ray observation of the comet Hyakutake (Lisse et al. (1996) & Cravens (1997)), it was almost immediately recognized as a potential foreground of the DXRB (Cox, 1997). Models predicted that charge exchange with interstellar neutral hydrogen and helium in the heliosphere could provide a significant fraction of the observed local 1/4 keV flux (e.g. Lallement (2004) & Koutroumpa et al. (2009)). Like thermal emission, SWCX emits lines corresponding to electron transitions in highly-charged ions—resulting in the same emission lines for the same ions. Separating thermal and SXCX contributions to the DXRB is therefore a major obstacle to studying the hot-ionized ISM.

Although SWCX and thermal emission mechanisms are capable of producing the same emission lines for a given ion, their spectra can in principle be distinguished by the

relative strengths of those lines. The different excitation mechanisms for the two emission processes (capture from a neutral for SWCX, electron collision for thermal) preferentially populate different quantum states. The resulting emission spectra, therefore, depend on the de-excitation path from that state governed by selection rules for electronic transitions. Unfortunately, SWCX emission spectra are known to have a strong velocity dependence over the range of interest (200–800 km/s) that is very difficult to accurately calculate. Recent laboratory measurements will serve to guide and verify these calculations, and provide known velocity-dependent spectra for select ions (e.g. Defay et al. (2013)).

In lieu of spectral information, there have been other efforts to quantify and account for the level of SWCX contamination in diffuse X-ray observations. Long exposures and repeated observations can be used to remove any time-variable signal, though this technique is most effective for geocoronal, rather than heliospheric contributions, and is insensitive to any DC contribution. Another approach has been to compare observations with viewing geometries sampling different interplanetary neutral densities. This method provides the current best estimates for the contribution of SWCX to the DXRB, finding sky-averaged values of $\sim 30\%$ at 1/4 keV, $\sim 7\%$ at 3/4 keV, and negligible contribution at higher energies (Uprety et al., 2016).

1.3 The X-ray Quantum Calorimeter

Studying the origins of the DXRB poses an instrumental and observational challenge. In order to resolve individual emission lines, spectral resolution on the order of a few electron volts is required. At the same time, however, large throughput is needed due to the low surface brightness of the DXRB. All of the instruments on the X-ray observatories currently in orbit lack these at least one of these requirements. Ionization detectors such

as CCDs have the required throughput, but their spectral resolution is fundamentally limited to a fraction of the required resolving power. In contrast, wavelength-dispersive spectrometers could in principle provide the required spectral resolution, but an entrance slit necessary for viewing extended sources would lack the required throughput (in fact, almost all X-ray spectrometers of this type are slit-less, limiting their use to point-like sources).

The University of Wisconsin-Madison/Goddard Space Flight Center's X-ray Quantum Calorimeter (XQC) is a sounding rocket payload specifically designed to meet the instrumental requirements of observing the DXRB. XQC couples a $\sim$1 sr field of view and thin filters with a microcalorimeter spectrometer to achieve large throughput and 6 eV FWHM resolution at energies below 1 keV. Microcalorimeters resolve photon energy by measuring the heat deposited in a pixel, and therefore do not suffer the same fundamental limitations as ionization detectors. Moreover, being energy-dispersive, microcalorimeters are capable of observing extended sources without degrading performance. The operating principle of microcalorimeters is depicted in Figure 1.3. To date, XQC has completed four successful science missions: two observations at high Galactic latitude ($\ell, b = 90°, +60°$) and two observations at low Galactic latitude ($\ell, b = 165°, , -5°,$).

1.4 Thesis Outline

This thesis conducts a detailed analysis of the DXRB below 1 keV, utilizing all of the available XQC spectra to date. Chapter 2 provides a concise account of the observations, analysis, and conclusions, as they relate to the astrophysics of the ISM and the DXRB. The development of a novel technique for processing microcalorimeter data was also crucial to this work. Chapter 3 provides a thorough description of this technique as well as tests

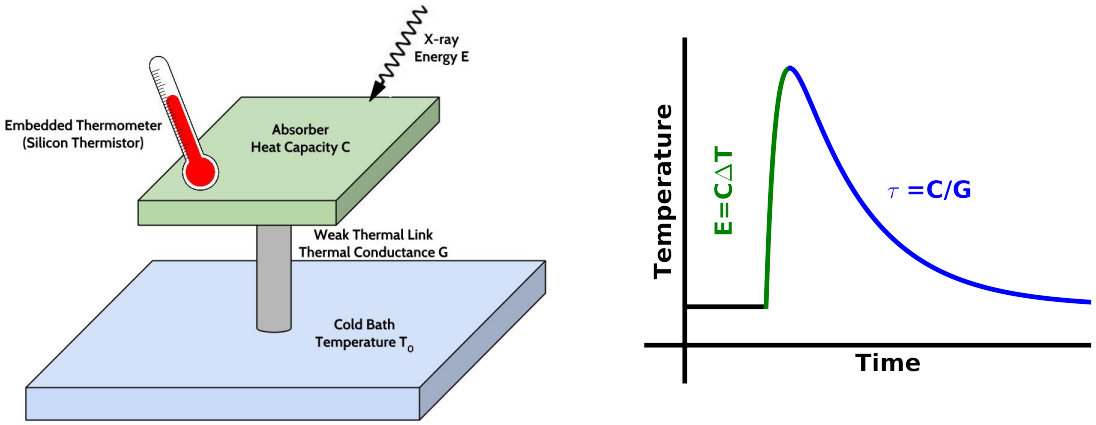

Figure 1.3: Operating principle of an ideal calorimeter. When an X-ray of energy E is stopped in an absorber with heat capacity C, the deposited energy causes an increase in the absorber temperature equal to E/C. This change in temperature is measured by a thermometer embedded in the absorber (e.g. ion-implanted silicon thermistor) before the energy dissipates and the temperature returns to the bath temperature.

characterizing its performance. Details of the analysis presented in Chapter 2 can be found in Appendix A. This includes payload and flight parameters for each observation, as well as details of data processing. Appendix B provides a description of *h5filter*, the software package used to implement the technique presented in Chapter 3. Here the reader can find descriptions of file formats and data directory structures.

Chapter 2

A High Spectral Resolution Study of the Soft X-ray Background

2.1 Introduction

The Diffuse X-ray Background (DXRB) below ∼1 keV originates from several sources ranging from solar to extragalactic distance scales. With the exception of the Cosmic X-ray Background (CXB) which can be modeled as a power law, most of the DXRB flux in this energy range is in lines from highly charged ions. Within the heliosphere, highly charged ions originating from the solar wind emit X-rays through an electron capture process known as charge exchange, or in this case specifically, Solar Wind Charge eXchange (SWCX). Throughout the Galaxy, thermal emission originates from $1–4 \times 10^6$K gas in the hot, ionized component of the interstellar medium (ISM) and the Galactic halo. Although different emission mechanisms will produce the same emission lines for a given ion, the line intensity ratios will vary. Therefore, observations with sufficient spectral resolution to resolve individual emission lines offer a valuable tool for studying the origins

of the DXRB.

To date, very few observations of the DXRB have had the spectral resolution needed to resolve individual emission lines. Some of the earliest observations made with gas proportional counters had almost no intrinsic energy resolution, but were able to derive some spectral information through the use of band-pass filters. CCD detectors provided a marked improvement, making it possible to identify individual ions (e.g. OVII, OVIII, and NeIX), but are incapable of resolving lines separated by less than ∼100eV, even when operating close to their fundamental limit.

Here we present a combined analysis of four high spectral resolution observations of the DXRB made by the University of Wisconsin-Madison/Goddard Space Flight Center sounding rocket payload, the X-ray Quantum Calorimeter (XQC). Section 2.2 outlines details of the payload hardware and sounding rocket observations. In Sections 2.3 and 2.4, the data reduction process and spectral analyses are described. A discussion of the inferred properties and components of the DXRB is presented in Section 2.5.

2.2 The X-ray Quantum Calorimeter

XQC is a 36-pixel silicon thermistor microcalorimeter X-ray spectrometer, achieving ∼6 eV FWHM resolution below 1 keV (the first generation detector used for Mission 27.041 achieved ∼9 eV FWHM). The low surface brightness of the DXRB, in combination with the limited observation times available on sounding rockets requires a large throughput. XQC meets this requirement by pairing a ∼1 sr non-imaging field of view with 1.44 cm^2 of detector area (0.36 cm^2 for the first generation detector). A series of five thin aluminum filters are used to shield the detector plane from 300 K infrared radiation ($\lambda \gtrsim 2\mu$m) while still having acceptable transmission to soft X-rays. On the most

recent flight, an outer sixth filter held at 300K was added to prevent contamination to the inner five filters which are held at temperatures ranging from 0.05–130K. The filter contamination issue is discussed more in Section 2.3.3. For a more detailed description of the payload, detector, and infrared-blocking filters, the reader is referred to McCammon et al. (2002) and McCammon et al. (2008).

2.2.1 Sounding Rocket Observations

All XQC missions to date have launched from White Sands Missile Range in New Mexico, USA (lat, long = 32:40:00N, 106:20:00W). Table 2.1 summarizes the observing details of these four missions. The first two missions, 27.041 and 36.223, observed a high latitude target ($\ell, b = 90^\circ, +60^\circ$), chosen to be typical of high latitudes while avoiding the bright enhancements of Loop I and the North Polar Spur. The second two missions observed at low latitude ($\ell, b = 165^\circ, -5^\circ$), avoiding emission associated with the Galactic center. Projections of the target areas onto all-sky maps are depicted in Figure 2.1.

Mission 27.041 featured the first generation detector format described in McCammon et al. (2002), while the following three missions featured the larger detector format described in McCammon et al. (2008). The new detector provides a factor of four improvement in collecting area, while thinner filters introduced at the same time provide greater transmission at low energies. A hard landing destroyed the detector from Mission 36.223, but a similar detector was used for Missions 36.264 and 36.294. A previous analysis of the data from 27.041 was presented in McCammon et al. (2002), and data from 36.223 has been presented in Crowder et al. (2012). Data from 36.264 and 36.294 are presented here for the first time.

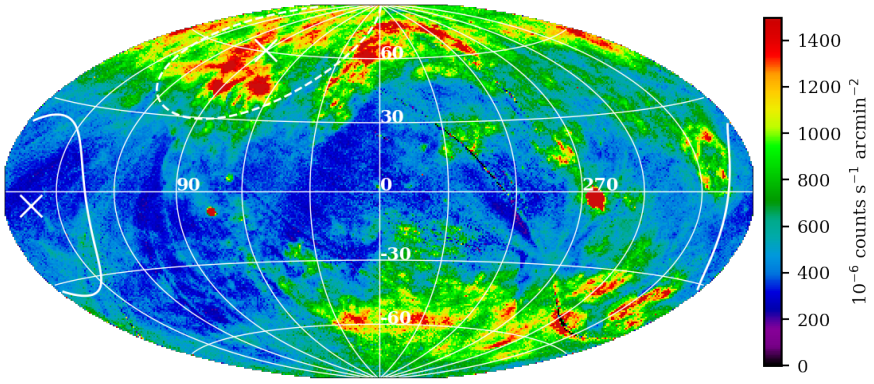

(a) 1/4 keV (R12)

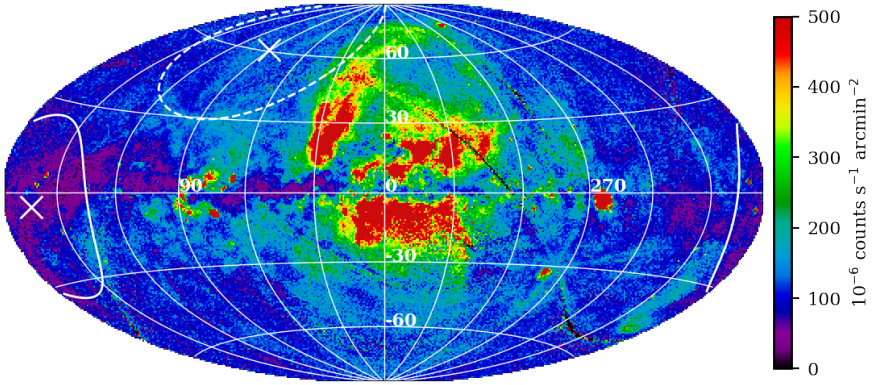

(b) 3/4 keV (R34)

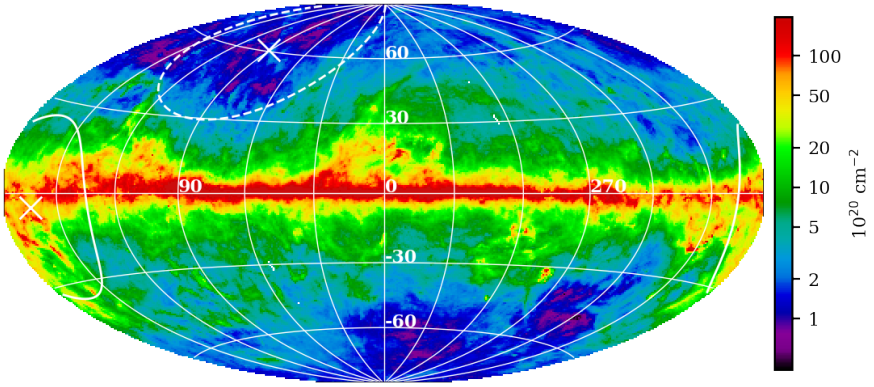

(c) IRAS-derived N_H

Figure 2.1: Target area of XQC observations projected on all-sky maps. Maps are in Galactic coordinates with the plane of the Galaxy located along the equator and the center of the Galaxy at the center of the map. $\sim$1 sr field of view is centered at $\ell, b = 90°, +60°$ (dashed line) and $\ell, b = 165°, -5°$ (solid line). ROSAT maps R12 and R34 reproduced from data presented in Snowden et al. (1997), and IRAS map reproduced from data presented in Schlegel et al. (1998).

Table 2.1: Summary of XQC sounding rocket observations.

Mission Number	Time (UT)	Date	Exposure (s)	Target (ℓ, b)
27.041	9:00	28 Mar 1999	100.7	90°,+60°
36.223	5:30	1 May 2008	176.9	90°,+60°
36.264	7:00	5 Nov 2011	205.4	165°,–5°
36.294	9:25	3 Nov 2013	150.0	165°,–5°

2.3 Data Reduction

Mission 36.223 was the first XQC flight to have on-board data recorders, enabling storage of the complete time-series data for all pixels throughout the observation. Prior to this, only triggered data segments containing pulses were telemetered to the ground station. The on-board data can be analyzed with the overlapped pulse-fitting algorithm presented in Chapter 3, which permits use of a longer filter template (for better energy resolution), while simultaneously improving processing live-time by a factor of 2–3 compared to the previous processing technique. No changes were made to the data reduction of 27.041 from McCammon et al. (2002).

As described in Crowder et al. (2012), data from 36.223 was low-pass filtered to remove microphonics due to gate-valve induced vibrations. Vibrational heating of the pixels also caused the detector gain to drift, but this could be corrected post-processing. Microphonics were significantly less severe on subsequent flights due to changes in the mechanical design of the pixels. 36.223 also suffered from 60.04 Hz pickup from a commutator in the housekeeping electronics. Because this pickup had a very stable frequency and amplitude, it could be fit and subtracted from the time-domain data. The source of this pickup was repaired for Missions 36.264 and 36.294.

2.3.1 Gain and Linearity Corrections

The gain of a microcalorimeter changes with temperature, partly because the heat capacity of the absorber increases with temperature, but primarily because the temperature response of the thermometers is non-linear. While the coldplate temperature can be precisely regulated at 50 mK, changes in thermal balance can affect the temperature offset between the pixels and the coldplate. For instance, although the inside of the gate-valve slide is highly polished to limit the emissivity, it still radiates more power on the detector than the night sky. As a result, there is a greater temperature offset between the pixels and the coldplate when the gate-valve is closed than when it is open. For a fixed coldplate temperature, this results in a higher pixel temperature and therefore lower gain. Likewise, heat deposited in the pixels by vibrations initiated when the gate-valve is operated can cause gain changes independent of coldplate temperature regulation. This was an issue for Mission 36.223. To track gain changes that do not follow the coldplate temperature, a ^{41}Ca source installed inside the cryostat provides a continuous source of K Kα and Kβ X-rays. The energies of these lines are above 3 keV and thus do not seriously contaminate energy range of interest below 1 keV. This source illuminates the entire detector array to enable tracking changes on a pixel-by-pixel basis.

During the launch of Mission 36.294, a burn instability in the second-stage motor caused significant vibrational heating that exhausted the cooling power of the refrigerator. This meant that we were unable to regulate the coldplate temperature throughout the observation, during which time it drifted from 57 down to 52 mK. In this case, the gain changes did follow the coldplate temperature, and recorded housekeeping data could be used to recover the gain. However, the magnitude of the gain change varied from pixel to pixel, so the internal calibration source was still used to normalize the gain correction across the detector array.

The same non-linear thermometer response that is responsible for temperature dependent gain is responsible for a non-linear energy response (i.e. as the temperature of the pixel changes during an X-ray event, the gain also changes). In order to convert measured pulse height to energy, a cubic polynomial is fit to calibration data containing emission lines of known energy from the gate-valve source. For each mission, this calibration data was collected during ∼30 minutes immediately preceding the launch.

2.3.2 Event Selection

As presented in more detail in Chapter 3, the overlapped pulse fitting algorithm is capable of processing nearly 100% of X-ray events. However, high energy events that saturate the electronics cannot be analyzed, and are therefore excised from the data streams prior to processing. Any other X-ray events that are within the filter length of a saturating event must also be removed. This dominates the processing dead-time for the three observations that are processed this way. Additionally, events are excluded from the final analysis if they arrive within 10 ms of another event on the same pixel or within 2.5 ms of an event on an adjacent pixel. The first condition ensures that selected events are sufficiently separated in the data stream to be recovered with optimal resolution. 10 ms is conservatively chosen to be several thermal rise-times of the XQC detectors (∼1.5 ms). The second condition is to remove cosmic ray events and cross-talk.

Finally, all the events from a single pixel are excluded if the arrival time distribution of events on that pixel is significantly different than uniform, or if the energy distribution (i.e. the spectrum) is significantly different than that of the entire detector array. The spectral comparison is made using the K-Sample Anderson-Darling test for equality (Scholz & Stephens, 1987), chosen because it is non-parametric (i.e. requires no assumptions about the underlying data distribution), and has sensitivity at the tails of the distribution. Pixels

were considered to fail this test if the null hypothesis was rejected at the 95% level. As discussed in Section 2.3.1, a number of factors can cause time- and pixel-dependent gain changes. This test is an added safeguard to ensure that gain corrections do not introduce any distortions to the final spectrum. Depending on the observation, this test affected between 3 and 5 pixels.

2.3.3 Filter Contamination

Beginning with Mission 36.223, water ice contamination to the infrared-blocking filters became an issue. Calibration data from immediately before and after the observation confirmed that the contamination formed during the time that the gate-valve was open. Between Missions 27.041 and 36.223, a hermetic bulkhead was removed between the experiment section and the other sections of the payload. Venting of the payload through the experiment section was initially attributed as the source of the contamination. For Mission 36.264 therefore, the payload was purged with dry nitrogen gas prior to launch. Unfortunately, this remedy failed to prevent water ice from accumulating on the filters. For Mission 36.294, a sixth filter held at 300K was added. While the additional filter reduced throughput slightly, it successfully prevented ice from forming on the inner filters, resulting in a net improvement to throughput. The exact source of water vapor responsible for the ice is still unknown.

For Missions 36.223 and 36.264, the correction for contaminated filter transmission is based on a combination of post-flight data (taken on the parachute after re-entry into the atmosphere) and in-flight data. There is a multi-target fluorescent calibration source installed on the gate-valve slide that emits a range of characteristic X-ray lines between 183 and 1740 eV which is viewed by the detector whenever the gave-valve is closed. This permits a relative measurement of filter transmission immediately before and after the

observation. This method has the advantage of measuring the transmission at well-defined energies. However, the gate-valve source samples a relatively small fraction of the filter area. To complement this measurement, the on-sky spectrum during the observation is compared to predictions made by the ROSAT All-Sky Survey (RASS) data. This method samples the entire filter area, but measures only the average transmission over the entire energy band pass. For both flights affected by ice contamination, the models used for filter transmission are consistent with both of these measurements.

2.3.4 Non X-ray Background

The primary source of background in the 100–1100 eV range are higher energy X-rays that have lost energy due to photoelectric escape. Whenever an X-ray is stopped near the surface of the pixel, there is some probability that it will eject an electron from the surface via the photoelectric effect. This happens $\sim$5% of the time. The energy carried away by the electron is lost, resulting in a lower observed energy. The observed spectrum can be approximated as flat between zero and the incident X-ray energy. The internal ^{41}Ca calibration source at 3.3 and 3.7 keV is responsible for most of this background, contributing 0.5 counts s^{-1} keV^{-1} to the 36.223 spectrum and 0.3 counts s^{-1} keV^{-1} to the other 3 spectra (the calibration source intensity was greater for 36.223).

Particle background is also possible. The minimum ionizing energy loss of primary cosmic ray protons in a pixel is $\gtrsim$5.5 keV, and thus does not contribute to the low energy spectrum directly. Cosmic rays can also eject electrons from surfaces surrounding the detector array which can produce events at lower energies. We have looked for evidence of background from such cosmic ray secondaries when the gate-valve is closed, but the observed spectrum has always been consistent with the spectrum taken with the gate-valve closed at sea level. Auroral electrons, on the other hand, have been a source of low

energy background on other thin-windowed sounding rocket observations. This has always been a concern for XQC, whose wide field of view prevents the use of magnetic brooms to keep electrons from entering the detector. Fortunately, only Mission 36.294 seems to have been affected by electrons. Based on comparisons with the previous observation of the same target, as well as with the ROSAT survey data, we estimate this background to be 0.6 counts s^{-1} keV^{-1} over the 100–2500 eV range.

2.3.5 Cross Check

As a consistency check on the absolute flux, the final event selection is compared to the band fluxes measured by the RASS. The energy resolution of XQC allows us to make a model-independent comparison by simply multiplying the observed spectra by the ratio of the ROSAT PSPC response to the XQC response at the energy observed by XQC. The relative fluxes are listed in Table 2.2. We find good general agreement, although R2 is systematically low for the three observations reported. A possible explanation may be differences in the contribution from heliospheric SWCX. Uprety et al. (2016) estimates that, on average, 30% of R2 flux observed by the RASS is due to heliospheric SWCX. However, as we will show, SWCX emission appears to account for only $\lesssim$10–20% of the R2 flux in our emission models for these three observations.

2.4 Spectral Analysis

Spectral fitting was performed with XSPEC version 12.9.0, using the PyXspec interface (Arnaud, 1996). Thermal emission was modeled with APEC version 2.0.2 (Smith et al., 2001) using solar abundances from Anders & Grevesse (1989). Depleted abundances from Savage & Sembach (1996) were also tried, but were not found to improve the fit. Charge

Table 2.2: Observed flux for all four missions relative to the RASS energy bands. R1 is not included due to large statistical uncertainties. Statistical uncertainties are given at the 1σ level.

	Relative ROSAT Band Flux				
Mission	R2	R4	R5	R6	R7
27.041[a]	-	96%	82%	70%	88%
36.223	84%±16%	97%±7%	91%±6%	99%±6%	107%±7%
36.264	75%±22%	88%±12%	89%±8%	101%±7%	101%±8%
36.294	72%±10%	95%±7%	96%±7%	110%±8%	113%±11%

[a] Values quoted from McCammon et al. (2002).

exchange emission was modeled with AtomDB's Charge eXchange model, ACX, version 1.0.1 (Smith et al., 2012). Following Smith et al. (2014), two SWCX components representing contributions from fast and slow solar wind were used, with relative abundances taken from Table 1 in von Steiger et al. (2000).

Interstellar absorption of the Cosmic X-ray Background (CXB) and Hot Galactic halo was taken to be the transmission-weighted average over the full XQC field of view. The column depth of the Galaxy in each direction was derived from the DIRBE-corrected IRAS 100 μm map (Schlegel et al., 1998). This infrared emission map is intended to approximate dust mass, which we can convert to total hydrogen column depth by assuming an average gas to dust ratio. This accounts fairly well for molecular as well as atomic gas. The equivalent hydrogen column depth is obtained by multiplying by a factor of 1.44×10^{20} cm^{-2} MJy^{-1}, as outlined in Snowden et al. (1997). Transmission was calculated assuming solar abundances from Anders & Ebihara (1982) and cross sections from Balucinska-Church & McCammon (1992). The resulting average interstellar transmission for each target area is shown in Figure 2.2.

All four spectra were fit with a minimum of four model components: 1) an unabsorbed, APEC plasma representing emission from the Local Hot Bubble (LHB); 2) an

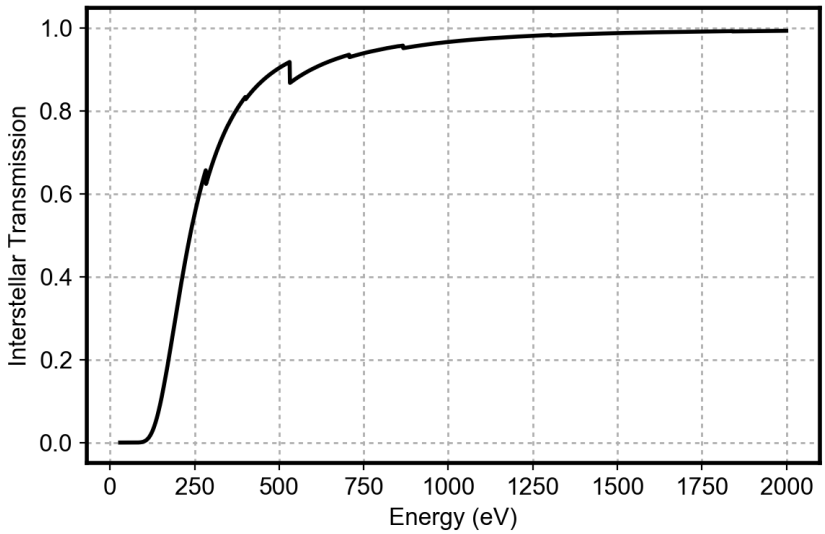

(a) High Latitude

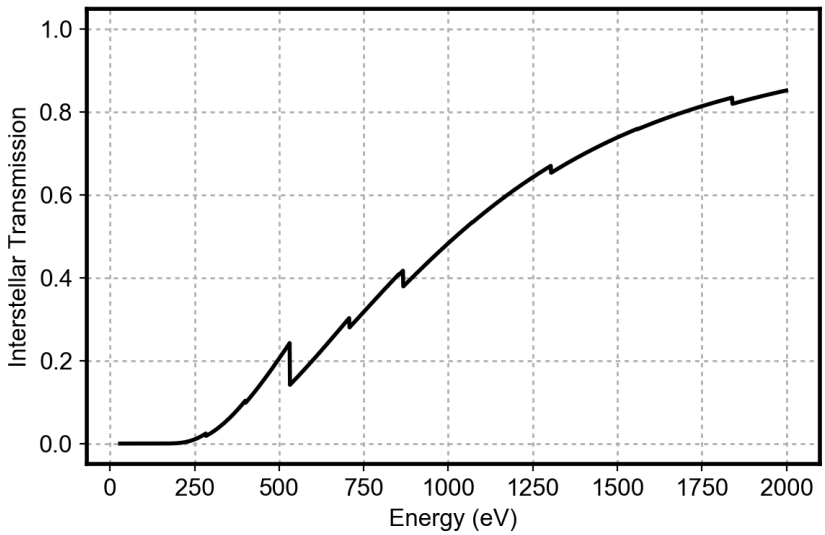

(b) Low Latitude

Figure 2.2: Interstellar transmission averaged over XQC field of view for the (a) high latitude observations and (b) low latitude observations. This transmission is used to account for interstellar absorption of the Cosmic X-ray Background (CXB) and hot Galactic halo emission in spectral models.

unabsorbed ACX component representing SWCX due to slow solar wind; 3) an absorbed APEC plasma representing emission from the hot Galactic halo; and 4) an absorbed, double broken power law representing the Cosmic X-ray Background (CXB) due to unresolved extragalactic Active Galactic Nuclei (AGN). For the low latitude observations, an additional absorbed power law component was included to represent the Crab nebula, which was within the field of view. The low latitude observations also required an additional thermal component at $\sim$0.9 keV, initially represented by an unabsorbed APEC plasma in what we will refer to as Low Latitude Model 1. Two other models were also fit to the low latitude observations, which will be discussed in more detail in Section 2.5. Lastly, the spectrum from Mission 27.041 required an additional unabsorbed ACX charge exchange component representing SWCX due to fast solar wind. The fast solar wind differs from the slow solar wind primarily in its ion temperature, but also differs slightly in elemental abundances.

A summary of model components and parameters are listed in Tables 2.3 to 2.6. Repeat observations of the same target were fit simultaneously with common model parameters, with the exception of the SWCX components, which have free normalizations for each observation. Parameters for the CXB and Crab were not fit, and instead were fixed to values from the literature. Figures 2.3 to 2.6 show each of the four spectra and corresponding models folded through the instrumental response. Individual model components are depicted in Figures 2.7 to 2.10, and their corresponding flux contributions to the ROSAT energy bands are listed in Table 2.7.

Table 2.3: Model parameters for simultaneous fit of high latitude observations. All uncertainties are reported at 90% confidence level. In instances where a component or parameter is unique to a single spectrum, the values are reported in the column of the corresponding spectrum.

			Value	
Component	Model	Parameter	27.041	36.223
Local Hot	APEC[a]	kT (keV)	$0.093^{+0.010}_{-0.019}$	
Bubble		Normalization[b]	$8.09^{+1.92}_{-1.79}$	
Fast SWCX	VACX[c]	kT (keV)	$0.055^{+0.011}_{-0.006}$	
		Normalization[d]	$4.17^{+3.54}_{-2.25}$	
Slow SWCX	VACX[c]	kT (keV)	$0.184^{+0.042}_{-0.045}$	
		Normalization[d]	$0.073^{+0.059}_{-0.042}$	$0.023^{+0.020}_{-0.017}$
Hot Halo	Absorbed	kT (keV)	$0.189^{+0.027}_{-0.018}$	
	APEC[a,e]	Normalization[b]	$4.99^{+2.04}_{-1.75}$	
CXB	Absorbed Double	Photon Index 1	1.54	
	Broken Power Law[e,f]	($E < E_{break}$)		
		Normalization 1[g]	5.7	
		Photon Index 2	1.96	
		($E < E_{break}$)		
		Normalization 2[g]	4.9	
		E_{break} (keV)	1.2	
		Photon Index	1.4	
		($E > E_{break}$)		
χ^2/dof			934.46/788 = 1.186	

[a] Solar abundances from Anders & Grevesse (1989).
[b] Units of 10^{14} cm^{-5} sr^{-1}.
[c] Relative abundances from Table 1 in von Steiger et al. (2000).
[d] Parameter has no intrinsic physical unit. See text for discussion.
[e] Transmission-weighted average interstellar absorption (Figure 2.2a).
[f] Double power law model given by Smith et al. (2007)
[g] Units of photons keV^{-1} s^{-1} cm^{-2} sr^{-1} at 1 keV.

Table 2.4: Same as Table 2.3 except for Low Latitude Model 1.

			Value	
Component	Model	Parameter	36.264	36.294
Local Hot	APEC	kT (keV)	0.097 (fixed)	
Bubble		Normalization	$3.43^{+0.69}_{-0.65}$	
Slow SWCX	VACX	kT (keV)	$0.185^{+0.065}_{-0.047}$	
		Normalization	$0.010^{+0.009}_{-0.008}$	$0.016^{+0.015}_{-0.012}$
Hot Halo	Absorbed	kT (keV)	$0.186^{+0.021}_{-0.026}$	
	APEC[a]	Normalization	$14.1^{+4.1}_{-3.9}$	
Unidentified	APEC	kT (keV)	$0.952^{+0.083}_{-0.143}$	
Hot Feature		Normalization	$1.91^{+0.48}_{-0.45}$	
CXB	Absorbed Double	Photon Index 1	1.54	
	Broken Power Law[a]	($E < E_{break}$)		
		Normalization 1	5.7	
		Photon Index 2	1.96	
		($E < E_{break}$)		
		Normalization 2	4.9	
		E_{break} (keV)	1.2	
		Photon Index	1.4	
		($E > E_{break}$)		
Crab	Absorbed	N_H (10^{22} cm^{-2})	0.43	
	Power Law[b]	Photon Index	2.04	
		Normalization[c]	8.8	
χ^2/dof			995.43/788 = 1.263	

[a] Transmission-weighted average interstellar absorption (Figure 2.2b).
[b] Absorbed power law given by Table 1 in Weisskopf et al. (2010).
[c] Units of photons keV^{-1} s^{-1} cm^{-2} at 1 keV.

Table 2.5: Same as Table 2.3 except for Low Latitude Model 2. Not listed are components modeling the CXB and Crab, which are identical to those in Table 2.4.

			Value	
Component	Model	Parameter	36.264	36.294
Local Hot Bubble	APEC	kT (keV)	0.099 (fixed)	
		Normalization	$3.67^{+0.67}_{-0.63}$	
Slow SWCX	VACX	kT (keV)	$0.191^{+0.120}_{-0.044}$	
		Normalization	$0.010^{+0.009}_{-0.008}$	$0.014^{+0.014}_{-0.011}$
Hot Halo	Absorbed APEC	kT (keV)	$0.188^{+0.028}_{-0.027}$	
		Normalization	$11.8^{+4.3}_{-4.2}$	
dM Stars	Absorbed Two-Temperature APEC	Normalization[a]	$1.19^{+0.25}_{-0.24}$	
χ^2/dof			1053.68/789 = 1.336	

[a] Dimensionless. See text for description of model and normalization parameter.

2.5 Discussion

2.5.1 Solar Wind Charge Exchange

The ACX model used to generate SWCX spectra is only approximate, since it lacks any dependence on the solar wind velocity. Laboratory measurements of charge exchange emission spectra have demonstrated that the relative strength of different emission lines changes significantly over the range of solar wind velocities (e.g. Defay et al. (2013)). Nevertheless, the two SWCX components used here improve the fits and help resolve differences between repeat observations of the same target—especially at high latitudes. As previously mentioned, 27.041 is the only observation to require a significant contribution from fast solar wind. Coincidentally, this is also the only observation to exhibit strong excess emission at 263 eV (See Figure 2.3). Perhaps this feature can also be attributed

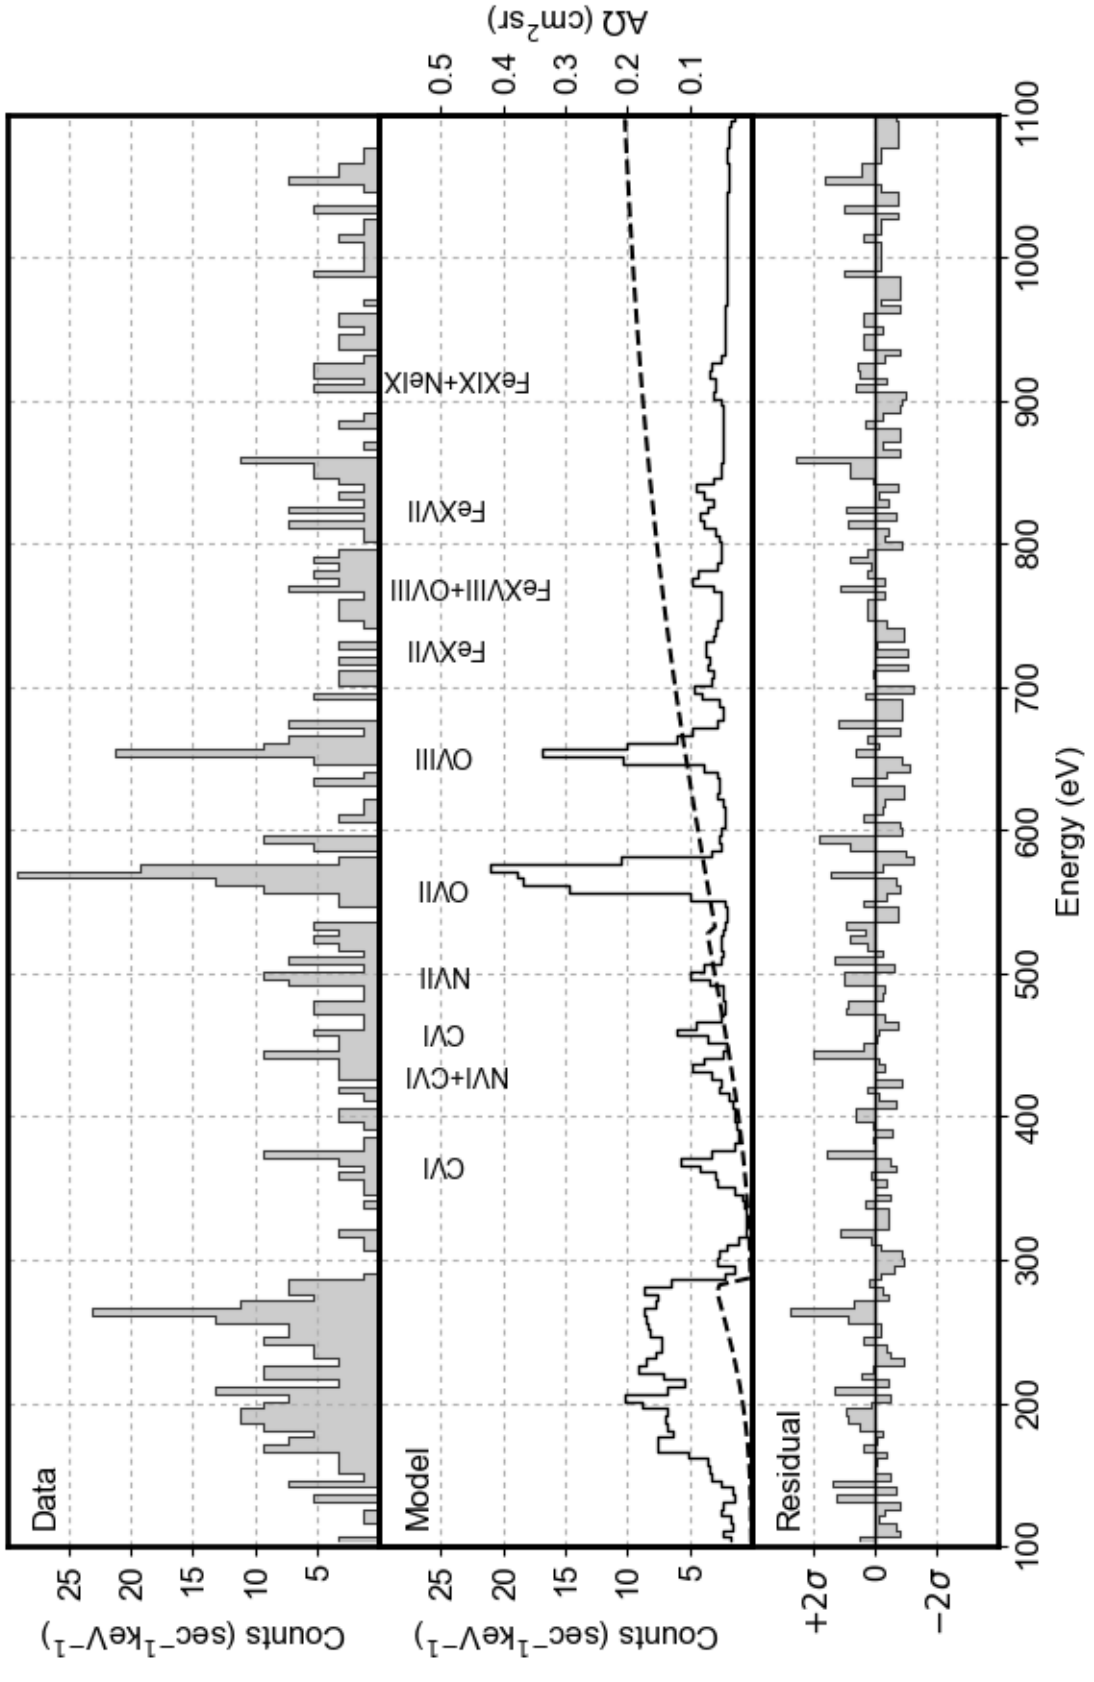

Figure 2.3: 27.041 High Latitude Spectrum. Data is shown in the top panel with the model from Table 2.3 shown in the middle panel. The dashed line represents the energy-dependent effective area. Residuals are shown in the bottom panel in units of sigma, estimated from the model.

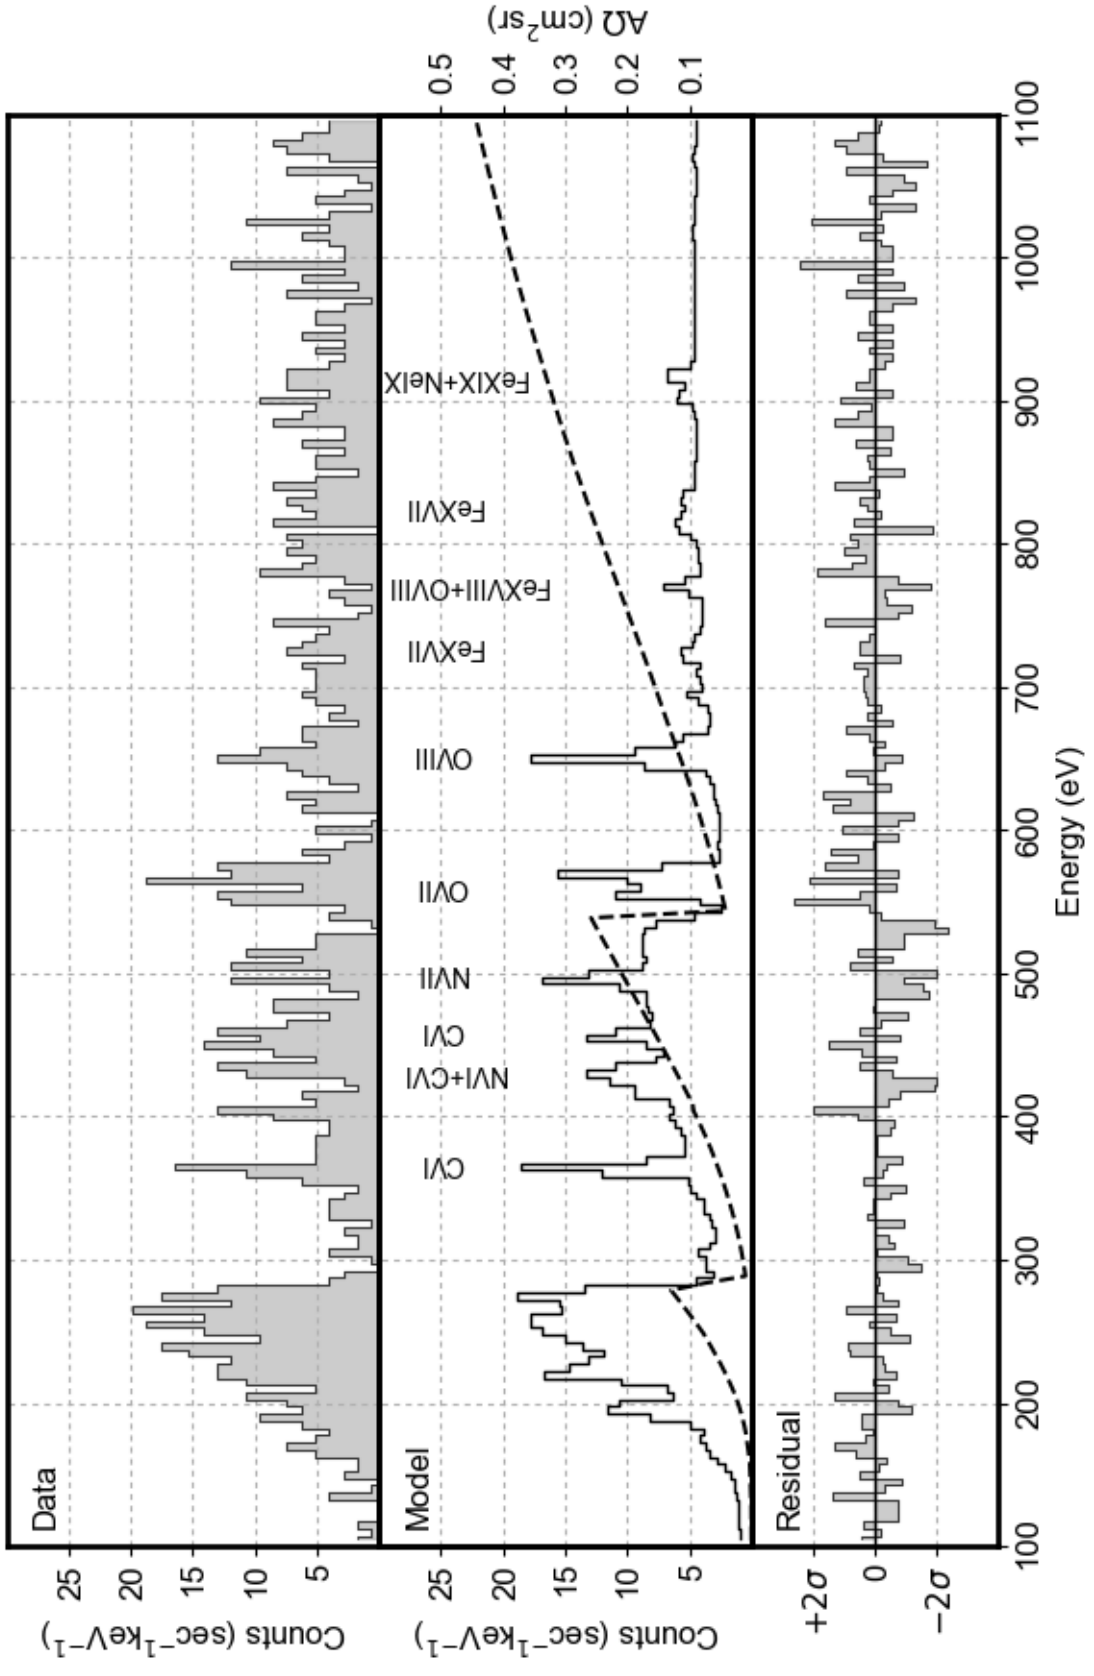

Figure 2.4: 36.223 High Latitude Spectrum. Data is shown in the top panel with the model from Table 2.3 shown in the middle panel. The dashed line represents the energy-dependent effective area. Residuals are shown in the bottom panel in units of sigma, estimated from the model.

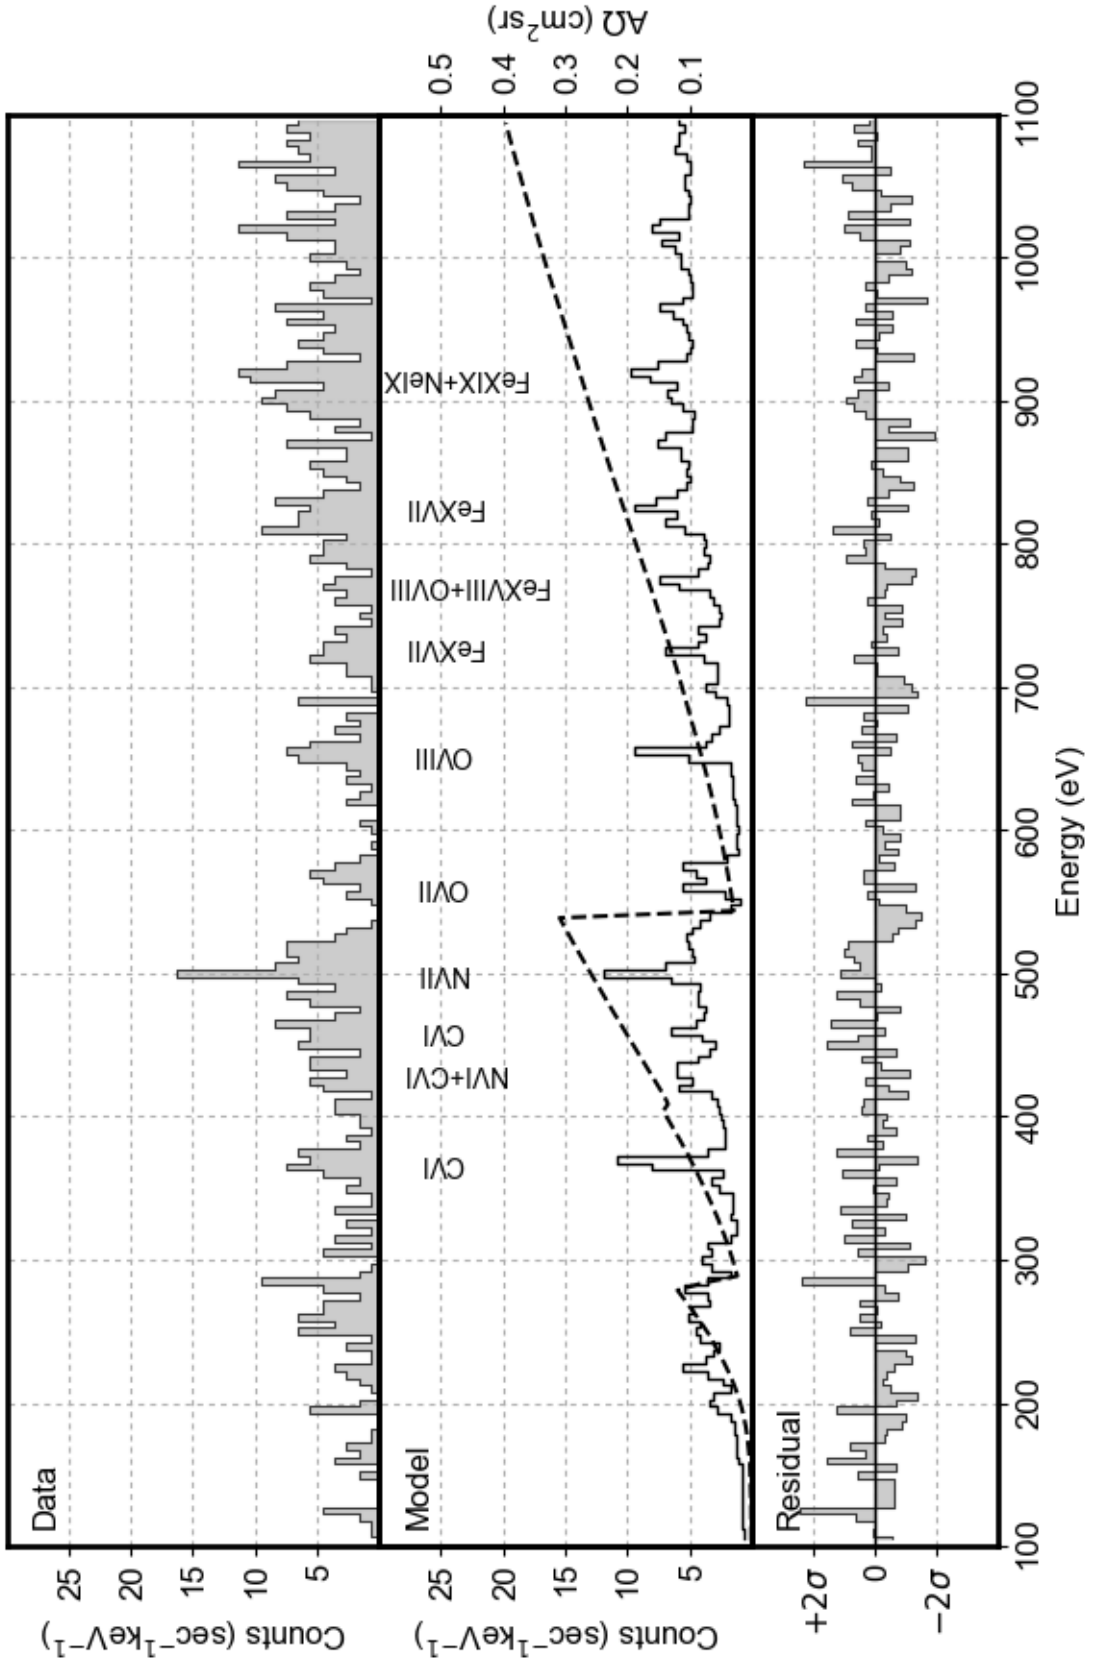

Figure 2.5: 36.264 Low Latitude Spectrum. Data is shown in the top panel with the model from Table 2.6 (Model 3) shown in the middle panel. The dashed line represents the energy-dependent effective area. Residuals are shown in the bottom panel in units of sigma, estimated from the model.

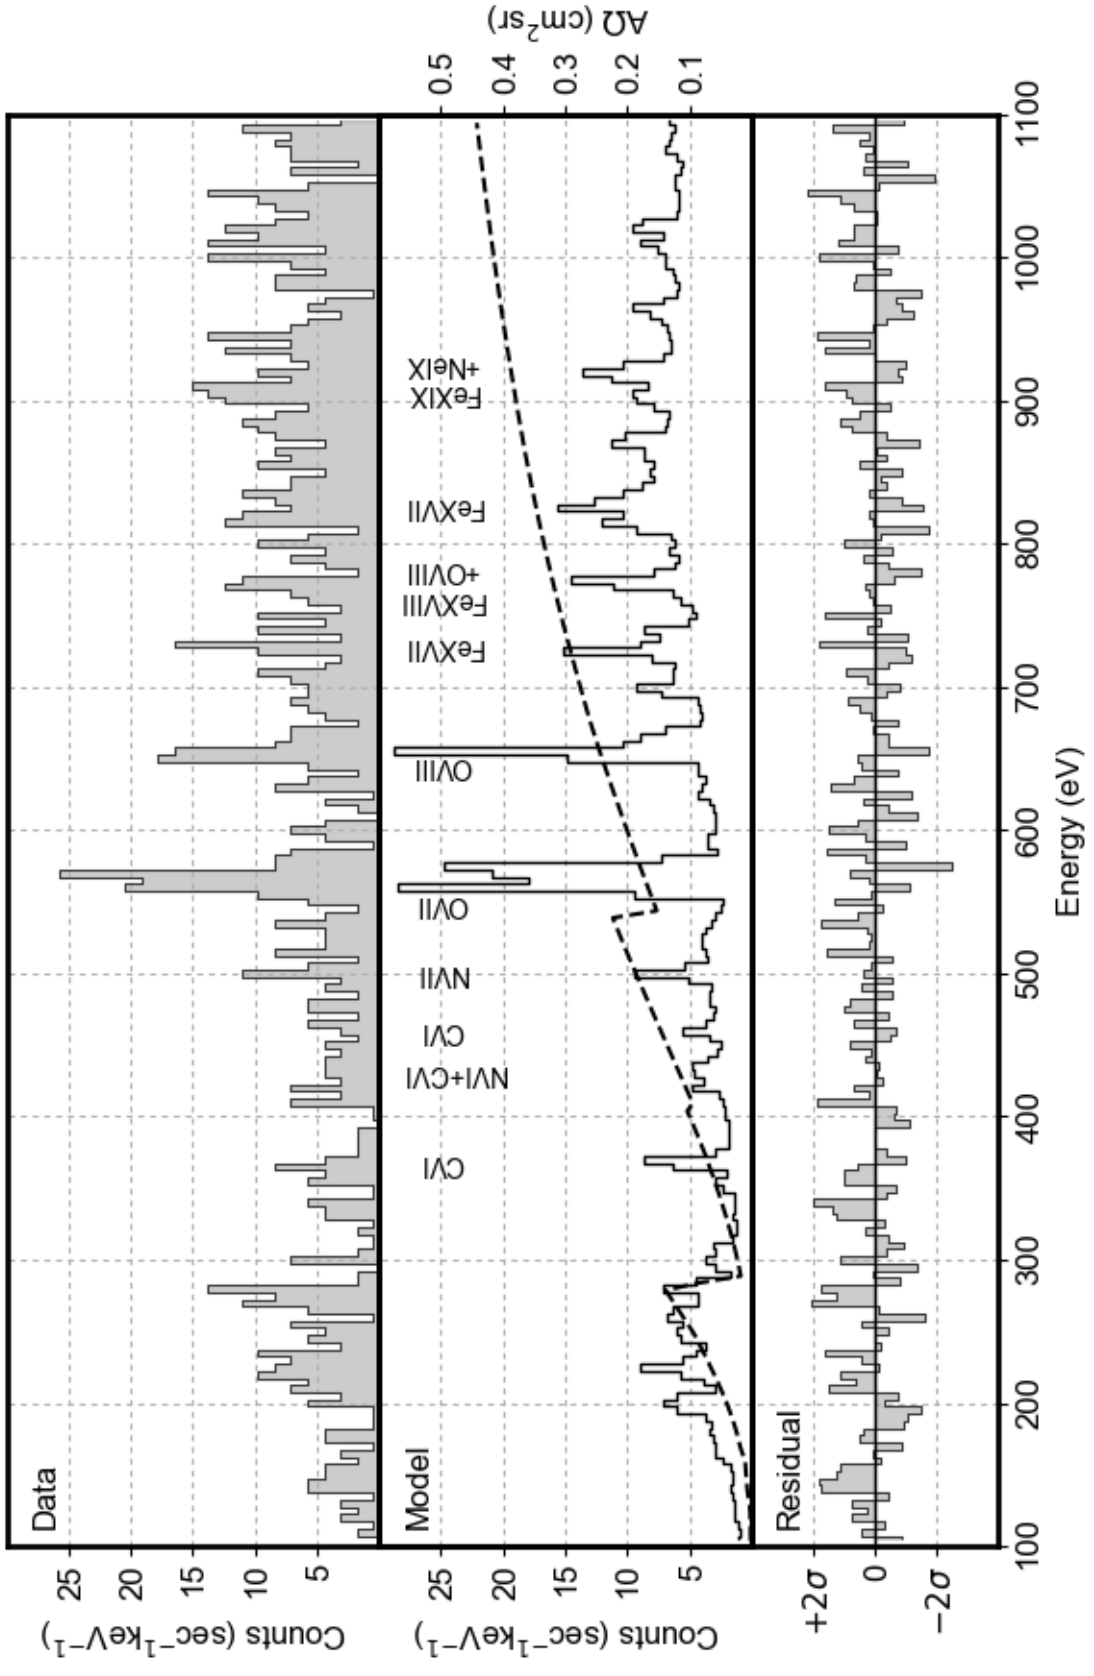

Figure 2.6: 36.294 Low Latitude Spectrum. Data is shown in the top panel with the model from Table 2.6 (Model 3) shown in the middle panel. The dashed line represents the energy-dependent effective area. Residuals are shown in the bottom panel in units of sigma, estimated from the model.

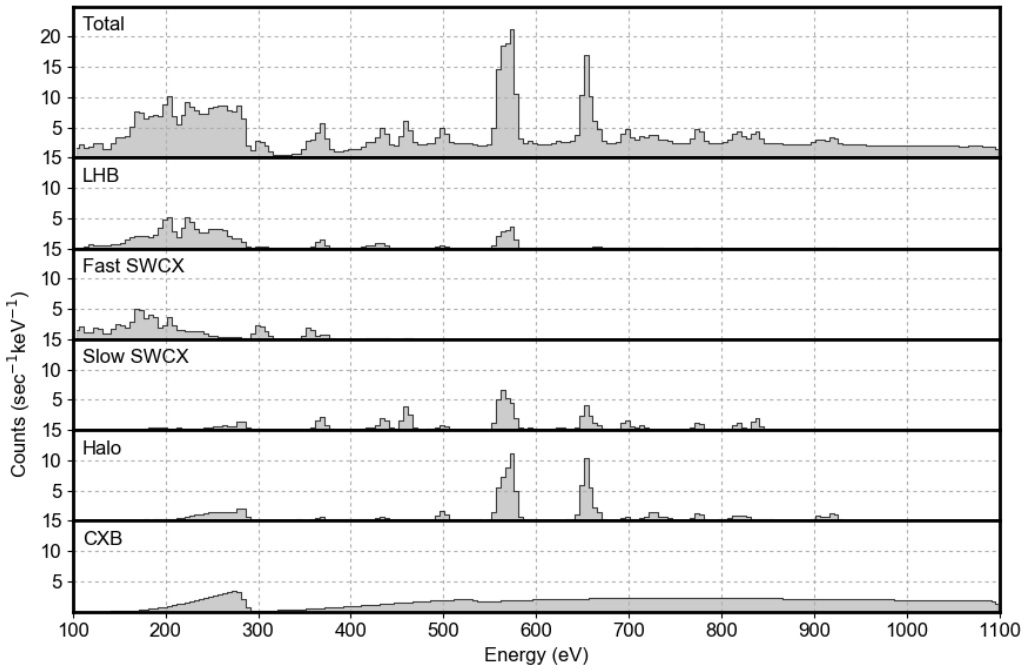

(a) 27.041

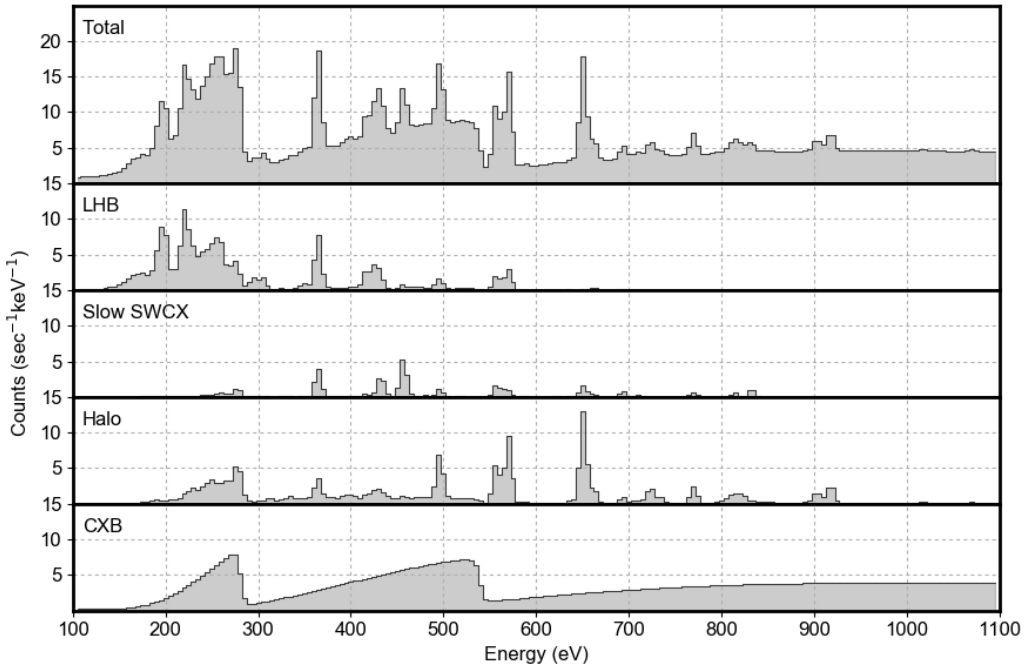

(b) 36.223

Figure 2.7: High latitude model components (Table 2.3).

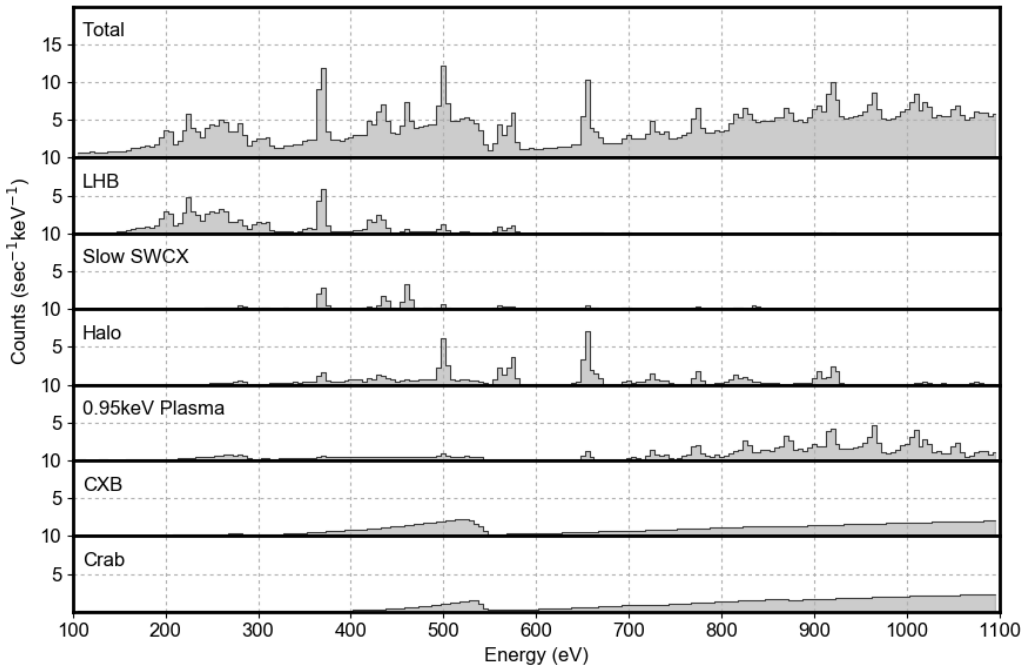

(a) 36.264

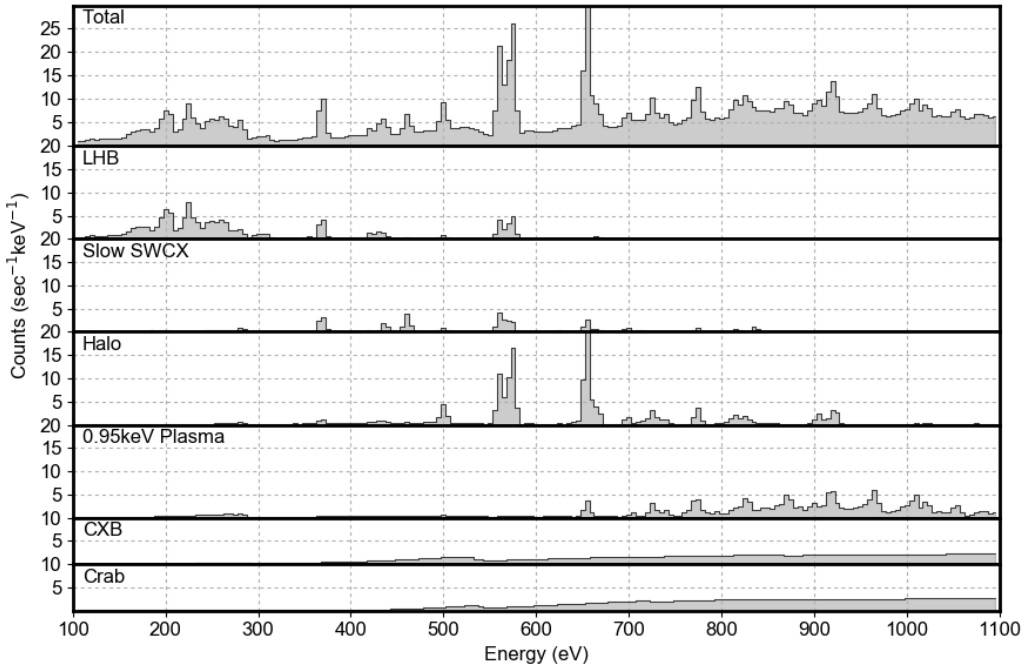

(b) 36.294

Figure 2.8: Low latitude model components (Model 1 from Table 2.4).

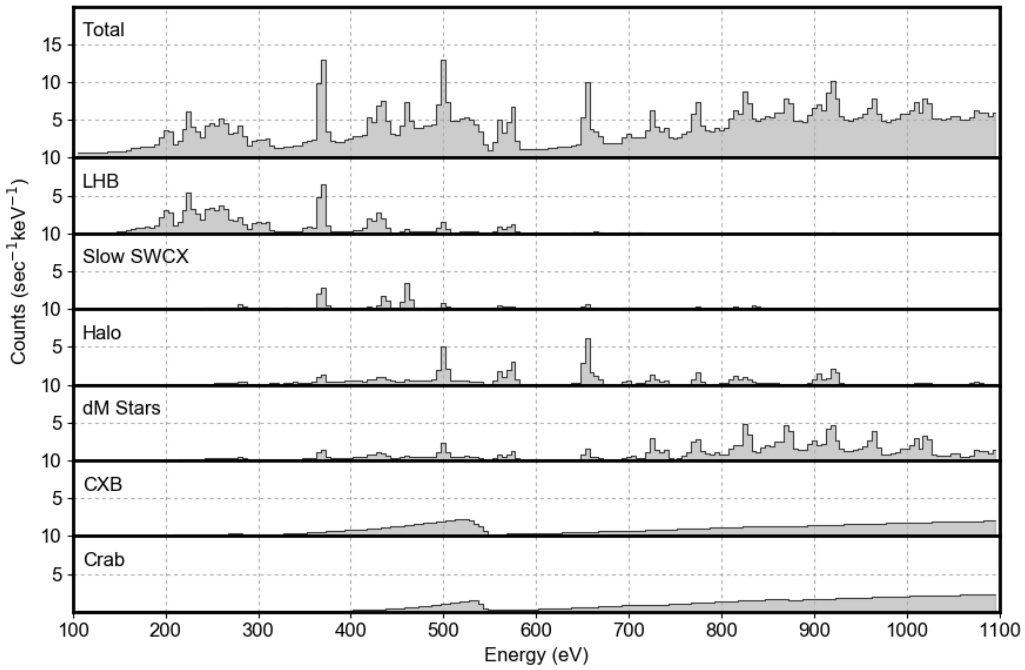

(a) 36.264

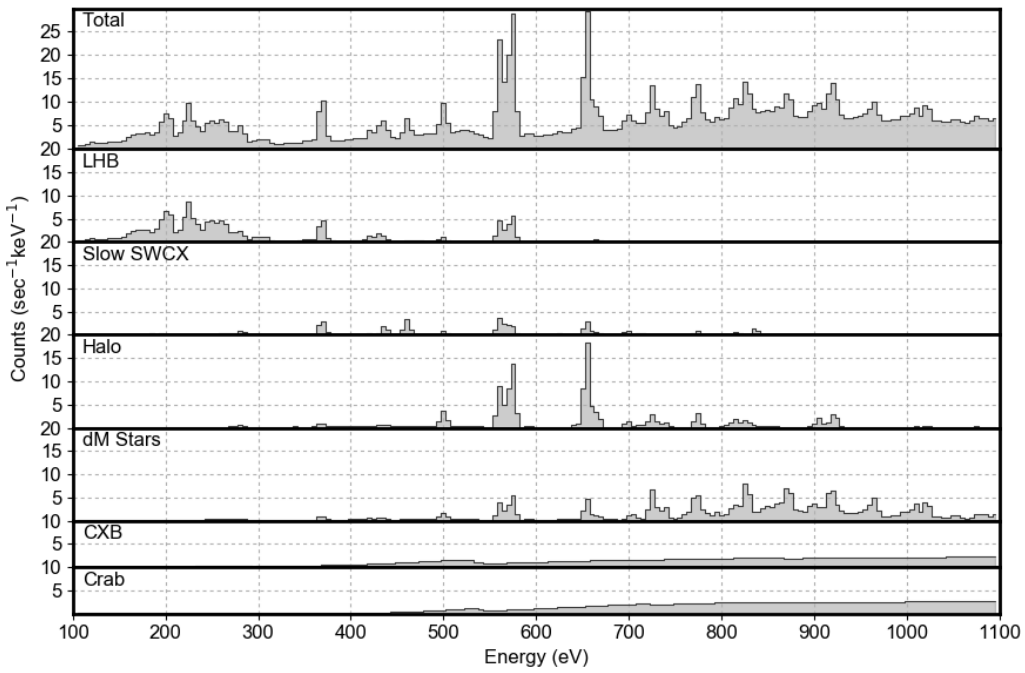

(b) 36.294

Figure 2.9: Low latitude model components (Model 2 from Table 2.5).

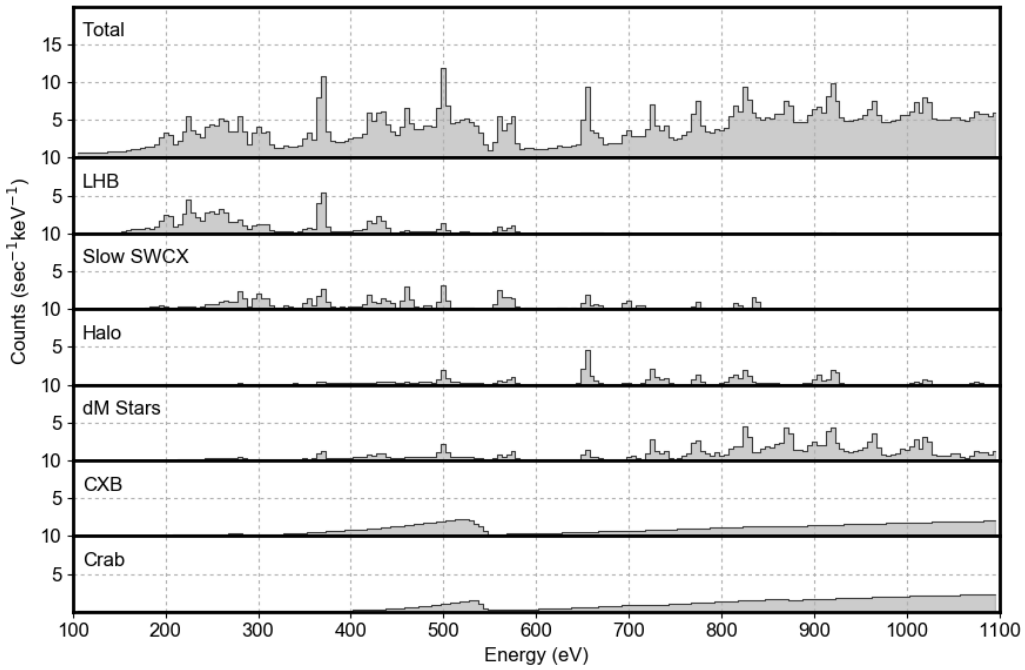

(a) 36.264

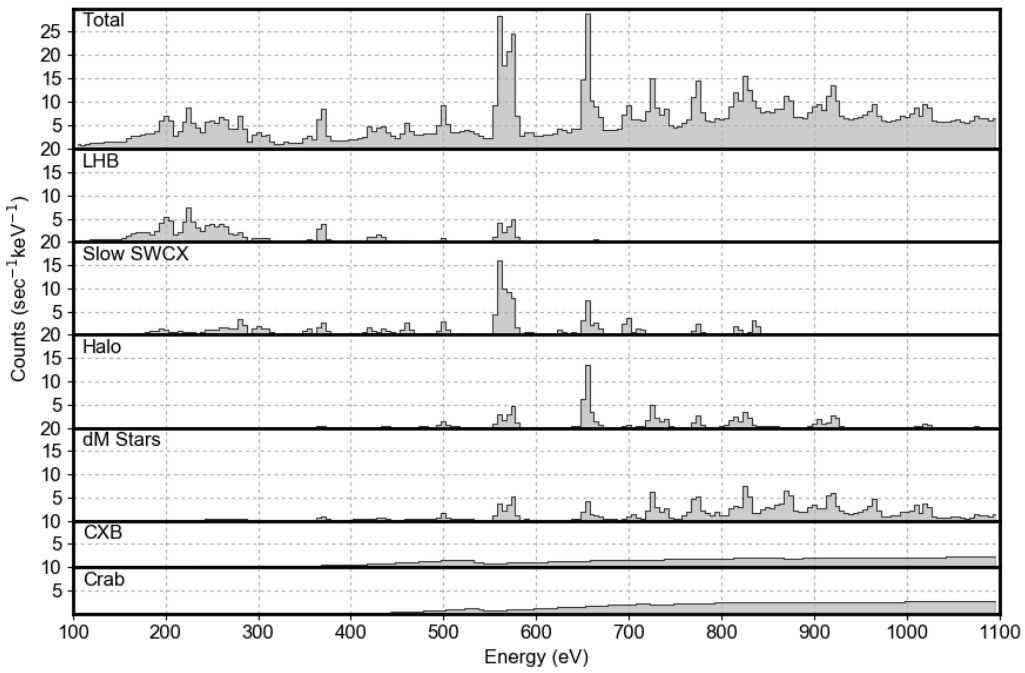

(b) 36.294

Figure 2.10: Low latitude model components (Model 3 from Table 2.6).

Table 2.6: Same as Table 2.3 except for Low Latitude Model 3. Not listed are components modeling the CXB and Crab, which are identical to those in Table 2.4.

			Value	
Component	Model	Parameter	36.264	36.294
Local Hot	APEC	kT (keV)	0.100 (fixed)	
Bubble		Normalization	$3.01^{+0.73}_{-0.69}$	
Slow SWCX	VACX	kT_C (keV)	$0.086^{+0.012}_{-0.086}$	
		kT (keV)	$0.177^{+0.016}_{-0.019}$	
		Normalization	$0.050^{+0.029}_{-0.025}$	$0.065^{+0.032}_{-0.029}$
Hot Halo	Absorbed	kT (keV)	$0.240^{+0.068}_{-0.056}$	
	APEC	Normalization	$5.34^{+4.3}_{-3.3}$	
dM Stars	Absorbed Two-Temperature APEC	Normalization[a]	$1.11^{+0.28}_{-0.27}$	
χ^2/dof			1030.97/788 = 1.308	

[a] Dimensionless. See text for description of model and normalization parameter.

to SWCX from fast solar wind, but it is not included in the current ACX model. In contrast to the fast solar wind, contribution from the slow solar wind is observed in all four spectra. Again, the SWCX contribution modeled in this way is only approximate.

A more robust measurement of the SWCX contribution can be made by analyzing individual emission lines, made possible by the high spectral resolution of these observations. After fitting a global model to each spectrum, the intensities of individual lines are measured by setting the abundance for a single element to zero in each of the model components and adding δ-functions for each of the strong lines of that element. Fluxes measured in this manner for select emission lines are presented in Table 2.8, with ratios of those lines given in Table 2.9. Of particular interest are lines from CVI and OVII, since they are relatively bright, are isolated from other lines, and can be used to unambiguously differentiate between thermal emission and SWCX.

Table 2.7: Contribution of the diffuse model components in Tables 2.3 to 2.6 to the ROSAT All-Sky Survey energy bands. Only contributions greater than 1 ROSAT Unit are reported. SWCX contributions are given as an upper bound. It is important to note that these are upper bounds for the SWCX contribution to these four observations only. All uncertainties and bounds are reported at 90% confidence level.

		ROSAT Band Flux (10^{-6} counts s^{-1} $arcmin^{-2}$)					
Target	Component	R1	R2	R4	R5	R6	R7
High Latitude	LHB	256^{+61}_{-57}	203^{+48}_{-45}	$3.8^{+0.9}_{-0.8}$	$1.3^{+0.3}_{-0.3}$	-	-
	SWCX	<855	<379	<17.9	<11.9	<3.3	-
	Hot Halo	$22.9^{+9.3}_{-8.0}$	$42.9^{+17.5}_{-15.0}$	$15.7^{+6.4}_{-5.5}$	$11.3^{+4.6}_{-4.0}$	$4.5^{+1.8}_{-1.6}$	-
	CXB	55.1	89.5	26.0	38.4	62.6	42.6
Low Latitude	LHB	104^{+21}_{-20}	$87.1^{+17.5}_{-16.5}$	$2.0^{+0.4}_{-0.4}$	-	-	-
Model 1	SWCX	<12.0	<12.5	<4.3	<2.9	-	-
	Hot Halo	$1.0^{+0.3}_{-0.3}$	$3.7^{+1.1}_{-1.0}$	$10.4^{+3.0}_{-2.9}$	$9.2^{+2.7}_{-2.5}$	$4.9^{+1.4}_{-1.4}$	-
	0.95 keV Plasma	$12.9^{+3.2}_{-3.0}$	$10.8^{+2.7}_{-2.54}$	$5.2^{+1.3}_{-1.2}$	$15.2^{+3.8}_{-3.6}$	$21.7^{+5.5}_{-5.1}$	$4.1^{+1.0}_{-1.0}$
	CXB	0.9	2.8	7.5	15.1	34.2	31.2
Low Latitude	LHB	102^{+20}_{-19}	$88.6^{+17.1}_{-16.1}$	$2.2^{+0.4}_{-0.4}$	-	-	-
Model 2	SWCX	<11.2	<11.5	<4.2	<3.0	-	-
	Hot Halo	$1.0^{+0.3}_{-0.3}$	$3.1^{+1.1}_{-1.0}$	$8.9^{+3.2}_{-3.2}$	$8.0^{+2.9}_{-2.8}$	$4.3^{+1.6}_{-1.5}$	-
	dM Stars	$1.5^{+0.3}_{-0.3}$	$3.8^{+0.8}_{-0.8}$	$8.1^{+1.7}_{-1.6}$	$18.8^{+4.0}_{-3.8}$	$22.6^{+4.8}_{-4.6}$	$4.4^{+0.9}_{-0.9}$
	CXB	0.9	2.8	7.5	15.1	34.2	31.2
Low Latitude	LHB	$88.0^{+21.3}_{-20.2}$	$77.11^{+18.7}_{-17.7}$	$2.0^{+0.5}_{-0.4}$	-	-	-
Model 3	SWCX	<34.8	<34.8	<10.8	<7.1	<1.9	-
	Hot Halo	-	$1.2^{+1.0}_{-0.8}$	$5.5^{+4.5}_{-3.4}$	$6.9^{+5.6}_{-4.3}$	$4.7^{+3.8}_{-2.9}$	-
	dM Stars	$1.4^{+0.3}_{-0.3}$	$3.6^{+0.9}_{-0.9}$	$7.6^{+1.9}_{-1.8}$	$17.5^{+4.4}_{-4.3}$	$21.1^{+5.3}_{-5.1}$	$4.1^{+1.0}_{-1.0}$
	CXB	0.9	2.8	7.5	15.1	34.2	31.2

Table 2.8: Line identification and fluxes. F, I, and R refer to the forbidden, intercombination, and resonance lines due to the fine structure splitting of the He-like K-α triplet. The OVII K-α fine structure splitting could not be resolved for mission 27.041, though a total flux of $4.74^{+1.52}_{-1.41}$ LU is observed for the triplet. All uncertainties are reported at 90% confidence level.

	Energy	Flux[a] (LU)			
Emission Line	(eV)	27.041	36.223	36.264	36.294
CVI Ly-α	368	$5.88^{+5.15}_{-3.55}$	$2.02^{+1.03}_{-0.84}$	$0.70^{+0.44}_{-0.34}$	$0.80^{+0.67}_{-0.48}$
CVI Ly-γ	459	$0.75^{+1.27}_{-0.75}$	$0.38^{+0.37}_{-0.29}$	$0.22^{+0.20}_{-0.15}$	$0.09^{+0.27}_{-0.09}$
OVII K-α F	561	-	$1.88^{+1.20}_{-0.93}$	$0.81^{+0.95}_{-0.67}$	$1.08^{+0.48}_{-0.39}$
OVII K-α I	569	-	$0.93^{+1.14}_{-0.93}$	$0.59^{+0.99}_{-0.59}$	$0.59^{+0.49}_{-0.41}$
OVII K-α R	575	-	$2.37^{+1.18}_{-0.96}$	$0.68^{+0.87}_{-0.58}$	$0.58^{+0.44}_{-0.35}$
OVIII Ly-α	654	$1.48^{+0.79}_{-0.61}$	$0.91^{+0.48}_{-0.39}$	$0.81^{+0.51}_{-0.38}$	$0.66^{+0.30}_{-0.25}$

[a] 1 LU $\equiv$ 1 photon sec^{-1} cm^{-2} sr^{-1}

Table 2.9: Ratios of lines identified in Table 2.8. Statistical uncertainties are given at the 1σ level.

	Mission Number				Thermal	SWCX
Ratio	27.041	36.223	36.264	36.294	Prediction	Prediction[a]
CVI Ly-γ/Ly-α	0.13±0.09	0.19±0.07	0.31±0.13	0.11±0.09	0.04	0.5
OVII (F+I)/R[b]	-	1.19±0.33	2.06±1.09	2.88±0.86	0.97	6.0

[a] Line ratios for SWCX are velocity dependent and poorly understood. Given values represent upper limits.
[b] Also known as the G-ratio.

As noted by a previous analysis of Mission 36.223, the Ly-γ/Ly-α ratio of H-like CVI and the G-ratio of He-like OVII can be used to constrain the allowed thermal contribution to the observed emission from these respective ions (Crowder et al., 2012). Although the ratios for SWCX spectra are velocity dependent and poorly understood, upper limits can be used to conservatively estimate the maximum likely contribution from thermal emission. Table 2.10 presents the results of this analysis. In agreement with the previous analysis, we find that thermal emission likely dominates the observed emission for both CVI and OVII at high latitudes, though the statistical uncertainties allow for a non-negligible

contribution from SWCX.

At low latitudes we find that SWCX is likely responsible for at least half of the observed OVII emission, though this seems to be due to a reduction in thermal emission, rather than an increase in SWCX. At all latitudes, we find that SWCX accounts for $\sim$1 LU[1] of OVII Kα emission. SWCX may also be dominating CVI emission in the Mission 36.264 observation, though this not statistically very significant. The strong relative contribution of SWCX to OVII is not reflected in Low Latitude Models 1 or 2, which motivated Model 3. In order to produce the OVII flux inferred from the G-ratio, the ion temperature of carbon in the ACX model (kT_C in Table 2.6) was allowed to vary independently of the temperature of the other ions. Without this additional free parameter, the normalization of the Slow SWCX component was limited by CVI emission. Compared to Model 2, this change significantly reduced in the halo emission measure, but only slightly improved the fit (See Tables 2.5 and 2.6 and Figures 2.9 and 2.10).

Table 2.10: Relative thermal contribution to observed emission lines from select ions, based on line ratios given in Table 2.9. Statistical uncertainties are given at the 1σ level.

	Mission Number			
Ion	27.041	36.223	36.264	36.294
$\mathrm{CVI}_{Thermal}/\mathrm{CVI}_{Total}$	$0.74^{+0.25}_{-0.22}$	$0.59^{+0.18}_{-0.16}$	$0.33^{+0.29}_{-0.24}$	$0.79^{+0.21}_{-0.22}$
$\mathrm{OVII}_{Thermal}/\mathrm{OVII}_{Total}$	-	$0.86^{+0.14}_{-0.17}$	$0.50^{+0.50}_{-0.24}$	$0.32^{+0.20}_{-0.13}$

2.5.2 The Local Hot Bubble

At high latitudes, we fit a local plasma temperature of $0.093^{+0.010}_{-0.019}$ keV ($1.08^{+0.12}_{-0.22}$ MK), assuming solar abundances. This value is within statistical uncertainty of several previously reported values. Snowden et al. (1998) reports 0.10 keV from an analysis of the

[1] 1 LU $\equiv$ 1 photon sec^{-1} cm^{-2} sr^{-1}

RASS data. More recently, Liu et al. (2017) reanalyzed the RASS data taking contributions from SWCX into account and reports a temperature of 0.097 keV. Smith et al. (2014) modeled SWCX contributions to the DXS HISM spectrum (Sanders et al., 2001) with the same ACX model used for this analysis and reports a LHB temperature of 0.097 keV with depleted abundances. With the exception of the DXS observation, the data presented here have significantly better energy resolution than these previous observations.

The aforementioned analyses rely on data below the carbon edge at 282 eV. Observations that include higher energy data including OVII Kα emission at $\sim$570 eV tend to fit higher temperatures that are inconsistent with lower energy observations (e.g. Yoshino et al. (2009)). For the observations analyzed here, we find that the temperature of the LHB emission is very sensitive to the intensity of OVII. If included in the fit, we derive a low latitude LHB temperature that is inconsistent with the high latitude observations. However, if the energy range is restricted to $<$500 eV during the fit, then the 90% confidence temperature ranges agree at both latitudes (See Tables 2.3 to 2.6). As mentioned in Section 2.5.1, OVII is likely dominated by SWCX at low latitudes, for which we do not have accurate spectral models. Therefore, in this analysis the LHB temperature at low latitudes is fixed to the value derived from the low energy spectrum before fitting the entire 100–1100 eV range.

The best-fit LHB emission measures correspond to ROSAT R12 band fluxes of 165 and 459 counts s^{-1} $arcmin^{-2}$ at low and high latitudes, respectively (Table 2.7). Correcting for SWCX, these values are comparable to those reported by Snowden et al. (1998), who gives a combined LHB+SWCX flux of 398 and 561 counts s^{-1} $arcmin^{-2}$ for these two targets. For reference, Uprety et al. (2016) estimates that SWCX contributes an average of $\sim$170 counts s^{-1} $arcmin^{-2}$ to R12 across the entire sky. Moreover, the factor of 2.8 relative change in intensity that we observe is consistent with the spatial extent of the

LHB varying by a factor of ∼3 and being greater at high Galactic latitudes (Snowden et al., 1998).

2.5.3 Hot Galactic Halo

A single temperature halo model was used for this analysis, with a best-fit temperature of ∼0.2 keV at both low and high latitudes, assuming solar abundances. This is consistent with other observations which fit single temperature halos with $kT \sim 0.2$ keV (Yoshino et al., 2009). However, previous observations have also established that the thermal and spatial distribution of the halo is complex (e.g. Kuntz & Snowden (2000), Lei et al. (2009), Yao et al. (2009), & Liu et al. (2016)), and a single temperature model is therefore likely overly simplistic. One should be careful not to ascribe too much significance to any physical interpretation. Moreover, the halo component at low latitudes has a strong coupling to other model components of comparable brightness at both lower and higher energies.

Motivated by Kuntz & Snowden (2000), we investigated a two temperature halo model with $T_C = 0.099$ keV and $T_H = 0.249$ keV. For the low latitude observations, interstellar absorption makes us insensitive to the cooler component, so the model is effectively reduced to a single temperature model once again. For this target and field of view, Snowden et al. (1998) estimates that 96% of the ROSAT R12 flux originates locally (LHB+SWCX). The high latitude observations, in contrast, should be sensitive to contributions from a distant ∼0.1 keV halo component. Despite the sensitivity however, attempts to include a cooler halo component result in zero fit emission measure and no improvement to the fit.

2.5.4 Stellar Contribution

At low latitudes, we observe emission that is fit by a 0.95 keV APEC plasma (Model 1). Since it is difficult to provide physical justification for a truly diffuse source at this temperature, we investigated stars as a possible source of the emission. Following Kuntz & Snowden (2001) and Masui et al. (2009), we constructed a spectral model of dwarf M stars (dM stars), taking into account Galactic absorption and spatial star distribution. Stellar spectra were approximated as a two temperature plasma with $T_C = 0.138$ keV and $T_H = 0.78$ keV, based on ROSAT observations of nearby dM stars (Giampapa et al., 1996). In this model, each temperature component has a line-of-sight emission measure density of 4.13×10^{14} cm^{-5} sr^{-1} kpc^{-1} in the plane of the Galaxy at the solar radius. Parameters detailing the spatial distributions and densities for dM stars of different ages are given in Table 2.11. Absorption was calculated from Equations 1 and 2 in Ferrire (2001). We refit the low latitude observations, replacing the 0.95 keV plasma with this stellar model scaled by a single free normalization parameter (Models 2–3). The best-fit normalizations are $1.19^{+0.25}_{-0.24}$ and $1.11^{+0.28}_{-0.27}$ for Model 2 and 3, respectively. These values are consistent with the value of 1.27 reported by Masui et al. (2009), which used a similar model adjusted for viewing geometry. Compared to Model 1, the halo emission measures in Low Latitude Models 2 and 3 are reduced due to emission from the cooler stellar component (the halo emission is further reduced in Model 3 due to stronger relative SWCX contribution as discussed in Section 2.5.1).

As an additional check, we also tested for a stellar contribution at high latitudes. Due to the small scale height of young dM stars (<0.15 Gyr) relative to the gas scale height, the stellar X-ray brightness is expected to drop sharply with increasing Galactic latitude. 71% of the X-ray luminosity in our model is due to young dM stars, whose vertical density distribution drops to half of the in-plane value in only 60 pc. For comparison,

Table 2.11: Parameters used in spectral emission model for dM stars.

Age (Gyr)	K_+^a (pc)	K_-^a (pc)	$c^{a,b}$	Midplane EM Density[c] (10^{14} cm^{-5} sr^{-1} kpc^{-1})
0–0.15	5000	3000	0.0140	2.93
0.15–1	2226	494.4	0.0279	0.85
1–2	2226	494.4	0.0457	0.10
2–3	2226	494.4	0.0662	0.07
3–5	2226	494.4	0.0867	0.11
5–10	2226	494.4	0.0958	0.07
Total				4.13

[a]Scale parameters for (1) and (2) of Appendix A in Bienayme et al. (1987).
[b]Dimensionless ellipticity from Notes of Table 2 in Kuntz & Snowden (2001).
[c]Adapted from Table 2 in Kuntz & Snowden (2001) and Table 3 in Masui et al. (2009).

the half-density height for neutral hydrogen is ∼110 pc at the solar radius (Dickey & Lockman, 1990). Recalculating the stellar spectral model for the high latitude target and assuming the Model 2 normalization of 1.19, we predict contributions to the ROSAT bands R45 and R67 of 12.7 and 11.2 counts s^{-1} arcmin^{-2}, respectively (9.8% and 9.1% of the total flux in these bands). Fitting for the high latitude normalization directly, we get an upper limit of 0.79 at 90% confidence, with corresponding band fluxes at this limit of 8.43 and 7.43 counts s^{-1} arcmin^{-2}. Alternatively, we can place an upper limit on the emission measure of a 0.95 keV plasma at 0.66×10^{14} cm^{-5} sr^{-1}, which contributes 6.9 and 8.8 counts s^{-1} arcmin^{-2} to R45 and R67, respectively. In both cases, the upper limits are below the predictions made by extrapolating the stellar model fit at low latitudes. Therefore we conclude that the source of this emission can be consistently explained by our stellar model, or else by a source whose surface brightness falls off at least as quickly as our model with increasing Galactic latitude.

2.6 Conclusions

Here we presented a combined analysis of the past four XQC sounding rocket observations. The three most recent observations have all been processed with an overlapped pulse fitting algorithm to improve spectral resolution and reduce processing dead-time compared to conventional processing techniques. The resulting spectra are consistent with fluxes measured by the RASS, with discrepancies comparable to the estimated contributions from SWCX. Likewise, repeated XQC observations of the same targets can be fit with identical spectral emission models, adjusting for differences in instrumental response and allowing for variable SWCX intensity. Analyses of individual emission lines of CVI and OVII permit a more robust measurement of the SWCX contribution that is less sensitive to inaccuracies in the spectral models for charge exchange. The derived parameters for the LHB and halo are consistent between high and low latitude observations, and compatible with prior values reported in the literature. Emission fit by a 0.95 eV plasma observed at low latitudes can be consistently explained by unresolved dM stars. Ultimately, the significance of many of the results is limited by large statistical uncertainty. An orbital mission of an XQC-like instrument could greatly refine our understanding of the DXRB and its origins.

Chapter 3

Optimal Filtering for Microcalorimeter Data at High Count Rates

3.1 Introduction

Optimal filtering for microcalorimeter detectors is well understood. For a signal shape $S(t)$ with Fourier transfrom s_j, this technique constructs the pulse-amplitude estimator

$$E = \sum_{j=1}^{\infty} w_j s_j \tag{3.1}$$

with weights w_j chosen to maximize the signal-to-noise ratio. Under the assumptions that the statistical properties of the noise are not changing with time and the signal shape is not changing with amplitude, it can readily be shown that these weights are given by

$$w_j = s_j^*/n_j^2 \tag{3.2}$$

where s_j^* and n_j^2 are the complex conjugate of the signal and the mean-square of the noise in the j^{th} frequency bin, respectively (for complete derivation, see e.g. Szymkowiak et al. (1993)). Since the best estimate of E is given by a single point in the time-domain, it is usually computationally efficient to do the filtering there via the Convolution Theorem. That is, the inverse Fourier transform of w_j, $W(t)$, is convolved with the time-domain signal to output the pulse-amplitude E at time $t = 0$.

$$E = W(t) \otimes S(t)|_{t=0} \tag{3.3}$$

In practice, the detector output is sampled and digitized at a rate sufficient to avoid aliasing the useful terms in Equation 3.1. The useful bandwidth extends until the frequency f_{max}, beyond which additional terms will not improve the signal to noise ratio. According to Equation 3.2, the value of f_{max} will depend on the relative behavior of the signal and noise with respect to frequency, which will depend on the details of a given detector. However, in the interest of maximizing the useful bandwidth and thus the energy resolution, detectors are conventionally designed such that their useful bandwidth is limited by reaching a noise floor, rather than by loss of signal. To achieve this goal, it can be shown that the primary decay time τ must be of order α times longer than any internal time constants of the device, where α is the dimensionless thermometer sensitivity defined by

$$\alpha \equiv \frac{d\log R}{d\log T} = \frac{T}{R}\frac{dR}{dT} \tag{3.4}$$

where R and T are resistance and temperature, respectively (for a detailed discussion,

the reader may refer to McCammon (2005)). For large values of α, this can place strong requirements on τ, limiting the speed of the detector.

The time-domain length of $W(t)$ determines the frequency resolution of the filter, with longer filters having finer frequency-space binning. Longer filters are thus more efficient for optimizing the signal-to-noise ratio (particularly when the noise spectrum contains features), providing improved spectral resolution. One major drawback of conventional optimal filtering is its requirement that pulses produced by individual photons in a single pixel have a minimum separation of at least the time-domain length of the optimized filter, $W(t)$. This can be responsible for significant processing dead-time even at modest count rates (Figure 3.1). Consequently, a compromise must usually be made between the desired spectral resolution and the corresponding dead-time when choosing the filter length. On Hitomi, for example, filters of different lengths could be chosen to produce spectra graded as "high-resolution" (longest filter) "mid-resolution" (shorter filter) or "low-resolution" (no filter)(Ishisaki et al., 2018).

In response to this problem, others have developed techniques for recovering piled-up pulses with minimal loss of resolution compared to conventional optimal filtering. Crowder et al. (2012) began the work on this technique. It required significant human intervention and was tailored to a single, relatively small data set. Fowler et al. (2015) presents a more generalized approach, but which sacrifices performance in both resolution and live-time in favor of computational efficiency. In contrast, the technique presented here makes no simplifying assumptions beyond those made by conventional optimal filtering (i.e. stationary noise and linear response), yet which is still automated and applicable to all calorimeter data sets satisfying these assumptions. The result is a general algorithm capable of recovering pulses separated by as little as the detector rise time without loss of resolution compared to conventional optimal filtering techniques.

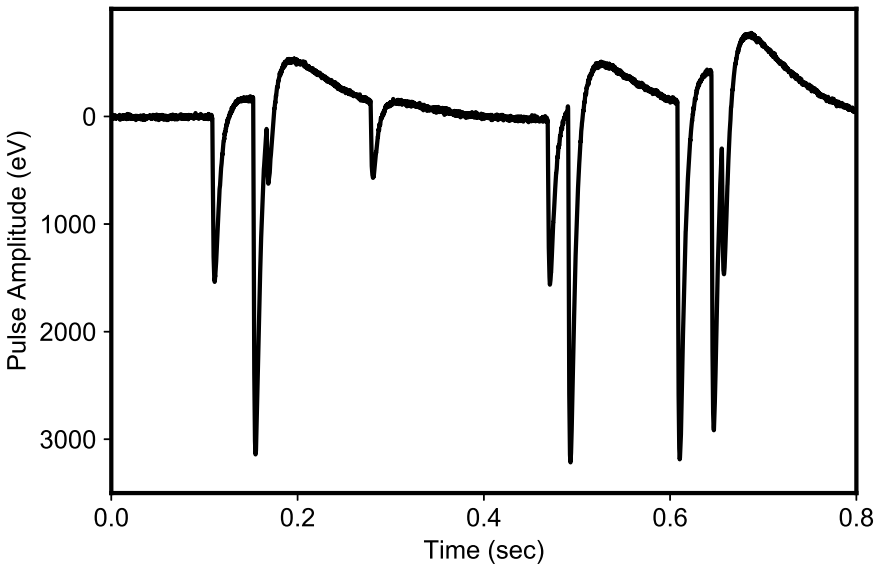

Figure 3.1: Example of pulse pile-up in XQC data ($\tau_{eff} = 9$ ms). Each of the nine pulses in this data segment have a minimum separation of less than the 200 ms time-domain filter length typically used for this detector, and therefore cannot be processed using conventional optimal filtering techniques. Using the techniques presented here, however, each of these pulses may be recovered without loss of spectral resolution.

This technique was largely developed using data from the University of Wisconsin-Madison/Goddard Space Flight Center X-ray Quantum Calorimeter sounding rocket payload, henceforth referred to as XQC. The reader may refer to McCammon et al. (2002) for a complete description of the payload. XQC is a 36-pixel, silicon thermistor microcalorimeter detector optimized for X-ray energies below 1 keV. Most pixels achieve a baseline resolution of ~ 6 eV FWHM under typical noise conditions. The pixels have an effective decay time $\tau_{eff} \sim 9$ ms, and are about 3% non-linear at 3 keV. Each pixel is approximately current biased, and the pixel voltages are digitized to 12 bits at a sampling rate of 10.4 kHz.

3.2 Technique

3.2.1 Filtering

To begin, we construct the usual filter optimized for pulse amplitude described by Equation 3.2. In practice, the s_j and n_j are determined by averaging pulses and noise spectra from calibration data (Figures 3.2a–3.2b). Since the dead-time of our technique does not depend on filter length, this length can be as long as necessary to maximize the spectral resolution, and/or as long as is necessary to contain the entire signal shape. The latter requirement ensures that the signal shape is zero at the endpoints and not meaningfully altered by windowing.

In addition to the usual amplitude filter, a filter optimized for pulse arrival-time is also constructed. To understand this filter, it is instructive to revisit the amplitude filter in slightly more detail. The complex phase of the weights in Equation 3.2 are chosen to make each term in Equation 3.1 entirely real, and thus maximize the sum by ensuring constructive phase addition of each term. Another consequence of this choice is that the

time-domain convolution of the filter with the signal shape can be expressed as a pure cosine sum, with each term constructively adding to attain a global maximum equal to E at time $t = 0$ (Equation 3.3). Therefore, the output of the amplitude filter for an isolated pulse will be maximum at the pulse arrival-time. We can take advantage of this fact to construct an arrival-time filter which maximizes the slope of the signal relative to the noise and whose time-domain convolution with the signal shape will have a zero-crossing at $t = 0$ (i.e. the pulse arrival-time). Rather than deriving such a filter from scratch, we can invoke the Derivative Theorem for Fourier Transforms, simply multiplying the amplitude filter a factor of $2\pi if$. Based on the construction of Equation 3.1, the corresponding estimator for the slope can be written as

$$E' = \sum_{j=1}^{\infty} 2\pi i f_j w_j s_j = \sum_{j=1}^{\infty} w'_j s_j \tag{3.5}$$

where the w_j are the same as in Equation 3.2. Here f_j is the central frequency of bin j, and i is the complex number $\sqrt{-1}$.

Rather than convolving these filters with short, individually triggered data segments containing a single pulse, as is done in conventional pulse processing, the filters are each convolved continuously with the entire pixel data stream. In practice, these continuous, large convolutions can be done efficiently in the frequency-domain and account for a small fraction of the total computational cost. For ease of processing, these two filtered data streams are split into short segments typically a few times the length of the filter and containing several pulses. For an isolated pulse, the zero-crossing time in the arrival-time-filtered segment and the value of the amplitude-filtered segment at this time give the best estimates of the pulse arrival-time and amplitude. There is no additional information in any of the other points. However, there is substantial ringing at times up to the filter

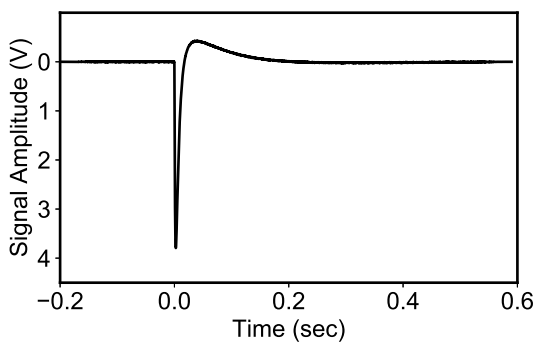

(a) Average Pulse Shape

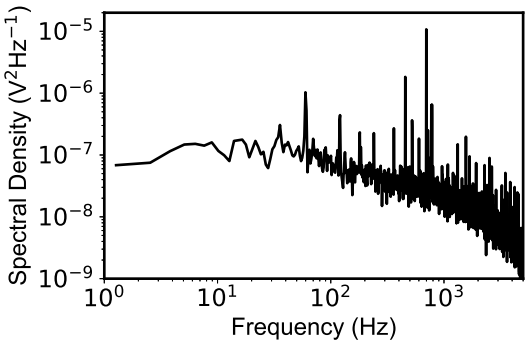

(b) Average Noise

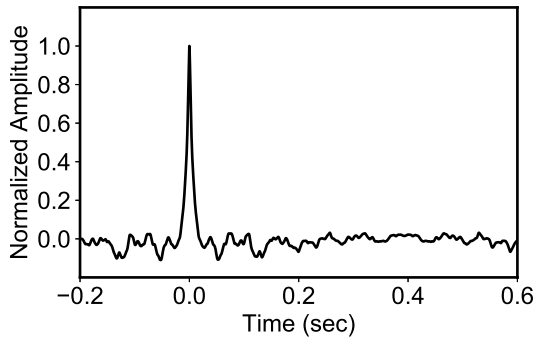

(c) Pulse Filtered for Amplitude

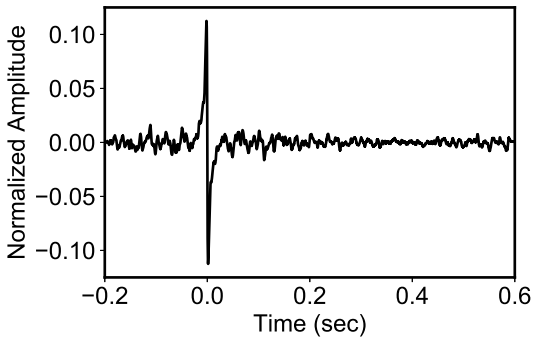

(d) Pulse Filtered for Arrival-Time

Figure 3.2: (a) Average pulse shape made by averaging several isolated 3.3 keV calibration pulses. (b) Average detector noise spectral density, also made from calibration data. (c) Average pulse shape after being optimally filtered for amplitude and normalized to unity amplitude. As in conventional optimal filtering, the value at the pulse arrival-time gives the best (highest signal-to-noise) estimate of pulse amplitude. (d) Average pulse shape after being optimally filtered for arrival-time and normalized such that the mean-square fluctuations of the filtered noise are the same as the pulse-amplitude-filtered data. The zero-crossing near $t = 0$ gives the best estimate of pulse arrival-time.

length away from the pulse which can distort the signal at the arrival-times of other pulses. This ringing is entirely deterministic given the amplitude and arrival-time (Figures 3.2c–3.2d), and can be subtracted. Taking advantage of the linear properties of the convolution, and assuming approximate linearity of the detector, a filtered data segment containing overlapped pulses will simply be the sum of the appropriately scaled and shifted copies of an isolated pulse (Figure 3.3). The values of the filtered segments at a pulse arrival-times are no longer determined by the amplitude and arrival-time of a single pulse, but instead depend on the amplitudes and arrival-times of every pulse within the filter convolution length.

3.2.2 Fitting

Templates for fitting pulses in the filtered data streams are made by convolving the same filters with the average signal shape. Because pulses in the data stream arrive randomly with respect to the sampling times, a dictionary of fitting templates is made for different inter-sample arrival-times. Provided that all of the signal power is below the Nyquist frequency, an exact interpolation can be done by advancing the phases in frequency space, most conveniently by using the Shift Theorem for Fourier Transforms.

Pulses in each segment of the data stream are detected as local maxima in the amplitude-filtered segment, since this provides the best signal-to-noise. However, since the ringing caused by the largest pulses can be comparable in amplitude to smaller pulses, pulses must be detected iteratively, starting with the largest pulses. In the initial pass, the peaks of all pulses larger than the expected ringing are identified in the amplitude-filtered segment, and initial guesses for their amplitudes and arrival-times are made from the value and time of these peaks, respectively. These two numbers are used to scale and time-shift both amplitude- and time-filtered pulse templates, and to construct a model

which predicts the ringing due to each pulse at all times in the respective filtered segments.

After adding the templates of multiple pulses together, the ringing from every other pulse disturbs the peak amplitudes and zero-crossing times from the initial guess. So, the amplitudes and arrival-times must be adjusted until the sums of all the templates are consistent with the observed values in each of the two filtered segments at the pulse arrival-times, which are the only points with useful information. Note that, for N pulses, there are 2N data points (the deviation from zero in the arrival-time-filtered stream and the amplitude in the amplitude-filtered stream at the arrival-time of each pulse) and 2N parameters (arrival-time an amplitude for each pulse), so there will always be an exact solution, even in the presence of noise (though the noise will affect the accuracy of the derived solution). A model based on this best-fit solution can then be subtracted from the filtered segments, reducing the ringing and enabling detection of smaller pulses in the next iteration. With each iteration, newly detected pulses are added to the previous ones, and a new self-consistent set of amplitudes and arrival-times is found.

This process continues until all pulses above the detector noise level are detected and included in the fit. For our filters, the largest ringing peak is $< 1/10^{\text{th}}$ the value of the main peak, allowing a $10\times$ reduction in the pulse-detection threshold with each iteration. Note that, since we are working with pulse-amplitude filtered data, the noise-level is intrinsically related to the optimally filtered energy resolution. Based on our original formulation in Equation 3.1 and assuming that the noise is uncorrelated at different frequencies, the mean-square fluctuations in our estimation of energy is given by

$$\Delta E^2 = \sum_{j=1}^{\infty} (w_j n_j)^2 \tag{3.6}$$

This is equivalent, however, to the mean-square fluctuations in the amplitude-filtered

output at *all* times, not just at the pulse arrival-times. In practice, this means that the 5σ lower threshold for pulse detection is about two times the optimally filtered FWHM, which is considerably better than typical non-optimal triggering algorithms. Moreover, this means that any ringing due to undetected pulses will at most contribute $1/10^{\text{th}}$ of this lower threshold to the residuals, and have negligible impact on the accuracy of the fit.

Pulses that arrive within a filter-length of the end of each segment may be affected by ringing of pulses that arrive in the following segment but are yet undetected. For this reason, consecutively analyzed segments are chosen to overlap by at least the filter length. Pulses whose templates are completely contained in one segment are subtracted from the overlapped region in the following segment. The remaining pulses, which may be affected by pulses in the following segment and vice-versa, are passed to the following segment to be refit with this new information. This process ensures that the final fit of every pulse in the data stream includes the information of every other pulse that arrives within a filter-length before or after it.

3.2.3 Addressing Changing Pulse Shape

If the data stream contains pulses extending in energy beyond where the assumption of constant pulse shape holds, then the subtraction may leave significant residuals which impair the process of detecting smaller pulses. In this case, filtered fitting templates for pulses of multiple energies may be included in the dictionary of fitting templates (Figure 3.4), allowing one to select the appropriate shape for each pulse depending on its amplitude. These energy-dependent pulse shapes can be determined empirically with calibration data, or else modeled with knowledge of the particular detector. It is important to note that, as these shapes only depend on the energy of a single pulse, they do not

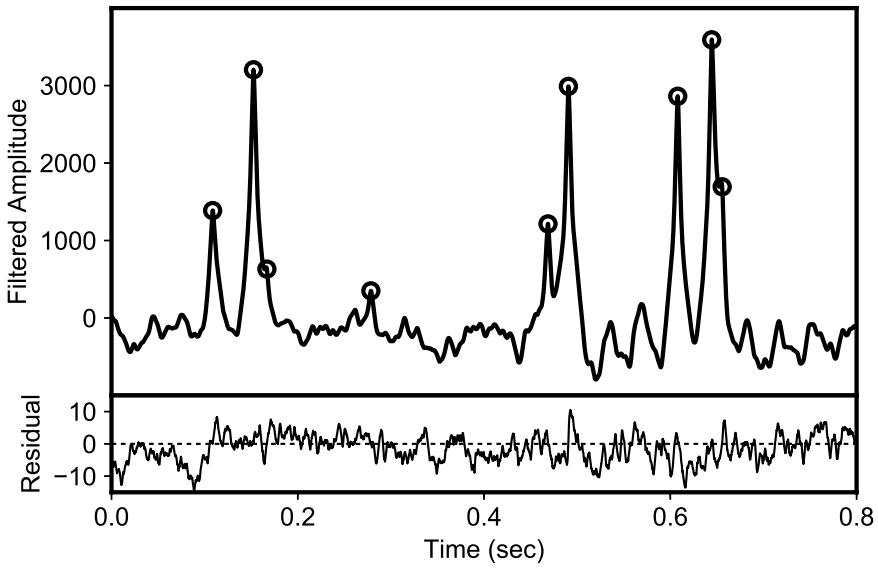

(a) Amplitude Filtered Data

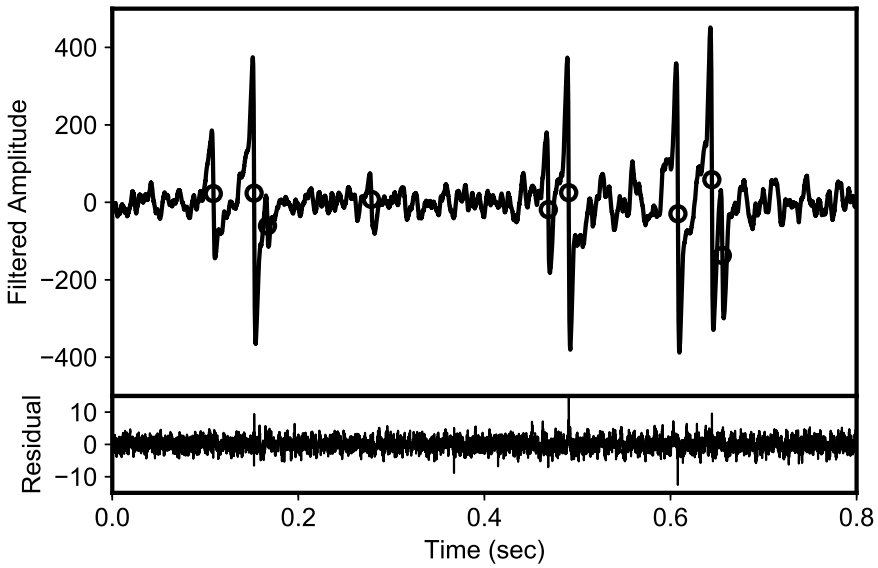

(b) Arrival-Time Filtered Data

Figure 3.3: Data segment from Figure 3.1 after optimally filtering for (a) amplitude and (b) arrival-time. Open circles mark the filter outputs at the pulse arrival-times, and are the only points considered in the solution. Residuals at all times are shown for the model constructed from the best-ft solution for pulse amplitudes and arrival-times.

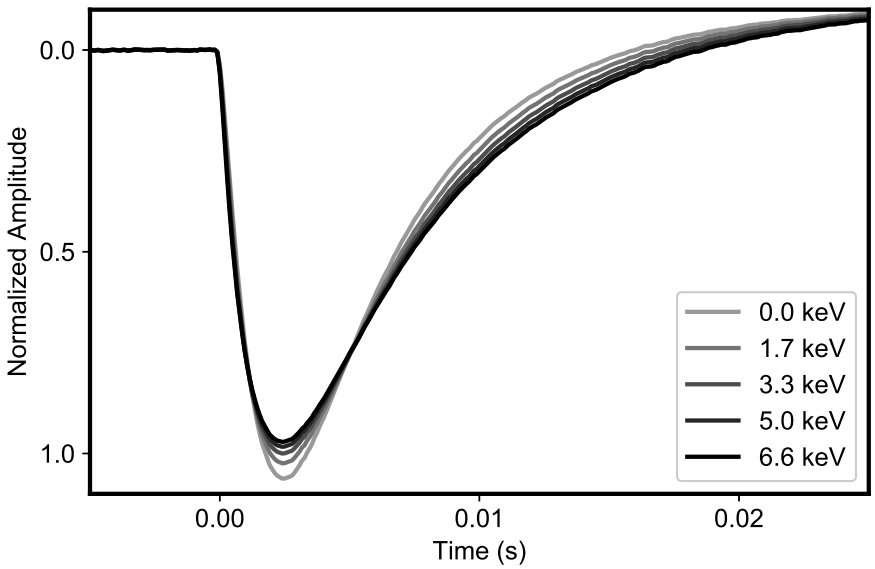

Figure 3.4: Pulse templates for pulses of different pulse energies. For illustrative purposes, the magnitude of the shape changes are exaggerated. Energy-dependent fitting templates are made by convolving each of these pulse templates with the optimal filters for pulse-amplitude and pulse-arrival-time.

account for shape changes that result due to overlap. For instance, a small pulse that arrives while the pixel is still cooling from a large pulse may have a different shape than an isolated small pulse of the same amplitude. However, because the fit amplitudes of the two small pulses in this scenario would be the same, the same pulse shape would be used. More elaborate algorithms could in principle account for this effect as well, but have not been implemented in this work. Moreover, in the limit of extreme non-linearity, the assumption of stationary noise breaks down and different techniques are needed for optimization altogether (Fixsen et al., 2002).

3.2.4 Gain Correction

The best-fit model parameters provide a "raw" pulse amplitude which needs to be converted to the photon energy. Since resolving powers ($E/\Delta E$) of these detectors can exceed several thousand, even small non-linearities that have negligible affect on pulse shape or noise must still be taken into account. Therefore, a correction function fit to an adequate amount of calibration data is used to convert the filtered pulse amplitudes to photon energy. This non-linearity is primarily due to reduced thermometer sensitivity at higher temperatures, and to a lesser extent to increased detector heat capacity at higher temperatures. For pulses sufficiently separated in time, only the raw pulse amplitude is needed to determine the corresponding energy from this correction function. However, pulses that arrive while the pixel is still cooling from the preceding event will begin at a higher temperature than isolated pulses, and have will therefore have less thermometer response. In this case the raw amplitudes of both pulses and their separation time is required to determine the energy of the second pulse (Figure 3.5). The appropriate correction can once again be determined from an adequate amount of calibration data, or else calculated from a detector model.

3.2.5 Livetime Determination

This technique can yield a live-time efficiency of nearly 100% up to high count rates. The only irreducible source of dead-time arises from pulses with time separations comparable to the rise time of the detector ($\sim$ 1.5 ms for XQC). As the time separation of two pulses approaches this limit, the combined shape of the pulses begins to approximate the shape of a single pulse, and only the sum of pulse amplitudes is constrained. For even closer spacings, the event will be indistinguishable from a single pulse with the total energy of

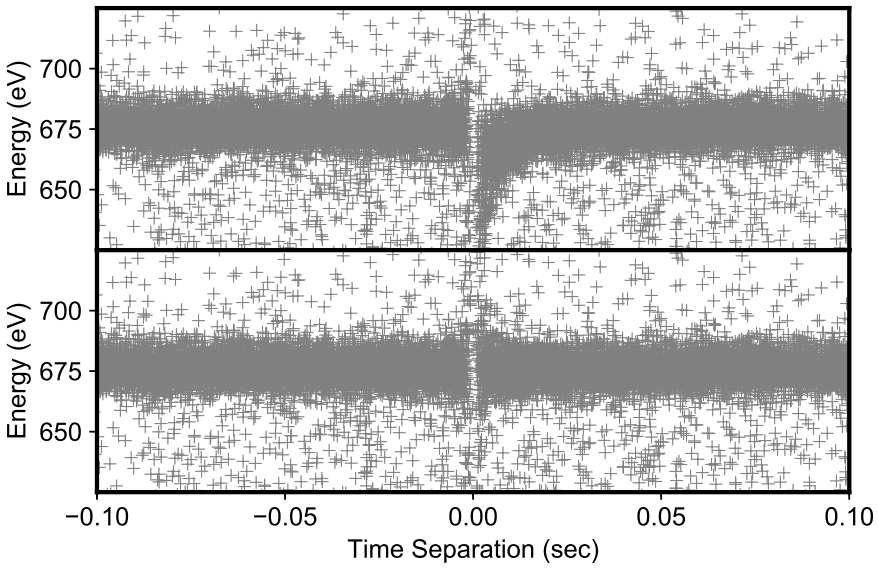

Figure 3.5: Scatter plot of pulse energy versus time separation from overlapping pulse at $t = 0$, shown for F Kα line at 676.8 eV. Pulses plotted at negative times arrive prior to the overlapped pulse, positive times after. Top: Pulses that arrive shortly after another pulse (within ~ 20 ms, or a couple decay times) experience low detector gain due to the pixel still being hot from the preceding event. Bottom: Same pulses, after correcting for this effect. Spectral resolution does not degrade until pulses are separated by as little as a few rise times, at which point the sum of pulses cannot be distinguished from a single pulse.

the two pulses. However, the fraction of pulses affected is very small for count rates much less than the inverse of the detector rise time.

Another source of dead-time encountered in practice is the inability to fit and subtract features in the convolved data streams with unexpected shapes. Such features include thermal and electrical cross-talk with other pixels, direct photon hits to the thermometer, and large pulses that saturate the electronics. Without templates for fitting these events, they cannot be included properly in the model, resulting in residuals that impact the solution for amplitudes and arrival-times of good pulses. While it is theoretically possible construct templates for such non-pulse features, this is difficult to implement in practice and is usually not worth the small improvement in live-time compared to simply skipping over affected segments. To minimize the live-time cost associated with these non-pulse features, affected segments can be masked from the data stream prior to convolution, to avoid "smearing out" the effects of the features over the filter length.

3.3 Performance Tests

To evaluate this technique, we have compared it to conventional optimal filtering of the same data. Here we use 3 hours of XQC data. There are two X-ray sources in this data: the first is a ^{41}Ca source producing K Kα and Kβ, and the second is an alpha-excited fluorescent source producing Kα lines from C, O, F, Al, and Si. The combined rate of these two sources averages 1.8 counts/second/pixel, and spans an order of magnitude in photon energy (277–3590 eV). Figure 3.6 depicts the spectra of the F Kα line that result from the two different processing techniques. Even at this modest count rate, the processing live-time of conventional optimal filtering is only 12% when using the 200 ms filter needed for near-optimum spectral resolution. In comparison, the technique

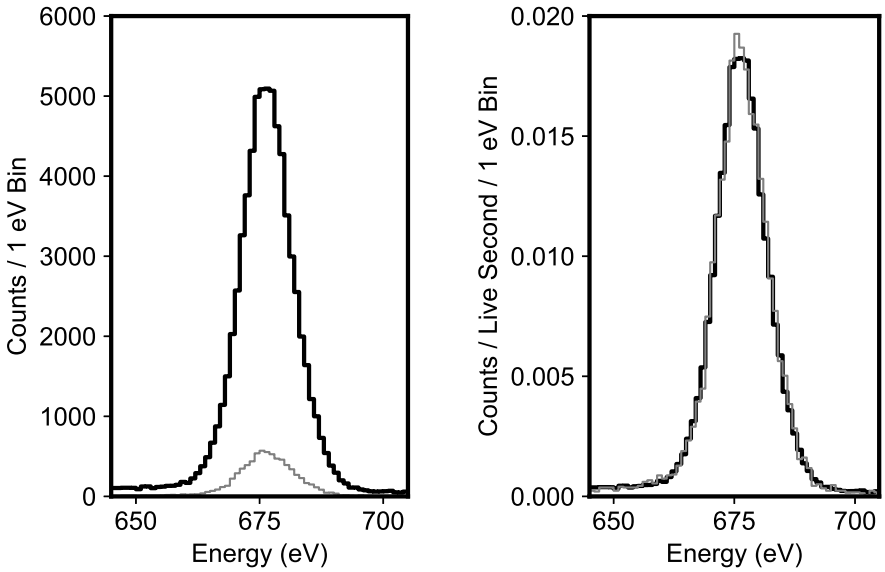

(a) Comparison of Total Counts (b) Comparison of Normalized Counts

Figure 3.6: XQC spectra of F Kα line at 676.8 eV processed with this technique (heavy black line) compared to conventional optimal filtering of the same data set using a 200 ms filter (light gray line) (a) Histograms of total counts illustrate the $> 8\times$ improvement in processing live-time with this technique. (b) Histograms normalized to live-time illustrate identical spectral resolution.

presented here achieves 98% live-time while achieving identical spectral resolution to the conventional technique (the 2% dead-time is almost entirely due to background events saturating the electronics). This comparison emphasizes the dramatic gains that can be achieved under normal operating conditions.

3.4 Computational Requirements

This purpose of this section is only to provide a general sense of the computational requirements for this technique as it is currently implemented. The process is coded entirely in Python, and is executed on a reasonably equipped desktop computer (4-core,

3.2 GHz processor; 12 GB of RAM). While some effort has been made to improve the efficiency of the algorithm, we have not conducted a detailed optimization. In practice, the 2N equation solution described in Section 3.2.2 is done with MINPACK's Levenberg-Marquardt least squares algorithm (More et al., 1980), since the estimation of arrival-time lacks an analytic solution (i.e. the solution cannot be found by matrix inversion). This is currently the most computationally intensive step of this process, accounting for $> 60\%$ computation time. Therefore, the computation time depends more on the number of pulses than it does on the length of the data record. A single 3.2 GHz processor core can process ~ 10 pulses/second using the current algorithm. This processing requirement puts the algorithm out of reach for realtime onboard analysis of a large array on a spacecraft. However, with some cleverness it may be possible to find an algorithm that gives an approximate solution of adequate accuracy with much smaller computational demands.

3.5 Conclusions

We have presented a technique for processing microcalorimeter data capable of recovering pulses separated by as little as the rise time of the detector without loss of resolution compared to conventional optimal filtering. As demonstrated here, this technique can dramatically improve instrument throughput even at modest data rates. As microcalorimeters have gained in popularity in a variety of applications, this technique may have wide-reaching applications across many scientific disciplines.

This approach is ideal for data sets like the one presented here, for which the assumptions of linear response and stationary noise are approximately satisfied. More tests are in progress to determine how this technique performs when these assumptions are less valid. It is likely possible to reduce computation time modestly with optimization of the

existing algorithm and its coding. Major gains, however, will require making some simplifying approximations to the overlapped pulse solution (e.g. linearizing the solution for pulse arrival-time), and it remains to be seen how performance would be affected.

Appendix A

Details of XQC Flight Analysis

Payload, flight, and analysis parameters are presented in this appendix.

Table A.1: Naming conventions and basic descriptions of flights. All flights launched from White Sands Missile Range, NM (lat, long = 32:40:00N, 106:20:00W)

Flight Number	Mission Number	File Name	Date	Galactic Target
Flight 4	36.223	h6r19	1 May 2008	High Latitude
Flight 5	36.264	j5r62	5 Nov 2011	Low Latitude
Flight 6	36.294	k8r61	3 Nov 2013	Low Latitude

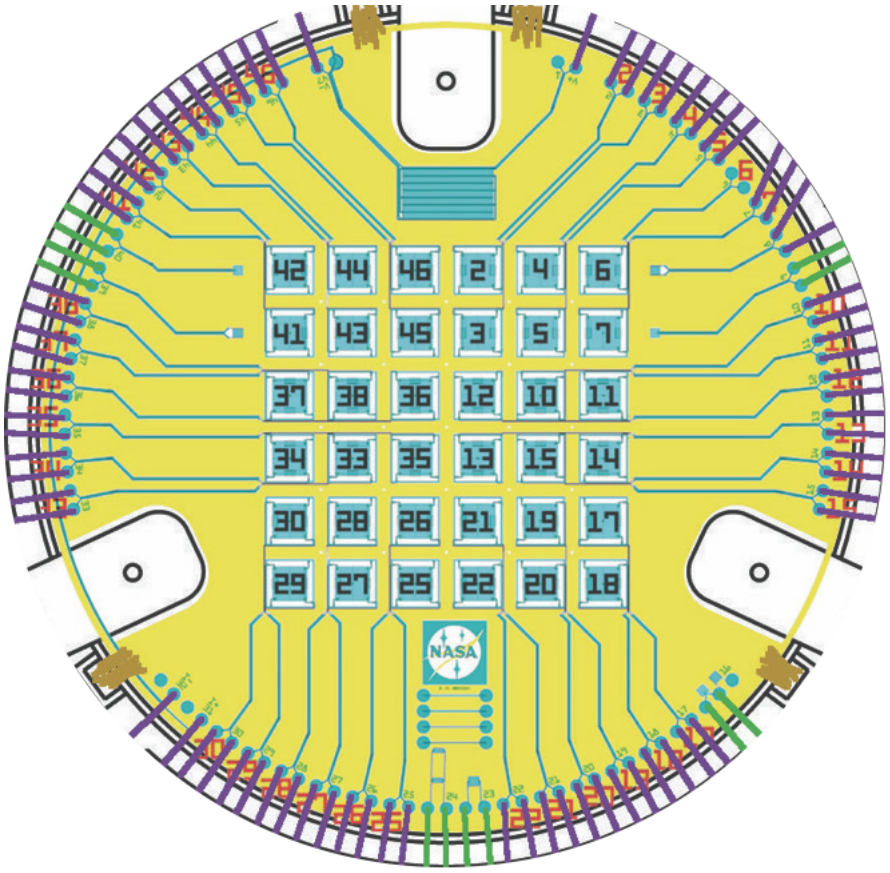

Figure A.1: Detector array layout and Goddard Pixel Numbers. Corresponding channel numbers are listed in Table A.2

Table A.2: Equivalent detector channel naming conventions. Only Goddard number refers to physical pixel on the array (See Figure A.1). Wisconsin number, readout channel, and cardfile channel all refer to readout electronics.

Wisconsin Number	Goddard Number	Readout Channel	Cardfile Channel
0	46	A0	a1_0
1	45	A1	a1_1
2	44	A2	a1_2
3	43	A3	a1_3
4	42	A4	a1_4
5	41[a]	A5	a1_5
6	38	A6	b1_0
7	37	A7	b1_1
8	36	A8	b1_2
9	35	A9	b1_3
10	34	A10	b1_4
11	33	A11	b1_5
12	15	B0	a2_0
13	14	B1	a2_1
14	13	B2	a2_2
15	12	B3	a2_3
16	11	B4	a2_4
17	10	B5	a2_5
18	7[b]	B6	b2_0
19	6[a]	B7	b2_1
20	5	B8	b2_2
21	4	B9	b2_3
22	3	B10	b2_4
23	2	B11	b2_5
24	30	C0	a3_0
25	29	C1	a3_1
26	28	C2	a3_2
27	27	C3	a3_3
28	26	C4	a3_4
29	25	C5	a3_5
30	22	C6	b3_0
31	21	C7	b3_1
32	20	C8	b3_2
33	19	C9	b3_3
34	18	C10	b3_4
35	17	C11	b3_5

[a] After damage incurred in July 2013 to Goddard pixel number 41, Goddard pixel number 6 was rerouted via readout A5.
[b] Goddard pixel number 7 is not bonded. Instead, readout B6 is used to read out the frame thermometer.

Table A.3: Timeline for Flight 4. T=0 occurs at 5:30 UT 1 May 2008.

Event	T-Time (s)
Terrier ignition	0.0
Rail release	0.6
Terrier burnout	5.2
Frequency control pulse (10 s delta)	10.0
Black Brant ignition	12.0
S-19 canard de-couple	18.0
Black Brant burnout	44.4
Magnet controller “R” ON	52.0
Pump valve OPEN	61.0
Magnet controller “X” ON	61.0
De-spin	64.0
Payload separation	68.0
Nose tip eject	72.0
ACS1 move to nadir	74.0
ACS1 complete	90.0
160 km up leg (nominal)	112.8
Exp. bias #1 OFF	114.0
ACS2 move to solar horizon	116.0
Gate valve OPEN	116.6
ACS2 complete	131.0
Magnet controller “R” OFF	131.0
ACS3 move to science target	145.0
ACS3 complete	160.0
Apogee (nominal)	269.3
Exp. bias #2 OFF	418.0
Gate valve CLOSE	420.0
160 km down leg	425.8
Pumping valve CLOSE	435.0
Reentry door CLOSE	450.0
Payload spin-up	453.0
300,000 ft down leg (nominal)	468.0
ACS vent	493.0
Ballistic impact	514.4
Chute deploy	618.5
ACS off	900.0
Payload impact	923.3
Frequency control OFF	1200.0
Magnet controller “X” OFF	1200.0

Table A.4: Timeline for Flights 5 and 6. Flight 5 T=0 occurs at 7:00 UT 5 November 2011. Flight 6 T=0 occurs at 9:25 UT 3 November 2013.

Event	T-Time (s)
Terrier ignition	0.0
Rail release	0.6
Terrier burnout	5.2
Frequency control pulse (10 s delta)	10.0
Black Brant ignition	12.0
S-19 canard de-couple	18.0
Black Brant burnout	46.0
Pump valve OPEN	61.0
Magnet controller “X” ON	62.0
De-spin	63.0
Payload separation	66.0
Nose tip eject	70.5
ACS#1 move to intermediate target	74.0
ACS#2 move to science target	96.0
ACS#2 complete	104.0
160 km up leg (nominal)	110.6
Exp. bias #1 OFF	114.0
Gate valve OPEN	125.0
Exp. bias #2 OFF	268.2
Apogee (nominal)	274.4
Exp. bias #3 OFF	426.0
Gate valve CLOSE	436.0
160 km down leg	444.8
Pumping valve CLOSE	445.0
Slammer door CLOSE	450.0
Payload spin-up	453.0
300,000 ft down leg (nominal)	473.3
ACS vent	493.0
Ballistic impact	555.0
Exp. bias #4 OFF	615.7
Chute deploy	650.0
ACS off	900.0
Payload impact	913.7
Frequency control OFF	1500.0
Magnet controller “X” OFF	1500.0

Table A.5: Flight 4 filter stack. Thicknesses of aluminum and polyimide are fit to Philippe transmission data. Oxide thickness is fixed. Total aluminum thickness includes oxidized layer (i.e. oxide model only contains oxygen).

	Thickness (Å)			Support	Luxel
Filter	Total Al	Oxide	Polyimide	Mesh	R/N
130 K	221.1	50.0	1629.6[a]	UW0809W1_130K	-
Outer 2 K	288.3	50.0	491.0	UW0749W4_2K	11828-1
Inner 2 K	251.5	50.0	1365.4 [a]	UW0625W2_2K	-
Outer Detector Box	206.8	50.0	336.4	-	12579-2
Inner Detector Box	179.6	50.0	431.8	-	12576-2

[a] Parylene

Table A.6: Flight 5 filter stack. See Table A.5 for description.

	Thickness (Å)			Support	Luxel
Filter	Total Al	Oxide	Polyimide	Mesh	R/N
130 K	167.3	50.0	795.7	UW1115W1_130K	13537-2
Outer 2 K	210.0	50.0	680.0	UW1134W3_2K	30061/30048.3
Inner 2 K	210.0	50.0	680.0	UW1134W1_2K	30061/30048.3
Outer Detector Box	201.1	50.0	718.8	-	11933-4
Inner Detector Box	214.5	50.0	659.2	-	12576-2

Table A.7: Flight 6 filter stack. See Table A.5 for description.

	Thickness (Å)			Support	Luxel
Filter	Total Al	Oxide	Polyimide	Mesh	R/N
RT	210.2	50.0	695.2	UW1317W1_RT	32124-2
130 K	188.7	50.0	679.8	UW1315W1_130K	32123
Outer 2 K	230.4	50.0	914.9	UW1315W2_2K	32127-2
Inner 2 K	210.0	50.0	680.0	UW1134W1_2K	30061/30048.3
Outer Detector Box	222.3	50.0	672.0	-	32127-1
Inner Detector Box	187.2	50.0	677.4	-	11895-2

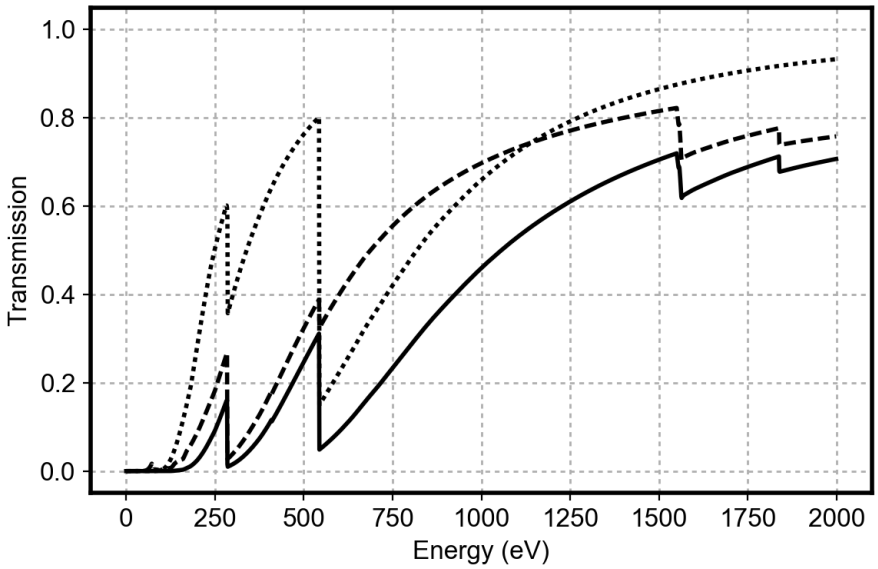

Figure A.2: X-ray transmission of clean filter stack (dashed line), filter contamination (dotted line), and contaminated filter stack for Flight 4 (solid line). Clean model parameters in given in Table A.5. The contamination transmission is modeled as 10.1 μg/cm^2 of carbon in series with 0.865 μm thick water ice, each covering 100% of the filter area. All transmission curves are averaged over the full field of view.

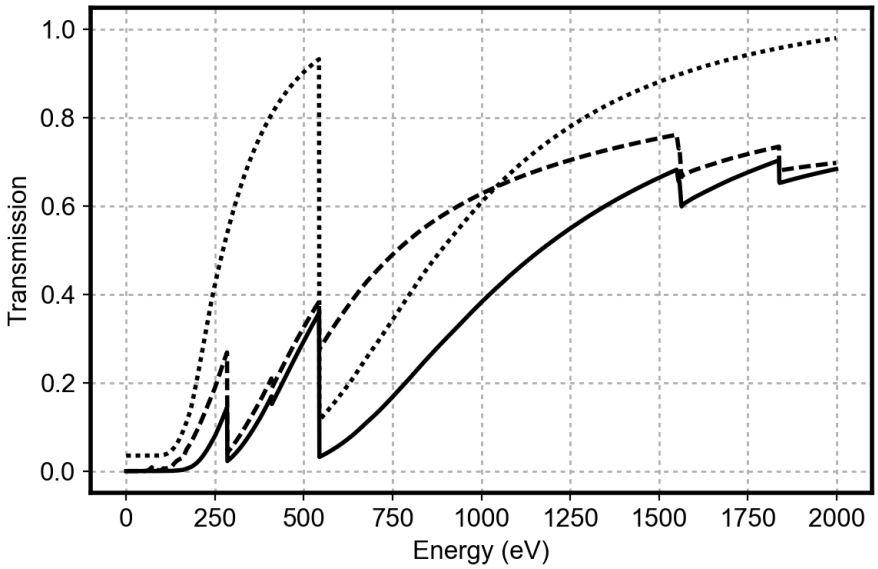

Figure A.3: X-ray transmission of clean filter stack (dashed line), water ice contamination (dotted line), and contaminated filter stack for Flight 5 (solid line). Clean model parameters in given in Table A.6. The contamination transmission is modeled as 0.785 μm thick ice instantly deposited at the time of gate-valve open (T=125), and growing at a rate of 3.8 nm/sec throughout the observation and covering 96% of the filter area (ice growth is assumed to stop when gate-valve closes at T=436). All transmission curves are averaged over the full field of view.

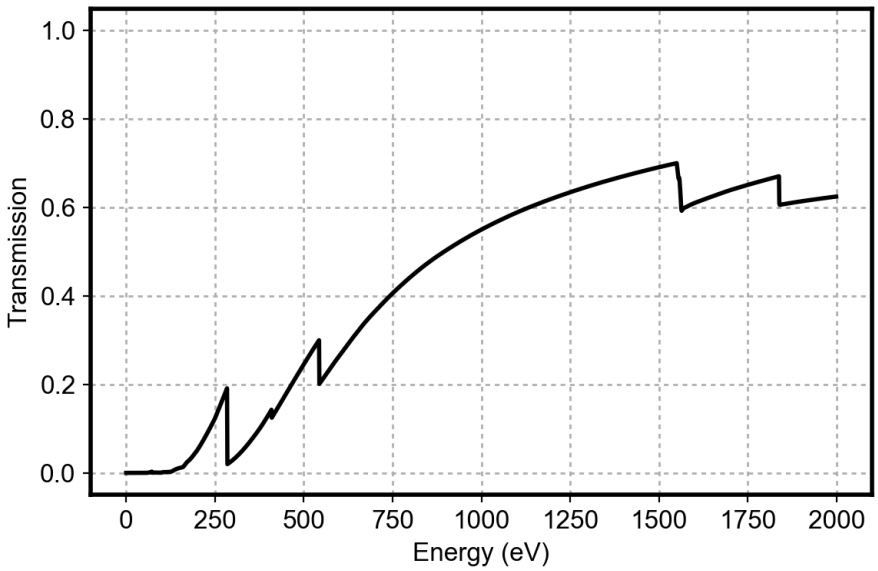

Figure A.4: X-ray transmission of Flight 6 filter stack using model parameters in Table A.7. Transmission is averaged over the full field of view.

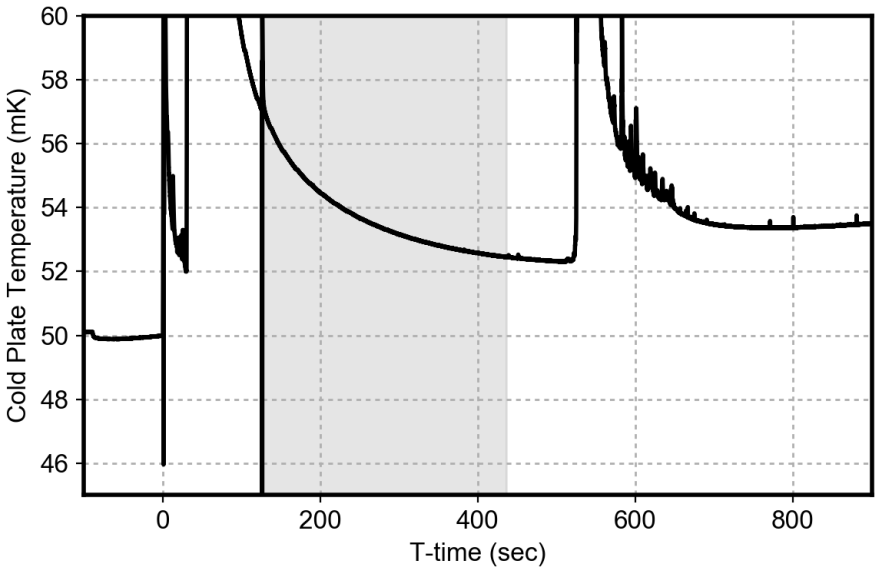

Figure A.5: Coldplate temperature during the Flight 6 observation used to correct gain drift. Shaded region corresponds to times that the gate-valve is open.

Table A.8: Baseline energy resolution as determined by choice of optimal filter

	Baseline Resolution (eV FWHM)		
Pixel	Flight 4	Flight 5	Flight 6
0	7.12	7.91	7.34
1	7.31	6.13	6.24
2	7.75	7.00	7.03
3	7.11	7.03	6.48
4	7.33	7.04	7.30
5	7.34	-	8.21
6	16.41	8.73	7.83
7	6.96	-	-
8	11.17	7.44	7.29
9	6.79	6.68	7.02
10	7.60	7.17	6.63
11	6.37	7.47	6.94
12	9.08	6.30	7.06
13	6.33	6.87	6.29
14	7.65	6.36	5.96
15	6.22	6.33	6.00
16	5.79	6.25	7.59
17	6.45	6.41	5.98
18[a]	-	-	-
19	6.95	-	-
20	7.72	7.79	7.06
21	7.66	6.55	6.94
22	7.51	6.47	6.16
23	5.73	6.55	6.12
24	7.31	7.19	-
25	7.72	6.31	6.71
26	7.72	8.29	10.39
27	15.01	6.61	7.65
28	8.43	8.29	7.34
29	6.17	6.42	7.89
30	7.43	6.61	6.30
31	6.64	7.46	-
32	6.42	6.46	6.98
33	6.69	7.21	8.76
34	6.81	6.47	7.06
35	6.92	6.39	6.14

[a] Frame Thermometer

Table A.9: Best fit parameters for pixel gain during the Flight 4 observation. Gain function has the form $g(t) = g_\infty(1 - \Delta g e^{\frac{t_0 - t}{\tau}})$, where t is T-time in seconds. Parameters fit to the pulse-height of the K Kα line at 3313.8 eV when the gate-valve is open. Pixels with only g_∞ given are fit to constant gain. Dividing the fit pulse heights by the gain will normalize the mean of the K Kα line to 3313.8 eV.

Pixel	g_∞	Δg	t_0 (s)	τ (s)
0	0.9985	-	-	-
1	1.1884	1.60e-01	148.5	33105.4
2	1.0008	1.19e-01	183.9	79.6
3	0.9983	1.90e-03	146.1	54.2
4	0.9972	9.99e-11	146.1	1386.1
5	1.0097	2.61e-03	146.1	4474.8
6	1.0017	1.00e-10	191.3	545.8
7	0.9990	9.45e-03	146.8	10.9
8	1.0010	4.98e-02	185.3	59.7
9	0.9939	9.99e-11	146.1	523.9
10	0.9976	8.28e-02	225.1	67.8
11	1.0014	-	-	-
12	0.9954	3.98e-02	213.6	64.2
13	0.9982	4.19e-04	146.5	57.0
14	1.0318	6.34e-03	146.2	84.5
15	1.0022	3.44e-04	147.5	27306.8
16	0.9993	-	-	-
17	0.9973	2.05e-03	149.2	77.8
19	0.9962	1.00e-10	146.1	1542.2
20	0.9918	9.31e-02	166.9	51.1
21	1.1447	1.21e-01	146.0	39662.1
22	0.9989	1.59e-01	146.3	86.9
23	0.9951	-	-	-
25	1.0049	-	-	-
24	1.0016	2.20e-03	150.6	47.0
26	0.9917	1.11e-01	163.6	38.0
27	0.4068	5.13e-02	178.4	47.4
28	0.9111	7.17e-02	146.3	75.2
29	0.8395	8.15e-02	146.2	13500.3
30	0.9974	-	-	-
31	0.9975	-	-	-
32	0.8750	-	-	-
33	1.0245	2.72e-02	152.5	23045.4
34	0.9989	3.93e-03	146.1	14.7
35	1.0969	9.28e-02	146.2	24483.8

Table A.10: Best fit parameters for pixel gain during the Flight 5 observation. Gain is constant throughout observation. Gain fit to the pulse-height of the K Kα line at 3313.8 eV when the gate-valve is open. Dividing the fit pulse heights by the gain will normalize the mean of the K Kα line to 3313.8 eV.

Pixel	Gain
0	0.9999
1	0.9994
2	0.9910
3	0.9989
4	0.9996
6	1.0000
8	0.9979
9	0.9987
10	0.9980
11	0.9988
12	0.9994
13	1.0002
14	0.9986
15	1.0000
16	0.9970
17	0.9984
20	0.9984
21	0.9985
22	0.9980
23	0.9989
24	0.9995
25	0.9828
26	0.9994
27	0.9988
28	0.9962
29	0.9954
30	0.9978
31	0.9953
32	0.9976
33	0.9986
34	0.9973
35	0.9992

Table A.11: Best fit parameters for pixel gain during the Flight 6 observation. Gain function has the form $\ln(g(T)) = a_{GV} + a_0 + a_1\Delta T + a_2\Delta T^2$, where ΔT is the difference between the cold plate temperature (Figure A.5) and 50mK. $a_{GV} = 0$ whenever the gate-valve is closed. Parameters fit to the pulse-height of the K Kα line at 3313.8 eV during time segments on the rail, throughout the observation, and on the parachute. Dividing the fit pulse heights by the gain will normalize the mean of the K Kα line to 3313.8 eV.

Pixel	a_{GV}	a_0	a_1	a_2
0	0.004655	0.002957	-0.035510	0.000580
1	0.005164	0.003793	-0.041823	0.001317
2	-0.001493	-0.000098	-0.043162	0.001896
3	0.004156	0.003731	-0.039918	0.000895
4	0.003613	0.002285	-0.037516	0.000609
5	0.023883	0.006236	-0.068593	0.008830
6	0.002366	0.003000	-0.035689	0.000848
8	0.007990	-0.089296	-0.043067	0.002372
9	0.006548	0.002565	-0.022610	-0.001090
10	0.001552	0.004030	-0.043452	0.001340
11	0.004252	0.002790	-0.038652	0.000690
12	0.005424	0.003306	-0.037024	0.000853
13	0.004554	0.002851	-0.035500	0.000594
14	0.005006	0.003209	-0.038329	0.000481
15	0.004158	0.002372	-0.036493	0.000320
16	0.002092	0.000469	-0.031136	-0.000421
17	0.002846	-0.067348	-0.037843	0.001020
20	0.004087	0.003382	-0.040597	0.001130
21	0.004685	0.003470	-0.036835	0.000675
22	0.003521	0.003987	-0.037574	0.000354
23	0.003842	0.003869	-0.038506	0.000852
25	0.005646	0.004025	-0.046810	0.002311
26	0.004890	0.000445	-0.032356	-0.000184
27	0.003709	0.005391	-0.039389	0.000936
28	0.003128	0.002062	-0.026428	-0.000465
29	0.012136	0.006440	-0.060122	0.003899
30	0.004603	0.002849	-0.040586	0.001029
32	0.004142	0.001703	-0.029351	-0.000391
33	0.002367	0.001678	-0.028953	0.000000
34	0.002736	0.004166	-0.039521	0.001203
35	-0.009083	-0.957500	-0.049251	0.005306

Table A.12: Non-linear Energy coefficients for Flight 4. Energy scale is a polynomial in gain-normalized pulse-height. I.e., $E(x) = a_1x + a_2x^2 + a_3x^3$, where x is the optimally filtered pulse-height divided by the pixel gain (Table A.9). Coefficients fit using characteristic lines of C (277.0 eV), F (676.8 eV), Si (1740.0 eV), and K (3313.8 eV) on the rail prior to flight.

Pixel	a_1	a_2	a_3
0	0.966711	3.25e-06	2.04e-09
1	0.966161	2.40e-06	2.33e-09
2	0.962569	4.64e-06	2.00e-09
3	0.965866	2.02e-06	2.50e-09
4	0.980411	-1.01e-05	4.82e-09
5	0.977847	-1.17e-05	5.45e-09
6	0.986843	-2.14e-05	7.63e-09
7	0.969787	-2.34e-06	3.44e-09
8	0.959804	7.48e-06	1.38e-09
9	0.968054	4.32e-06	1.58e-09
10	0.962013	2.79e-06	2.61e-09
11	0.964681	3.92e-06	2.03e-09
12	0.965353	3.46e-06	2.10e-09
13	0.971420	-4.30e-07	2.74e-09
14	0.968420	-6.36e-06	4.76e-09
15	0.971068	-1.84e-06	3.11e-09
16	0.969840	-2.20e-06	3.42e-09
17	0.962474	7.13e-06	1.24e-09
19	0.973352	-4.21e-06	3.70e-09
20	0.968095	1.18e-06	2.54e-09
21	0.965882	3.49e-06	2.03e-09
22	0.963261	5.15e-06	1.79e-09
23	0.970187	-8.72e-07	2.97e-09
24	0.968743	2.73e-06	2.02e-09
25	0.975603	1.29e-06	1.85e-09
26	0.955498	5.95e-06	2.26e-09
27	0.960214	3.73e-06	2.47e-09
28	0.957006	8.51e-06	1.31e-09
29	0.973883	-1.42e-06	2.78e-09
30	0.967029	7.61e-06	7.04e-10
31	0.965732	6.56e-07	2.90e-09
32	0.960811	4.47e-06	2.22e-09
33	0.960811	4.57e-06	2.17e-09
34	0.970164	5.50e-06	1.05e-09
35	0.972711	-9.07e-07	2.76e-09

Table A.13: Non-linear Energy coefficients for Flight 5. Energy scale is a polynomial in gain-normalized pulse-height. I.e., $E(x) = a_1x + a_2x^2 + a_3x^3$, where x is the optimally filtered pulse-height divided by the pixel gain (Table A.10). Coefficients fit using characteristic lines of C (277.0 eV), F (676.8 eV), Si (1740.0 eV), and K (3313.8 eV) on the rail prior to flight.

Pixel	a_1	a_2	a_3
0	0.982680	-9.68e-07	1.87e-09
1	0.971980	3.39e-06	1.53e-09
2	0.971960	5.96e-06	3.61e-10
3	0.972567	3.27e-06	1.51e-09
4	0.972269	4.80e-06	1.07e-09
6	0.986469	1.35e-06	8.30e-10
8	0.981212	-6.41e-07	1.91e-09
9	0.976438	4.46e-06	8.00e-10
10	0.971229	4.70e-06	1.20e-09
11	0.973449	3.89e-06	1.25e-09
12	0.974722	3.85e-06	1.13e-09
13	0.977093	1.88e-06	1.52e-09
14	0.975230	2.15e-06	1.61e-09
15	0.974775	3.38e-06	1.28e-09
16	0.976794	2.36e-06	1.40e-09
17	0.973273	4.84e-06	9.75e-10
20	0.971256	5.86e-06	8.53e-10
21	0.977224	1.09e-06	1.75e-09
22	0.974134	4.61e-06	9.61e-10
23	0.975246	2.16e-06	1.60e-09
24	0.973826	3.77e-06	1.24e-09
25	0.975289	7.98e-07	2.01e-09
26	0.977214	3.67e-07	1.96e-09
27	0.978006	-1.22e-06	2.38e-09
28	0.977421	1.69e-06	1.54e-09
29	0.960074	5.15e-06	2.08e-09
30	0.974709	2.84e-06	1.44e-09
31	0.981423	7.08e-07	1.48e-09
32	0.974503	2.10e-06	1.69e-09
33	0.980575	2.49e-06	1.01e-09
34	0.978812	3.17e-06	9.75e-10
35	0.974963	2.11e-06	1.64e-09

Table A.14: Non-linear Energy coefficients for Flight 6. Energy scale is a polynomial in gain-normalized pulse-height. I.e., $E(x) = a_1x + a_2x^2 + a_3x^3$, where x is the optimally filtered pulse-height divided by the pixel gain (Table A.11). Coefficients fit using characteristic lines of C (277.0 eV), F (676.8 eV), Si (1740.0 eV), and K (3313.8 eV) on the rail prior to flight.

Pixel	a_1	a_2	a_3
0	0.978117	9.30e-07	1.69e-09
1	0.970124	4.22e-06	1.41e-09
2	0.965395	9.61e-06	3.59e-10
3	0.967139	6.03e-06	1.13e-09
4	0.964917	7.98e-06	7.65e-10
5	0.975894	2.15e-06	1.53e-09
6	0.986980	-1.10e-06	1.53e-09
8	0.979801	-4.14e-06	3.05e-09
9	0.978199	-2.65e-07	2.08e-09
10	0.971092	1.78e-06	2.07e-09
11	0.972274	5.03e-06	1.01e-09
12	0.977843	-2.30e-06	2.71e-09
13	0.972800	5.38e-06	8.04e-10
14	0.971208	3.79e-06	1.46e-09
15	0.972075	2.88e-06	1.61e-09
16	0.971696	1.79e-06	2.00e-09
17	0.964653	1.06e-05	-4.24e-11
20	0.966633	8.33e-06	5.05e-10
21	0.970057	4.74e-06	1.33e-09
22	0.969282	5.89e-06	9.91e-10
23	0.973977	1.55e-06	1.88e-09
25	0.967786	5.41e-06	1.29e-09
26	0.980685	1.69e-06	1.28e-09
27	0.977545	8.05e-07	1.81e-09
28	0.973135	4.12e-06	1.20e-09
29	0.957656	-2.03e-07	4.07e-09
30	0.970310	4.63e-06	1.26e-09
32	0.970495	6.97e-06	5.24e-10
33	0.980585	-1.06e-06	2.11e-09
34	0.972588	4.13e-06	1.19e-09

Table A.15: Summary of flight exposures.

Flight	Good Pixels	Begin Observation (T-time)	End Observation (T-time)	Net Live-time (Pixel-Seconds)
Flight 4	26	160	420	4598.2
Flight 5	27	140	436	5546.7
Flight 6	24	185	436	3600.9

Table A.16: Pixels excluded from Flight 4 analysis.

Pixel	Reason
0	Non-uniform Pulse Arrival Times
3	Bad Fit to Non-linear Energy Scale
5	Bad Fit to Non-linear Energy Scale
10	Inconsistent Spectrum
11	Inconsistent Spectrum
14	Bad Fit to Non-linear Energy Scale
18	Frame Thermometer
23	Non-uniform Pulse Arrival Times
29	Inconsistent Spectrum
34	Non-uniform Pulse Arrival Times

Table A.17: Pixels excluded from Flight 5 analysis.

Pixel	Reason
5	Dead Pixel
7	Dead Pixel
10	Inconsistent Spectrum
18	Frame Thermometer
19	Dead Pixel
20	Non-uniform Pulse Arrival Times
23	Inconsistent Spectrum
31	Non-uniform Pulse Arrival Times
33	Non-uniform Pulse Arrival Times

Table A.18: Pixels excluded from Flight 6 analysis.

Pixel	Reason
7	Dead Pixel
8	Bad Non-linear Energy Scale
14	Inconsistent Spectrum
18	Frame Thermometer
19	Dead Pixel
20	Inconsistent Spectrum
24	Dead Pixel
26	Poor Resolution
28	Inconsistent Spectrum
29	Bad Non-linear Energy Scale
31	Dead Pixel
35	Bad Non-linear Energy Scale

Appendix B

User's Guide to *h5filter*

Documentation for the *h5filter* program used for the analysis in Chapter 3 is presented in this appendix. The program gets its name from the HDF5 file format that it uses, much like *fitsfilter* uses the FITS format.

B.1 Directory Structure

All scripts and data are on wisp5. Scripts are located in `/home/calorim/src/h5filter/`. Data files (henceforth BSN files, see Section B.3) are located in `/home/calorim/Runs/*/BSN/`, where the * is replaced by the name of the cool-down (e.g. "k8r"). The `BSN/` directory is a symbolic link to a directory in `/data_nobackup/bsn_directories/` with the same name as the cool-down. In addition to the BSN files, the `BSN/` directory contains the directory `plots/`, where any plots made during the analysis are saved.

B.2 Organization of Python Code

h5Filter.py is a wrapper which calls code from four other python scripts. Figure B.1 depicts these scripts and the python object classes they contain, in the approximate order that they are used. xqcHDF5.py contains code for creating, writing to, and reading from the BSN files. optimumFilter.py contains code for making optimum filters, fitting templates, and convolving the data streams with filters. pulseFitting.py contains the code for fitting pulses in the filtered data streams. eventCat.py contains code for determining the non-linear energy scale and correcting the energies of overlapped pulses. It also contains code for selecting good events and determining live-time.

B.3 The BSN File Format

Within xqcHDF5.py, there are two object classes: PackedDataBBB and singlePixel. PackedDataBBB is used for converting the 12-bit binary files written by the on-board data recorders to BSN files in the HDF5 format. In the future, the BSN files may be written directly by the streaming software, making this class obsolete. The singlePixel object class, on the other hand, is integral to almost all the other processing tasks, since it provides access to the data for a single pixel group in the BSN file. The singlePixel object class defines what attributes, datasets, and formats are used in the BSN file. The current format of the BSN file is outlined in Figure B.2.

As with *fitsfilter*, the data for all 36 pixels are stored in a single file, known as the BSN file. The name of the BSN file follows the format "*.bsn.hdf5," where the * is replaced with the name of the run. This will usually be a 3-character string referring to a cool-down, followed by a 2- or 3-digit number indicating the file number of that cool-down, e.g. "k8r61." The first three characters should refer to an existing directory

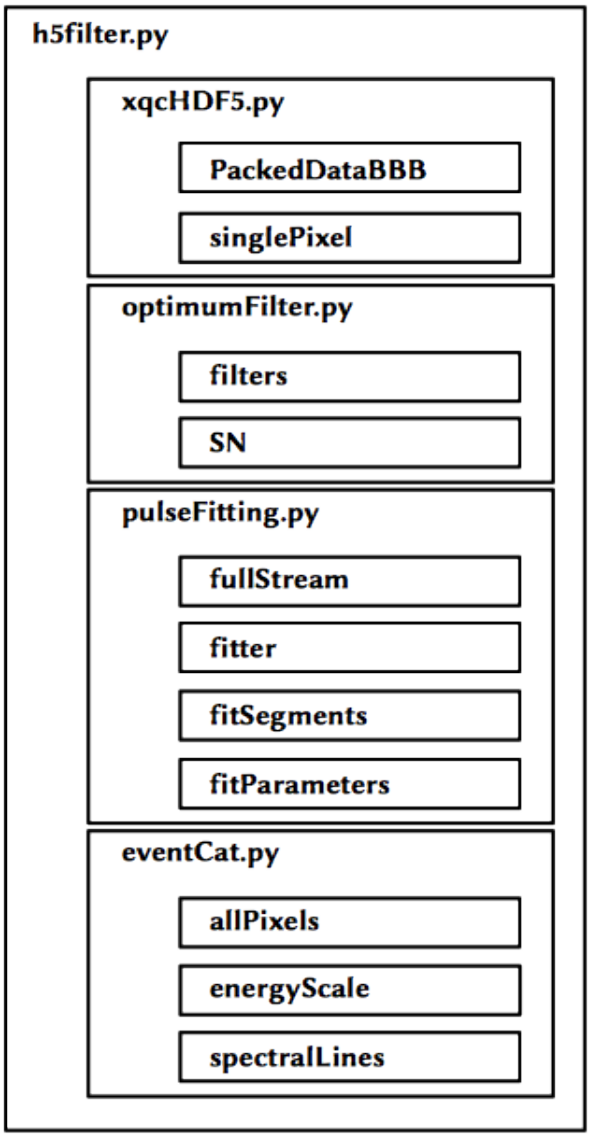

Figure B.1: *h5filter* organizational structure. Boxes with the extension ".py" refer to scripts, while those without an extension are object classes within those scripts.

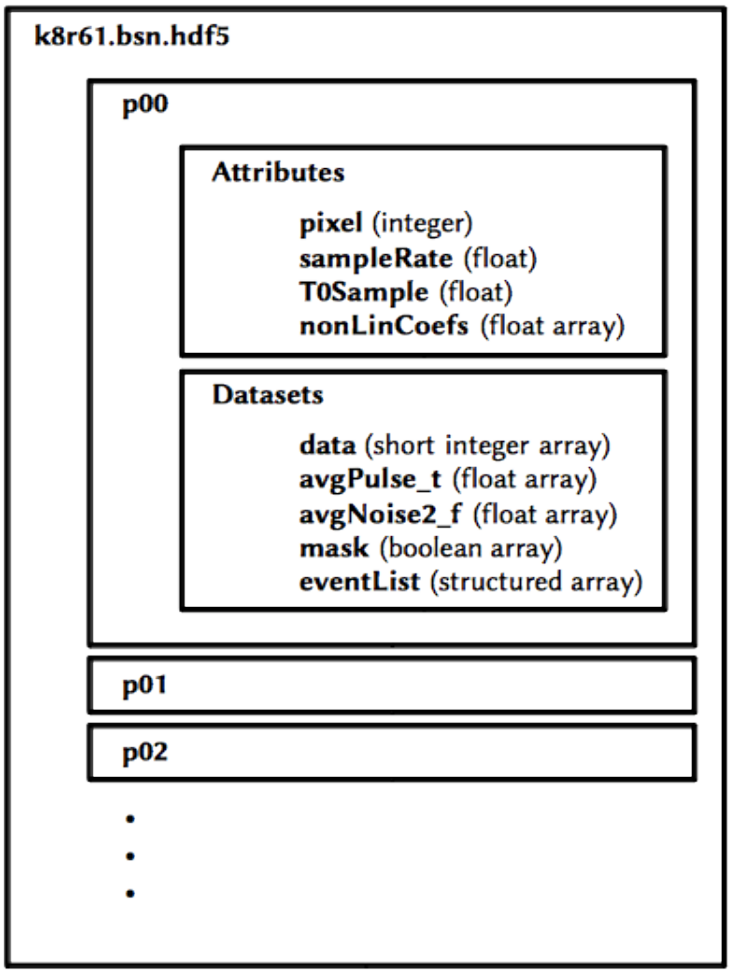

Figure B.2: BSN file format. Each file ("k8r61.bsn.hdf5" in this example) has groups "p00," "p01," "p02," etc. Within each group, there are attributes and datasets. Attributes are small pieces of metadata, and datasets contain larger pieces of data.

within `/home/calorim/Runs/` on wisp5, where the program will look for the file (See Section B.1).

Adopting the language of h5py (`www.h5py.org`), the BSN file contains a group for every data stream. By default, each group contains the data stream from a single pixel, and groups have names like "p00," "p01," "p02," etc. Group names must be unique (think of them like directories). Within each group, there are attributes and datasets. Attributes are small pieces of metadata, and datasets contain larger pieces of data.

For XQC data, the pixel attribute will almost always be the same number that is in the group name (e.g. 5 for group "p05"). The sampleRate attribute is the sample rate of the data stream in Hertz (∼10416.7 Hz for XQC). The T0Sample attribute is the relative sample number of the data stream that coincides with zero time. For rocket flights, zero time is T = 0. For laboratory data, zero time is midnight UT. Note that this number will be negative if zero time occurred before the beginning of the file. It will also generally be a floating point number. The nonLinCoefs attribute is an array that stores the polynomial coefficients a_i for the non-linear energy scale $E(x) = \sum_{i=1}^{n} a_i x^i$, where x is the gain-normalized pulse-height. The order of polynomial n is inferred by the number of coefficients and there is no zeroth order term (i.e. zero pulse-height⇒zero energy). The first coefficient in the array is a_1, followed by a_2, etc.

The data stream for each group is stored in the dataset called data. By default, this is a 1D array of short integers. The average signal shape for the optimal filter is stored in the dataset avgPulse_t as a 1D array of floats. The signal shape should be in the time-domain and in units of Volts. The mean-square noise for the optimal filter is stored in the dataset avgNoise2_f as a 1D array of floats. The mean-square noise should be in the frequency-domain and in units of V^2/Hz^{-1}. Note that there should be agreement between the lengths of the average signal and mean-square noise—i.e. for a average signal

length $2N$ in the time-domain, the mean-square noise should be length $N + 1$ in the frequency-domain. The mask dataset is a boolean array with the same shape as the data dataset, and is used to keep track of where the fitting completed successfully (1 = fit, 0 = not fit). Under this convention, the live-time for the data stream is equal to the number of ones in the mask dataset divided by the sampleRate attribute. The eventList dataset is a structured array, with a sequence of numbers for every event in the data stream. The format is flexible, but generally contains fields such as "pixel," "sample," "ph" (i.e. pulse-height), "time," "energy," "offset," and various other flags or metrics for each event.

B.4 High Level Functions

B.4.1 fullStream Object Class

The fullStream object class handles the high level aspects of overlapped pulse fitting. It creates dictionaries of fitting templates from the saved signal and noise templates in the BSN file. It also segments, filters, and fits the data streams, and saves the fit information to the BSN file.

`pulseFitting.fullStream(run,pixel,grpName=None)`

Initiate instance of fullStream object class.

- Parameters
  - `run` (str): Run base name preceding .bsn.hdf5 extension of BSN file.
  - `pixel` (int): Pixel number
  - `grpName` (str, optional): Group name in BSN file. By default, grpName is determined from pixel (e.g. "p05" for pixel 5).

`fullStreamObject.makeTemplates(preTriggerFrac=0.25,calEnergy=3313.8, scale=np.linspace(0,1,5),gainCoefs=[0.97,1e-5],LCF=None,UCF=None,LCP=2, UCP=2)`

Create dictionaries of fitting templates.

- Parameters
  - `preTriggerFrac` (float, optional): Fraction of average signal trace length before pulse arrives.
  - `calEnergy` (float, optional): Energy in eV of pulses used to make average signal template.
  - `scale` (1d array, optional): Values to scale calEnergy when making pulse-height dependent shapes. If no pulse shape changes are desired, set to [1.].
  - `gainCoefs` (1d array, optional): Polynomial coefficients used to make pulse-height dependent shapes. First order coefficient given first, followed by higher order coefficients. No zeroth order coefficient.

- `LCF` (float, optional): Lower cutoff frequency used to deconvolve amplifier response. By default, values are taken from `/home/calorim/Runs/parameterfiles/xqcgains.dat`.
- `UCF` (float, optional): Upper cutoff frequency used to deconvolve amplifier response. By default, values are taken from `/home/calorim/Runs/parameterfiles/xqcgains.dat`.
- `LCP` (int, optional): Lower cutoff pole used to deconvolve amplifier response.
- `UCP` (int, optional): Upper cutoff pole used to deconvolve amplifier response.

`fullStreamObject.fitFilteredData(loRes=True,interactive=False, useExisting=True,adjustOffset=True,initThrshld=3000.0,nIter=3,redFctr=0.1, nTmplt=3,initOffset=0,maxPls=10,deltaS=0.1,deltaX=20)`

Segment, filter, and fit entire data stream.

- Parameters
  - `loRes` (bool, optional): If True, pulses that cannot be fit in filtered data streams will have low-pass filtered pulse heights saved with resolution flag set to 1.
  - `interactive` (bool, optional): If True, user will have the chance to interact with each segment graphically (See Figure B.3).
  - `useExisting` (bool, optional): If True, and eventList dataset already exists in BSN group, then data stream will not be refit.
  - `adjustOffset` (bool, optional): If True, program will attempt to adjust each segment for varying DC offset.

- `initThrshld` (float, optional): Initial threshold in units of eV to look for pulses in amplitude filtered data stream.
- `nIter` (int, optional): Minimum number of times to find pulses, fit, subtract, and reduce threshold.
- `redFctr` (float, optional): Amount to lower threshold in each iteration.
- `initOffset` (float, optional): Initial amount to adjust offset of unfiltered data in units of eV.
- `maxPls` (int, optional): Maximum number of pulses per template length that program will attempt to fit. If number of pulses exceeds this number, the fit will be cleared and the program will move on to the next segment.
- `deltaS` (float, optional): Maximum allowable difference between a pulse's shape parameter and one. Pulses exceeding this will be removed form fit.
- `deltaX` (int, optional): Maximum allowable difference between a pulse's fit location and the initial guess in units of samples. Pulses exceeding this will be removed form fit.

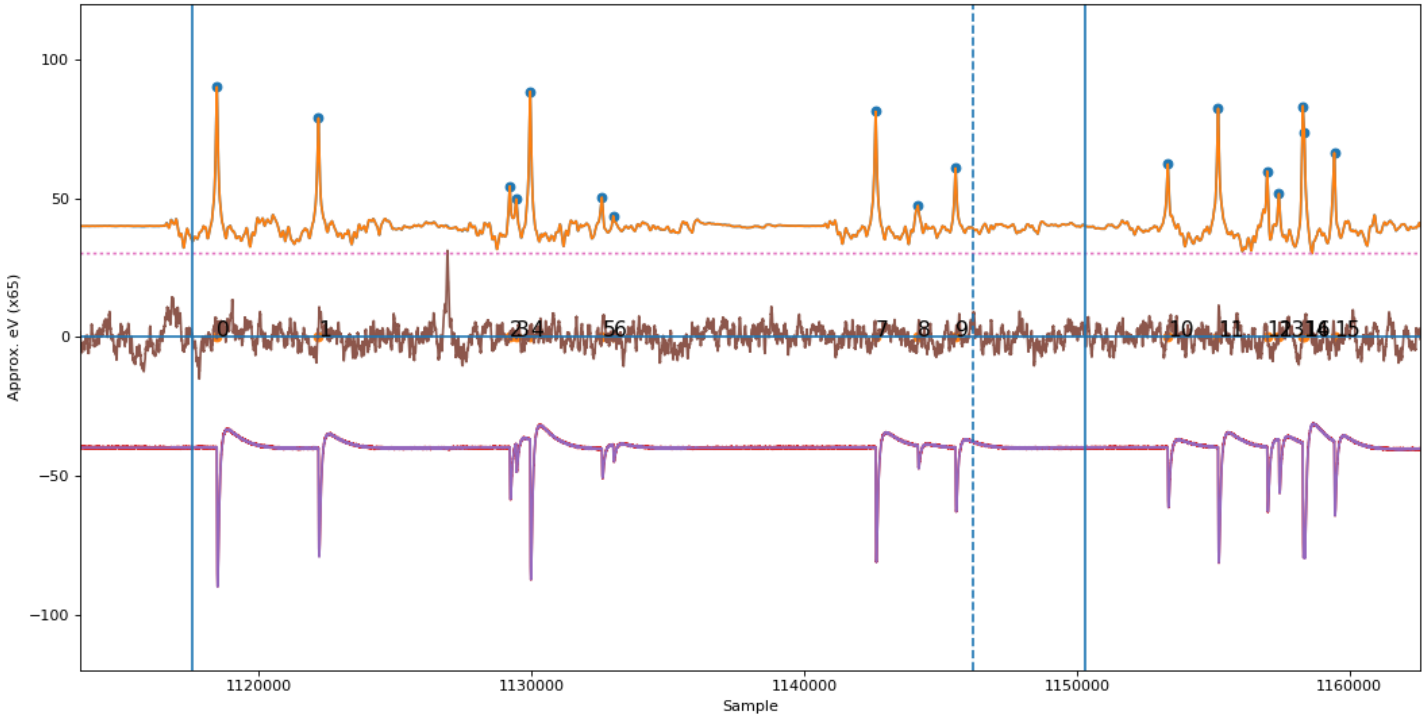

Figure B.3: *h5filter* interactive fitting interface. The top curves are the pulse-amplitude filtered data and the same model constructed from fitting templates. Circles mark the data points used in the fit. The middle curve is the residual of the of the top two curves, magnified to show detail. The locations of pulses in the residuals are annotated with identification numbers. The horizontal dashed line shows the lowest threshold for pulse detection. The bottom curves are the unfiltered data and model. If fitting of the previous segment was successful, there should be no pulses before the solid vertical line on the left, as they will already have been fit and subtracted. The next segment will begin at the dashed vertical line. Any pulses that arrive before the solid vertical line on the right will be subtracted before moving on the next segment.

B.4.2 energyScale Object Class

The energyScale object class is used for setting the non-linear energy scale.

`eventCat.energyScale(run,pixels=numpy.arange(36,dtype=np.int32))`

Initiate instance of energyScale object class.

- Parameters
  - `run` (str): Run base name preceding .bsn.hdf5 extension of BSN file.
  - `pixels` (array, optional): Pixel numbers.

`energyScaleObject.findNonLin(order=3,goodLines=["C","F","Si","Ka"], doPlot=False,write=True,calTimeLimits=None,grpName=None,interactive=False, maxE=6000.0,minE=0.0,binWidth=2.0)`

Determine polynomial coefficients for non-linear energy scale.

- Parameters
  - `order` (int ,optional): Order of polynomial
  - `goodLines` (list, optional): Calibration lines to use.
  - `doPlot` (bool, optional): If True, a plot of polynomial fit residuals will be shown.
  - `write` (bool, optional): If True, polynomial coefficients will be written to the nonLinCoefs attribute for each pixel group in the BSN file.
  - `calTimeLimits` (tuple, optional): Start and stop times to use for calibration

- `grpName` (str, optional): Group name in BSN file. By default, grpName is determined from pixel (e.g. "p05" for pixel 5). Currently, a non-default grpName only works for one pixel at a time.
- `interactive` (bool, optional): If True, user will be asked to manually select line locations.
- `maxE` (float, optional): Maximum energy in eV to histogram data.
- `minE` (float, optional): Minimum energy in eV to histogram data.
- `binWidth` (float, optional): Bin width in eV to histogram data

`energyScaleObject.eScale(closelySpacedCorrection=True,doPlot=False, ignoreGain=False,grpName=None,calEnergy=3313.8,LCF=None,UCF=None,LCP=2, UCP=2)`

Apply non-linear energy scale to events in BSN file.

- Parameters
  - `closelySpacedCorrection` (bool, optional): If True, event energies are corrected for reduced gain due to previous pulse.
  - `doPlot` (bool, optional): If True, a scatter plot of energy versus time will be shown.
  - `ignoreGain` (bool, optional): If True, the gain of every event is assumed to be one.
  - `grpName` (str, optional): Group name in BSN file. By default, grpName is determined from pixel (e.g. "p05" for pixel 5). Currently, a non-default grpName only works for one pixel at a time.

- `calEnergy` (float, optional): Energy in eV of pulses used to make average signal template.
- `LCF` (float, optional): Lower cutoff frequency used to deconvolve amplifier response. By default, values are taken from `/home/calorim/Runs/parameterfiles/xqcgains.dat`.
- `UCF` (float, optional): Upper cutoff frequency used to deconvolve amplifier response. By default, values are taken from `/home/calorim/Runs/parameterfiles/xqcgains.dat`.
- `LCP` (int, optional): Lower cutoff pole used to deconvolve amplifier response.
- `UCP` (int, optional): Upper cutoff pole used to deconvolve amplifier response.

B.4.3 allPixels Object Class

The allPixels object class is used for handling events from all pixels at once, for plotting all the events, and for selecting events.

`eventCat.allPixels(run,pixels=numpy.arange(36,dtype=np.int32),timeBounds=None, grpName=None)`

Initiate instance of allPixels object class.

- Parameters
  - `run` (str): Run base name preceding .bsn.hdf5 extension of BSN file.
  - `pixels` (array, optional): Pixel numbers.
  - `timeBounds` (tuple, optional): Start and stop event times to use. By default, events at all times are loaded from BSN file.

  - `grpName` (str, optional): Group name in BSN file. By default, grpName is determined from pixel (e.g. "p05" for pixel 5). Currently, a non-default grpName only works for one pixel at a time.

`allPixelsObject.selectEvents(dt1=0.01,dt2=0.0025,goodPixels=None,doPlot=False,fitRate=True,grpName=None,adjacent=True)`

- Parameters
  - `dt1` (float, optional): Minimum allowable time separation in seconds between two events on the same pixel.
  - `dt2` (float, optional): Minimum allowable time separation in seconds between two events on different pixels.
  - `goodPixels` (array, optional): Pixels to select events from. By default, all pixels are good.
  - `doPlot` (bool, optional): If True, a plot of the event arrival-time cumulative data distribution is shown for each pixel.
  - `fitRate` (bool, optional): If True, the rate of events on a pixel is fit to the slope of the event arrival-time cumulative data distribution. If False, the rate is calculated as the number of events divided by the live time.
  - `grpName` (str, optional): Group name in BSN file. By default, grpName is determined from pixel (e.g. "p05" for pixel 5). Currently, a non-default grpName only works for one pixel at a time.
  - `adjacent` (bool, optional): If True, time separation of pulses on different pixels only considers adjacent pixels.

References

Anders, E., & Ebihara, M. 1982, Geochimica et Cosmochimica Acta, 46, 2363

Anders, E., & Grevesse, N. 1989, Geochimica et Cosmochimica Acta, 53, 197

Arnaud, K. A. 1996, in Astronomical Society of the Pacific Conference Series, Vol. 101, 17

Balucinska-Church, M., & McCammon, D. 1992, The Astrophysical Journal, 400, 699

Bekhti, N. B., Fler, L., Keller, R., et al. 2016, Astronomy & Astrophysics, 594, A116

Bienayme, O., Robin, A. C., & Creze, M. 1987, Astronomy and Astrophysics, 180, 94

Bowyer, S., Byram, E. T., Chubb, T. A., & Friedman, H. 1964, Nature, 201, 1307

Burrows, D. N., & Mendenhall, J. A. 1991, Nature, 351, 629

Cowsik, R., & Kobetich, E. J. 1972, The Astrophysical Journal, 177, 585

Cox, D. P. 1997, International Astronomical Union Colloquium, 166, 121

Cox, D. P., & Smith, B. W. 1974, The Astrophysical Journal, 189, L105

Cravens, T. E. 1997, Geophysical Research Letters, 24, 105

Crowder, S. G., Barger, K. A., Brandl, D. E., et al. 2012, The Astrophysical Journal, 758, 143

Dame, T. M., Hartmann, D., & Thaddeus, P. 2001, The Astrophysical Journal, 547, 792

Defay, X., Morgan, K., McCammon, D., et al. 2013, Physical Review A, 88, 052702

Dickey, J. M., & Lockman, F. J. 1990, Annual Review of Astronomy and Astrophysics, 28, 215

Ferrire, K. M. 2001, Reviews of Modern Physics, 73, 1031

Finkbeiner, D. P. 2003, The Astrophysical Journal Supplement Series, 146, 407

Fixsen, D. J., Moseley, S. H., Cabrera, B., & Figueroa-Feliciano, E. 2002, AIP Conference Proceedings, 605, 339

Fowler, J. W., Alpert, B. K., Doriese, W. B., et al. 2015, The Astrophysical Journal Supplement Series, 219, 35

Giacconi, R., Gursky, H., Paolini, F. R., & Rossi, B. B. 1962, Physical Review Letters, 9, 439

Giampapa, M. S., Rosner, R., Kashyap, V., et al. 1996, The Astrophysical Journal, 463, 707

Gould, R. J., & Burbidge, G. R. 1963, The Astrophysical Journal, 138, 969

Henry, P., Bowyer, S., Lampton, M., Paresce, F., & Cruddace, R. 1976, The Astrophysical Journal, 205, 426

Hickox, R. C., & Markevitch, M. 2006, The Astrophysical Journal, 645, 95

Ishisaki, Y., Yamada, S., Seta, H., et al. 2018, Journal of Astronomical Telescopes, Instruments, and Systems, 4, 011217

Jenkins, E. B., & Meloy, D. A. 1974, The Astrophysical Journal Letters, 193, L121

Koutroumpa, D., Lallement, R., Raymond, J. C., & Kharchenko, V. 2009, The Astrophysical Journal, 696, 1517

Kuntz, K. D., & Snowden, S. L. 2000, The Astrophysical Journal, 543, 195

—. 2001, The Astrophysical Journal, 554, 684

Lallement, R. 2004, Astronomy & Astrophysics, 418, 143

Lei, S., Shelton, R. L., & Henley, D. B. 2009, The Astrophysical Journal, 699, 1891

Lisse, C. M., Dennerl, K., Englhauser, J., et al. 1996, Science, 274, 205

Liu, W., Galeazzi, M., & Ursino, E. 2016, The Astrophysical Journal, 816, 82

Liu, W., Chiao, M., Collier, M. R., et al. 2017, The Astrophysical Journal, 834, 33

Marshall, F. E., Boldt, E. A., Holt, S. S., et al. 1980, The Astrophysical Journal, 235, 4

Masui, K., Mitsuda, K., Yamasaki, N. Y., et al. 2009, Publications of the Astronomical Society of Japan, 61, S115

McCammon, D. 2005, in Cryogenic Particle Detection (Springer Science & Business Media)

McCammon, D., Burrows, D. N., Sanders, W. T., & Kraushaar, W. L. 1983, The Astrophysical Journal, 269, 107

McCammon, D., Almy, R., Apodaca, E., et al. 2002, The Astrophysical Journal, 576, 188

McCammon, D., Barger, K., Brandl, D. E., et al. 2008, Journal of Low Temperature Physics, 151, 715

Metzger, A. E., Anderson, E. C., Dilla, M. a. V., & Arnold, J. R. 1964, Nature, 204, 766

More, J. J., Garbow, B. S., & Hillstrom, K. E. 1980, User guide for MINPACK-1, Tech. Rep. ANL-80-74, Argonne National Lab., IL (USA)

Rosner, R., Avni, Y., Bookbinder, J., et al. 1981, The Astrophysical Journal, 249, L5

Sanders, W. T., Edgar, R. J., Kraushaar, W. L., McCammon, D., & Morgenthaler, J. P. 2001, The Astrophysical Journal, 554, 694

Sanders, W. T., Kraushaar, W. L., Nousek, J. A., & Fried, P. M. 1977, The Astrophysical Journal Letters, 217, L87

Savage, B. D., & Sembach, K. R. 1996, Annual Review of Astronomy and Astrophysics, 34, 279

Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, The Astrophysical Journal, 500, 525

Scholz, F. W., & Stephens, M. A. 1987, Journal of the American Statistical Association, 82, 918

Smith, R. K., Brickhouse, N. S., Liedahl, D. A., & Raymond, J. C. 2001, The Astrophysical Journal Letters, 556, L91

Smith, R. K., Foster, A. R., & Brickhouse, N. S. 2012, Astronomische Nachrichten, 333, 301

Smith, R. K., Foster, A. R., Edgar, R. J., & Brickhouse, N. S. 2014, The Astrophysical Journal, 787, 77

Smith, R. K., Bautz, M. W., Edgar, R. J., et al. 2007, Publications of the Astronomical Society of Japan, 59, S141

Snowden, S. L., Egger, R., Finkbeiner, D. P., Freyberg, M. J., & Plucinsky, P. P. 1998, The Astrophysical Journal, 493, 715

Snowden, S. L., Mebold, U., Hirth, W., Herbstmeier, U., & Schmitt, J. H. M. M. 1991, Science, 252, 1529

Snowden, S. L., Egger, R., Freyberg, M. J., et al. 1997, The Astrophysical Journal, 485, 125

Spitzer, Jr., L. 1956, The Astrophysical Journal, 124, 20

Szymkowiak, A. E., Kelley, R. L., Moseley, S. H., & Stahle, C. K. 1993, Journal of Low Temperature Physics, 93, 281

Uprety, Y., Chiao, M., Collier, M. R., et al. 2016, The Astrophysical Journal, 829, 83

von Steiger, R., Schwadron, N. A., Fisk, L. A., et al. 2000, Journal of Geophysical Research: Space Physics, 105, 27217

Weisskopf, M. C., Guainazzi, M., Jahoda, K., et al. 2010, The Astrophysical Journal, 713, 912

Yao, Y., Wang, Q. D., Hagihara, T., et al. 2009, The Astrophysical Journal, 690, 143

Yoshino, T., Mitsuda, K., Yamasaki, N. Y., et al. 2009, Publications of the Astronomical Society of Japan, 61, 805