Imaging X-ray Spectroscopy with *Micro-X* and *Chandra*

by

John Rutherford

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Physics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2013

Author .....
Department of Physics
August 30, 2013

Certified by .....
Enectali Figueroa-Feliciano
Associate Professor of Physics
Thesis Supervisor

Accepted by .....
Krishna Rajagopal
Professor of Physics
Associate Department Head for Education

Imaging X-ray Spectroscopy with *Micro-X* and *Chandra*

by

John Rutherford

Submitted to the Department of Physics
on August 30, 2013, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Physics

Abstract

High spectral resolution observations of X-ray phenomena have the potential to uncover new physics. Currently, only point sources can be probed with high resolution spectra, using gratings. Extended objects like supernova remnants cannot be dispersed, leaving the dense forest of emission lines blended by the moderate resolution of modern instruments.

In the first half of this thesis, I undertake two investigations of the supernova remnant Cassiopeia A using the flagship X-ray observatory, *Chandra*. The first study combines the spatial resolution of the ACIS instrument with the spectral resolution of the dispersive HETG to investigate the evolution of ejecta knots. The second improves on statistical limits of radioactive ejecta, and simulates what higher resolution instruments may observe.

Micro-X, a new high resolution X-ray telescope, is the focus of the second half. I detail the commissioning of this novel sounding rocket payload, which uses a focal plane of microcalorimeters to achieve high spectral resolution in each pixel. The flight hardware is in a final state of testing and integration before the launch, which is anticipated in 2014.

Thesis Supervisor: Enectali Figueroa-Feliciano
Title: Associate Professor of Physics

Acknowledgments

This work represents a true collaborative effort; any use of 'we' in the text should not be construed as the royal/scientific 'we'. My contributions are dwarfed by my collaborators', making these thank-yous a gross underpayment.

Tali, thanks for convincing me to work on *Micro-X*. I probably would have left MIT if it weren't for your leadership and boundless optimism. The freedom you provided gave me a sense of proud ownership, rekindling my love of experimental physics. Patrick, the once and future *Micro-X* engineer: I have learned so much under your guidance. I continually try to emulate your design sense and work ethic. Sarah, I loved working side-by-side with you. I would be hard pressed to find another colleague with such indefatigability and shared appreciation for trashy pop music.

The combined expertise of Steve Leman, Kosuke Sato, Phil Oakley, and Rick Foster helped me through many problems. I thank Maurice Leutenegger, Kelsey Morgan, Megan Eckart, and Jan-Patrick Porst, for their time and knowledge, which enabled the filter calibrations. Thanks to Dan Dewey for his HETG mastery and Cas A intuition. Jean Papagianopoulos saved me many a time, and always had a smile to spare. It's been an honor to work with the all the exceedingly smart MIT astrograds.

My parents have provided continual support and unconditional encouragement in my career aspirations, whether garbage man (age 4), storm chaser (age 12), or physicist (age 22). Finally, Te, you're pretty awesome.

Contents

Chapter 1

X-Ray Astrophysics

Nature so often leaves the most daring imagination of man far behind.

– Bruno Rossi

The tranquil night sky belies the ongoing violent turmoil surrounding us in space, wherever we point our telescopes. Stars fuse elements at extreme temperatures and churn plasma into the environment around them. Galaxy clusters collide, heating great volumes of gas. Large stars explode, leaving an unfathomably dense core spinning many times per second. Enormous black holes at the centers of galaxies send jets flying, beacons that can be seen at cosmological distances.

Energy scales of these phenomena dwarf the processes we know on Earth. High energy processes typically emit X-ray light, and are observable with X-ray telescopes only above Earth's absorbent atmosphere. Notions of hot intra-cluster gas, accretion-powered spinning neutron stars, and many other X-ray objects were not conceptualized until technology made it possible to observe above the atmosphere. Our understanding of the high energy sky today is the result, largely, of serendipity made possible by experimenters, looking in seemingly dark parts of the sky. I will recapitulate the history of the field first, summarizing our current understanding of the X-ray sky along the way, before focusing on supernova remnants, a primary subject of this thesis.

1.1 A History of X-ray Astronomy

The field of X-ray astronomy is merely five decades old, but has advanced to maturity at a rapid pace. From Galileo's first telescope to the *Hubble Space Telescope*, 380 years elapsed. From the first blip of an extrasolar X-ray source to the launch of the *Hubble*-equivalent X-ray flagship, *Chandra*, only 38 years passed. The quick advancement is a testament to the compelling physics revealed in the X-ray sky.

The heritage of X-ray astrophysics at MIT is strong, especially for rocket-based experiments. The field was essentially birthed in the university's neighborhood, with critical help from an MIT professor. *Micro-X* represents the latest small contribution to this heritage. Throughout this brief history, I will try to highlight MIT's involvement in advancing the field, so the reader should realize that this account is biased.

Scientists had studied solar X-ray emission with German V-2 sounding rockets since the end of World War II, but the field truly started in the summer of 1962, when Riccardo Giacconi and his American Science and Engineering (AS&E)[1] team serendipitously discovered a bright source, later named Sco X-1 [47]. Launched from White Sands Missile Range, their Aerobee rocket payload carried Geiger counters, pointing up at a fixed angle, without any optics to focus the light. The 2 Hz spin of the rocket (pointed toward zenith) provided coverage of the sky, and the data counts were binned into azimuth angle. A peak, both higher and shifted in angle from the moon's expected fluorescence contribution, provided the first evidence of an X-ray source outside of the solar system (Figure 1-1). This source's luminosity was 1000× brighter in X-rays than in optical, suggesting they had found a truly strange object, one not yet predicted by astronomers. In addition, the data revealed a background level of X-ray emission. Many future experiments were designed to uncover the nature of this background – whether it was diffuse or made of point sources, and whether it originated in our galaxy.

Hale Bradt and other researchers at MIT's newly founded Center for Space Research

[1] AS&E used to be located in Cambridge, MA, near MIT, but now is headquartered in Billerica, MA. Their primary business is X-ray scanning technology to combat narcotics, smuggling, and terrorism, but was originally founded to develop scientific instruments for NASA.

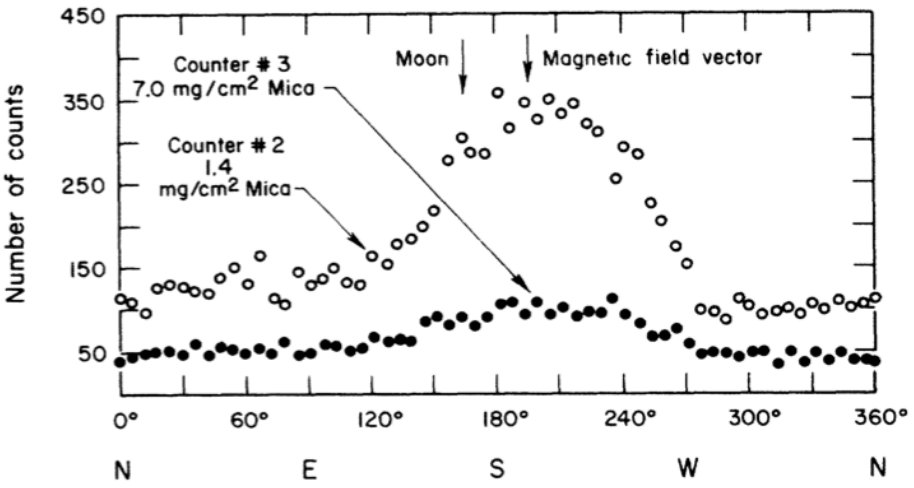

Figure 1-1: The first extrasolar X-ray source was detected as strong excess emission in the direction of the constellation Scorpius by Giacconi et al. [47]. The x-axis represents degrees around the rocket payload's azimuth. Fluorescence from the moon (labeled) could not explain the signal, nor could the coincidental alignment with the earth's magnetic field (also labeled). Sco X-1 was later understood to be an entirely new kind of phenomenon, a star accreting onto its collapsed binary companion. (Reprinted from Giacconi et al. [47].)

(CSR) collaborated with AS&E on another X-ray rocket to pin down the location of Sco X-1, in 1966 [50]. The payload utilized two modulation collimators at different angles to constrain the source location to $1'$. A companion analysis identified the optical counterpart to be a faint, variable, and massive star [117]. The ability to link objects across wavelengths was essential to advance the science and the growing field of X-ray astronomy. Without counterparts or images, X-ray astronomy in this early era was driven by serendipity, and struggled to gain legitimacy among communities of other wavelengths. The sources, too, were few, with little spectral information, leaving any physical theories of the emission which *could* be formed unconstrained by further observations. However, the uniqueness of the sources – some being highly variable, which suggested compact objects – proved compelling enough to fly more rockets.

Sounding rockets became an accessible and productive platform for X-ray astronomy. MIT alone flew 16 rocket payloads in the 1960s and 1970s[2]. In 1967, Bradt and his student

[2] `http://space.mit.edu/sites/default/files/Overview_CSR_1977.pdf`

Saul Rappaport led another CSR payload[3] which discovered extragalactic X-ray emission from M87 and constrained 6 other sources to $20'$ [10, 11, 111]. The British Skylark design, launched from Australia, began the search for extrasolar X-ray objects in the Southern sky. Dual launches pinned down the location of (low mass X-ray binary) GX3+1 to within an arcsecond by using the remarkable technique of moon occultation [31, 109]. Unfortunately, no optical companion was found – the case, also, for less precisely located sources.

With existence proofs of cosmic sources in hand, Giacconi set out to bring X-ray objects into focus. He and MIT's own Bruno Rossi were the first to conceive of an imaging X-ray telescope, in 1960 [46]. They adapted the decade-old designs of Hans Wolter's X-ray microscopes that utilized multiple low-angle reflections to focus the high energy light without coma or spherical aberration. Giacconi and Rossi thought to leverage the larger, inverted scale of telescopes[4], which didn't require machining tolerances as precise as microscopes. Their original conception (Figure 1-2) nested paraboloidal shells to increase the collecting area of the mirror without vignetting.

Focusing light yields two immediate advantages: first, imaging becomes possible; second, the signal-to-noise increases. Simple geometrical arguments show that the signal-to-noise scales as the ratio of the area of the mirror (which collects the signal) to the area of the collector (which is subject to readout noise and stray light or particles). For X-ray astronomy, the latter was required for the characteristically dim sources, but the former was also necessary to validate the field, as I will discuss below. The first true X-ray image came in 1965, of the Sun [43].

The first orbiting X-ray observatory, *Uhuru*, launched in 1970, another first for Giacconi's group at AS&E. The survey mission operated for three years, mapping out the sky in the 2–20 keV band [36]. As with the Aerobee payload, the instrument utilized gas-filled proportional

[3]A gripping account of the launch can be found online: `http://space.mit.edu/sites/default/files/Rocket_Diary_1967_0.pdf`.

[4]Curiously, the authors used the word 'telescope' lightly, always enclosing it in quotation marks, seemingly uncomfortable with comparing this new method of focusing light to the established methods from optical astronomy. By now, however, astronomers have embraced diverse forms of telescopes, such as huge radio dishes, photomultiplier tubes arrayed deep in ice, and coded aperture masks, and would not hesitate one bit to call a grazing-incidence mirror part of a telescope.

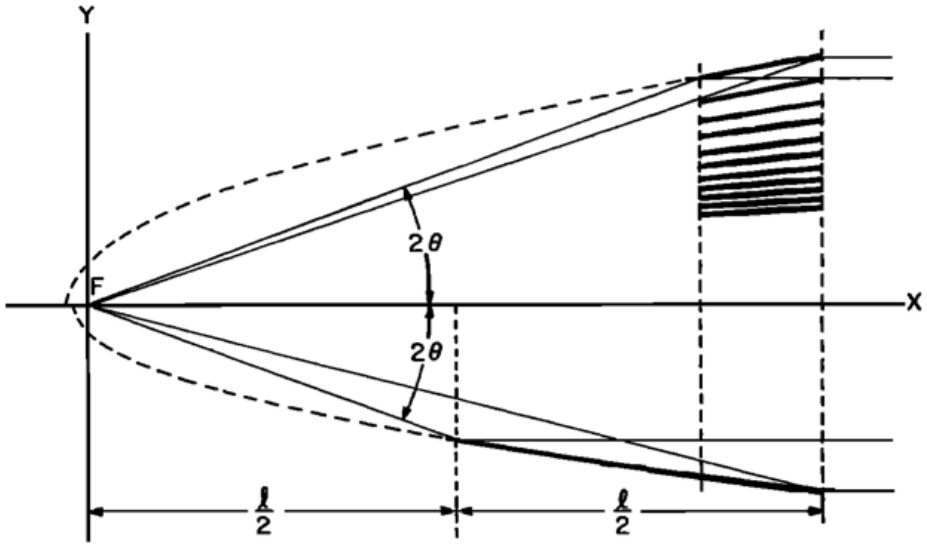

Figure 1-2: Giacconi and Rossi's original grazing incidence telescope design proposed both nested shells (top) and longer paraboloidal segments (bottom). Nested shells provide greater collecting area. Modern designs nest shells for greater collecting area, and approximate the parabolic sections with two easier-to-machine shapes. (Reprinted from Giacconi and Rossi [46].)

counters without any focusing optics, rotating on its axis to provide sky coverage. *Uhuru* discovered many other X-ray sources similar to Sco X-1, helping to develop the theory for a class of objects called X-ray binaries, which consist of stars accreting onto their compact companion, a neutron star or black hole. *Uhuru* also discovered emission from galaxy clusters, inferring that enormous volumes and masses of gas had been heated to X-ray temperatures – aproximately ten times the mass of the cluster member galaxies themselves.

Ariel 5, *SAS-3*, and *Hakucho* succeeded *Uhuru*, with similar technology and science results. *Ariel 5* and *SAS-3* utilized their low orbits to process data quickly, ushering in an era of rapid response both by space and ground telescopes. (Notifications of a flaring source were announced much quicker than the six months it took Giacconi to publish the original discovery.) Positional accuracy improved with these satellites, aiding counterpart identification. *Ariel 5* helped Seyfert galaxies and active galactic nuclei (AGNs) become recognized as important extragalactic X-ray sources [109]. *SAS-3* found objects as disparate as a quasar, white dwarfs, and the (confirmed extragalactic) diffuse background. *Hakucho*

discovered the first quasi-periodic oscillation in a burster [118].

SAS-3, the MIT-led satellite launched in 1975, was headed by George Clark, Hale Bradt, and Walter Lewin[5]. Like *Ariel 5*, data from a modem link[6] was analyzed in real time, 24 hours a day, to look for targets of opportunity. Responsive steering ('pointed Y-axis observations') proved to be the satellite's greatest asset, though required great finesse and intuition to prevent the instrument from slewing to empty sky[7]. Over ten highly variable objects were discovered and studied[8], such as the aptly named Rapid Burster [86]. Though also a 'small astronomy satellite' like the single-instrument *Uhuru* (*SAS-1*), *SAS-3* carried four instruments each tuned for specific purposes: astrometry to 15″, transient source discovery, spectral characterization, and wide field low-energy surveying.

A steady stream of science results and growing collaborations ushered in observatory-class X-ray missions in the 1970's. NASA oversaw three satellites for its High Energy Astronomy Observatory Program. *HEAO 1* carried four X-ray instruments, *HEAO 2* – later renamed the *Einstein Observatory* – featured an X-ray mirror with interchangeable instruments at the focal plane, and *HEAO 3* studied gamma rays and cosmic rays.

HEAO 1 offered improved spectral, astrometric, and timing resolution of X-ray objects, but no images. The four instruments onboard were provided by different teams with different goals. The Large-Area Sky Survey nearly tripled the number of known point sources. The Cosmic X-ray Experiment, from NASA Goddard, collimated 2 – 60 keV light into thin window proportional counters to characterize the cosmic X-ray background. MIT's Hale Bradt helped build the Modulation Collimator [98], which refined point source positions for follow-up in other wavebands. The Hard X-ray/Low Energy Gamma-ray Experiment

[5]Other future Building 37 researchers such as Claude Canizares, Saul Rappaport, and George Ricker contributed significantly to the project.

[6]Unlike *Chandra* to follow, *SAS-3* was operated off campus at a Goddard Space Flight Center control room.

[7]John Doty recalled, 'Walter Lewin loved [real-time steering], and demanded its use for all his observations. SAS Control hated it, and used two controllers whenever we were trying it: one to control SAS and one to try to deal with the phone calls by the frantic duty scientist.' [125]

[8]'We had discovered a pulsar in 4U0115 + 63 which later became the first transient X-ray source to be proven to be a member of a binary system,' wrote Lynn Cominsky. 'Being duty scientist, however, is the one thing *SAS-3* allowed me to do that no other woman had done, and I always found it exciting and challenging.' [125]

employed solid state scintillation counters to look at the sky above 20 keV. Walter Lewin of MIT headed the data reduction.

The *Einstein X-ray Observatory*, or *HEAO 2*, launched in 1978 under Giacconi's leadership. The satellite originally started out under a different name[9], the *Large Orbiting X-ray Telescope*, and with a more ambitious design. Due to cost overruns in the Viking program, *LOXT* was cancelled, then restarted as *HEAO 2*, with a smaller mirror and half the budget [44]. A single high resolution mirror focused 0.2 – 4 keV X-rays onto a focal plan with swappable instruments. In addition to two proportional counters, several types of spectrometers ensured that objects could be studied with high spectral resolution ($E/dE \sim 10 - 1000$). Among these, the Focal Place Crystal Spectrometer (FPCS) is remarkable for producing Bragg-diffracted spectra of sources, some even extended like Puppis A [88]. Finally, the High Resolution Imager produced pictures at $\sim 2''$ resolution in addition to timing resolution down to milliseconds, with a new technology of channeltron multipliers.

Thanks to the grazing-incidence mirror, the first images of extrasolar objects in the X-ray band from *Einstein* secured the mission its place in history. The young supernova remnant Cassiopeia A showed structure in X-rays different than in the visible band (Figure 1-3), and temperatures and abundances could be roughly mapped [94]. Phase-folded freezeframes showed the pulsar within the Crab Nebula blinking at 30 Hz, many point sources in Andromeda were resolved from each other, and a jet could just be seen emanating from the center of M87 [44]. A long exposure of 780 square degrees found just over one source per degree, a remarkable improvement in sensitivity (though 1000× fewer objects than optical telescopes of the time could resolve) [118]. These images inspired the public and astronomers alike.

Despite the quality science and inspiring images, the most lasting contribution of *Einstein*, according to Giacconi himself, has been its open data policy. Large Principal Investigator projects had typically afforded their teams proprietary data rights. *Einstein*, however, made

[9]In fact, according to Heilbron [55], 'Giacconi had wanted to name this satellite *Pequod*, from Melville's *Moby Dick*, as a reminder of Massachusetts and the American Indian. The inevitable comparisons of Giacconi and the egomaniacal Captain Ahab indicate the qualities needed to drive a large and complex scientific project to completion. NASA declined to associate its satellite with a white whale.'

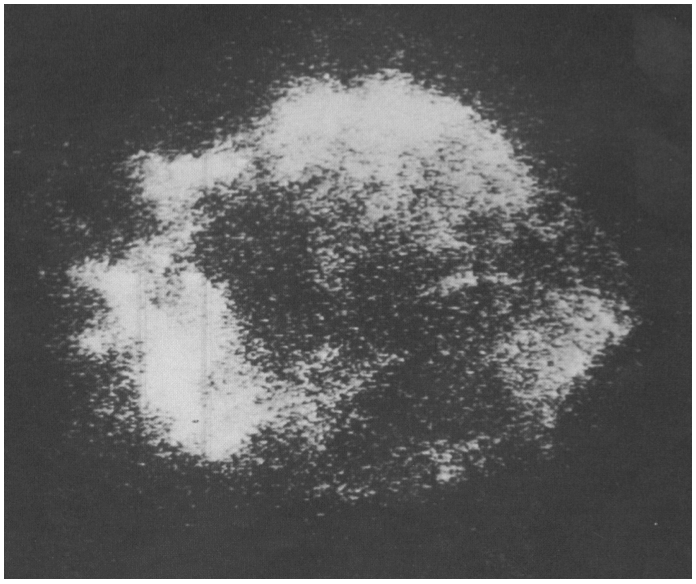

Figure 1-3: A striking first X-ray image of Cas A demonstrated the potential of grazing X-ray mirrors for astrophysical exploration. This 36 ks observation by the HRI onboard *Einstein* has 2″ resolution. (Reprinted from Murray et al. [94].)

data available immediately, in nicely reduced form. This access facilitated time-sensitive multi-wavelength studies and even helped to 'legitimize' X-ray astronomy among the other sub-disciplines. Online data reduction, calibration, archival, and retrieval were all put into place for the first time on a large scale. This open data policy is now part of any observatory, such as *Hubble*, the Very Large Array, the Sloan Digital Sky Survey, *Kepler*, and *Chandra*.

Giacconi additionally envisioned a guest investigator program to allow scientists outside the collaboration to avail themselves of the telescope's unique facilities. Astronomers could apply for time to observe objects of their own interests, and be judged by peer review. Over its lifetime, the satellite hosted 150 guest observers, producing a remarkable publication rate of over 100 publications per year.

During the time between the conclusion of the *HEAO* observations and *Chandra*'s first light – over 15 years – several projects made important advances. European missions included *EXOSAT*, *ROSAT*, and *BeppoSAX*, while Japan matured its program with *ASCA*, and the US continued apace with *RXTE*.

European groups developed X-ray expertise while *Chandra* was being meticulously built

in the United States. *EXOSAT* jumpstarted the European Space Agency, allowing the nascent program to eventually match pace with NASA and launch *XMM Newton* in the same year as *Chandra*. To its own credit, *EXOSAT* continued to further understanding of point-like sources like low mass X-ray binaries, pulsars, and AGNs, due in part to its unique orbit which could use moon occultation for source positioning, like its Skylark rocket predecessors [109]. *ROSAT*, a German satellite, began its famous all sky survey campaign in 1990, cataloguing 19,000 sources in the soft X-ray band to an accuracy just below $10''$ at best [135, 105]. *ROSAT* also provided compelling images of supernova remnants, galaxy clusters, and the moon (providing a nice punctuation to the original goal of Giacconi's Aerobee). The primary strength of *BeppoSAX*, an Italian-Dutch collaboration, was its quick determination of gamma ray burst locations, which were followed up in X-rays. The satellite performed well far beyond its target lifespan, like *ROSAT*, from 1996 – 2002.

Japan, too, found a foothold in the field in the 1990s. *ASCA*, the *Advanced Satellite for Cosmology and Astrophysics*, combined good imaging and spectral resolution, using new technologies. The lightweight nested shell mirrors provided usable effective area up to 10 keV, with a FWHM of $20'$ [121]. MIT's CCD Laboratory, led by George Ricker, provided the first ever CCD (charge coupled device) X-ray camera for the satellite [13]. CCDs, which are more commonly found in digital cameras, are solid state (i.e., not gaseous like proportional counters) devices that have become the modern focal plane workhorse since their use on *ASCA*. The new technology allowed individual spectroscopic features to be mapped across extended objects, such as ionized element emission lines in supernova remnants. The Japanese space agency ISAS granted US observers 40% of the telescope time, as thanks for the significant technology contributions.

The *Rossi X-ray Timing Explorer* (*RXTE*), named for MIT's Rossi, proved to be a specialized instrument capable of unique, revolutionary science. Launched at the end of 1995, the satellite provided a wealth of data on characteristically fast phenomena, such as spinning neutron stars, for a remarkable 16 years. MIT provided one of three instruments, the All Sky Monitor (ASM), which was headed by Hale Bradt, then Alan Levine[10] and Ron

[10]Levine completed his PhD at MIT 20 years earlier, having launched two sounding rocket payloads to

Remillard. As opposed to the hour to month long timescales of the ASM, the Proportional Counter Array and HEXTE afforded 1 and 8 microsecond timing resolution, respectively. This sub-millisecond data was processed on-board by the Experiment Data System (EDS), another MIT contribution headed by Ed Morgan, which reduced the large amount of data which could not all be telemetered down. NASA Goddard built and tested the satellite and managed the flight operations. *RXTE* shed light on some of the most extreme environments in the universe, near so-called compact objects. Analyses showed how X-ray pulsars – neutron stars spinning hundreds of times per second – could be spun up by accretion. The instruments could also track accretion flows onto black holes, spreading bursting conflagration across the surface of an accreting neutron star, and other highly variable X-ray sources across the sky.

The impact of *Chandra*, launched in 1999, cannot be overstated. *Chandra*'s imaging capabilities alone are unsurpassed: it can distinguish very fine details in extended objects at $0.5''$ resolution. In addition, with the insertion of dispersive gratings, the instrument achieves high spectral resolution for point-like objects. Due to its importance, and because I used the instrument for this thesis work, I will describe it in more detail than its predecessors.

Giacconi originally conceived of a *Chandra*-like mission only a year after the discovery of Sco X-1, to study the diffuse background [137]. He and Harvey Tananbaum proposed in 1976 for technology development to begin in earnest for mirrors and detectors that could resolve sub-arcsecond sources which made up this background. The project was awarded to the Marshall Space Flight Center/Smithsonian Astrophysical Observatory team, and called *AXAF*, the *Advanced X-ray Astrophysics Facility*[11]. The mission was truly science-driven, so the technology had to improve greatly to match the science goals. Though given a high priority, the technology developed slowly (some blame *Hubble* as a resource drain). In the early 1990s, the orbit was changed from low-earth to more eccentric, thereby abandoning the possibility of astronaut servicing while also requiring a lighter instrument. The change in orbit and budget cut two instruments from the observatory, as well as inner mirror segments,

study the soft X-ray background [85].

[11]The word 'telescope' was omitted to avoid two telescopes (*Hubble* being the other) as line items on Congress' budget [137].

but also simplified thermal management and increased continuous observing time [137].

The final observatory design consisted of the mirror (4 nested iridium-coated paraboloid-hyperboloid pairs), the Advanced Camera for Imaging Spectroscopy (ACIS, an array of CCDs), the High Resolution Camera (HRC, made of microchannel plates, which operate like tiny photomultiplier tubes), and the two transmission gratings for low energy and high energy, LETG and HETG. The 10 CCDs which comprise ACIS were developed, manufactured, and calibrated at MIT [41]. At 2 keV, the sensors demonstrate 80 eV resolution (full width half max) when the best 'grades' of events are selected. Some of the chips are back-illuminated to achieve better low-energy sensitivity (at the cost of worse energy resolution), as the photons are unimpeded by the semiconductor gate layer. The HETG, developed at MIT [15], can be inserted in the focal path after the mirror, dispersing photons onto a line of ACIS CCDs. The assembly is comprised of thin gold bars electroplated on many individual squares of polyimide, with bar dimensions chosen to highlight different orders for medium and high energy photons. Additionally, being CCDs, the detectors themselves can untangle any order overlap with their own broader energy resolution. The SAO established the operations control center for *Chandra*, which schedules observations and handles data, on MIT campus, where it functions to this day.

Chandra launched aboard the space shuttle Columbia, and has operated without incident or significant degradation for 14 years now. The first light image was dotted with point sources, including a $z = 0.32$ active galactic nucleus, validating the mirror performance and the original science goal of resolving background sources. The team revealed the public first light image of Cassiopeia A, the striking supernova remnant, which for the first time resolved the central neutron star, in addition to its finely knotted structure.

The telescope has uncovered many phenomena, resulting in over 5000 publications across of a variety of fields [45]. The superior angular resolution revealed point sources which account for 75% of the diffuse background originally discovered in 1962. The *Chandra* Deep Field surveys (similar to the *Hubble* campaigns) found sources 10 orders of magnitude fainter than Sco X-1, with count rates of approximately 1 per day. Shocks and bubbles in galaxy

clusters are now evident, thanks to deep observations, offering a glimpse of how the central massive black hole interacts with its environs out to great distances. X-ray flares from star forming regions were discovered, constraining stellar formation scenarios. *Chandra* resolved the jet and nucleus of Cen A, the nearest active galaxy, complementing high resolution radio images. The fine wispy structures of pulsar wind nebulae have been monitored on timescales of weeks to years, showing changes as the plasma gets whipped around at half the speed of light. Resolving some 1500 X-ray sources in globular clusters has showed how these objects are formed dynamically, not in situ. The HETG revealed changing temperatures and different bulk velocities for various ions in the blast of supernova 1987A. (See [45] for references.) Clearly, as more high quality observations continue every day, and with over a decade of archival data, the science reach of *Chandra* has not yet been fully realized.

Since it launched just after *Chandra*, the European mission *XMM-Newton* draws only comparisons to the NASA observatory. *XMM*, like *Chandra*, utilizes CCDs on its focal plane and gratings after the mirror, but the gratings are reflective, not transmissive [73]. *XMM* sports a much higher effective area, as well as better sensitivity at low energies, though it cannot match the spatial resolution of *Chandra*.

Quick on the heels of both *Chandra* and *XMM*, Japan launched *Astro-E* in 2000, but the payload crashed into the ocean. The payload was rebuilt as *Astro-EII*, and christened as *Suzaku* upon successful orbit in 2005. This instrument, too, though, had a fatal problem: the microcalorimeter instrument XRS – the first of its kind in space – lost its necessary cryogens due to a design flaw. The liquid helium vented into the internal vacuum space, causing run-away evaporation and rendering the low-temperature instrument useless soon after launch. While the loss of XRS and its high spectral resolution thwarted the hopes of the community, two instruments remained intact: a *Chandra*- and *XMM*-like CCD imaging spectrometer, and a hard X-ray detector. Like *XMM*, *Suzaku* complements *Chandra* with slightly different capabilities, but also with observing time. With three comparable instruments orbiting at the same time, the backlog of observations can be alleviated.

NuSTAR is the latest major X-ray payload, and has started releasing results from its

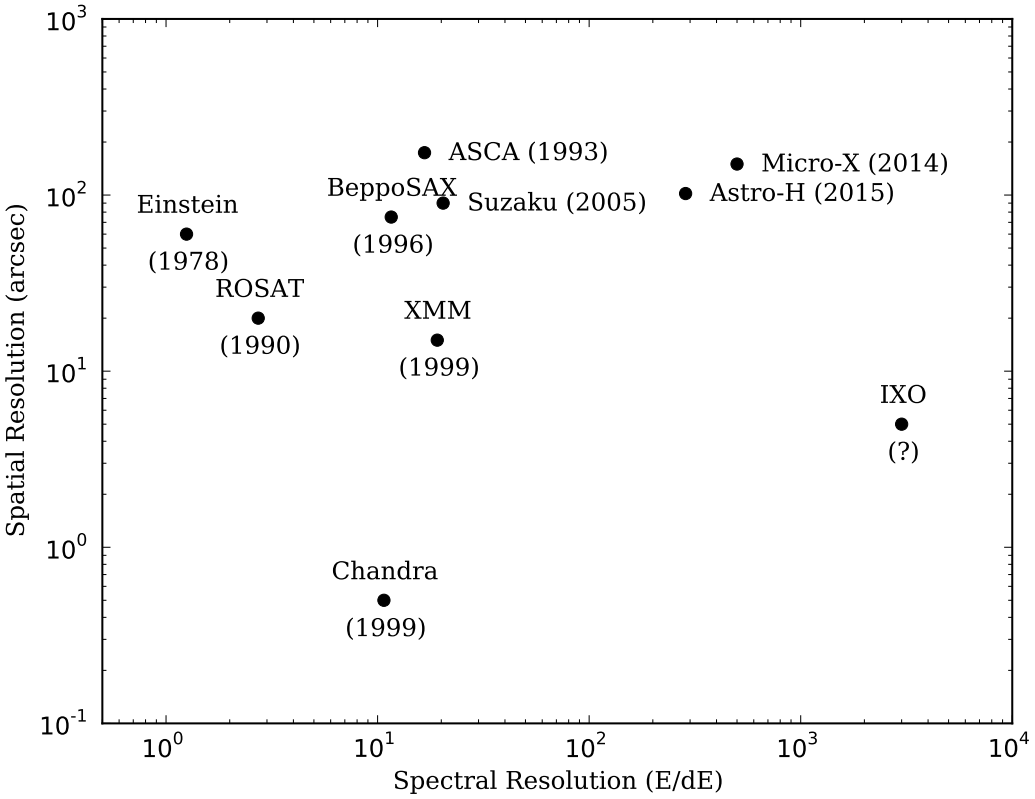

Figure 1-4: Imaging X-ray spectrometers typically must trade off between high spatial resolution and high spectral resolution. The characteristics shown here for important instruments are fiducial values over pertinent energy bands, and on-axis imaging. For multi-instrument missions, the most comparable instrument to *Micro-X* is used: e.g. ACIS on *Chandra*, or SXS for *Astro-H*.

first observations this year (2013) [52]. The NASA small explorer features two redundant, co-aligned hard X-ray (6–80 keV) detectors. The 10 m focal length mirrors were deployed on a self-assembling mast after launch on a small rocket. The telescope boasts good energy resolution (despite the lack of many lines above 6 keV) and sub-arcminute imaging. Initial discoveries have included radioactive scandium in Cas A, flaring from the Milky Way's central supermassive black hole, and, in concert with *XMM-Newton*, discovering the high spin of a nearby galaxy's central black hole.

In the near future, the Japanese/American collaboration *Astro-H* will carry an upgrade to the failed XRS aboard *Astro-EII*, to provide 7 eV spectral resolution across 0.3–10 keV.

The multi-mirror observatory will house additional instruments to detect photons up to 600 keV.

Until this point, I have described the history of X-ray astronomy as a series of successful spacecrafts, but the current era is defined by the lack of one: the *International X-ray Observatory* (*IXO*). *IXO* coalesced from two similar proposed observatories from Europe and the United States in 2008. The planned successor to *Chandra* featured a very large mirror with long focal length and switchable instruments on the focal plane, including a microcalorimeter array. The Decadal Survey, a report from the US National Academy of Sciences which strongly influences astronomy funding, relegated *IXO* to lower priority in 2010. NASA, subsequently, terminated the project, but the European Space Agency is still considering a curtailed version of the mission. Many have blamed the large budget of *Hubble*'s successor, the *James Webb Space Telescope*, for *IXO*'s cancellation, echoing the discontent with *Hubble* decades earlier as the possible cause of *Chandra*'s delay. Regardless, the X-ray community is no longer galvanized by the advent of the next great observatory.

Figure 1-4 gives a graphical overview of the spatial/spectroscopic resolution phase space that X-ray instruments cover.

1.2 This Work in Context

This thesis focuses on the use of high resolution spectroscopy to understand supernova remnants (SNRs), the bright, transitory remains of violent stellar explosions. These two domains, one from astronomical instrumentation, the other from observational astrophysics, are interesting in their own rights. The application of spectroscopy *to* SNR, though, promises to push both fields further, leading to an exceptionally rich understanding of new physics.

Spectroscopy complements imaging to provide a full physical picture of astrophysical sources, for reasons I will elaborate later. High resolution X-ray spectroscopy has been available for some time, but utilizes dispersion to sort the signal by energy, like a prism. Dispersive elements, however, only work for point sources, otherwise the signal overlaps on the detecting elements, smearing the energy resolution. Thus, ideal spectroscopic targets

are unresolved points of X-rays, like the central supermassive black hole of a distant galaxy, an accreting neutron star in the Milky Way, or the corona of a nearby star. The technology has been unavailable to extended objects like Galactic SNRs, leaving many open questions about their physical states.

The maturation of microcalorimeter technology will allow high resolution spectroscopy for all X-ray sources, extended or not, since each pixel can determine the energy of individual photons to high precision. *Micro-X*, the novel rocket payload detailed in the second half of this thesis, is helping to lead this sea change in X-ray astronomy.

SNRs are ideal targets for this young X-ray imaging spectroscopy technology: they are extended on the sky, display many atomic transition lines, and are often bright. For these reasons, Puppis A, a 4000 year old Galactic remnant, is the target of *Micro-X*'s first flight. Ahead of launch, the *Micro-X* team began several studies of SNRs, to become familiar with the science and current analytical and spectroscopic techniques. The first half of this work details two investigations with *Chandra* into the ejecta properties of the Cas A remnant. These studies both feature high resolution spectroscopy data or simulations to glean more understanding than the broad energy resolution of CCDs can provide.

Supernova remnants, of course, are interesting in their own right, besides being ideal imaging spectroscopy targets. The Big Bang did not produce any elements heavier than Li (i.e., nearly all of the periodic table), leaving the remainder to be produced in stars and distributed far and wide via their explosive deaths. Our solar system coalesced from old SNRs, forming thanks not only to the passing shock wave, but the enrichment of metals from the remnants, which allowed the protostar to cool. Studying SNRs is akin to a star autopsy: much can be learned about the life and death of the star from a post hoc investigation of circumstantial evidence. Moreover, the plasma formed in the wake of the explosion shock is of much lower density than any produced in a laboratory, offering unique opportunities to study extreme conditions.

While other objects and phenomena – galaxy clusters and the diffuse X-ray background, for instance – could benefit from spatial spectroscopy, I will not discuss them here. Likewise,

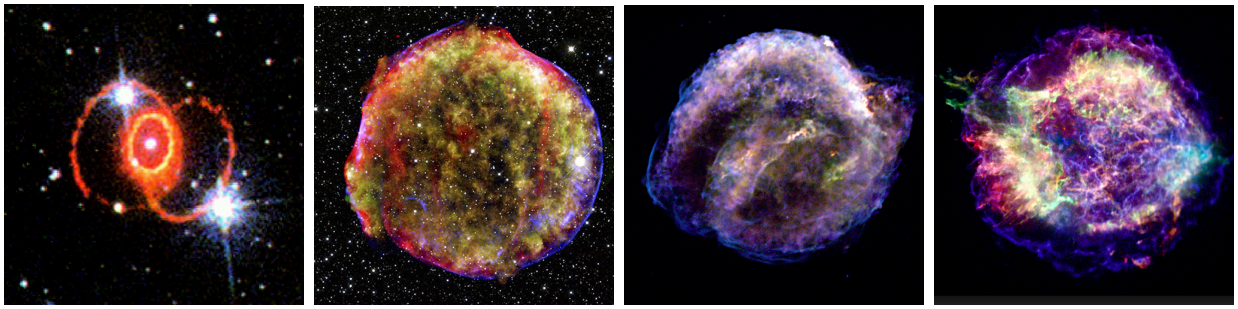

(a) (b) (c) (d)

Figure 1-5: The complex morphologies and spectral features of supernova remnants yield a great amount of information about their explosive origins. SN 1987A displays three rings in optical/infrared light (a), and will provide a close-up view of evolution for decades to come. Tycho's remnant (b), named for the man who studied SN 1572, and Kepler's remnant (c), likewise named after the astronomer who investigated SN 1604, represent the evolution of two historical events during the birth of modern astronomy. Cas A (d), the subject of this work, shows a prominent jet and broken symmetries. The three X-ray images (b–d) display shells of nonthermal emission in blue encircling thermally emitting shocked ejecta.

to keep this document concise, I will omit any broader survey of X-ray objects besides SNRs, referring the reader to more comprehensive sources [31, 118, 45].

1.3 Supernovae

Supernovae are extremely energetic stellar explosions, reaching luminosities 10^{10} times brighter than the Sun, on par with most galaxies. The energy released in a supernova, around 10^{53} ergs , is carried off mostly by neutrinos, but approximately 1% is dumped into the surrounding medium as the remains of the star get ejected outward (roughly the energy output of the Sun in its 10^{10} year lifetime). Photons carry away about only 1% of the kinetic energy. The shock wave interacts with the ejecta and interstellar medium (ISM) to create a remnant of the supernova, short-lived evidence of the blast like a puff of smoke suspended near a fired gun.

Two separate physical processes define this single observationally-defined class of 'supernova': the core collapse of a massive, dying star, and the thermonuclear detonation of an

accreting white dwarf. The first process occurs when a star's core can no longer balance the inward force of gravity. Fusing any elements heavier than Fe requires – instead of releases – energy, so as burning of lighter elements proceeds in a star, Fe 'ash' rains down on the core and forms a hot, dense Fe core. The core is supported against runaway collapse by the degeneracy pressure of the free electrons in the core, until the mass of the core grows beyond the Chandrasekhar limit. At that point, the core streams inward with a free fall time less than a second, quickly heating the material and trapping neutrinos, halting only when the density reaches the nuclear regime ($\sim 10^{14}$ g/cm^3). The quick stop sends a shock wave through the infalling material, expelling it outwards with great speed ($\sim 0.1c$). The collapse of lighter progenitors ($M < 50M_\odot$) will force electrons to combine with protons, forming an entire 'star' of neutrons, supported by neutron degeneracy pressure, while a larger collapsing mass will form a black hole. The explosion physics is still an area of active research, with several competing models for how the shock pushes through to the outer layers. These models must feature asymmetry, as many remnants display a preferred axis, and neutron stars have been observed recoiling from the center of the explosion at ~ 1000 km/s.

The second supernova process occurs when a white dwarf – the result of a less massive star's death – accretes enough matter from a binary companion star to initiate another round of burning and a subsequent explosion. During a lower mass star's ($M < 8M_\odot$) lengthy lifetime, the cooler temperatures in the core prevent burning of elements beyond helium, and at some point, gravity overcomes the radiation pressure from the core. The He or C core collapses until electron degeneracy supports it (again, the same large-scale quantum mechanical phenomenon seen in neutron stars and the Fe core). Meanwhile, winds slough off the outer layers of the star, leaving the degenerate core to cool on its own. If the white dwarf, as the bare core is called, has a binary companion, or captures one, it too can achieve a dramatic supernova exit, despite its humbler low mass beginnings. To explode, the white dwarf must accrete enough mass to initiate C burning. As the object is degenerate, it cannot puff up to regulate the high temperatures from the nuclear reactions. The runaway reactions release enough energy to unbind the star, either by detonation or

deflagration, and the supernova explosion ensues. Material can accrete either from the outer layers of a companion star, or from the coalescence of two white dwarfs. A number of recent observations have pointed to the latter 'double degenerate' model as the favored – or at least dominant – method, as signatures of disrupted companions have not been seen in X-ray or UV, nearby galaxies are underluminous in accretion-produced X-rays, and large white dwarf progenitor masses have been inferred for some supernovae.

Classes of supernovae are divided not according to their explosion mechanisms, but rather their optical observational signatures, reflecting how observation outstripped theory. Thermonuclear explosions resulting from accretion onto a white dwarf are classified as Type Ia (no H lines, with Si absorption). Core-collapse supernovae comprise the rest of the subclasses. Type Ib (no H or Si) and Ic (no H, He or Si) are thought to be massive stars that have shed their outer H layer, or H and He in the case of very massive Ic's. The Type II core collapse supernovae exhibit H absorption lines, and are divided further according to their lightcurves: Type IIP lightcurves plateau, while Type IIL decline linearly. A final class, Type IIb supernovae initially show Type II characteristics, but their spectra evolve into that of a Type Ib, presumably due to a thin H outer layer. If Type IIb's are any indication, these classes represent a continuum of progenitor masses and environmental conditions. Supernova types can be determined even well after the explosion, from the spectrum of scattered light 'echoes' off a dust cloud; Cas A is now believed to be of Type IIb [79]. Inference of the original explosion type does not have to rely on spectra, however. Since Type Ia's are presumed to obliterate the white dwarf progenitor, the existence of a compact object within the remnant greatly favors a core-collapse explosion model.

The latest generation of telescopes discover and monitor hundreds of supernovae each year. However, in the modern era, every one of these events has occurred outside our galaxy. After their bright appearance, the point sources fade in a few months, precluding any more follow-up spectra, while their distance prevents morphological studies. The estimated rate of supernovae for the Milky Way is two per century, which either means we are in a Poisson dry spell, or all the action is happening on the other side of the galaxy, along obscured sightlines.

No Galactic supernovae have appeared since the birth of the observational astronomy, in the 16th and 17th centuries, the time of Galileo Galilei, Johannes Kepler, and Tycho Brahe. (The progenitors of Cas A and G1.9+0.3, the two youngest known remnants, exploded around 1680 and 1870, respectively, but went largely unnoticed.) An explosion in 1987 in a satellite dwarf galaxy, the Large Magellanic Cloud, has borne the most scrutiny of any supernova, thanks to modern instrumentation and its proximity (50 kpc). SN 1987A, as it is called, has yielded fascinating data in the optical band, in X-rays, and from neutrinos, and has only whetted astronomers' appetites even more for a local event.

1.4 Supernova Remnants

The supernova sends hot ejecta streaming outward supersonically, parcels of metals which will compose the structure of a supernova remnant (SNR). This ejecta cools quickly as it expands, but is reheated by the shock that forms at the expanding boundary between the ejecta and circumstellar material (CSM). A secondary shock, the *reverse shock*, arises once enough CSM has been swept up to push back on the ejecta, and proceeds inward (in the frame of the forward shock), heating the interior ejecta as well to X-ray emitting temperatures. The SNR ejecta ionizes as it heats, forming a diffuse plasma. Observations with precise imaging over long timescales have confirmed the high expansion velocities of thousands of km/s by measuring proper motions, which Doppler measurements corroborate for line-of-sight motion. Both forward and reverse shock fronts can be delineated in some SNRs like E0102-72 [39] and Cas A [48] by looking for discontinuities in emissivity, temperature, or ionization state across some radial boundary.

The SNRs in our galaxy provide us with a closer look at the aftermath of these highly energetic explosions. The explosive remains of a dead star stay visible for only a few thousand years, $\sim 1000\times$ shorter than the star's life. Figure 1-5 shows a few example SNR images, hinting at the variety of morphologies possible from different explosions and surroundings. SNRs are unique laboratories for studying astrophysical plasmas, shocks, nuclear physics, and chemical evolution. The hot, low-density plasma formed in the wake of the supernova

explosion shock wave cannot be created in laboratories, or even in the coronae of nearby stars. The spectral signatures of elements can constrain explosive nucleosynthesis models, and the distribution of the metals throughout the remnant can inform explosion models. And, most pertinent to our existence on Earth, the ejecta from the supernova provide the metals for the enrichment and chemical evolution of the galaxy, while the passing shock seeds the formation of new stars. If it weren't for nearby supernovae, we wouldn't be here to study them.

Approximately 300 SNRs have been observed in the Milky Way, and nearly all of them are visible in the radio band. Of those, some 110 emit X-rays – 20% as point sources, and the rest as extended emission. The radio emission and 'non-thermal' X-ray emission both arise from electrons emitting synchrotron emission at sites of magnetic field amplification. Magnetic field strength cannot be easily be inferred from a powerlaw spectrum without some assumptions, but combined radio and X-ray observations imply magnetic fields of 10's of μG in various SNR, $\sim 10\times$ greater than the ambient field of the Galaxy [114].

During the early stage of SNR evolution, the shock proceeds unimpeded at a constant speed of roughly one parsec per century, sweeping up material ahead of it. The supersonic growth sets up a contact discontinuity between the hot expanding interior and the swept-up ISM, behind the shock. The material between these two boundaries gets heated by a factor proportional to the Mach number squared, so can reach X-ray-emitting temperatures quite efficiently.

Once the mass of the swept-up ISM approaches that of the ejecta, the SNR enters the Sedov stage, characterized by a constant energy condition: $E \propto \rho r^3 v^2$. As $v = dr/dt$, we can integrate the equation to find that $r \propto (Et^2/\rho)^{1/5}$, slower than the linear free expansion of the earlier stage, $r \propto t$. The temperatures behind the shock characteristically emit in the X-ray band as the shock velocity slows from its initial value of $\sim 10,000$ km/s:

$$k_B T = 3/16\, m_p v^2 \approx 2 \text{ keV} \left(\frac{v}{1000 \text{ km/s}} \right)^2 .$$

The Sedov phase lasts approximately 30,000 years, and encompasses all of the Galactic X-ray

SNRs. (For reference, the Cygnus loop is a soft X-ray source ($kT \sim 200$ eV) at most 10,000 years old, with a shock velocity now of ~ 200 km/s.)

At the point where radiative cooling becomes significant, the assumption of constant energy no longer holds, and the SNR enters its third stage of evolution. During this 'snow plow' phase, momentum is conserved: $\rho r^3 \dot{r} = \text{const}$, so that $r \propto t^{1/4}$. Lower mass ions like C, N, and O recombine with electrons as the temperature behind the shock drops below 10^6 K. The thin shell expands outward with a velocity of 100–200 km/s, which is only weakly dependent on the explosion energy, density, and time.

After approximately 10^5 years, the SNR comes to the end of its brief existence, as the shock velocity approaches the sound velocity, a few tens of km/s. At this point, radial expansion stops, and the remnant thermalizes with the ISM around it. Having spread its metals over a sphere of radius 40 pc, the remnant fades from view. As the Galactic disk is only about 200 pc thick, 15 kpc in radius, and over 10^{10} years old, nearly every point in the galaxy has interacted with an SNR in the past. Stars that form later, seeded by SNR shocks and metals, display higher metallicity in the spectra of their coronae. Metal-rich clouds also coalesce easier than the earlier generation of protostars, as metals cool a protostar efficiently, allowing it to compact and begin fusion.

The simple Sedov-Taylor solution outlined above neglects the effect of the reverse shock, which arises in the early free expansion phase. The growing pressure of the ISM at the contact discontinuity eventually exceeds the ejecta's declining thermal pressure, sending a shock wave backward towards the explosion origin (in the rest frame of the expanding shock). To the inside of this shock, 'cold' ejecta continues to expand, while outside of its boundary, but within the contact discontinuity, ejecta is heated to high temperatures. Eventually the reverse shock reaches back to the center, heating all of the ejecta. Models which take into account the velocity distribution of the ejecta produce this double-shock scenario, and indicate full reheating after a few thousand years [132]. The Sedov-Taylor model can be modified further if the CSM has been blown out into a bubble, which may produce high X-ray luminosity when the shock reaches the dense outer region. Pan et al. [100] showed

that such superluminous X-ray supernovae can occur with typical explosion energies when the CSM is both massive and blown far out, a regime which results in an optically thin shock with high enough density to produce strong emission.

In addition to enriching it with metals, SNRs impart energy to the ISM via cosmic rays. SNRs are thought to be responsible for the creation of cosmic rays at their shock fronts, potentially accounting for a third of the ISM energy budget [132]. The magnetic field amplification, evident from synchrotron emission, provides indirect proof of cosmic ray acceleration at the shock. The cosmic ray spectrum is a broken power law, suggesting a common acceleration method for particles up to the 'knee' at $\sim 10^{16}$ eV. If SNRs were this source of constant cosmic ray energy density, 10–20% of the kinetic energy of SNe would need to be converted to cosmic rays [56].

Shocks in SNRs have high Mach number, but are collisionless, relying instead on plasma wave excitations to 'communicate' the presence of the shock over scales much smaller than the collisional, Coulombic mean free path. The ISM has a density of only a few particles per cubic centimeter, many orders of magnitude lower than, say, Earth's atmosphere, where sonic-boom-causing shocks are generated by collisions between particles. The characteristic diameters of SNRs are *smaller* than the average distance a particle will travel before hitting another. This phenomenon requires an existing magnetic field. The magnitude of such a field can be estimated from the physical volume of X-ray synchrotron emission. Collisionless shocks are not well understood, so SNRs represent a good testbed for various theories. Synchrotron emission from accelerated electrons has been observed at the shock boundaries over a wide range of bands, from radio through X-ray.

The primary elements formed in supernovae are O, Ne, Mg, Si, S, Ar, and Ca, dubbed the 'alpha-elements'. Fe, Ni, and other 'Fe-group elements' (including radioactive ^{44}Ti), also help to provide constraints on the progenitor mass. O, Ne and Mg dominate the yields of core-collapse supernovae, and the O yield is a sensitive function of the progenitor mass for most models; O-rich SNRs are assumed to derive from high-mass progenitors [132]. For this reason, core-collapse SNe are believed to be responsible for creating most of the oxygen in

the universe. Thermonuclear SNe produce 10 times more Fe than core collapse SNe, so Fe-rich SNRs likely result from Type Ia explosions. Since X-ray emitting plasmas are optically thin, abundances can be tabulated relatively easily, constraining nucleosynthesis models, in principle. The only complication is if elements in some volume of space have not been heated to X-ray emitting temperatures.

Spatial spectroscopy has become a standard technique in the age of *Chandra*, producing maps of temperature, ionization and elemental abundances across individual SNRs. Practitioners tolerate the moderate energy resolution of CCDs, leaning heavily on the spatial resolution to paint a finer picture of the remnants' states. Hwang and Laming [66], for instance, used observations totaling 1 megasecond (11.5 days) to map Cas A's properties, showing gradients of temperature and abundances, inferring the distribution of ionization ages, and carefully summing the total ejecta mass. Future spatio-spectral studies with microcalorimeters will need to start in the opposite regime, as pixels will be large. (For bright sources, dithering the telescope may provide sub-resolution imaging.)

Spatio-spectroscopic studies have revealed that a wide range of ionization conditions can exist across a remnant. This fact may not be too surprising, given the large volume and long mean free paths of the plasma, yet it highlights the important fact that SNRs are inevitably more complex than the textbook examples.

The chemical makeup of SNRs also diverges from textbook spherical symmetry; it is well accepted now that the core collapse process is chaotic. Elements are distributed through remnants randomly, with coherent bullets of pure ejecta observed streaming outward at high velocity. Such bullets persist even in old SNRs like Puppis A and Vela. 'Overturning' is observed: low-lying Fe found farther out than outer-layer elements like Si or O [61] (though overturning is also disputed [18]).

ASCA and *Chandra* observations, in conjunction with the γ-ray instrument Fermi, have confirmed that SNRs are factories for cosmic ray production [131, 30, 136]. Thin wisps of hard X-rays are observed on the outskirts of young SNRs like Cas A and SN 1006, outside the line-emitting ejecta. These wisps are the site of magnetic field amplification, near the edge

of the shock, which exhibit synchrotron emission, and are coincident with γ-ray emission.

Morphological studies of remnants can also inform the explosion type. Type Ia remnants are generally more symmetric than their core-collapse counterparts. The thermonuclear explosion proceeds from a quickly expanding, regular burning front, which is more likely to imprint the ejecta with symmetry. Core collapse may occur asymetrically, however, resulting in an explosion regardless. Furthermore, the winds of massive core-collapse progenitors may distribute material without regard to spherical symmetry. Lopez et al. [87] confirmed these notions quantitatively by computing multipole moments of many SNRs, discovering a clean segregation between Type Ia and core-collapse remnants along the axes of second- and third-order multipoles. Type Ia's occupy the lower corner of this phase space, with low values for higher-order multipoles.

1.5 Plasmas in Supernova Remnants

Plasma, the state of matter least familiar in daily life, consisting of ionized atoms and their stripped electrons, is common in the cosmos where temperatures can often reach above 10^6 K. The environments of stellar coronae, compact accretion disks, supernova remnants, and galaxy clusters all heat gas to high temperatures, turning it into X-ray emitting plasma.

Astrophysical plasmas are generally either collisional or photoionized. Collisional plasmas are found in explosive situations where shocks develop, such as in SNRs or interacting binary star systems, or in hot environments like stellar coronae or the intra-cluster medium. Gas can also be ionized to a plasma state by high energy photons; accretion-powered sources often form photoionized plasmas. For a collisional plasma, the hot electrons hit the ions, exciting or further ionizing them, causing the ions to emit spectral lines as they relax. In a photoionized plasma, the spectral lines arise via excitation from the bath of energetic photons, as the electrons are too cool to impart enough thermal energy. The electrons and ions often exhibit different temperatures, and the role of the *electron temperature* strongly influences the different properties of collisional and photoionized plasmas. For instance, recombination lines, emitted as an electron is captured and falls down to a lower atomic

shell, are narrow in photoionized plasmas, as the kinetic energy of the electron prior to capture is low.

X-ray emission from plasma in SNRs is split broadly into two categories: thermal and non-thermal emission. Thermal emission, which arises from thermal excitations within and between the ions, can inform us about the temperature of the plasma, its progress towards equilibrium, and the various metals in the ejecta. Plasmas in young SNRs, unlike more dense plasmas, are most often found in the state of *non-equilibrium ionization* (NEI). Non-thermal emission, by construction, is everything else and includes synchrotron emission at the shock fronts, as well as radioactive decay.

The low density of the collisional plasmas in SNRs results in NEI, a state which affects spectroscopic signatures. The eventual, energetically preferred state is *collisional ionization equilibrium* (CIE), where the population fractions of various ionization states are stationary at a given temperature. CIE occurs once enough collisions per ion have transpired to reach a steady state, which is governed by the temperature and density. This situation can be expressed as system of coupled differential equations for the number fraction F_j of ions (of a given species) with j electrons stripped [70]:

$$dF_j/dt = n_e(I_{j-1}F_{j-1} + R_jF_{j+1} - I_jF_j - R_{j-1}F_j).$$

The ionization rates I_j and recombination rates R_j out of and into state j, respectively, are functions of temperature. (This dependence has been suppressed to illustrate the simplicity of this ionization balance sheet. Second-order modifications have also been neglected, like inner shell ionization, which will eject a second Auger electron with some probability.) With this formalism, NEI can be defined by $dF_j/dt \neq 0$, while CIE represents $dF_j/dt = 0$. Modern spectroscopic modeling software solves this system of equations efficiently by finding the eigenvalues of $d\mathbf{F}/dt = n_e\mathbf{A} \cdot \mathbf{F}$ to decouple the differential equations.

At a given temperature, an NEI plasma will be ionized to a lesser degree than a CIE plasma, due to fewer collisions. Thus, inference of a plasma's electron temperature will differ wildly, depending on the assumption of NEI or CIE. A CIE plasma will exhibit strong O VIII

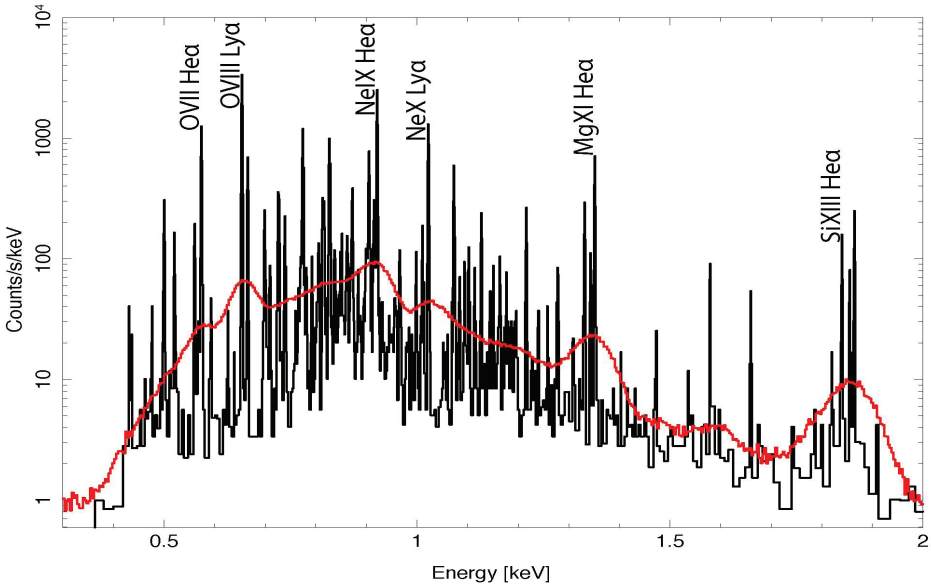

Figure 1-6: Thermal emission results in a thick forest of emission lines. The true underlying spectral features (black) are blended by the medium energy resolution of *Chandra* (red). High spectral resolution instruments will help uncover a wealth of new information, encoded in the emission lines and widths.

emission at only $kT_e \sim 200$ eV, while an NEI plasma could display the same emission at 1.5 keV.

The relevant model parameter, then, for an NEI plasma must be a measure of progress towards CIE which does not have to keep track of number of interactions per ion. The appropriate choice is the *ionization age*, $\tau = n_e t$, a product of the electron number density and time since shock, with units s/cm^3. Young SNRs have ionization ages around $10^{10} - 10^{11}$ s/cm^3, and the rule of thumb is that CIE is achieved after $\tau \sim 10^{12}$ s/cm^3.

1.5.1 Thermal Emission

When ejecta encounter a shock, a great deal of energy is imparted to the atoms, allowing electrons to escape their bonds. Generally, different temperatures will characterize the two

new populations of ions and free electrons. Simulations are required to determine the effect of shock instabilities and cosmic ray production on these two temperatures, but observations suggest these two temperatures segregate for shocks faster than $\sim$ 500 km/s. A fiducial picture to keep in mind for a strong shock is a population of ions immediately heated to 10's or 100's of keV (depending on the atomic mass), cooling slowly over time to equilibrate with the electrons which heat from 100 eV to meet the ions at 10 keV [132]. These timescales depend on the density of the plasma, so the appropriate quantity again is the ionization age. Coulomb interactions are responsible for the slow equilibration of the ion temperature and electron temperature, which could take 10,000 years. The electron temperature is the observable, imprinting on the continuum shape and emission line ratios of X-ray spectra. The ion temperature affects line broadening, an effect often too small to observe.

The smooth hump of continuum emission is produced primarily via bremsstrahlung ('braking' or 'free-free' emission), whereby electrons scatter and emit photons as they lose energy. Simulations support the assumption that local electron populations will relax quickly to a Maxwellian distribution of energies, so the emissivity can be written as

$$\varepsilon \propto \sqrt{T_e}\, e^{-h\nu/kT_e} \sum_i n_e n_i Z_i^2 \text{ erg/s/Hz/cm}^3.$$

(For commonly used symbols and abbreviations, see Appendix A.) The normalization term above simplifies for solar abundance plasmas, which are dominated in number by H [1]. Some default software plasma models therefore truncate the sum after $i = 1$, so the normalization becomes

$$1/(4\pi r^2) \int n_e n_H dV \approx n_e^2 V/(4\pi r^2),$$

where V is the volume of the emitting region and r the distance to the remnant. In the case of ejecta, which will be highly enriched with metals, the higher-Z terms must be kept, not only to account for the higher absolute abundances n_i, but also because these ions contribute more electrons. These considerations become important, for example, in our investigation of individual ejecta knots in Cas A (Chapter 2).

In some situations, 'free-bound' emission will contribute a significant continuum component [132]. As its name implies, a free electron emits free-bound emission when it falls into a partially ionized ion, releasing its kinetic energy and the ionization potential via a photon. The resulting spectrum is piecewise continuous (as the kinetic energies are continuous) between absorption edges of the atomic species.

Line emission – the forest of individual atomic transitions that sits on the hill of continuum emission – arises from collisional excitation of partially ionized atoms. Figure 1-6, a simulation of Puppis A emission with the `vnei` model, shows the variety and density of lines expected from a SNR. This collisional process is inherently thermal. The energy imparted to the ion, usually by an electron, bumps a bound electron into a higher shell, relaxing eventually by emitting a photon with energy characteristic of the transition. (The chance of a second collision to *de*-excite the ion before emission is small, and, for similar reasons, the atom can be assumed to be in the ground state initially. The extreme situation of infrequent interactions allows the plasma to exist outside of ionization equilibrium, explained in the next section.) Two main line diagnostics can help infer the electron temperature: the strength of inner shell ionization, and the ratio of finely-spaced lines of 'He-like' ions. The latter is primarily available only with high-resolution spectra.

When an ion loses in an inner shell electron due to a collision, an electron in a higher shell may fill the vacancy. The electron can fall down to the K shell in exchange for the energy difference, which may emit a photon (*fluorescence*) or eject a fellow electron (*Auger process*). (Here, the terminology is K shell for the $n = 1$ level, L for $n = 2$, M for $n = 3$, etc.) Ions with higher Z boast higher probability of fluorescing (*fluorescence yield*), and the fluorescence energies increase with ionization.

Fe is a common target for K shell emission, around 6.4 keV, which arises for electron temperatures above 2 keV. CCD instruments can detect the centroid shift of the 'blend' of Fe I (neutral), Fe II (singly ionized), Fe III (doubly ionized), and so on, to determine the most prevalent ionization state. Fe-L emission, where the ejected inner shell electron came from the L shell, arises at much lower temperatures, above 100 eV. The combination of

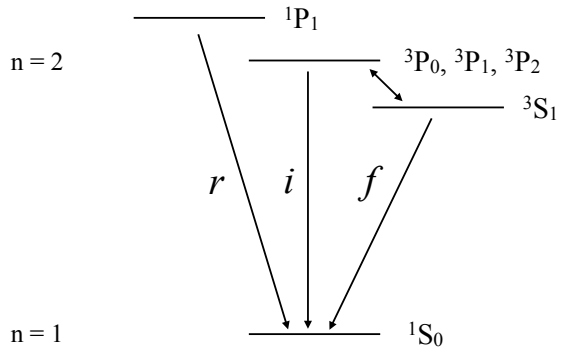

Figure 1-7: The transitions of He-like ions, stripped of all but two electrons, are easily enumerated and their energies calculated.

both Fe-K and Fe-L information can pin down the ionization state of the plasma, even with CCD instruments. High resolution instruments can provide an even more detailed picture, of course.

He-like ions, stripped of all but two electrons, provide a tractable and important diagnostic for high resolution spectra. The He-like ion is the most common species over a wide range of elements and temperatures because the ground state is a closed shell [31]. Excited states are easily enumerated (Figure 1-7) and their energy differences can be readily calculated. Transitions between the states are revealed by a triplet of lines: forbidden (f, magnetic dipole), intercombination (i, a 'blend' of two transitions of very similar energy), and resonance (r, electric dipole). The ratio between these lines can reveal the temperature and ionization state. For instance, the resonance line level continues to be populated as temperature increases, due to its electric dipole transition, but the other level populations flatten. The ionization state can be revealed via the f line, as it indicates inner shell ionization of Li-like ions, which leaves the He-like ion in a state which can only relax via forbidden magnetic dipole radiation.

Chapter 2 utilizes the high resolution capabilities of the *Chandra* HETG instrument to investigate the plasma characteristics of ejecta via the Si He-like triplet.

1.5.2 Non-thermal Emission

Relativistic electrons emit synchrotron emission as they spiral around magnetic field lines. Electrons attain such high velocities via repeated head-on interactions with magnetic ripples in the shock front, in a process called Fermi acceleration[12] The aggregated increase in individual particle energies results in a non-Maxwellian distribution. Rather a power law distribution arises, $N(E) \propto E^p$, resulting in characteristically *non-thermal* emission. Radiation from this population spans a remarkably large range, from radio frequencies (10^7 Hz) to the X-ray band (10^{18} Hz). As mentioned earlier, most SNRs are visible to radio telescopes due to synchrotron radiation.

High spatial resolution telescopes like ASCA and *Chandra* confirmed the presence of non-thermal X-ray emission in SNRs, in thin shells ahead of the bright thermal emission. Prior work had not assumed synchrotron emission at the shock front could extend into the X-ray band; instead, the non-thermal component was attributed to a central pulsar wind nebula [78]. The strength of the magnetic field can be inferred fairly easily from the synchrotron spectrum, implying electrons of 10's of TeV traveling in SNR fields of 100's of μG. These high electron energies make SNRs a prime candidate for the production of electron cosmic rays observed from Earth. They may also explain γ-ray emission from SNRs, although pion-production from ions could also account for such signatures [58].

X-ray synchrotron spectra are characterized by the photon index Γ: $F(E) \propto E^{-\Gamma}$. Most spectra are 'steep', with Γ ranging from 2 to 3.5, characteristic of radiation from electrons close to the high-energy cutoff, where radiative losses regulate the maximum energy. The steepness could also be characteristic of a very young shock, since the acceleration process requires time to produce high energy electrons.

Radioactive decay of newly synthesized unstable elements, an entirely different process than synchrotron radiation, is also classified as non-thermal emission. This classification is sensible, though: such emission depends entirely on the decay times and amount of the

[12]First order Fermi acceleration, which applies to shocks, produces energy gains $\Delta E \propto \frac{v}{c} E$. Second order acceleration is due to the random motion of magnetic interstellar clouds, providing electrons reflecting off these 'magnetic mirrors' with an energy gain, as the name suggests, of $\Delta E \propto \left(\frac{v}{c}\right)^2 E$.

material, not its temperature. (The decay lines may be broadened by thermal processes, of course, but the radiation mechanism is not inherently thermal.)

Radioactivity is important during the very early life of an SNR. The explosion forms significant quantities of ^{56}Ni ($\sim 0.5 M_\odot$), which has a 6 day half life before decaying to ^{56}Co, then further to ^{56}Fe [132]. The decay products heat the surrounding ejecta, significantly affecting the light curve in the first year.

Evidence of radioactive decay can also be found at later stages of SNR evolution, and can constrain nucleosynthesis models [84]. The trace element ^{44}Ti – also formed during a supernova, but with much smaller yield ($\sim 10^{-5} M_\odot$) – has a much longer half life of 60 years, admitting the possibility of detection in young SNRs. The yield of this specific isotope can be a sensitive gauge of the progenitor and explosion characteristics, but is strongly correlated with ^{56}Ni [8]. On its way to stable ^{44}Ca, the isotope emits photons in the γ-ray band as nuclear relaxations, and around 4 keV. The latter signal results from electron capture, where the nucleus absorbs an inner shell electron, and results in fluorescence exactly imitating inner shell ionization, discussed above. (Stable Sc is not common, so normal inner shell ionization is not expected to pollute the radioactive signal.) Specific details of the ^{44}Ti decay chain are discussed in Chapter 3. Radioactive non-thermal emission, which occurs regardless of the isotope's ionization state, reveals both shocked and unshocked ejecta, an advantage over thermal emission, which is insensitive to 'cold' ejecta.

Unfortunately, not many other decay chains emit in the X-ray band with appropriate decay times and sufficient yields. Besides the two above examples, we are left with

$$\begin{aligned} ^{55}\mathrm{Fe} &\rightarrow {}^{55}\mathrm{Mn},\ \tau_{1/2} = 2.7 \text{ yr} \\ ^{59}\mathrm{Ni} &\rightarrow {}^{59}\mathrm{Co},\ \tau_{1/2} = 75 \text{ kyr} \\ ^{53}\mathrm{Mn} &\rightarrow {}^{53}\mathrm{Cr},\ \tau_{1/2} = 3.7 \text{ Myr}, \end{aligned}$$

which all emit between 4 and 7 keV [84].

With such a low rate of local SNe, though, astronomers must avail themselves of even these improbable shreds of evidence, and radioactive line detection may become feasible and

important for future imaging spectroscopy.

Chapter 3 details the search for a radioactive non-thermal signal in Cas A.

1.6 The Need for High Resolution Spectroscopy

Spectra encode a large amount of information. Ion abundances, temperature, magnetic field strength, velocity, foreground absorbers, and ionization state can all be imprinted by photons of different energies which emanate from astrophysical sources. Including spatial or timing information enriches understanding even further.

This wealth of data from spectroscopy speeds the cycle of observation and prediction. Modern spectral models of plasma, for instance, need only reproduce broad features observed under the 'blend' of *Chandra*'s or *XMM*'s resolution. This is not to say theorists do not take known physical processes into careful consideration; but until unexpected features hidden in the blend are brought into sharp relief, they have no guidance to refine their models. Figure 1-6 illustrates the amount of information loss from blended spectra from medium-resolution instruments.

Finer spectra will also make it easier to read off the physical state. The forest of inner fluorescent lines or He-like triplet ratios will pin down the electron temperatures for multiple elements, across the field of view. The resulting picture will be rich with physical detail, and will inevitably include surprises.

The remainder of this thesis describes the pursuit of high resolution X-ray spectra. The first half explores what the best current generation instrument can reveal about SNRs. The commissioning of *Micro-X*, a next-generation high resolution telescope, follows.

Chapter 2

Ejecta Knots in Cassiopeia A: A Spatio-spectral Case Study

The bright X-ray knots of Cas A present a unique opportunity for high resolution spatio-spectroscopy. First, these nodular features – conspicuously clear in the first short *Chandra* observation – are most likely discrete parcels of ejecta, admitting easier physical interpretation of their spectra. The knots, only a few pixels across on the *Chandra* ACIS instrument, display bright Si emission but still require modeling of the broad background/foreground for accurate inference of the plasma state. Second, their brightness and tiny size allow us to obtain high resolution dispersed spectra. This rare use of *Chandra*'s transmission gratings (HETG) on an extended object has been employed for only a few other remnants, notably SN1987A [23] and E01012 [14]. Third, a number of well spaced observations over *Chandra*'s 13 year lifetime allow us to track evolution in these knots over a long baseline. The *lack* of evolution we observe leads to a more consistent picture of the knots' compositions as richer in metals and less dense than previously inferred.

This chapter details the analysis of 21 bright X-ray knots in the Cas A supernova remnant from observations spanning ten years. Dan Dewey provided the motivation for the investigation, and much of the expertise for the gratings-dispersed spectra, while the *Micro-X* team – primarily myself, Tali Figueroa, and Sarah Heine – led the ACIS analysis. I

led the ACIS parameter inference effort, providing the statistical framework, quantified our evolution results, and synthesized the various conclusions for publication.

We performed a comprehensive set of measurements to reveal the kinematic and thermal state of the plasma in each knot, using a combined analysis of 2 high energy resolution HETG and 4 medium energy resolution ACIS sets of spectra. The ACIS electron temperature estimates agree with the HETG-derived values for approximately half of the knots studied, yielding one of the first comparisons between high resolution temperature estimates and ACIS-derived temperatures. We did not observe the expected spectral evolution – predicted from the ionization age and density estimates for each knot – in all but 3 of the knots studied. The incompatibility of these measurements with our assumptions has led us to propose a dissociated ejecta model, with the metals unmixed inside the knots, which could place strong constraints on supernova mixing models.

This analysis pushes this state-of-the-art instrument to its limits. It required long observations with two instruments, operating the gratings spectroscope beyond its designed purpose. The difficulty of the analysis alone makes the case for microcalorimeter-based telescopes. Moreover, only the state of Si – the brightest element in the knots – can be probed with the dispersed spectra, while the high quantum efficiency of a microcalorimeter would produce quality spectra for more species. No microcalorimeter focal plane will match the spatial resolution of *Chandra* anytime soon, though. Such trade-offs make it clear that X-ray astronomers will need to combine different observations from high-spatial and high-spectral resolution instruments for some time to come.

This chapter was published in volume 769 of *The Astrophysical Journal* [116].

2.1 The Supernova Remnant Cassiopeia A

The SNR Cassiopeia A (Cas A), distinguished by its filamentary bright X-ray features, has been well studied in the X-ray band. From the *Chandra* first light observation [61], these features have been identified as Si-rich ejecta knots, coherent material expelled from the deeper layers of the progenitor star. As Laming and Hwang [81] first reasoned, the

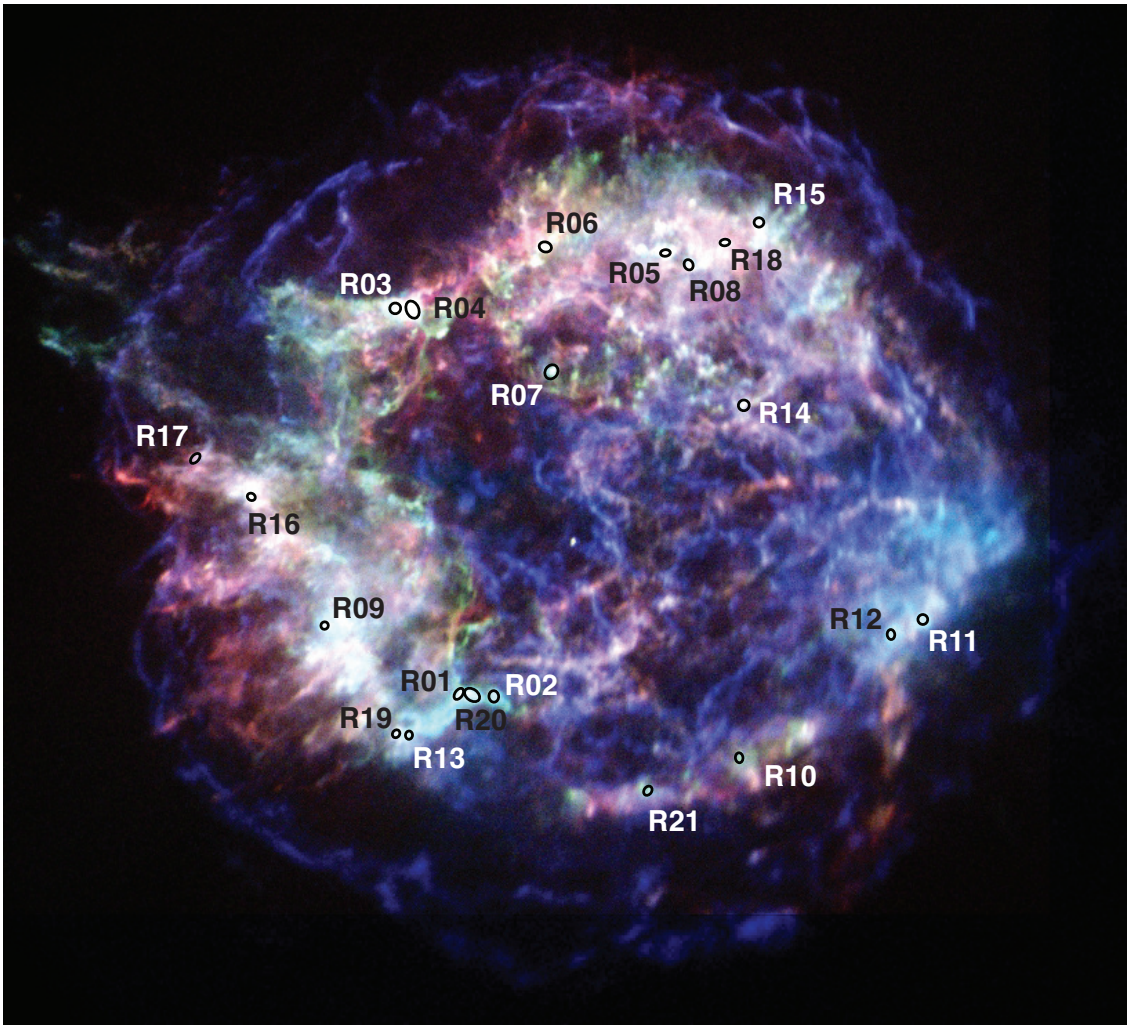

Figure 2-1: We investigated 21 bright Si X-ray knots in Cas A for this study. The first 17 coincide with those in Lazendic et al. [82]. R18 and R21 have been added, as they have brightened, and R19 and R20 were split off of R13 and R01, as they have separated on the sky since the first observation.

term 'knot' here refers to the bright X-ray features having densities of tens to hundreds of electrons per cubic centimeter — which is only a factor of a few above the density of the surrounding material — in contrast to the much denser 'optical knots' of Fesen et al. [32]. Laming and Hwang, along with the companion paper [62], systematically investigated the plasma properties of X-ray knots first, using the regions as tracers to infer the surrounding envelope's density profile. Lazendic et al. [82] (Paper A hereafter) performed the first high spatial-spectral resolution analysis of the bright X-ray knots, deriving temperatures, Doppler velocities, and abundances. That analysis utilized dispersed spectra of the extended remnant, a difficult and uncommon technique that our current work also employs. The entire 3-dimensional structure of Cas A was fleshed out beyond the skeleton of these bright knots by DeLaney et al. [18]. The exhaustive investigation suggests a flattened explosion and ejecta 'pistons'. Recently, Hwang and Laming [66] finely mapped abundances and plasma states across the entire remnant for the 1 Ms *Chandra* observation in an analysis tour de force.

These thermal X-ray knots are the subject of our investigation. We sought to characterize the plasma state of a set of localized ejecta regions in Cas A and investigate their evolution over a decade in the remnant's approximately 330 year lifetime. We collected both imaging and dispersed data of Cas A with the Advanced CCD Imaging Spectrometer (ACIS) and High Energy Transmission Grating Spectrometer (HETGS) instruments, respectively, on board *Chandra*. Analysis of the dispersed data yields ratios of the strongest lines (we only look at Si), but these ratios alone cannot fully describe the plasma state, including the electron temperature and ionization age. We therefore use the broadband spectra from the imaging data — with a model — to provide a fuller picture of the plasma state. We address two primary questions regarding the physical state of these knots. First, over the decade of observations of Cas A, can we see these knots evolving spectrally? Second, how do the two pictures of the knots' plasma states — the broadband plasma model parameters from the ACIS analysis and the Si line ratios from the HETGS data — compare and agree with predictions?

Cas A has been the subject of several other long-term investigations recently. Patnaude

and Fesen [101] took a similar approach to this work and looked at the evolution of four bright regions over four years, observing flux variation and some changes in plasma parameters. A full decade's worth of high quality *Chandra* data on Cas A was utilized by Patnaude et al. [102]. That study found a 1.5% per year decrease in the nonthermal X-ray flux across the remnant, while the thermal component stayed fairly constant. In the optical band, Fesen et al. [32] studied the flickering phenomenon of outer ejecta knots over half a decade, which can reveal the initial enrichment of the interstellar medium.

Due to the low densities of these knots and relatively short timescales, the ions and electrons have not interacted enough to achieve collisional ionization equilibrium (CIE), remaining instead in non-equilibrium ionization (NEI). Ionization age is a measure of the evolution of a plasma from NEI toward CIE; $\tau = n_e t$, where n_e is the electron density in cm^{-3} and t the time in seconds since encountering the reverse shock. As τ increases over time, the electron temperature increases to meet the the proton and ion temperatures. In NEI, the fractions of the different ionization species for a given element evolve with time as the ionization and electron capture processes balance. The CIE timescale depends on the plasma temperature and its elemental abundances; for a Si-dominated plasma at 2 keV the timescale is $\tau \sim 10^{12}\,\mathrm{cm}^{-3}\,\mathrm{s}$ [123]. The ionization ages of around $10^{10}\,\mathrm{cm}^{-3}\,\mathrm{s}$ derived in Paper A for several knots suggested the knots were in various states of NEI. The derived electron densities $n_e \sim 100\,\mathrm{cm}^{-3}$ implied that significant signs of plasma evolution should be evident over a ten-year period, through a measurable increase of the ACIS-derived ionization age and an increased Si XIV to Si XIII emission ratio in the HETG data.

We report results from analyses of observations over a ten year baseline, both with a new HETG observation (Table 2.1), and existing ACIS data (Table 2.2). We focus on 21 knots, which include the 17 knots studied in Paper A. A comprehensive set of measurements of the thermal and kinematic properties of these knots is presented. Unlike Paper A, wherein plasma parameters like temperature were inferred from line ratios, our analysis directly models and fits these parameters with ACIS data. We thereby obtain independent high-resolution line ratios and plasma parameters, ensuring cross-validation for our results.

ObsID	Start Date	Exposure (ks)	RA	DEC	Roll Angle
1046	2001 May 25	69.93	$23^{\rm h}23^{\rm m}28^{\rm s}.00$	$+58°48'42''.50$	85.96
10703	2010 Apr 27	35.11	$23^{\rm h}23^{\rm m}27^{\rm s}.90$	$+58°48'42''.50$	60.16
12206	2010 May 2	35.05	$23^{\rm h}23^{\rm m}27^{\rm s}.90$	$+58°48'42''.50$	60.16

Table 2.1: Summary of HETG observation parameters.

This study finds significantly higher temperature estimates and lower ionization ages than Paper A for most knots. The rapid evolution predicted from these new values was largely not seen. These data present strong evidence that one cannot model these knots as single-temperature non-equilibrium systems evolving without interactions with their surroundings. We propose a dissociated model of the metals inside the knots, which could constrain supernova explosion models.

The observations and analysis are described in Section 2.2, with the results presented in Section 2.3 and discussed in Section 2.4. The details of the analysis are presented in this chapter's Appendices 2.B (for HETG) and 2.C (for ACIS). The tabulated results and several plots for all the knots are presented in Appendix 2.A.

2.2 A Decade Baseline Investigation

We chose 21 knots for our analysis. These are shown in Figure 2-1 and listed in Table 2.3. The first 17 were defined to coincide with those in Paper A, which were selected for their bright Si features. Four more were added: R18 and R21 because they have brightened; R19 and R20 were split off of R13 and R01 because they have separated on the sky since the first observation.

The data we collected for this work are of two types: gratings-dispersed images, which have higher spectral resolution, and non-dispersed images, with moderate spectral resolution. For the dispersed HETG analysis, a new Monte Carlo method called Event-2D was used to extract the Si lines of the knots from the very complex dispersed datasets. For the non-dispersed ACIS analysis, we compared the derived plasma model parameters of these knots over the decade. The fine-grained spectral detail of the HETG work and the broad-brush

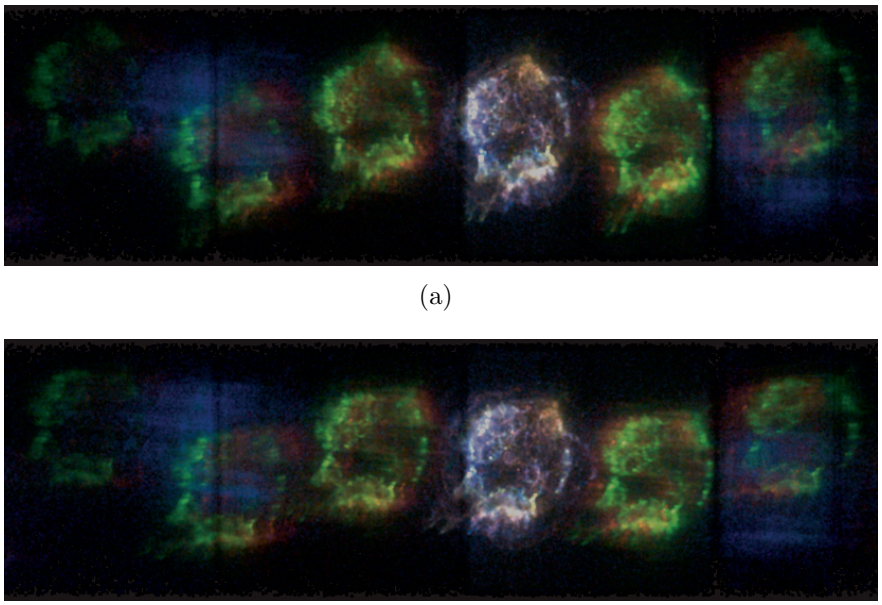

(a)

(b)

Figure 2-2: The bright knots of Cas A can be analyzed using dispersive gratings. The two HETG observations of Cas A — 2001.4 in (a), 2010.3 in (b) — were taken at different roll angles. Left to right, the images show orders MEG-3, HEG-1, MEG-1, zeroth, MEG+1, HEG+1. The false color image shows the Si XIV line (red), the Si XIII blend (green), and lines below 1 keV (blue).

nature of the ACIS analysis complement each other to provide a better picture of the physical processes in Cas A.

2.2.1 High Spectral Resolution: HETG

The prospect of directly measuring temperature or ionization state evolution for X-ray knots – based on the results of Paper A in 2001 – motivated a second HETG observation of Cas A in 2010. The observations span a nine-year baseline, and were taken as part of the *Chandra* GTO program (Table 2.1). The latter observations were taken at a different roll angle than the Paper A dataset, affording the ability to cross-check the data reduction process, as the knots would be dispersed across different slices of Cas A. Figure 2-2 shows how the bright features of Cas A can be tracked over multiple dispersive orders.

For the HETG analysis we used — for the first time — the Event-2D technique [21].

ObsID	Start Date	Exposure (ks)
114	2000 Jan 30	50.56
4637	2004 Apr 22	165.66
9117	2007 Dec 5	25.18
10935	2009 Nov 2	23.58

Table 2.2: Summary of ACIS observation parameters.

This complex high-energy-resolution analysis incorporates a 2D model of each knot and its background, then folds it through the response function of the HETG to generate a model image that includes the 2D morphology, the spectral model, and the dispersion response. This model image is then compared to the actual data image collected with the ACIS-S CCD, and fit with a Monte Carlo technique. The method yields increased energy resolution over naive approaches by modeling complex shapes and utilizing the CCD spectral resolution to disentangle the source from the background.

The Event-2D analysis uses a 4-Gaussian model corresponding to the Si XIII (He-like) recombination, inter-combination and forbidden (r, i and f, respectively) lines and the Si XIV (H-like) Ly α line. The f/i ratio is fixed at the low density value of 2.45 and the remaining free parameters are then the overall flux, the f/r ratio, the H-like/He-like ratio and the Doppler velocity. A thorough description of this analysis is detailed in Appendix 2.B.

2.2.2 High Spatial Resolution: ACIS

To complement the new HETG data, four archival ACIS observations spanning a ten-year baseline were analyzed (Table 2.2), from early 2000, mid-2004, late 2007, and late 2009.

We assume these small, bright knots are individual clumps of similar material, so we therefore defined their boundaries spectrally. Arbitrarily bounding these regions by hand, using only brightness information, could unintentionally adulterate the knot's spectrum with non-knot material of similar surface brightness, thereby skewing results of the analysis. Our automated, systematic method found where the surrounding spectrum differed significantly from that of the central knot core, and drew the knot boundary accordingly.

We used the fainter regions surrounding the knots as the 'background', choosing diffuse

plasma showing little variation with position, under the assumption that the spectral core we extracted from the knot region contained this diffuse material either in front of or behind it.

With the regions and backgrounds defined, we extracted the data using a custom pipeline based on the `specextract` CIAO script that correctly takes into account the dithering of the telescope during the observation, which is important at the small scales of our knot regions.

Each knot was fit with the conventional model: a non-equilibrium ionization (NEI) plasma with variable abundances (`vnei` [9]), with the equivalent absorption by neutral hydrogen (n_H) between the source and observer accounted for (`phabs`, abundances from Balucinska-Church and McCammon [4]). Figure 2-3 shows a typical fit to the data for a single epoch. All fitting was performed with the ISIS software [60]. As we assume the foreground and knot compositions do not change over our observation period, we performed fits to the four epochs of ACIS data for each knot simultaneously, tying n_H and the abundances across all datasets. The redshift was fixed after fitting it to the bright Si lines. (For our radial velocity measurements, we used the more accurate HETG Doppler shift.) The final values of the tied parameters were dominated by the 2004.4 dataset due to its longer integration time. The overall scaling factor (the norm), the electron temperature kT, and ionization age τ were allowed to float independently for each epoch in the combined fit. Pileup corrections were also used. The details of the ACIS analysis are presented in Appendix 2.C.

In contrast to other analyses (Paper A, Hwang and Laming), we include the contributions of low-Z elements other than oxygen to the spectral continuum. H and He are held at solar values, and C, N, and O are fixed at 5 times solar. As will be discussed in Section 2.4.1, the model and data cannot distinguish between which elements prop up the continuum, so these values are merely fiducial. We include $Z < 8$ for two reasons: (1) nucleosynthesis models predict yields of C and N on order with O, and (2) Dewey et al. [22] argue that O cannot account for all of the continuum in Cas A, otherwise the line emission would be too strong.

Region	RA	Dec
R01	$23^{\mathrm{h}}23^{\mathrm{m}}33^{\mathrm{s}}.176$	$+58°47'47''.34$
R02	$23^{\mathrm{h}}23^{\mathrm{m}}31^{\mathrm{s}}.566$	$+58°47'46''.73$
R03	$23^{\mathrm{h}}23^{\mathrm{m}}35^{\mathrm{s}}.934$	$+58°50'03''.80$
R04	$23^{\mathrm{h}}23^{\mathrm{m}}35^{\mathrm{s}}.199$	$+58°50'04''.59$
R05	$23^{\mathrm{h}}23^{\mathrm{m}}23^{\mathrm{s}}.626$	$+58°50'23''.88$
R06	$23^{\mathrm{h}}23^{\mathrm{m}}29^{\mathrm{s}}.187$	$+58°50'26''.64$
R07	$23^{\mathrm{h}}23^{\mathrm{m}}28^{\mathrm{s}}.843$	$+58°49'43''.34$
R08	$23^{\mathrm{h}}23^{\mathrm{m}}22^{\mathrm{s}}.554$	$+58°50'20''.83$
R09	$23^{\mathrm{h}}23^{\mathrm{m}}39^{\mathrm{s}}.236$	$+58°48'12''.36$
R10	$23^{\mathrm{h}}23^{\mathrm{m}}20^{\mathrm{s}}.273$	$+58°47'25''.01$
R11	$23^{\mathrm{h}}23^{\mathrm{m}}11^{\mathrm{s}}.871$	$+58°48'14''.00$
R12	$23^{\mathrm{h}}23^{\mathrm{m}}13^{\mathrm{s}}.388$	$+58°48'09''.34$
R13	$23^{\mathrm{h}}23^{\mathrm{m}}35^{\mathrm{s}}.506$	$+58°47'33''.70$
R14	$23^{\mathrm{h}}23^{\mathrm{m}}20^{\mathrm{s}}.101$	$+58°49'31''.14$
R15	$23^{\mathrm{h}}23^{\mathrm{m}}19^{\mathrm{s}}.376$	$+58°50'35''.02$
R16	$23^{\mathrm{h}}23^{\mathrm{m}}42^{\mathrm{s}}.746$	$+58°48'58''.16$
R17	$23^{\mathrm{h}}23^{\mathrm{m}}45^{\mathrm{s}}.045$	$+58°49'12''.76$
R18	$23^{\mathrm{h}}23^{\mathrm{m}}20^{\mathrm{s}}.914$	$+58°50'28''.55$
R19	$23^{\mathrm{h}}23^{\mathrm{m}}35^{\mathrm{s}}.935$	$+58°47'34''.50$
R20	$23^{\mathrm{h}}23^{\mathrm{m}}32^{\mathrm{s}}.519$	$+58°47'47''.61$
R21	$23^{\mathrm{h}}23^{\mathrm{m}}24^{\mathrm{s}}.498$	$+58°47'13''.35$

Table 2.3: The coordinates of the knot centers in the 2004.4 epoch.

2.3 Results

The two analysis pipelines produced several independent plasma diagnostics: primarily, the f/r ratio in Si XIII, the Si XIII/XIV line ratio, the ionization age from `vnei`, and the electron temperature, also from the `vnei` model. For considerations of flow and space, we present the full tabular results of our analyses in appendices.

The HETG analysis was performed for the observations in Table 2.1, and the results of those fits are shown in Appendix 2.A, Figure 2-11 and Table 2.6. Uncertainties are given for all but the flux, which has a statistical uncertainty generally less than 2% and so will be dominated by systematic errors, such as the calibration of the effective area at a knot's specific location.

The ACIS analysis was performed for the observations in Table 2.2, and the results of these fits are shown in Appendix 2.A, Figures 2-12 – 2-14, Tables 2.7, and 2.8. The reported parameter ranges represent bounding edges of confidence contours (see Figure 2-19 and Appendix 2.C).

2.3.1 Comparison with Paper A and Other Results

The values of the HETG line ratios from the first epoch are consistent with those obtained in the Paper A analysis, validating the Event-2D technique. The latest HETG observation and new multi-epoch fits to the ACIS data extend their work, but challenge their conclusions.

Our temperature and ionization age results differ substantially from Paper A. Paper A used Si line ratios to infer kT and τ values by comparing to the XSPEC model ratios. We believe our ACIS analysis of four combined epochs has much smaller systematics to directly determine kT and τ. Perhaps this tension is not too surprising, though; Paper A derived the parameters from Si only, while we infer kT and τ from a broadband spectrum composed of several ion species and a continuum component. More specifically, while all but three — R08, R10, and R17 — of Paper A's knot temperatures are *lower* than $kT = 1.6$ keV and are clustered around $kT \simeq 1$ keV, all of our results are *higher* than $kT = 1.6$ keV and are clustered around $kT \simeq 2.2$ keV (see Figures 2-12 – 2-14). With respect to ionization age,

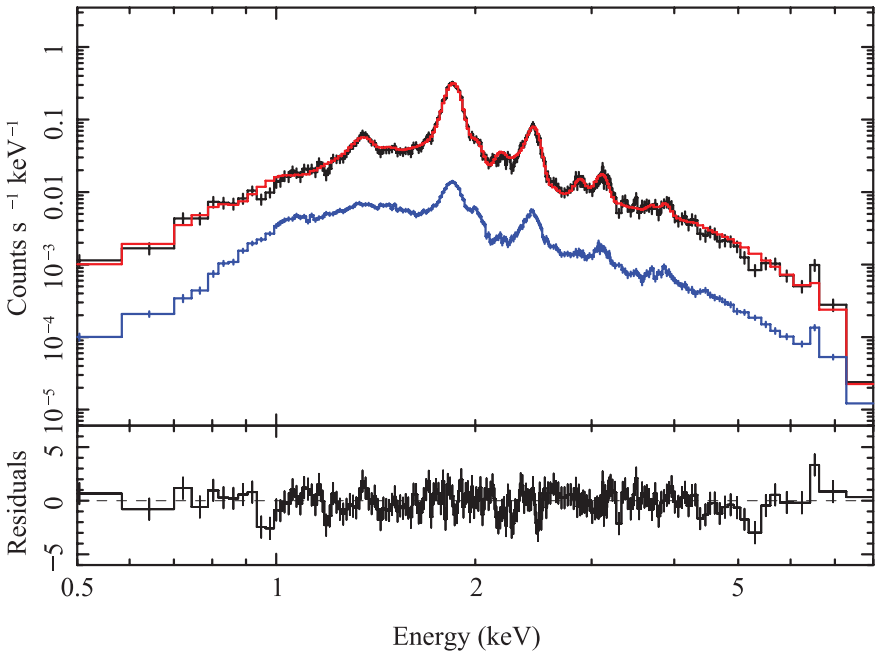

Figure 2-3: The `vnei` model fits the ACIS data well; shown is the 2004.4 epoch for R12. The data appear in black, the background in blue, and the combined model in red.

Table 2.7 lists systematically lower τ values than Paper A. All but two of our knots — R09 and R16 — have τ lower than $10^{11}\,\mathrm{cm}^{-3}\,\mathrm{s}$, while all but four — R07, R08, R10, and R17 — of Paper A's values are above $10^{11}\,\mathrm{cm}^{-3}\,\mathrm{s}$.

This work also portrays a slightly different ionization state than investigations by Hwang and Laming. Their mosaic of Cas A for the 2004.4 epoch gave distributions of kT, peaked at 1.4 keV, and τ, peaked at $2.5 \times 10^{11}\,\mathrm{cm}^{-3}\,\mathrm{s}$. This handful of knots, however, grouped around $kT = 2.2$ keV and $\tau = 4 \times 10^{10}\,\mathrm{cm}^{-3}\,\mathrm{s}$. The discrepancy could be due to choice of model (`vpshock` vs `vnei`), or simply show these knots to be a different population. The Si and S abundances are higher than in the mosaic regions, which is to be expected, but the other elements are consistent. Laming and Hwang [81] analyzed several of the same knots: R03 = NE8, R04 = NE6 + NE7, R08 = NNW4. We see general agreement, except for R08, which exhibits double the temperature and half the ionization age. The Si/O ratios are consistently twice as high in this analysis, though we fix the O abundance. None of these discrepancies are cause for concern, especially given the different analysis schemes.

The abundances of Si and S are consistent with nucleosynthesis models, though cannot constrain them. Even though these knots are local, Si-rich features, it is instructive to compare the prominent elements with fiducial global abundance predictions. We turn the abundance point estimates from Table 2.8 into a mass ratio with the solar abundances from Anders and Grevesse [1], resulting in a large spread: $1.4 \leq M_{Si}/M_S \leq 2.4$. Rauscher et al. [112] consider several masses for a Type II progenitor, and two sets of reaction rates, resulting in mass ratios $1.9 \leq M_{Si}/M_S \leq 2.5$.

2.3.2 Theory Approaching Observation: Two Measurements of Temperature

This collection of knots forms a unique set of isolated low density plasmas, the characteristics of which can be compared against the theoretical predictions of atomic transition codes. Table 2.6 shows the He-like f/r and H/He ratios of Si for this sample of knots from the HETG analysis. The 90% confidence bounding edges for kT and τ are shown in Table 2.7.

The knots exhibit a variety of ionization states. The Si XIII f/r ratios run from 0.18 to to 0.88, while the Si H/He ratios similarly cover a wide range, from to 0.035 to 0.60. The ACIS analysis revealed knots covering much of kT–τ space, with temperatures from below 1 to above 7 keV, and ionization ages from 2 to above $50 \times 10^{10}\,\mathrm{cm}^{-3}\,\mathrm{s}$, as can be seen in Figure 2-4.

He-like transition ratios have long been used as temperature diagnostics. Gabriel and Jordan [38] applied the technique to derive temperatures in the corona of the Sun, where electron densities reach 10^{11} cm^{-3}. With an eye toward extrasolar *Chandra* observations, Porquet et al. [107] significantly improved calculations for ions out to Si XIII in collisional plasmas. Most recently, Smith et al. [124] applied new fully relativistic code to Ne IX, demonstrating significant corrections to previous calculations. (More complete references on the utility of He-like transitions can be found therein.)

The HETG and ACIS analyses yield two different measurements of temperature, which we may compare against each other. The ratios of Si XIII lines from the HETG analysis can

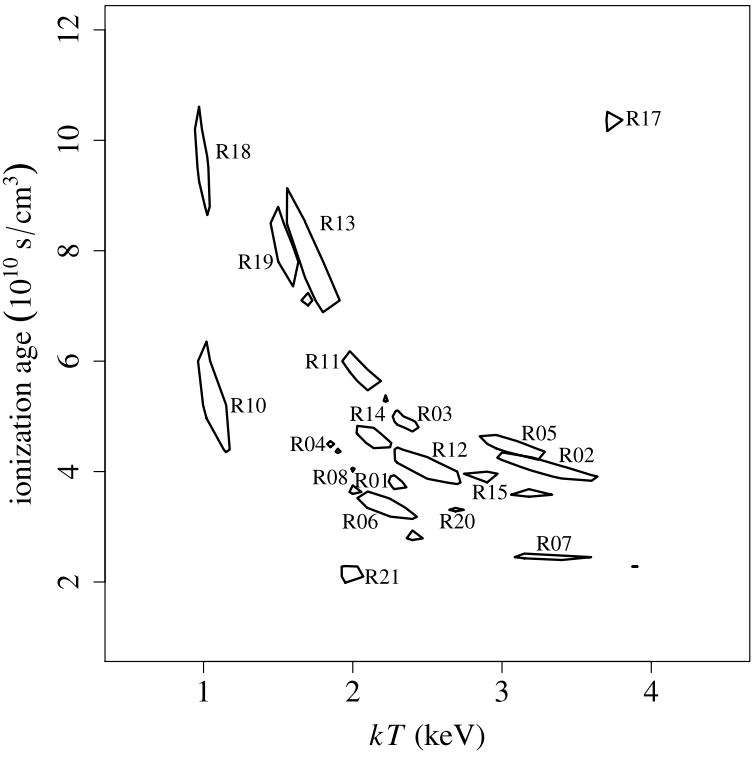

Figure 2-4: The knots exhibit a variety of plasma states. 68% confidence contours in the kT–τ plane for knots R01–R20 in the 2004.4 epoch display systematically larger values than Paper A. A large anti-correlation between kT and τ can be seen. This correlation widens the confidence intervals, compared to the 1D confidence interval obtained from ISIS or XSPEC, and it highlights the pitfall of reporting 1D confidence levels without accounting for correlations in the fit parameters.

be modeled with atomic transition codes to infer the temperature of the emitting plasma, while the ACIS analysis produces a temperature which is interpreted within the `vnei` model. We note that emission models in such atomic transition codes assume CIE, while our ACIS model assumes NEI.

The ionization age, τ, provides a measure of the stage of evolution of the plasma from NEI to CIE. Since at large ionization age an NEI plasma will come to equilibrium and asymptote to CIE Smith and Hughes [123], we would expect good correlation between the line-ratio-derived temperature and the `vnei`-derived temperature for plasmas with large τ.

Figure 2-5 shows a comparison of the emission line-derived and model-derived temperatures. The y-axis shows the average f/r value of the two HETG observations, and the x-axis

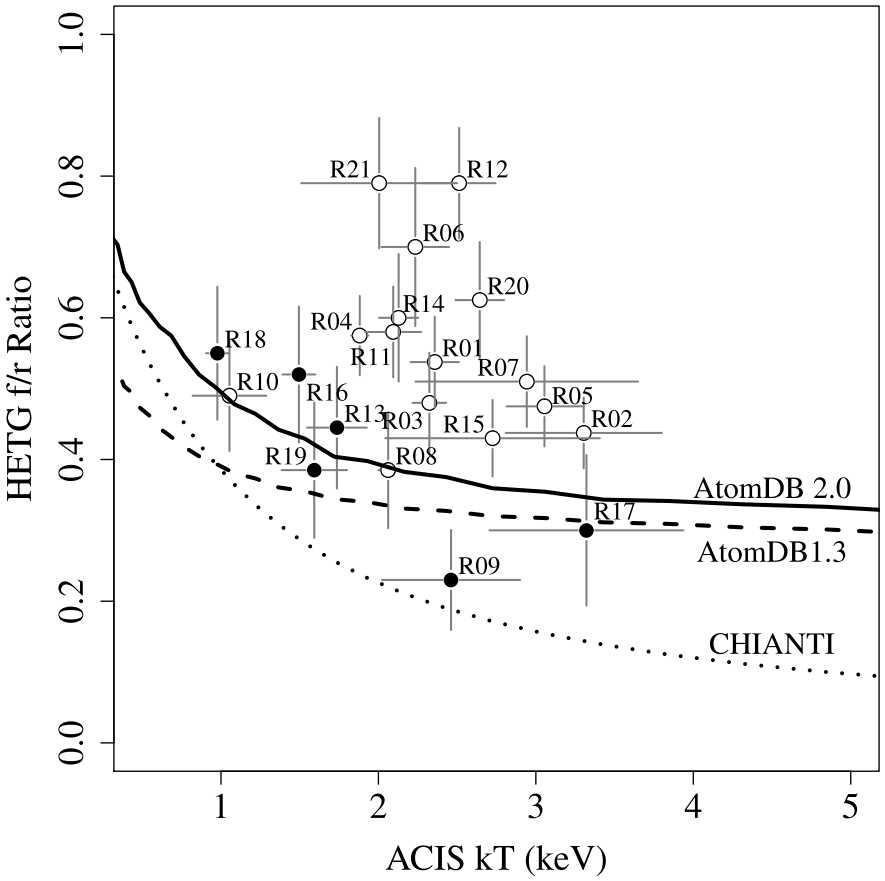

Figure 2-5: Modern plasma emission codes (curves) cannot fully describe the ACIS-derived temperatures (x-axis) and HETG-derived f/r ratios (y-axis) of the 21 ejecta knots. The 'NEI outliers' (white dots), which differ most significantly from the plasma emission codes, have systematically lower ionization ages ($\tau < 6 \times 10^{10}$ s cm^{-3}). These knots therefore break the collisional ionization equilibrium assumption implicit in the plasma database calculations. 68% confidence and 1-σ error bars are shown for the knots along the x- and y-axes, respectively. The knots are identified by the knot number from Figure 2-1. The *Chandra* measurements are compared against the CHIANTI 6.0 database (dotted), AtomDB 1.3.1 (dashed), and AtomDB 2.0.0 (solid).

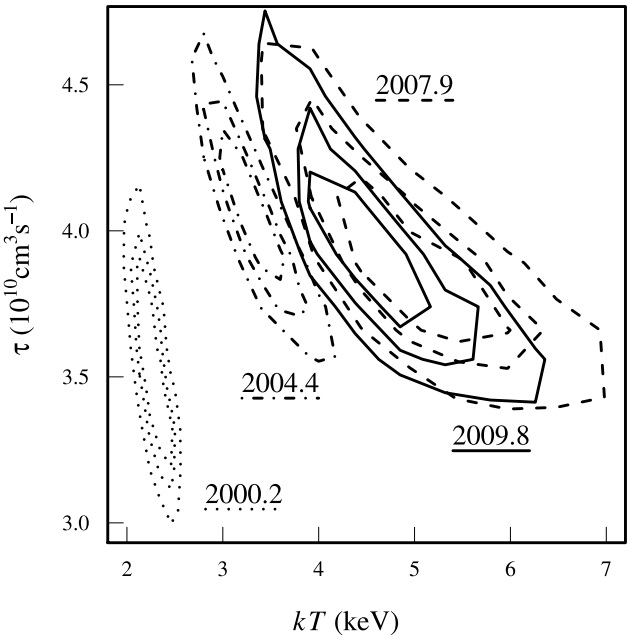

Figure 2-6: R02 shows distinct signs of evolving temperature.The confidence contours for the four epochs do not overlap in kT-τ space, indicating that the ACIS parameters have evolved. The three rings for each epoch indicate 68%, 95%, and 99% confidence levels.

shows the temperature for the 2004.4 ACIS observation. As most knots did not show much evolution, the average and the mid-decade values are appropriate.

Nearly half of the knots are consistent with the latest plasma emission code, AtomDB 2.0 [37]. Two other plasma emission codes, an earlier version of AtomDB and the latest CHIANTI database [19], predict lower f/r ratios over the same temperature range. All databases are evaluated in the low density limit. The new calculations of excitation and recombination rates have brought the new AtomDB into greater agreement with these knot observations.

The 'NEI outliers' above the curve exhibit too much forbidden Si XIII line radiation for their derived temperatures, by a factor of two in some cases. These outliers exhibit low, tightly grouped ionization ages. Figure 2-5 shows the two different distributions for the 2004.4 epoch maximum likelihood values. The outlier distribution is clustered around an ionization age of $\tau = 5 \times 10^{10}\,\mathrm{cm}^{-3}\,\mathrm{s}$, indicating a plasma fairly far from CIE [123]. These knots have been caught early in the evolution process, so they deviate most from the

AtomDB ionization equilibrium calculations. The knots closer to the AtomDB 2.0 line vary more widely in derived τ values, though some show equally low values as the NEI outliers. The populations do not differentiate as clearly in kT, H/He, or f/r.

The results in Figure 2-5 also show that the f/r ratio is a poor predictor of temperature for these low density, evolving plasmas. Even the theoretical curves for AtomDB 1.3 and 2.0 flatten and lose their predictive power above 2 keV.

2.3.3 Plasma Evolution

On the whole, the selected knots show few signs of evolution, either in the derived plasma model parameters or the dispersed Si line ratios.

The HETG results (Figure 2-11 and Table 2.6) yielded mixed trends. Six knots showed increasing H-like/He-like Si ratios (the plasmas ionized), while the ratio decreased significantly for only one knot (the plasma recombined). The f/r ratios for most knots are consistent between the two epochs, within the errors.

The ACIS results (Figures 2-12 – 2-14 and Tables 2.7 and 2.8) revealed only five knots with evolving (i.e., non-overlapping) kT–τ contours. All evolving knots trended to higher temperatures over the decade of observation. Three knots in this evolving group showed evidence of aging (growing τ), while one in fact cut its ionization age by a factor of four. As an example, Figure 2-6 exhibits distinctly non-overlapping confidence contours, which we identify as knot evolution.

Evolution Does Not Correlate with Radius

It is natural to ask if the few evolving knots are closer to the reverse shock than the others. To answer this question, we need first to locate the knots within the 3D structure of Cas A, then we must define a quantitative measure of evolution.

In order to place the knots in 3D, the HETG measurements are essential: both the precise velocity measurements and the proper motions are needed to transform to physical coordinates with the appropriate scaling. We apply the prescription of DeLaney et al. [18]

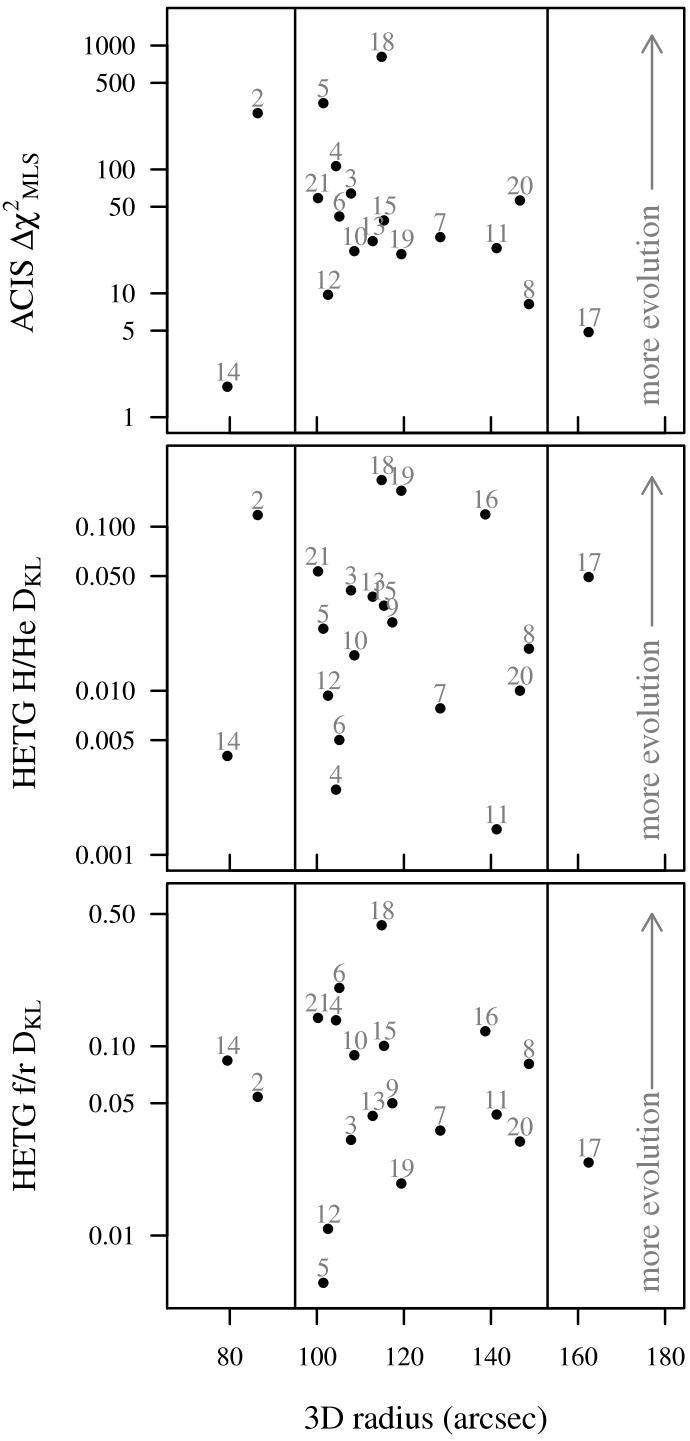

Figure 2-7: The evolution of a knot does not correlate with its distance to the reverse shock. As a function of derived 3D radius (the x–axis), we plot three measures of evolution (the y–axis of the three separate plots). For the HETG parameters, we apply the K-L divergence, and for the ACIS values we employ a maximum likelihood swap. In both cases, higher values indicate more evolution. The reverse and forward shock (black vertical lines) are shown for comparison.

Region	Doppler velocity (km s^{-1})	Radial proper motion (milliarcseconds yr^{-1})
R01	-2340	38
R02	-1360	221
R03	0	167
R04	-150	134
R05	-1450	210
R06	-1110	192
R07	3780	125
R08	2740	140
R09	-540	146
R10	1670	213
R11	750	66
R12	-90	265
R13	-810	249
R14	4300	279
R15	1070	279
R16	-1360	313
R17	-1710	307
R18	-850	112
R19	-880	166
R20	-2300	109
R21	680	229

Table 2.4: Doppler velocities and average radial proper motions of the knots.

to effect the transformation from velocity space to position space, refining their technique to use individual knot expansion rates instead of the average. The results are presented in Table 2.4.

Evolution is quantified differently for the HETG and ACIS results, due to the different statistics employed. Since the error bars in the HETG analysis are interpreted as widths of probability density functions, the Kullback-Leibler divergence (D_{KL}) can be used [80]. The KL divergence measures the level of similarity between two distributions in an information theoretic manner. In Figure 2-7, higher D_{KL} values indicate that the probability distributions of the 2010.3 parameters differ more from those in 2001.4. The ACIS confidence contours are not probability distributions, so we use a different measure of evolution: what we call a 'maximum likelihood swap' (MLS). We substitute the best fit values of the 2009.8 epoch for the 2000.2 parameters, then evaluate the χ^2. The resulting difference, $\Delta\chi^2_{MLS}$, corresponds

to the probability of the true 2000.2 parameters being as extreme as the 2009.8 ones.

Our study does not reveal any correlation between evolution and distance from the reverse shock, which is shown in Figure 2-7. No correlations with norm or τ were detected, either (not shown). A complicated shock morphology may underlie this observed lack of correspondence. That is, any smaller scale secondary shocks could prematurely age some areas. The 3D morphology findings of DeLaney et al. [18] reveal many asymmetries, for instance, which may affect shock propagation.

2.3.4 ACIS and HETG Cross-check

The results of the two different analyses are consistent with each other. The ACIS spectral response cannot resolve the He-like triplet lines, but it can separate the Si XIII triplet (as a single blend) from the Si XIV Ly-α. As a cross-check of the ACIS and HETG results, we looked at the measured Si XIV / Si XIII ratio from HETG, and the predicted ratio from the ACIS `vnei` model (Figure 2-8). The results, in Figure 2-9, show consistency for most knots within 30%. The effective area uncertainties for the *Chandra* instruments are on order of 10%.

2.4 Discussion

2.4.1 The Lack of Measured Evolution in the Knots' Plasma States

The new HETG observation was intended to catch the knots during evolution, motivated by the high electron density of the knots inferred from Paper A. Combining densities on the order of 100 cm^{-3} with a baseline of $\sim$9 years predicts a change of $\Delta\tau \approx 2.8 \times 10^{10}\,\mathrm{cm}^{-3}\,\mathrm{s}$, which would produce observable changes in the plasma through its line ratios and τ values. As demonstrated above, neither the `vnei`-modeled temperatures and ionization ages nor the HETG line ratios show changes of this magnitude. (Cf. Figure 2-6, which shows a spread

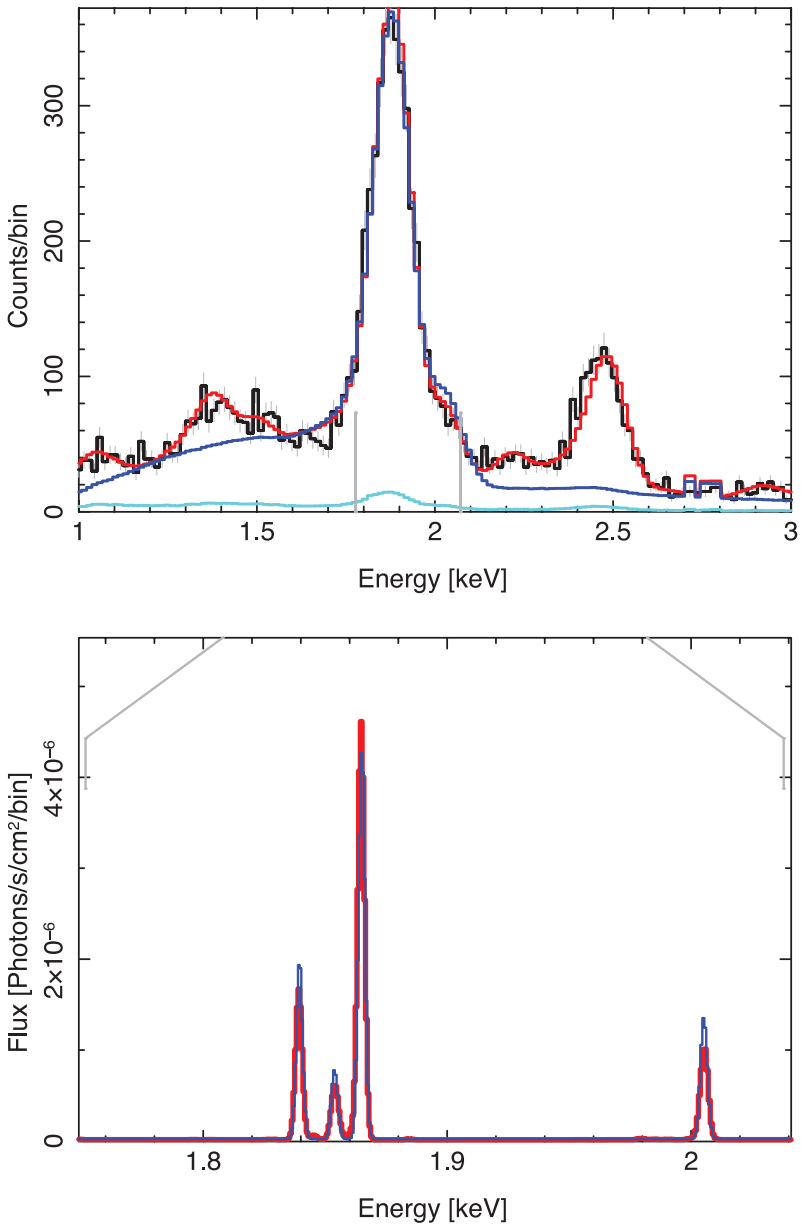

Figure 2-8: The measured HETG H/He-like ratios for the Si line agree well with the prediction from the best fit `vnei` model from the ACIS analysis. Top: the ACIS data (black), `vnei` model (red), and HETG 4-gaussian model, scaled, folded through the ACIS response, and with a bremsstrahlung component added (blue) agree in the Si region for R13. (ACIS background model is in light blue.) Bottom: a zoom-in of the underlying Si lines shows remarkable consistency between the HETG data (black), HETG model (blue), and scaled `vnei` model (red). The lines are, left to right: Si XIII f, Si XIII i, Si XIII r, and Si XIV.

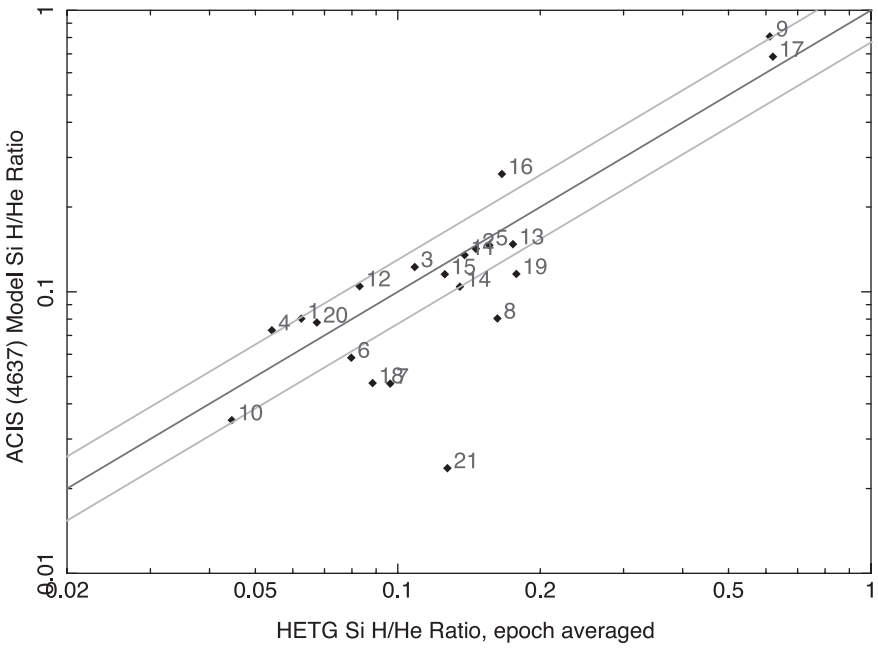

Figure 2-9: The quantified comparison of Figure 2-8 shows good agreement between the HETG and ACIS models for most knots. The 30% assumed systematic errors in the HETG measurement are shown with the top and bottom solid lines.

of only $1.5 \times 10^{10}\,\mathrm{cm}^{-3}\,\mathrm{s}$. For even the low resolution ACIS spectral model, forcing τ to a higher value results in obvious disagreement with the data, due to an increased Si XIV/XIII ratio.) We must, therefore, look into the various assumptions in our modeling to explain this lack of evolution.

Explaining the Non-evolution with Chemical Makeup

As we will show, under our original model assumptions, the most optimistically low estimates for the knots' electron densities cannot explain their lack of spectral evolution. However, when we relax these assumptions about the elemental makeup of the knot, we find a physical state that could explain the observed spectral stasis.

We can turn our limits on $\Delta\tau$ into an upper limit on the electron density, since the plasma evolution depends on n_e, for a given time baseline. A very generous upper limit is given by considering a τ change given by the difference between the epoch 2009.8 90% upper-limit τ

Region	Volume (10^{50} cm^3)	τ (2004.4) (10^{10} s cm^{-3})	n_e (90% limits) (cm^{-3})	n_e (nominal) (cm^{-3})	n_e (CN-rich) (cm^{-3})	n_e (Si-rich) (cm^{-3})
R01	21.7	3.81	35.0	286.6	118.1	15.5
R02	12.4	4.05	51.2	280.9	116.0	17.2
R03	15.1	4.97	21.8	338.0	139.1	16.1
R04	102.3	4.60	50.8	239.8	98.7	11.7
R05	5.8	4.53	57.1	413.1	169.9	19.0
R06	7.7	3.49	65.7	263.1	108.5	15.1
R07	24.5	2.93	30.7	155.8	64.8	12.8
R08	7.5	4.03	45.2	248.4	102.8	17.1
R09	7.2	20.56	1105.	360.1	149.3	26.7
R10	9.6	3.58	181.	191.5	79.3	13.3
R11	17.7	5.79	46.2	274.3	113.0	14.5
R12	14.1	4.02	54.8	260.8	107.4	13.7
R13	5.2	7.39	609.	323.6	133.6	19.7
R14	9.1	4.73	57.8	313.1	129.0	16.7
R15	7.0	3.73	211.	244.0	100.7	14.8
R16	12.7	12.55	173.	267.9	110.4	14.2
R17	16.2	10.37	333.	107.1	45.0	11.2
R18	4.3	10.20	-499*	615.4	253.5	32.4
R19	6.1	8.59	903.	327.3	135.1	19.7
R20	60.1	3.31	28.7	193.8	80.0	11.4
R21	8.1	2.40	38.6	105.5	43.6	7.2

Table 2.5: Electron densities derived under various assumptions of the knot components, discussed in Section 2.4.1. (*) R18 underwent 'recombination', with its τ values decreasing over the decade. Therefore, the simple aging model can provide no information about the density.

value and the epoch 2000.2 90% lower-limit value. Dividing this $\Delta\tau$ by the time baseline gives the 'n_e (90% limits)' values in Table 2.5, which are in the 10's to 100's of cm^{-3}.

The conversion of `vnei` parameters to plasma density is straightforward, although not typically implemented in the modeling software. Paper A outlines the technique (and we detail the calculation in Appendix 2.C): approximate the average number of electrons stripped per ion, then convert that to an absolute density through factors of the `vnei` abundances and emission norm, by approximating the knot as a uniform density spheroid. The 2004.4 epoch model parameters are used to produce the densities 'n_e (nominal)' given in Table 2.5. The values are too high above the upper limit densities.

The straightforward calculations of n_e yielded values too high, so we must find out how to reduce the density while maintaining spectral consistency. Our derivation of the electron density leaves us with only a few slightly adjustable parameters to tune:

$$n_e \propto \sqrt{\frac{X_H \times n_e/n_H}{f}},$$

namely X_H – the `vnei` abundance of hydrogen relative to solar, n_e/n_H – the number of electrons per hydrogen ion in the plasma, and f – the filling fraction of the material in the emitting volume. (The uncertainties of the distance measurements to Cas A translate to only a 4% error in the density, so we neglect that parameter.) We can start by changing the first two parameters with a different chemical makeup for the knot.

The low-Z elements provide the continuum, but the model cannot distinguish between their contributions. During fitting, therefore, we froze C, N, and O to 5 times solar (Appendix 2.C), and H and He to 1. We can investigate alternate abundance sets of low-Z elements that produce the same continuum, yet lead to different values of n_e. For instance, higher Z elements have more electrons to give up, so can compensate for a lower abundance of H.

As an example of this continuum degeneracy, we consider a predominantly C and N makeup. (Cas A exhibits little emission from C or N, but we care about the lack of H and He in this straw man model.) Changing the {H, He, C, N, O} abundances from {1, 1, 5,

5, 5} to a C- and N-rich plasma with {0.002, 0.20, 45, 45, 5} will produce a spectral model with the same norm, but a reduced n_e. (We increase C and N, but not O, based on the results of Dewey et al. [22], though a N-rich plasma could do the same job.) Tweaking the abundances this way increases n_e/n_H by about 100, partially undercutting the benefit of a lower X_H. The resulting n_e values ('n_e (CN-rich)' column in Table 2.5) are reduced by over a factor of 2, but remain comfortably above the upper limits. This chemical makeup cannot explain the lack of evolution.

We discard the CN-rich model, but the above example of an enriched metal plasma motivates an extreme scenario: a plasma constituted only of elements with $Z > 8$. Such a composition places a lower limit on the electron density that can still reproduce the observed line fluxes. Densities calculated with each knot's individual best-fit abundances – but with $Z \leq 8 = 0$ – fall below the estimated '90% limits' densities (Table 2.5, 'Si-rich'), in the range of 10–20 cm^{-3}. The lack of evolution for many knots could therefore be explained by a low density, Si-rich plasma component.

To make this metal-rich knot model more physical – and spectrally consistent – we need only split our nominal knot into two components. For both components, we use the same parameters as the nominal knot: the 2004.4 maximum likelihood kT, τ, and abundances (with the fiducial continuum set {H, He, C, N, O} = {1, 1, 5, 5, 5}). The metal-rich component is given an 85% fill fraction, and the elements with $Z \leq 8$ are assigned contamination-level abundances of 10^{-2} solar. Complementarily, the low-Z component fills 15% of the knot to bolster the continuum, and has $Z > 8$ abundances set to contamination level.

The two models with complementary abundance sets recover the same well fitting spectral model as the nominal case (Figure 2-10). Unlike the CN-rich case, the altered chemical composition results in slightly higher n_e/n_H, while X_H is cut by a factor of 100. The resulting n_e's for the metal-rich components – which contain the Si we have been using as a tracker of evolution – are below the upper limit values (Table 2.5), consistent with the non-evolution we observe.

Evidence for Metal Dissociation

A model that dissociates high-Z metals from the other elements can produce low n_e, thereby explaining the lack of evolution. Though we found no other simple solution to regain self-consistency, we acknowledge the model presented here may not be the only answer.

This dissociated plasma is made up of low density pure ejecta and a low-Z component. The low-Z ($Z \leq 8$) plasma provides the continuum emission we see and has higher density: the n_e's range from 300–1600 cm^{-3}, which correspond to radiative cooling times greater than 1900 years. This low-Z plasma would be physically separated from a less dense, ejecta-rich ($Z > 8$) plasma that accounts for the observed line emission. These plasmas share the same kT and τ, by construction, so the dual model sums to single-component `vnei` model, reproducing the good spectral fits (Figure 2-10). Only the interpretation of the physical picture is changed: the components are unmixed.

The data in this work cannot constrain the specific physical model for knot dissociation, but hydrodynamical simulations offer indications. We envision a low density parcel of pure high-Z ejecta pushing into a denser layer of low-Z material ahead of it, much like the tip of a Rayleigh-Taylor (RT) finger. The high-Z component, though, could manifest as an expanding clump implanted into the low-Z plasma, or be turbulently mixed on finer scales.

The scales for this dissociation would be small, if the components are not separated along the line of sight. [87] found evidence for well mixed ejecta in Cas A, down to 3″ ($\sim 10^{17}$ cm). Likewise, our knot-defining algorithm did not see large spectral differences at the pixel scale. Investigating infrared knots, Isensee et al. [67] found a larger scale (0.1 pc, 5″) line-of-sight separation between the O and Si layers in one 40″ by 40″ region of Cas A. These knots are more recently shocked than the X-ray knots, and do not admit direct comparison. Our proposed model would require dissociation as a line-of-sight effect, or at scales unresolved by the *Chandra* optics, which have a PSF matched to the pixel size, 0.5″ ($\sim 2.5 \times 10^{16}$ cm).

Early-times explosion models consistently show large scale RT mixing between burning layers, though just barely resolve the small scales of our knots. Hammer et al. [51] found strong mixing on large scales, with high-Z bullets probing deep into the outer hydrogen

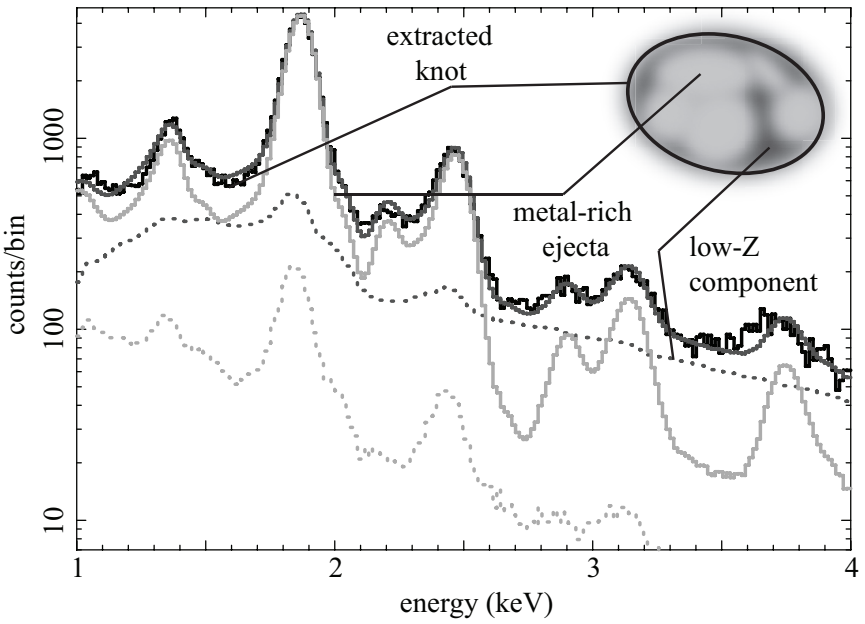

Figure 2-10: The simple model of a knot split into two spectral components – high metallicity ejecta expanding into a denser low metallicity plasma – reproduces the observed spectrum. The observed knot spectrum (black) is a sum (dark grey) of the low metallicity component (dotted dark grey) and the high metallicity component (light grey). The background spectrum (dotted light grey) does not resemble the low metallicity component. Metal dissociation within a knot could explain why so little evolution was observed over the decade. The cartoon is only a schematic; the data and model cannot distinguish the actual morphology.

layer. The smallest structures at 2.5 hours span 10^{11} cm, roughly the size of these X-ray knots at that time. The work of Joggerst et al. [74] found that mixing can persist longer ($\sim$ 1 day after the explosion). The rapid expansion of the remnant may leave elemental mixing incomplete, though to date there have been no studies to determine on what scale (private communication, H. Janka, 2012).

In the late-stage mixing simulations, several enhancements to the RT mixing have been suggested to force the fingers through the whole intershock region, where our Si-rich knots lie. RT instabilities could reach to the forward shock with the help of high enough compression ratios [6], large, underdense Fe bubbles colliding with the reverse shock [7] or high density contrasts in the originating ejecta clumps [99]. Orlando et al. observe density contrasts of 3 to 4 orders of magnitude over 0.5 pc (for a 1000 yr old simulated remnant), while we claim

contrasts of just over 1 order of magnitude, spanning 10^{-3} pc in Cas A.

The recent work of Ellinger et al. [29] aims to bridge the early and late epochs. Their simulations have evolved Cas A to 30 years, but have revealed little mixing between the Si and O layers, in contrast with Hammer et al. [51] and Joggerst et al. [74], who found more Si in the RT fingers. We await their extended simulations to compare with our dissociated metal hypothesis.

It is worth noting that our dissociated model remains simple, with a shared temperature and ionization age. If the ejecta are not fully mixed at these small scales, we would expect the two components to have different temperatures and ionization ages, though we achieve good spectral agreement without turning these additional knobs. This additional freedom could reconcile our differences with Paper A's kT and τ values. (Initial investigations show that, if freed, kT and τ for the high-Z component would not change much, while the dissociated model permits a much wider range of low-Z τ values.)

If this dissociation is the answer to the non-evolution riddle, it could inform elemental mixing during supernova explosions.

Flux Evolution

The brightened knot, R21, which we added to the set from Paper A, was also investigated in [101]. They too claimed spectral evolution in R21, with kT increasing from 1.1 to 1.5 keV from 2002 to 2004, and τ increasing from 4.6×10^{10} to 53.0×10^{10} s cm^{-3}. From 2000 to 2004, our results show a similar increase in temperature, but τ is not only lower ($\sim 2 \times 10^{10}$ s cm^{-3}), but fairly constant over that period. The disparities may arise from different model assumptions: both analyses use an NEI model with variable abundances, but [101] fit the spectrum over a smaller energy range, did not tie parameters across epochs, and derived only single-parameter confidence intervals.

In general, for many of the knots, the brightness of the Si lines changed dramatically – up to a factor of 3 – over the decade, as Figure 2-11 shows. However, a large-area, broadband investigation of Cas A by [102] showed only the slightest decrease in 1.5 – 3.0 keV flux in the

same period. Evidently, the uncorrelated variations in the individual knots average out over time and position, giving the impression of steady thermal emission that Patnaude et al. observe.

2.4.2 Model Degeneracies

We have found that the `vnei` parameters kT and τ are often anti-correlated. A higher kT and lower τ can produce a spectrum with equal statistical likelihood to one with lower kT and higher τ. The characteristic of this degeneracy is a banana-shaped confidence contour, such as in Figure 2-6.

There is a physical justification for the strong anti-correlation between kT and τ. The equilibrium ratio of the Si H-like Ly-α to the He-like triplet for a plasma is monotonically increasing with temperature. As an NEI plasma evolves toward CIE, this ratio also monotonically increases with τ. Thus, for a given spectrum, models with higher kT and lower τ or models with lower kT and higher τ can result in the same likelihood.

This degeneracy explains why the discrepancies in the kT and τ values derived in Paper A and in this analysis are anti-correlated.

Had we not used 2-dimensional confidence contours for the 2 `vnei` parameters of interest, we would have reached different conclusions about the data, due to problems that arise with the use of 1-dimensional confidence intervals for multidimensional data. First, during the interval generation for one parameter, the other parameters of interest are treated as nuisance parameters, which is wrong from a mathematical perspective (the incorrect χ^2 distribution is used). Second, the use of a χ^2 distribution with fewer degrees of freedom underestimates the parameter uncertainty. Moreover, any underlying correlations are lost, leaving the analyst susceptible to false trends in best fit values, when in fact those points may live in the same 'banana'.

2.5 Conclusions

We have performed a detailed analysis of two HETG and four ACIS observations of Cas A spanning a 10-year baseline. A new set of detailed plasma kinematic and temperatures measurements has been presented. Good agreement was found for the HETG– and ACIS–derived temperatures for about half of the knots. The outliers are well described by a plasma far from collisional ionization equilibrium, the state which is assumed for the HETG line-ratio temperature estimates.

The low τ values and the high electron densities derived from this analysis predict a significant amount of plasma evolution in the 10-year baseline, in stark contrast with the observations. We propose a physical model of two plasmas – one high metallicity, one low metallicity – that remain unmixed at small spatial scales. The Si emission comes from the pure heavy metal ejecta (Ne through Fe), while the continuum is provided by the low density component, rich in the elements C, N, and O. This model fits the data well and explains the lack of evolution, since it requires a much lower electron density for the Si plasma and, therefore, a longer timescale for evolution. If validated, this model could place strong constraints on turbulence in supernova explosion models.

2.A Appendix A: Results Tables and Figures

The tables and figures in the following pages present the results of our kinematic and thermodynamic analysis of 21 knots in Cas A.

Table 2.6 lists the values extracted from the high-energy-resolution HETG analysis. We extract the velocity of each knot, two line ratios, and the photon flux in the fitted line complex. The two line ratios are the ratio of the Si He-like triplet forbidden and recombination lines (f/r), which can only be obtained through the high-energy-resolution analysis, and the Si H-like to He-like ratio, which can be compared to the ACIS model fit value (shown in Figure 2-9).

Region –Epoch	Diam. (")	Velocity ($\mathrm{km\,s^{-1}}$)	Si XIII (f/r)	Si XIV (H/He ratio)	Flux
R01–I	3.00	-2320 ± 90	0.56 ± 0.12	0.040 ± 0.025	0.77
" –II	4.80	-2360 ± 100	0.52 ± 0.05	0.075 ± 0.020	0.14
R02–I	4.80	-1430 ± 90	0.48 ± 0.06	0.095 ± 0.020	1.1
" –II	5.04	-1300 ± 70	0.40 ± 0.08	0.165 ± 0.030	0.99
R03–I	3.72	$+30 \pm 100$	0.44 ± 0.10	0.120 ± 0.030	0.64
" –II	4.20	-30 ± 120	0.52 ± 0.10	0.080 ± 0.025	0.82
R04–I	5.76	-290 ± 100	0.53 ± 0.05	0.055 ± 0.020	1.4
" –II	5.04	-20 ± 100	0.62 ± 0.10	0.045 ± 0.020	0.90
R05–I	3.48	-1490 ± 100	0.46 ± 0.08	0.130 ± 0.030	0.80
" –II	3.36	-1420 ± 140	0.49 ± 0.08	0.150 ± 0.040	0.84
R06–I	3.60	-1200 ± 140	0.75 ± 0.20	0.070 ± 0.040	0.42
" –II	3.60	-1030 ± 140	0.65 ± 0.10	0.075 ± 0.035	0.39
R07–I	6.00	$+3920 \pm 100$	0.54 ± 0.10	0.085 ± 0.025	0.54
" –II	5.04	$+3650 \pm 100$	0.48 ± 0.08	0.085 ± 0.030	0.68
R08–I	3.36	$+2810 \pm 200$	0.32 ± 0.10	0.125 ± 0.050	0.35
" –II	3.60	$+2680 \pm 200$	0.45 ± 0.13	0.160 ± 0.060	0.32
R09–I	3.84	-490 ± 140	0.28 ± 0.10	0.560 ± 0.090	0.61
" –II	3.24	-610 ± 150	0.18 ± 0.10	0.500 ± 0.080	0.33
R10–I	4.80	$+1760 \pm 350$	0.56 ± 0.10	0.035 ± 0.034	0.22
" –II	4.80	$+1600 \pm 200$	0.42 ± 0.12	0.045 ± 0.044	0.18
R11–I	4.92	$+760 \pm 120$	0.54 ± 0.08	0.125 ± 0.035	0.58
" –II	5.04	$+760 \pm 120$	0.62 ± 0.10	0.115 ± 0.035	0.54

Table 2.6: Measured Si line parameters for 21 Cas A knots from spatial-spectral modeling of the HETG dispersed data. These ratios and fluxes are 'un-absorbed' based on the N_H determined by ACIS fitting. The last column is the un-absorbed flux is the sum of the four Si lines in units of 10^{-3} ph $\mathrm{s^{-1}}$ $\mathrm{cm^{-2}}$. (Continued on next page.)

R12–I	3.48	-90 ± 200	0.80 ± 0.12	0.070 ± 0.030	0.31
” –II	5.88	-90 ± 120	0.78 ± 0.10	0.075 ± 0.025	0.79
R13–I	3.60	-740 ± 120	0.41 ± 0.10	0.135 ± 0.040	0.40
” –II	4.20	-890 ± 200	0.48 ± 0.14	0.175 ± 0.055	0.43
R14–I	4.32	$+4300 \pm 250$	0.54 ± 0.10	0.130 ± 0.050	0.32
” –II	4.08	$+4320 \pm 120$	0.66 ± 0.15	0.110 ± 0.050	0.28
R15–I	3.72	$+1200 \pm 140$	0.47 ± 0.09	0.095 ± 0.035	0.43
” –II	4.20	$+960 \pm 120$	0.39 ± 0.06	0.135 ± 0.030	0.52
R16–I	4.56	-1710 ± 200	0.60 ± 0.15	0.130 ± 0.040	0.29
” –II	3.36	-1020 ± 250	0.44 ± 0.12	0.190 ± 0.080	0.16
R17–I	3.60	-2110 ± 300	0.26 ± 0.14	0.600 ± 0.150	0.11
” –II	4.20	-1310 ± 250	0.34 ± 0.16	0.580 ± 0.100	0.095
R18–I	3.00	-1100 ± 350	0.65 ± 0.17	0.045 ± 0.044	0.18
” –II	3.60	-610 ± 120	0.45 ± 0.08	0.140 ± 0.030	0.50
R19–I	3.60	-530 ± 240	0.35 ± 0.13	0.115 ± 0.040	0.33
” –II	3.60	-1240 ± 220	0.42 ± 0.14	0.220 ± 0.080	0.25
R20–I	5.40	-2300 ± 80	0.65 ± 0.13	0.050 ± 0.020	1.7
” –II	6.60	-2310 ± 100	0.60 ± 0.10	0.070 ± 0.020	1.6
R21–I	3.36	$+780 \pm 150$	0.88 ± 0.14	0.130 ± 0.050	0.088
” –II	3.36	$+590 \pm 100$	0.70 ± 0.12	0.095 ± 0.035	0.26

Table 2.6: Continued.

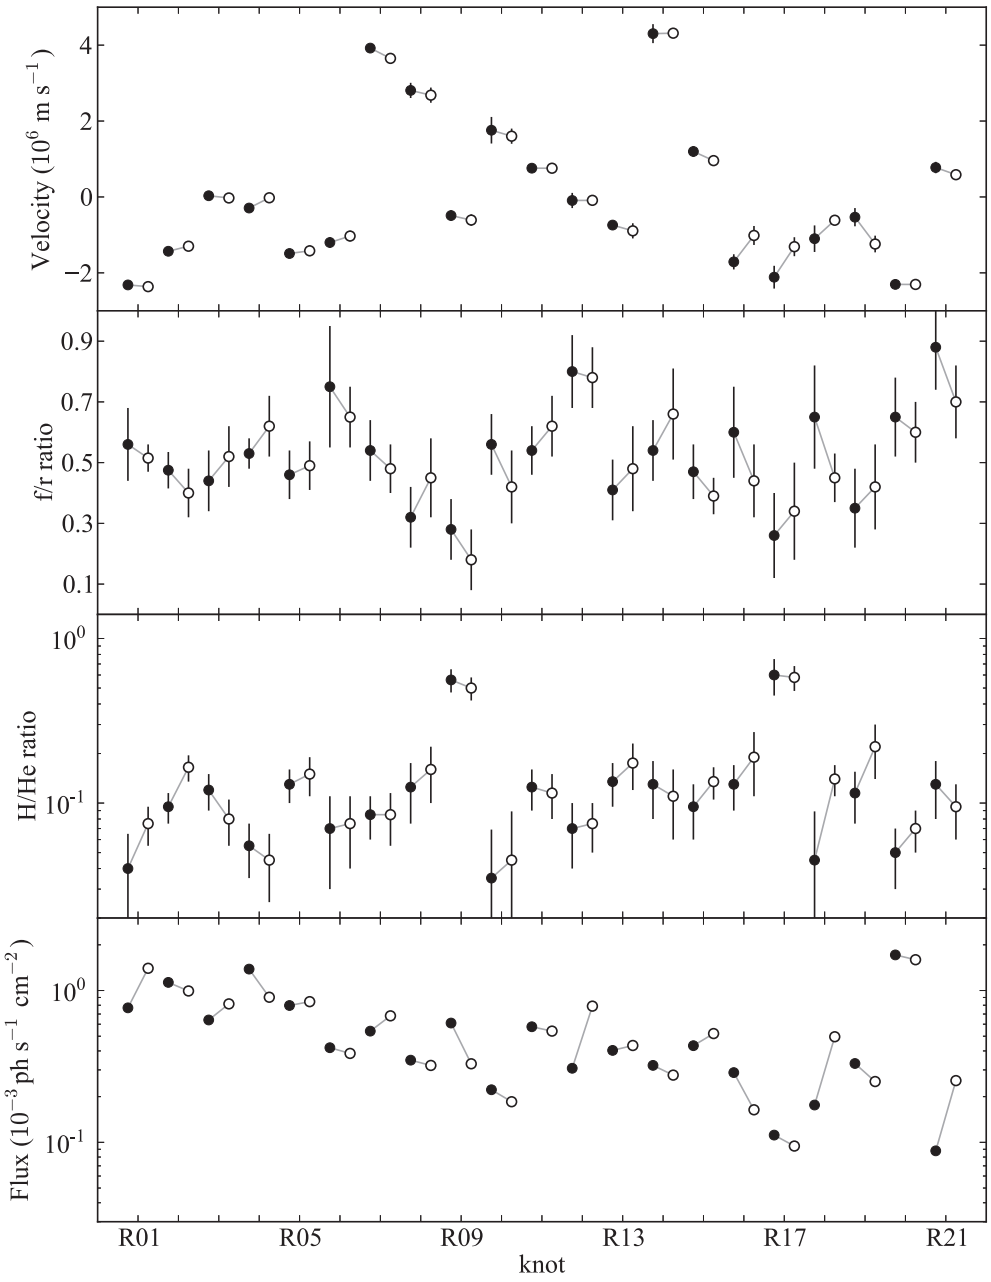

Figure 2-11: The HETG measurements did not show much change over the decade, besides the norm values. The left-most point for each knot (black) is the result for the first epoch (2001) and the right-most point (white) the result for the second epoch (2010). The 'f/r' label refers to the forbidden/recombination line ratio for Si XIII. The 'H/He' label refers to the (Si Ly-α)/(Si He-like triplet) line ratio. The norm errors were not computed.

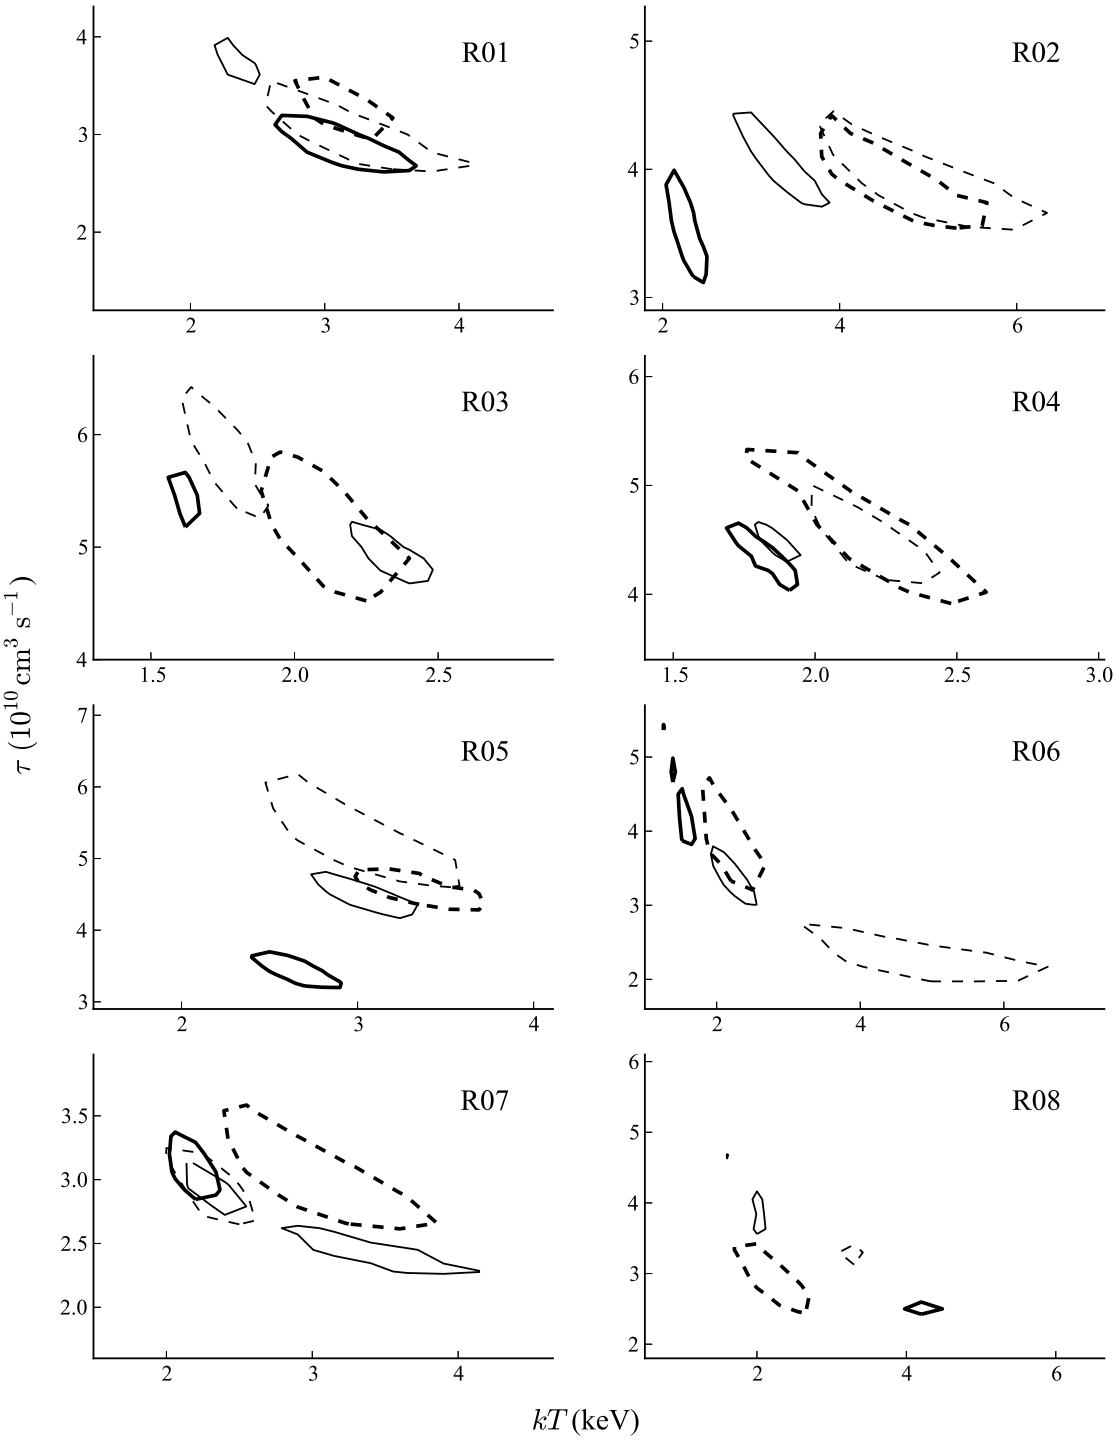

Figure 2-12: The ACIS data likewise showed little evolution. The 95% confidence contours in the kT–τ plane are shown for each epoch: 2000.2 (thick), 2004.4 (thin), 2007.9 (thin dashed), and 2009.8 (thick dashed). The joint fits between all four epochs are described in Appendix 2.C. Note the strong correlation between the two parameters.

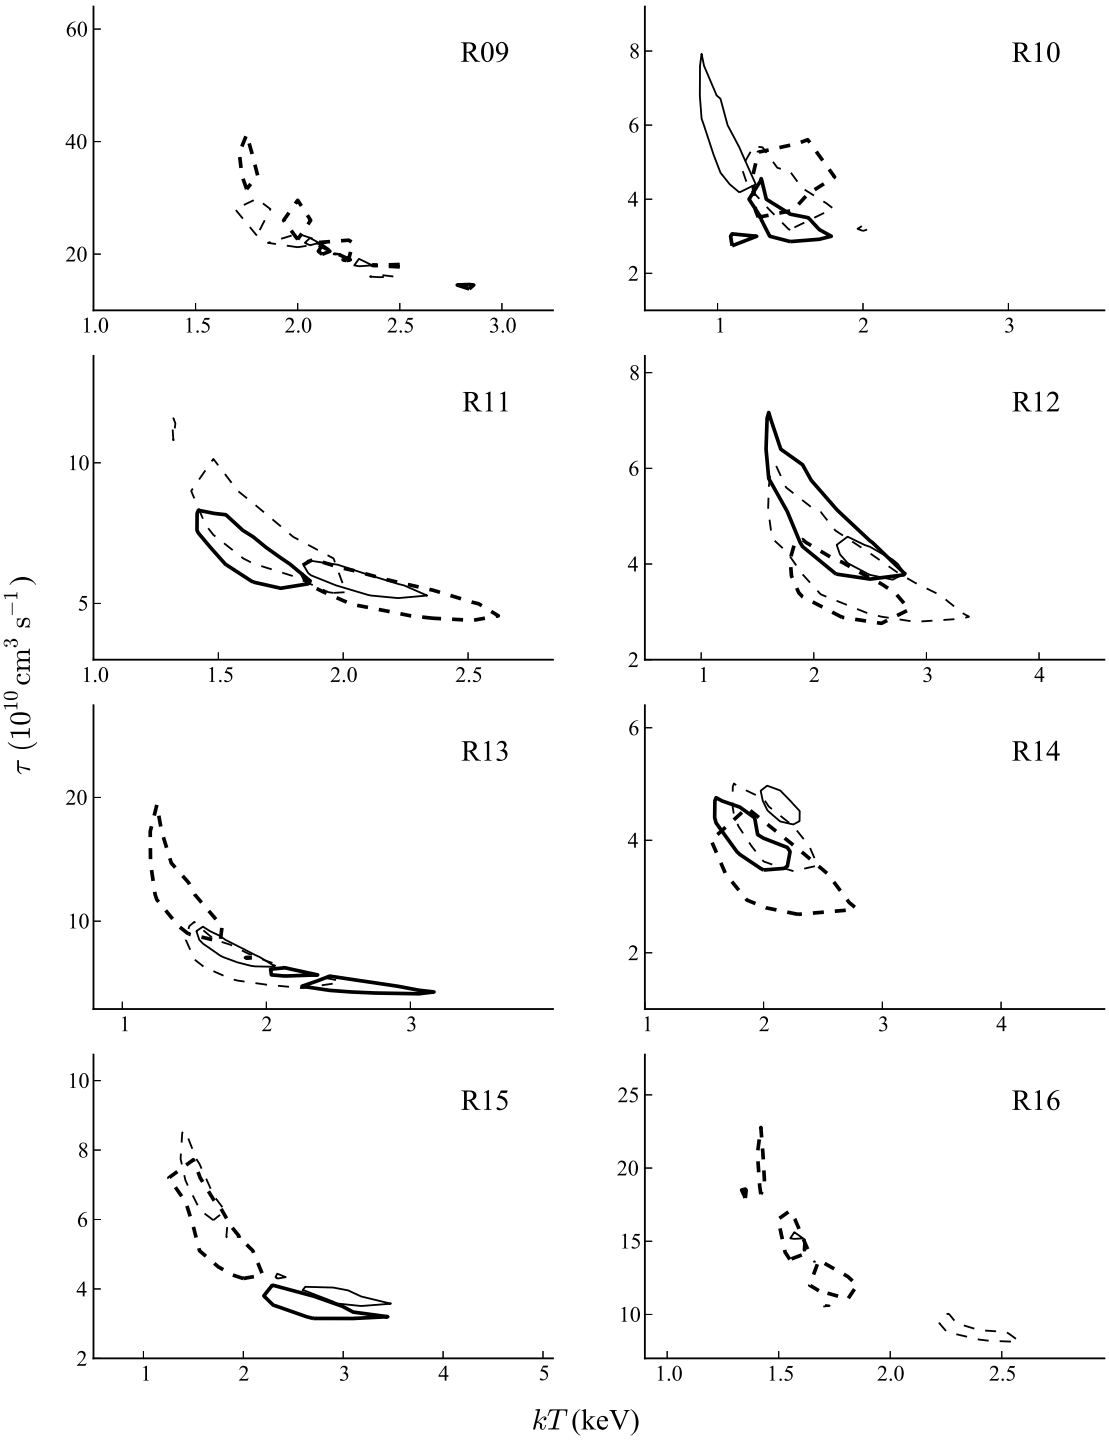

Figure 2-13: Continuation of Figure 2-12. The 95% confidence contours in the kT–τ plane are shown for each epoch: 2000.2 (thick), 2004.4 (thin), 2007.9 (thin dashed), and 2009.8 (thick dashed).

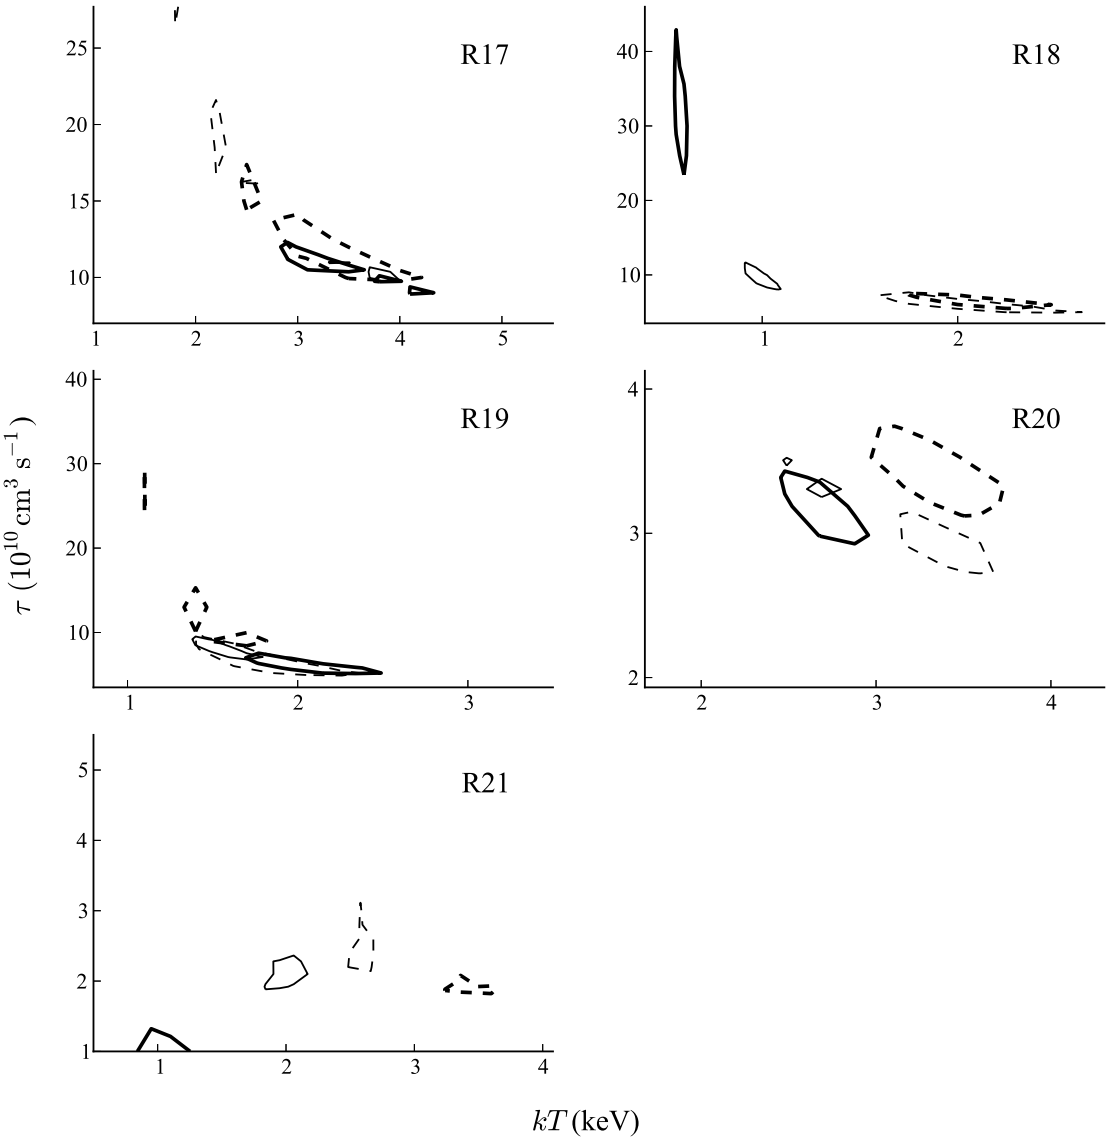

Figure 2-14: Continuation of Figure 2-12. The 95% confidence contours in the kT–τ plane are shown for each epoch: 2000.2 (thick), 2004.4 (thin), 2007.9 (thin dashed), and 2009.8 (thick dashed).

Region	Epoch	kT (keV)	τ (10^{10} s cm^{-3})	norm ($\times 10^{-4}$)
R01	2000.2	2.62 – 3.68	2.6 – 3.24	7.04
	2004.4	2.16 – 2.57	3.51 – 4.01	10.32
	2007.9	2.5 – 4.1	2.58 – 3.6	9.58
	2009.8	2.7 – 3.54	2.95 – 3.66	9.95
R02	2000.2	2.02 – 2.5	3.09 – 4.0	10.24
	2004.4	2.75 – 3.89	3.67 – 4.59	5.66
	2007.9	3.33 – 6.19	3.5 – 4.58	7.54
	2009.8	3.37 – 5.87	3.53 – 4.64	9.65
R03	2000.2	1.55 – 1.7	5.16 – 5.67	32.10
	2004.4	2.19 – 2.5	4.66 – 5.2	10.00
	2007.9	1.6 – 1.92	5.25 – 6.42	21.29
	2009.8	1.85 – 2.4	4.51 – 5.82	8.31
R04	2000.2	1.7 – 1.94	4.03 – 4.71	30.65
	2004.4	1.8 – 1.95	4.26 – 4.67	34.09
	2007.9	1.73 – 2.9	3.6 – 5.6	14.45
	2009.8	1.75 – 2.61	3.87 – 5.57	10.96
R05	2000.2	2.4 – 2.9	3.17 – 3.76	15.91
	2004.4	2.69 – 3.36	4.14 – 4.82	5.73
	2007.9	2.13 – 3.65	4.54 – 6.78	7.13
	2009.8	2.98 – 3.69	4.22 – 4.9	15.26
R06	2000.2	1.25 – 1.78	3.66 – 5.4	6.87
	2004.4	1.89 – 2.56	2.97 – 3.81	3.09
	2007.9	2.99 – 6.8	1.95 – 2.93	2.10
	2009.8	1.52 – 2.96	3.14 – 5.65	3.57

Table 2.7: 2D confidence interval bounds. Continued on next page...

R07	2000.2	2.03 – 2.4	2.83 – 3.38	2.54
	2004.4	2.15 – 4.4	2.27 – 3.3	3.43
	2007.9	1.94 – 2.64	2.65 – 3.36	5.67
	2009.8	2.14 – 3.95	2.52 – 3.76	4.11
R08	2000.2	3.89 – 4.56	2.41 – 2.65	1.42
	2004.4	1.6 – 2.19	3.57 – 4.59	2.66
	2007.9	1.11 – 4.0	2.5 – 6.1	1.88
	2009.8	1.1 – 3.5	2.24 – 3.78	4.37
R09	2000.2	2.0 – 3.08	12.51 – 25.0	11.13
	2004.4	1.9 – 2.9	14.0 – 26.0	5.38
	2007.9	1.41 – 2.48	16.0 – 58.0	7.37
	2009.8	1.48 – 2.53	15.92 – 46.0	4.49
R10	2000.2	0.89 – 1.75	2.69 – 5.5	2.04
	2004.4	0.79 – 1.54	3.39 – 9.2	2.05
	2007.9	0.94 – 1.95	3.02 – 5.4	2.58
	2009.8	1.18 – 1.96	2.78 – 8.2	1.41
R11	2000.2	1.4 – 1.89	5.31 – 8.44	13.38
	2004.4	1.81 – 2.34	5.11 – 6.59	7.72
	2007.9	1.28 – 2.12	5.06 – 11.59	15.68
	2009.8	1.77 – 2.64	4.33 – 6.71	15.08
R12	2000.2	1.46 – 2.87	3.49 – 7.7	3.80
	2004.4	2.17 – 2.83	3.64 – 4.58	5.56
	2007.9	1.46 – 3.58	2.45 – 6.31	5.57
	2009.8	1.15 – 4.0	2.35 – 5.15	6.73

Table 2.7: 2D confidence interval bounds. Continued on next page...

R13	2000.2	1.73 – 3.37	3.88 – 7.04	3.03
	2004.4	1.45 – 2.06	6.11 – 9.9	3.17
	2007.9	1.39 – 2.69	4.1 – 10.7	3.04
	2009.8	1.11 – 1.93	7.05 – 22.35	4.24
R14	2000.2	1.52 – 2.24	3.45 – 4.78	3.93
	2004.4	1.95 – 2.31	4.26 – 4.99	5.15
	2007.9	1.68 – 2.54	3.4 – 5.0	4.92
	2009.8	1.57 – 3.15	2.39 – 5.2	2.23
R15	2000.2	2.07 – 3.88	2.9 – 4.1	2.08
	2004.4	2.02 – 3.5	3.54 – 4.72	2.40
	2007.9	1.24 – 1.98	5.23 – 9.25	5.68
	2009.8	0.97 – 3.0	3.7 – 9.3	2.58
R16	2000.2	1.32 – 1.5	18.0 – 18.91	5.87
	2004.4	1.39 – 1.6	15.2 – 20.0	5.29
	2007.9	1.7 – 2.63	8.07 – 11.2	3.37
	2009.8	1.31 – 1.94	10.52 – 23.25	4.69
R17	2000.2	2.0 – 4.4	8.65 – 15.0	1.08
	2004.4	2.71 – 3.99	9.6 – 13.37	1.07
	2007.9	1.8 – 5.0	9.3 – 27.15	1.21
	2009.8	1.92 – 4.5	9.46 – 18.75	1.18
R18	2000.2	0.55 – 0.62	23.12 – 43.08	27.03
	2004.4	0.89 – 1.1	8.07 – 11.6	9.47
	2007.9	0.5 – 3.0	3.82 – 8.48	4.53
	2009.8	1.75 – 3.0	5.5 – 8.0	6.78

Table 2.7: 2D confidence interval bounds. Continued on next page...

R19	2000.2	1.53 – 2.53	5.04 – 8.2	3.35
	2004.4	1.35 – 1.9	6.24 – 9.9	3.77
	2007.9	1.34 – 2.4	4.71 – 11.15	3.41
	2009.8	1.1 – 2.61	5.0 – 32.4	4.52
R20	2000.2	2.48 – 2.96	2.92 – 3.46	9.19
	2004.4	2.49 – 2.85	3.22 – 3.57	13.07
	2007.9	3.09 – 3.75	2.73 – 3.16	16.30
	2009.8	2.91 – 3.8	3.13 – 3.79	12.25
R21	2000.2	0.76 – 1.28	1.0 – 5.5	0.66
	2004.4	1.51 – 2.5	1.2 – 2.29	0.52
	2007.9	2.37 – 2.71	2.0 – 3.11	0.80
	2009.8	3.31 – 3.6	1.76 – 2.17	1.26

Table 2.7: Limiting bounds of the 2-dimensional confidence kT–τ contours (90%) from the ACIS analysis. Bolded knots had significantly non-overlapping confidence contours, and we consider them to be evolving. Norm values were not included in the confidence contour parameter estimation, so no uncertainties can be provided. For the sake of estimation, a limited study yielded one-dimensional confidence intervals ranging between ±3% for brighter knots and ±50% for dimmer ones.

Region	N_H (10^{22} cm^{-2})	Mg (solar)	Si (solar)	S (solar)	Ar (solar)	Ca (solar)	Fe, Ni (solar)
R01	1.52	0.32	3.97	3.36	3.21	2.20	0.05
R02	1.63	0.12	5.04	4.41	4.26	3.87	0.00
R03	1.22	0.23	2.84	2.11	2.18	1.99	0.27
R04	1.23	0.34	2.67	2.06	2.84	1.42	0.27
R05	1.23	0.11	2.93	2.13	2.21	3.15	0.16
R06	1.08	0.22	4.08	3.44	3.94	4.48	0.12
R07	1.31	0.00	6.52	7.51	7.47	12.20	0.48
R08	1.41	5.16	4.25	5.68	7.17	9.72	1.02
R09	1.53	0.58	4.88	3.92	3.54	2.47	0.79
R10	1.45	0.47	5.21	5.43	5.48	3.58	1.10
R11	1.63	0.45	3.13	2.58	2.09	2.74	0.60
R12	1.58	0.41	3.30	2.65	2.34	2.16	0.34
R13	1.49	0.48	3.95	4.11	4.51	3.97	0.64
R14	1.41	0.20	2.72	2.38	3.37	2.77	0.82
R15	1.12	0.18	4.60	4.09	4.65	3.72	0.27
R16	0.91	0.41	2.74	2.92	2.32	1.47	0.74
R17	0.93	1.00	6.12	5.58	4.95	2.86	4.70
R18	1.42	0.31	1.44	1.62	1.63	1.22	0.00
R19	1.38	0.27	3.22	4.12	4.29	5.30	1.11
R20	1.47	0.31	4.65	4.12	4.02	3.25	0.07
R21	1.45	0.41	5.47	5.13	4.09	0.71	0.32

Table 2.8: The point estimates for the parameters of the ACIS analysis that were tied across all epochs. No confidence contours were made for these parameters, so any apparent trends should be treated with reservations. Moreover, the general Si-rich, Fe-poor trends are biased, as the HETG analysis requires regions with bright Si features. The hydrogen-equivalent column density values are consistent with previous analyses (e.g. Hwang and Laming [63]). Bolded knots exhibit significantly non-overlapping confidence contours, and we consider them to be evolving.

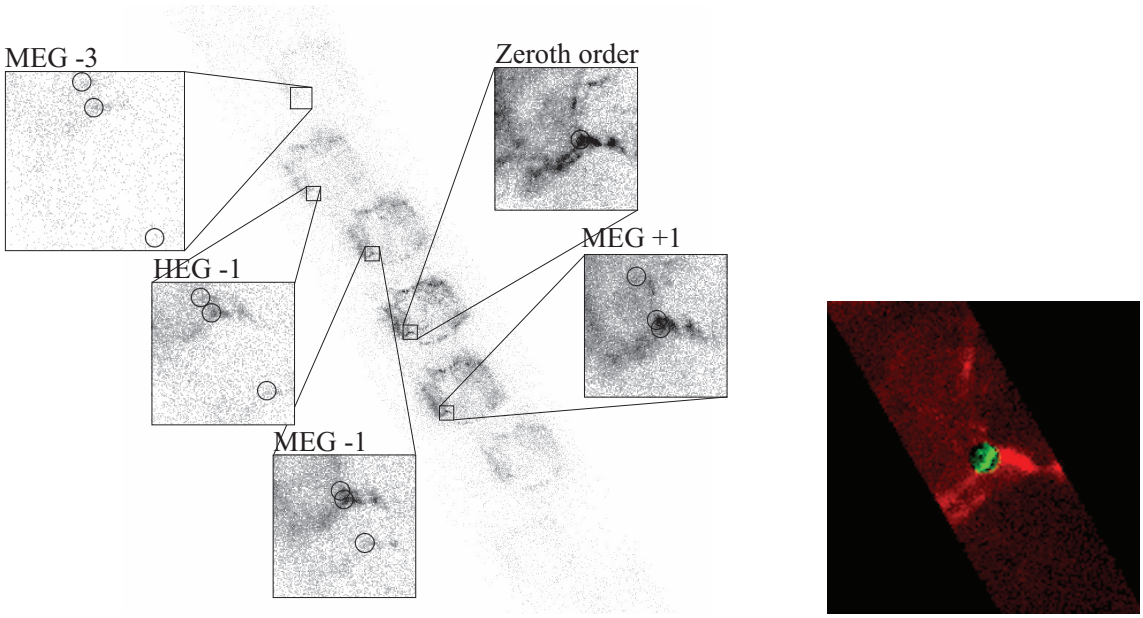

(a) (b)

Figure 2-15: The Event-2D software models and fits HETG dispersed data directly in two dimensions. (a) A knot (R1, in this case) can be found dispersed out to several orders. The circles show the expected dispersed locations of the H-like line and the He-like r and f lines for the best-fit velocity. (b) The zeroth-order spatial model for the same knot includes a 4-Gaussian spectral model (green, central region) and a `vnei` component for the surroundings (red), based on the ACIS results. Including the surrounding component along the dispersion direction (upper-left to lower-right) better models the spatial-spectral overlap.

2.B Appendix B: HETG Analysis Details

In general, extended sources cannot effectually be observed with slitless dispersive spectrometers such as the gratings on *Chandra* and *XMM-Newton*. However, for cases like the thermal knots of Cas A, their small size and bright line emission allow line information to be extracted from the dispersed data [20]. On CCD-readout instruments, the dispersed images (Figure 2-15(a)) carry useful information not only in the direction of dispersion, but also in the cross-dispersion axis, as well as in energy, which can provide order sorting.

The previous analyses of Paper A adapted the usual 1D spectral fitting machinery of `pha`, `arf`, and `rmf` files to the Cas A knots by collapsing the dispersed events along the cross-dispersion direction. To achieve higher resolution, the features were bent along the

cross-dispersion axis before creating that order's 1D spectrum. The companion `rmf` for each order encoded the spectral broadening introduced by the spatial distribution of the sheared zeroth-order events. A spectral model consisting of continuum plus several lines was fit to the `pha` data to determine Doppler shifts and line ratios.

For the current (re-)analyses of the HETG datasets, we improve the analysis technique by modeling and fitting the 2D dispersed events directly. The Event-2D software [1], briefly described in Dewey and Noble [21], is written in S-Lang [2] to provide an extension of the ISIS software. The software is an example of X-ray analysis that goes beyond the usual 1D fitting approach [97]. Event-2D removes the need to define a filament, allows different spectral models to be assigned to the knot and its surroundings, and, through its instrument simulation, utilizes a narrow order-sorting range to reduce background.

Knot Extractions

Spectral extraction consists of generating the usual spectral products (`pha`, `arf`, and `rmf` files) for a source observed with the HETG. The standard extraction process also adds grating-specific data (columns) to the `evt2` file based on the location of the source center. These values include an event's assigned diffraction order and 2-dimensional grating coordinates, `TG_R` and `TG_D`. (Note that many events will not be part of an extraction for a given source location; these events are flagged by a diffraction order set to 99.) It is this `evt2` file that Event-2D uses as input data in place of the `pha` file. Because of this, the standard HETG extraction steps are sufficient preparation for the Event-2D analyses.

The extractions for the Cas A knots were carried out with CIAO tools using standard HETG-appropriate steps. For convenience, scripts from TGCat[3] were used to simplify this process, yet still allow customization. The location for the extraction of each knot was manually set by comparing the regions of ACIS epoch 2004.4 with the HETG zeroth-order image to effectively select the same knot. We made some adjustments to the usual extraction

[1] Event-2D web page: http://space.mit.edu/home/dd/Event2D/
[2] S-Lang web page: http://www.jedsoft.org/slang/
[3] TGCat web page: http://tgcat.mit.edu/

parameters to improve the subsequent 2D analyses: (1) the cross-dispersion width was set to twice the knot diameter, (2) a large zeroth-order radius of 60 arcseconds was defined, (3) a narrow order-sorting range of $\pm\,0.06$ was set, and (4) `arf` files were made for orders $m = \{0, \pm1, \pm2, \pm3\}$. The large zeroth-order range was chosen to include (and hence model) off-knot spatial features that can produce spatial-spectral overlap with the knot. The narrow order sorting range still includes most of the Si-range flux while helping to reduce cross-talk from other spatial regions.

2D Spatial-spectral Modeling

The spatial model of each knot is self-described using the zeroth-order events within a given radius of the extraction center. The events surrounding the knot region are used to define a second spatial component of the model – see Figure 2-15(b). Modeling this surrounding region is an important part of generating simulated dispersed events in the wavelength range of interest.

Spectral models are assigned to the two spatial components. The knot spectrum is described by 4 Gaussians corresponding to the He-like r, i and f lines and the H-like Ly α line. The f/i ratio is fixed at the low density value of 2.45 and the remaining free parameters are then the overall flux, the f/r ratio, the H/He ratio and a common Doppler shift. For the surroundings' spectrum we start with the best-fit ACIS (Epoch 2004.4) `vnei` model and allow its τ to be adjusted. This can give a reasonable approximation of the surroundings, especially for the closest features that are most important. (Of course one could explicitly measure and use the spectrum of the surrounding region itself, perhaps even in multiple zones. However, we have not seen a clear need for this level of model fidelity in the analyses so far.) Because we are not very sensitive to the continuum shape, the temperature is held fixed. Although it need not be, the velocity of the surroundings is fixed at the knot velocity; we have not seen a clear indication in the data for a need to change this.

The 'Source-3D' routines (companion to Event-2D) are used to organize and access the spatial-spectral model described above. In particular Monte Carlo source events can be gen-

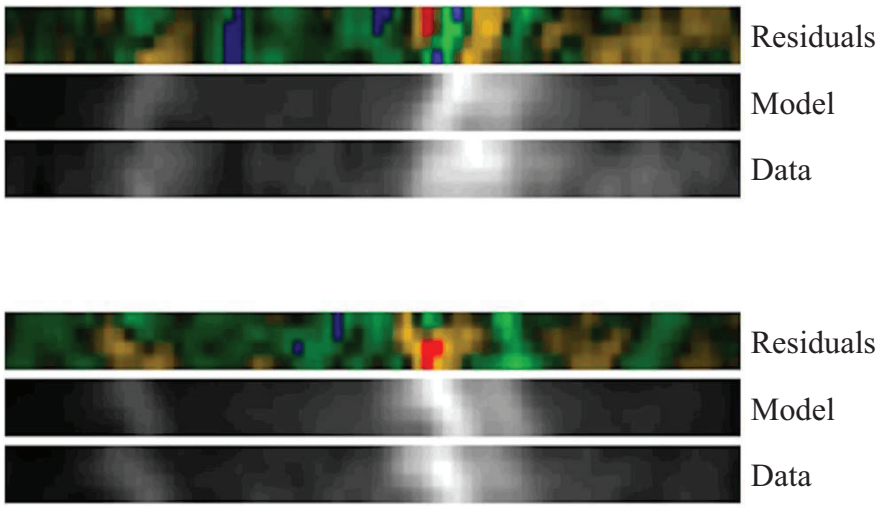

Figure 2-16: The Event-2D software shows remarkable agreement between data and the model. Shown are the epoch 2010.3 R5 MEG ±1 order 2D images. The Si-line region is shown with the wavelength range going from shortward of the Si XIV line at left to longward of the Si XIII triplet at right. Note how the apparent angle of the knot differs between orders.

erated from the defined model and then passed through an approximate HETG instrument simulator in Event-2D. In this way simulated model data are created that can be compared with the actual data (Figure 2-16). The inclusion of order-sorting effects in the instrument model is notable; the effects play a large role in shaping the observed and modeled events.

Parameter Estimation

Parameter estimation proceeds by forward-folding the spatial-spectral model to generate model events, which are then compared to the spectra. The data and model events are binned in defined 2D histograms, and the chi-squared statistic is calculated: $\chi^2 = \sum_i ((D_i - M_i)/\sigma_i)^2$, where $\sigma_i = \sqrt{0.63\,M_i + 0.37\,D_i}$. This modified definition of σ_i reduces the bias in the statistic compared to either $\sigma_i = M_i$ or $\sigma_i = D_i$.

The randomized nature of the model generation forces us to use noise-tolerant fitting algorithms. A small change in a parameter value can be masked by the 'model noise' of Markov Chain Monte Carlo (MCMC) samples (even though the model counts are over-

simulated by a factor of 10). For instance, around the best-fit value of parameter p, we have that $\Delta\chi^2 \approx a\,\Delta p^{\,2} + \sigma_{\chi^2}$ where a is a constant depending on the particular data and model and the σ_{χ^2} term represents the noise in the model computation. Hence changes in p are not 'noticeable' above the model noise unless $\Delta p > \sqrt{\sigma_{\chi^2}/a}$, effectively setting a minimum scale size for parameter p. Because of this, we use minimization schemes that include a minimum scale for noticeable changes in a parameter and that are tolerant of the model noise. For example, to determine a single parameter like the Doppler velocity, multiple evaluations of the statistic are performed at each step of 100 km s^{-1} to estimate σ_{χ^2}, which aids in determining the best parameter value. For multi-parameter fitting, a noise-tolerant conjugate-gradient (C-G) scheme can be constructed to minimize a parameter in 1D and include the noise level in its convergence diagnostics.

Another approach to dealing with noisy models is to use a MCMC method: the noise of the model generation is only a slight addition to the larger variation expected as the parameter space is randomly explored and so the technique is not sensitive to the model noise.

To analyze these X-ray knots, we combine the two methods. The C-G fitting coarsely sets the surrounding region norm and Tau values, fixing the Tau value. A 1D scan in velocity space is then performed and the velocity is fixed. A C-G scan is then done to fit the two line ratios and the ratio of flux from the surrounding region to that of the knot. MCMC exploration is carried out for these 3 parameters starting from the C-G best-fit values; in this way there is generally little settling time for the MCMC. The spread in the MCMC draws provides the error bars for the parameters.

As mentioned above, the modeling is sensitive to the CCD gain calibration through the narrow order-sorting range. After a knot's analysis is complete, the simulated and real CCD energy distributions are compared for the bright He-like Si line in each order. Where these differ by more than 0.8%, a custom gain correction is included for that epoch-knot-order (applied when the data are read in) and the knot is re-analyzed.

The measured values for a knot are its Doppler velocity, the two line ratios (f/r and

He-like/H-like), and the overall knot flux normalization; these are given in Table 2.6 for both epochs of HETG observations. Uncertainties are given for all but the flux, which has a statistical uncertainty generally less than 2% and so will be dominated by systematic errors, e.g., calibration of the effective area at a knot's specific location.

2.C Appendix C: ACIS Analysis Details

In this appendix we provide the details of the data reduction and analysis of the ACIS data. The observation IDs and exposure times are listed in Table 2.2, Section 2.2. The ACIS analysis used four archival datasets from early 2000, mid-2004, late 2007, and late 2009. All four of these observations were taken on the back illuminated ACIS-S3 chip in GRADED data mode. They were targeted at a right ascension of $23^{\rm h}23^{\rm m}26^{\rm s}.70$ and a declination of $+58°49'03''.00$.

Spectral Definition of the Knots

We assumed these small, bright knots are individual clumps of similar material, so we therefore defined their boundaries spectrally. Once defined, all of the enclosed material shares the same spectral features. Arbitrarily bounding these regions by hand, using only brightness information and an analyst's intuition, could unintentionally include non-knot material of similar surface brightness, thereby skewing results of the analysis. We found this to be the case, so developed a more systematic method than hand-selection.

Our region-determining algorithm finds the radius at which nearby material starts to contaminate the knot. For each knot we select an inner region consisting of the brightest central pixels – typically 4 to 6 in number; this region defines the spectrum of the knot. We then compare this spectrum to the spectra of annuli of increasing radius (Figure 2-17), evaluating the Poisson likelihood of the annulus data, given the central region's spectrum as the model. The knot boundary is defined at the annulus where the likelihood drops sharply, or where the normalized counts per unit area drop abruptly. In the absence of sharp changes, we visually compared the spectra to delineate the boundary.

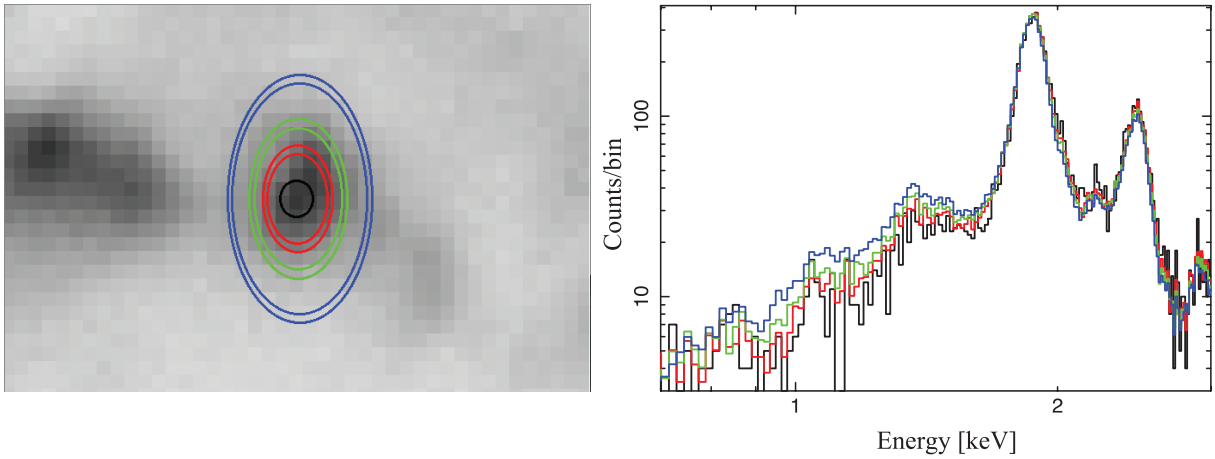

Figure 2-17: The knots are defined spectrally. Left: Anulii and spectra used for the knot region definition for knot R02. The two same-color ovals define an annulus, and the spectra from inside this annulus has been extracted and compared to the central black oval's spectrum. Right: normalized spectra from the central part of R02 (black) and three annuli. A soft-energy excess can be seen in the spectral shape of the outer annuli. These spectra were used to determine the boundary of the knot ensuring spectral homogeneity within the defined knot.

The Use of Close-in Backgrounds

The bright X-ray knots are expanding amidst a diffuse web of emitting material. We must therefore carefully consider how to define the 'background' (which could very well be the foreground) for spectral modeling, as the knot is seen *through* this web. We systematically examined the characteristics of the surrounding material to determine the most faithful way to define the background for each knot.

We decided to assign multiple diffuse, nearby regions as the background for each knot. Figure 2-18 shows a typical background assignment. These diffuse regions were characterized by low surface brightness (over the *Chandra* passband) and variations over larger angular scales. We typically chose several regions for each knot for two reasons. First, these regions were more spectrally consistent (and therefore yielded similar plasma parameters for the knot model) than background regions that contained brighter features varying on smaller angular

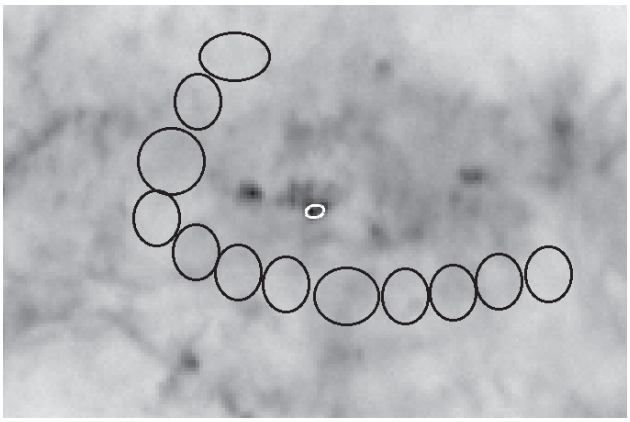

Figure 2-18: We assume that the low-brightness surrounding material fills the front or rear of the knot, at about the same level as in our background region. The background regions are shown in black with the R05 region shown in white. By taking many regions around the knot, we effectively average out the local small-scale variations to approximate the actual backgrounds and foregrounds in the spectral extraction 'core'.

scales. Second, since this diffuse material varied on longer scales, it was more probable that it contaminated the knot spectrum. Choosing multiple regions around the knot effectively interpolated the spectrum to the position of the knot. The regions remained the same across epochs, save for slight nudging required because of the ejecta expansion.

The background regions were chosen near to the knots but not directly adjacent. We considered using the discarded non-knot annuli from the region-determining algorithm. However, the material in these adjacent rings showed spectral features not present in the knot on order of the same surface brightness, suggesting different compositions.

In general, we find that the diffuse material we select as backgrounds tends to be cooler and to move more slowly than the ejecta knots. Including background spectra in the model generally results in higher derived plasma temperatures and Doppler velocities for the knots.

Data Extraction

The development of this algorithm yielded a unique image extraction technique. For an extended source like Cas A, the `specextract` CIAO script is typically used. The area-weighted `arf` (`warf`) created by `specextract` does not include the effect of angular dither,

since the region area is assumed to be large compared to the dither. The annuli usually have a size of 6-15 pixels, large enough to warrant a `warf` and `wrmf`, yet small enough to be affected by the dither of the instrument (David P. Huenemoerder, private communication, 2011). To create, then, a `warf` that includes dither, we first make the `pi` file and `wrmf` with `specextract`. Running the `sky2tdet` function outputs a `wmap` file describing the dither, which can then be used a second time in `mkwarf`, rewriting the `warf`. The tweaked calibration product is now appropriate to our peculiarly sized regions.

This extraction was done using CIAO version 4.3, with CALDB version 3.2.2.

Spectral Fitting

We applied a non-equilibrium ionization model with variable abundances (`vnei` version 2.0) to the data, also accounting for interstellar absorption (`phabs`) and pileup [17]. We chose not to employ the similar `vpshock` model due to the additional free parameter of the lower τ value. Subsequent investigations on several knots yielded maximum likelihood `vpshock` values of τ and kT within 10% of those for the `vnei` model, where we compared the upper τ for `vpshock` to twice that of the `vnei` τ. Therefore, we do not think our choice of model unduly biased the results.

For each knot, we expect several physical parameters to stay constant across the decade of observations: the equivalent neutral hydrogen column density and the elemental abundances. To this end we performed the fits to the four epochs of ACIS data for each knot simultaneously, while tying n_H and the abundances from the 2004.4, 2007.9 and 2009.8 datasets to that of the 2000.2 dataset. Values for C, N and O abundances were frozen to 5 times the solar value. (As discussed in Section 2.4.1, the actual values of the abundances used can be degenerate with the norm of the continuum fit. We thus chose a value of 5 to account for the expected C, N, and O dominance in the continuum. Vink et al. [133] argues for O, while Dewey et al. [22] argues for H/He, or C/N.) Fe and Ni abundances were tied together.

For the `pileup` component of the model, all parameters but n—the number of regions where pileup occurs—are fixed across epochs: $\bar{g_0} = 1$, $\alpha = 0.5$ and $f = 1$. The $\bar{g_0}$ and α

values are simply fiducial estimates for grade-migration processes. The parameter f was set to 1 because all of the knots are larger than *Chandra*'s PSF. Letting n alone vary captured the effects of pileup well, reducing the χ^2/d.o.f.

The redshift parameter in the model was treated only as a calibration gain factor, and it was allowed to differ from epoch to epoch. Redshifts were point-estimated at the beginning of the fitting process, then frozen throughout the confidence contour generation. This initial point estimate was fit with the He-like r, i, f, and H-like Ly α Si lines folded through the ACIS `rmf`. By fixing the redshifts to these maximum likelihood values, the confidence contour computation time was greatly reduced. (We attribute the added computational complexity to the discontinuous χ^2 space along the redshift axis, a persistent feature we found with this model.)

The combined four-epoch fit for each knot involved 88 parameters, with 25 of these free. With such a complicated parameter space, it was easy for the fit to fall into—and get hung up in—local minima. In order to better find the global minimum for each fit, we used the native, parallelized ISIS function `conf_loop`. This function performs confidence limit searches on each parameter individually about the current best fit value, often climbing out of a local minimum to find a better fit.

Generating the Confidence Contours

Once the spectral fits have found the maximum likelihood point for each epoch of each knot, we proceed to generate confidence contours for our fit parameters. A grid is laid down around this best estimate parameter vector and the S statistic (defined below) is minimized at each point to map out the confidence levels. The confidence contours are made with `conf_map_counts` in ISIS.

The results of these fits are shown in Appendix 2.A, Figures 2-12 – 2-14 and Tables 2.7 and 2.8.

The reported parameter ranges represent confidence levels. That is, if the analysis could be repeated, the true parameters would lie within the 90% contours for 90% of the re-analyses.

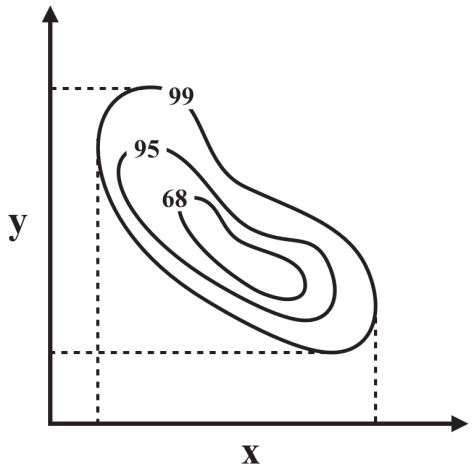

Figure 2-19: We report our uncertainties for the ACIS analysis as bounding edges of the confidence contours. Most of the knots exhibit correlated 'bananas' in kT–τ space, so the summary is sufficient.

Moreover, we only report the bounding edges (see Figure 2-19), even though the confidence contours are never rectangular. We do this for conciseness and because the contours all have the same anti-correlated 'banana' shape seen in Figure 2-12.

Assumptions and Choice of Fit Statistic

We make several assumptions to warrant our use of confidence contours for parameter estimation. We assume the counts in each energy bin are independent from each other and that they are Poisson-distributed, with each energy bin E_i having a mean λ_i given by the model. We group counts into energy bins of at least 30 counts, so that we may approximate the count distribution in each bin as normal, with mean λ_i and standard deviation $\sqrt{\lambda_i}$. (Our inclusion of pileup both undermines and supports this assumption. We admit that a count in one bin may originate from two in another, yet since we account for pileup effects, the bins of the model spectra will have less correlation.)

This large minimum number of counts per energy bin allows us to employ the S statistic, popularized by Cash for nonlinear X-ray models. (We did not use Cash's more appropriate C statistic solely because of the extra computation time within the ISIS framework.) The S

statistic is merely –2 times the log likelihood:

$$S = \sum_{i=0}^{N-1} \frac{(d_i - \lambda_i)^2}{\lambda_i},$$

where d_i are the data counts in each of the N bins and, as above, λ_i are the predicted model counts. We are interested in only a subspace of the whole parameter space, primarily the kTs and τs over the four epochs. Cash's result says that the difference between the minimized values of S over this q-dimensional subspace and the overall best fit S follows a χ^2_q distribution. Thus, with probability $1-\alpha$, the true values of the parameters reside within the contour(s) where this difference equals $\chi^2_q(\alpha)$.

As a final approximation, during confidence contour generation we assume that the epochs are independent, despite explicitly tying N_H and elemental abundances between epochs. That is, for one epoch of one knot, we grid up kT and τ values and minimize S at those points, with N_H, the abundances, and the kTs and τs from other epochs free to vary. We do this merely for computational tractability: under this approximation we only have to minimize S on $4 \times N^2$ grid points, instead of $N^{2\times4}$. The maximum likelihood parameters for each epoch differed little from the global best parameters, suggesting this assumption is sound.

Electron Density Derivation

To calculate n_e, we use a combination of model parameters and inferred values. Parameter estimation with the `vnei` model yields the norm, X_{norm}, and the abundances, $X_A(Z)$. We estimate the knot volume, V, as an oblate spheroid, with axes taken from our region-finding algorithm. We set the distance to Cas A to $d = 3.4$ kpc. We approximate the number of electrons stripped with a functional form:

$$Q(Z,T) = 0.41Z(\log(T(28/Z)^2) - 5.2).$$

(The function is constrained to the interval [0, Z].) Finally, the solar abundances, $A(Z)$, are drawn from Anders and Grevesse [1].

Using the conventional definition of X_{norm}, we can find the electron density.

$$\begin{aligned} n_e &= \sqrt{(n_e n_H)\left(\frac{n_e}{n_H}\right)} \\ &= \sqrt{\frac{4\pi d^2 10^{14} X_{norm} X_A(Z=1)}{V}\left(\frac{n_e}{n_H}\right)} \\ &= \sqrt{\frac{4\pi d^2 10^{14} X_{norm} X_A(Z=1)}{V}\frac{\sum Q(Z,T) X_A(Z) A(Z)}{X_A(Z=1)A(Z=1)}} \end{aligned}$$

The factor of $X_A(Z=1)$, called X_H in the main text, properly scales the norm, allowing us to set the hydrogen abundance in `vnei` to something other than 1.

Finally, if we deem than a `vnei`-modeled plasma component does not take up the full volume of the knot, then $V \to fV$, where f is the fill fraction.

Chapter 3

Searching for Scandium in Cas A

In contrast to the knots analysis, this second investigation of the ejecta of Cas A is a search for a *diffuse, non-thermal* signal. Recall from Section 1.5, 'non-thermal' means the radiation mechanism is not inherently thermal: the radioactive decay line emission depends entirely on the decay times and amount of the material, not its temperature. Kinematics and (thermal) line-broadening are taken into account, though. *Chandra*, again, suits the search well, primarily for the large amount of archival data. The diffuse signal, spread over the interior of the remnant, prohibits the use of the HETG, but I conclude this chapter with predictions of how higher spectral resolution can improve the search.

Due to extreme conditions, supernova explosions form entirely new elements that cannot be made during a star's lifetime of fusion or in the Big Bang. The relative mass fraction of these synthesized elements can depend heavily not only on the theoretical model, but also on physical characteristics such as the progenitor's mass and its explosion type. Observations of these trace elements will help constrain explosion models and progenitor characteristics. In this chapter, I describe my work in progress to unearth the faint signal of one such tracer, scandium, from the supernova remnant Cassiopeia A.

^{44}Ti, the original product of the ^{44}Sc signal we seek, is made predominantly at the exact moment of a massive star's death. Once the stellar core collapses, the surrounding layers stream inward down the gravitational potential well. A shock wave develops as the material

accelerates past the sound speed. Aided by neutrinos escaping the core, the shock heats material to temperatures above 10^9 K, partially dissociating nuclei into protons, neutrons, and helium nuclei (α particles). The abundance of such particles speeds the production of heavier nuclei from Si, such as ^{44}Ti, ^{56}Co, and ^{57}Co [75]. All the while, the neutrino flux ejects the material outward, causing it to expand and cool. The nuclear reactions quickly stop and leave an abundance of He nuclei, giving this important nucleosynthetic process the name α-rich freeze-out.

The presence of ^{44}Ti in the ejecta of a supernova can uncover a great deal about the explosion [128, 24]. As indicated above, the isotope is produced deep in the star, very near the border between the soon-to-be ejecta and compact object. The location of this border – called the mass cut – sets the yield for ^{44}Ti. Various reaction rates, the unknown ionization state of ^{44}Ti, and large scale turbulence near the mass cut also effect the yield, though astrophysical observations are far from constraining these second order effects.

Radioactive ^{44}Ti reveals itself via several emission lines as it decays to stable ^{44}Ca (Figure 3-1). ^{44}Ti transforms to ^{44}Sc via electron capture (a proton absorbs an inner shell electron, converting to a neutron) over a characteristic timescale of 85 years. The nucleus relaxes, emitting 68 keV and 78 keV lines, but an electron vacancy remains, usually in the K shell. The excited ^{44}Sc atom fills the vacancy with an electron, emitting the extra energy as a 4.09 keV photon 17% of the time; thus, the X-ray signal is fainter than the hard X-ray nuclear signatures. The radioactive isotope decays quickly ($\tau \sim 8$ hr) to the stable ^{44}Ca isotope via β^+ decay or electron capture, whereupon it emits a 1.16 MeV γ-ray to de-excite the nucleus.

Our signal of interest, the 4.09 keV photon, is emitted via fluorescence, where an electron transitions to a vacant inner shell. The atom fills the vacant inner shell with one of its own outer shell electrons, releasing energy, either as a photon or by ejecting an electron in a process known as the Auger effect. This process has a fluorescence yield – the probability that a photon, not an Auger electron, is emitted – of 19.9% for Sc [119]; thus, 4/5 of the signal is lost in the X-ray band.

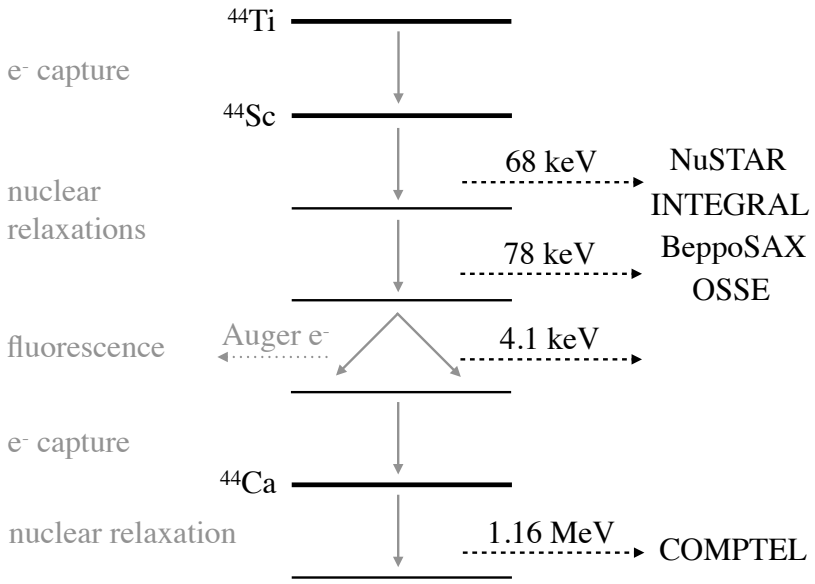

Figure 3-1: The ^{44}Ti decay chain exhibits several electromagnetic signatures in the X-ray and γ-ray bands. The decay time for radioactive ^{44}Ti is 85 years, well matched to young SNRs like Cas A. A selection of previous observations for the various decay lines is listed on the right.

The fluorescence yield[1] is divided among transitions to various atomic shells: 0.5% for the L shell ($\sim$ 0.4 keV), 17.2% for Kα (4.09 keV), and 2.3% for Kβ (4.46 keV). The 4.09 keV photon is our primary interest, since the Kα transition is most probable. However, the environment of the explosion may leave the ^{44}Ti in an excited or ionized state, changing these percentages from the laboratory-derived values.

Several hard X-ray/γ-ray instruments have already detected ^{44}Ti byproducts in Cas A, but an X-ray search remains relevant. The superior spatial resolution of X-ray instruments like *Chandra* presents the strongest case for relevance. Whereas the ^{44}Ti distribution is blurred for γ-ray instruments, *Chandra* has the opportunity to map the spatial distribution of the deep-lying ejecta. Additionally, cross-checks between different wave bands will confirm reported yields and validate decay chain models for this radioactive isotope.

[1]The fluorescence yield can be derived from both laboratory experiments and theory; it is a smooth, monotonically increasing function of atomic number, reaching 50% probability with Zn.

3.1 Observational Campaigns

Iyudin et al. [71] first identified the 1.16 MeV γ-ray line in the vicinity of Cas A with the COMPTEL instrument on the *Compton Gamma-Ray Observatory* (*CGRO*) in 1994. The ^{44}Ti decay line flux of 70 f_6, where $f_6 = 10^{-6}$ photons/s/cm^2, was later revised downward to 34 f_6 with more data from *CGRO* [26]. Another instrument on board *CGRO*, the Oriented Scintillation Spectrometer Experiment (OSSE), weakly detected flux in all three high energy lines, at 18 ± 15 f_6 [127]. *BeppoSAX* confirmed the two lower energy nuclear de-excitation lines in 2001 with total flux 19 f_6 [134]. The IBIS telescope on *INTEGRAL* provided a more recent glimpse of Cas A's ^{44}Ti decay lines (the coded aperture mask achieves $12'$ resolution, for sub-arcminute point source localization) and the tightest constraints on the 68 and 78 keV flux, 25 ± 3 f_6 [113].

Dupraz et al. [26] also carried out an all sky survey for the ^{44}Ti line with *CGRO*, in the hope of serendipitously discovering more SNRs. Based on statistical grounds, the authors expected to find 3 total, aided in part by the γ-ray transparency of the Galactic disk, but Cas A was the sole confident source. (The second highest significance source GRO J0852-4642 was later confirmed coincident with *ROSAT* X-ray source RX J0852.0-4622, and has been imaged with *XMM-Newton* [3, 72]. No X-ray ^{44}Sc line has been seen.) This source number discrepancy has caused some consternation, and the inconsistency has only grown with a more sophisticated Monte Carlo-based prediction of 4–7 detectable sources by The et al. [128]. Clearly, the mechanisms for ^{44}Ti production and distribution are still in dispute, and Cas A presents a great target of opportunity for probing ^{44}Ti emission via different avenues.

Most pertinent to this analysis, however, is the recently launched *NuSTAR* instrument [53]. The satellite mission features dual mirrors at the end of a 10 m deployable mast, which focus hard X-rays (3–79 keV) onto two 64×64 pixel cadmium zinc telluride semiconductor arrays. The optics attain $18''$ resolution (FWHM), a remarkable improvement over the previous generation of instruments. The detectors are surrounded by an anticoincidence shield, to prevent stray hard X-rays and cosmic rays from polluting the focused signal, a

common problem for hard X-ray instruments. *NuSTAR* achieves ~ 1 keV FWHM resolution at the 68 keV ^{44}Ti line, but only ~ 400 eV resolution at the 4.09 keV ^{44}Sc line. Such poor low-energy resolution renders *Chandra* a much more suitable instrument to detect the 4.09 keV line.

The *NuSTAR* team has recently reported preliminary results on observations of Cas A [49] (Paper B). A de-convolved image of the 68 keV ^{44}Sc signal — unpublished at the time of writing — shows emission clumped on large scales, distributed near the center of the remnant, and does not track with the X-ray emitting Fe. Two peaks in brightness appear on either side of the central compact object, along the axis of the jet, a feature identified and investigated in many works. The data, taken in several observations starting in the second half of 2012, is still proprietary, and no flux values have been reported. The analysis presented here will be revised once these results are known, but headway can still be made on the 4.09 keV signal search with the available information.

Theiling and Leising [129] (Paper C, hereafter) performed the first and only search for the scandium X-ray line in Cas A, and placed upper limits on the flux. They extracted spectra from large regions (Figure 3-2) of two 50 ks *Chandra* observations to increase signal-to-noise. With no obvious 4.09 keV line visible, the spectra were modeled empirically to account for the shape, ignoring the ionization state of surrounding nuisance lines. Upper limits for the ^{44}Sc line flux were derived from the error bars of a maximum likelihood search. The search produced remnant-wide '2σ' flux upper limit[2] of the 4.09 keV ^{44}Sc line at 38 f_6.

With only 100 ks of data and no observational constraints on the distribution or dynamics of the ^{44}Sc, Paper C was obligated to investigate large regions, chosen under weak motivations. One composite Fe-rich region and one continuum-poor area served to cover the hypotheses that the ^{44}Ti followed the Fe distribution or remained unshocked towards the center, respectively. The ^{44}Sc flux upper limit of Paper C is nearly 10 times the predicted value from γ-ray observations, pointing to the dimmest regions of the remnant for a successful follow-up search, where the faint line could poke above the continuum. A renewed ^{44}Sc search of Cas A could improve upon the initial efforts of Paper C via several avenues:

[2]Paper C does not specify the method used to determine this value.

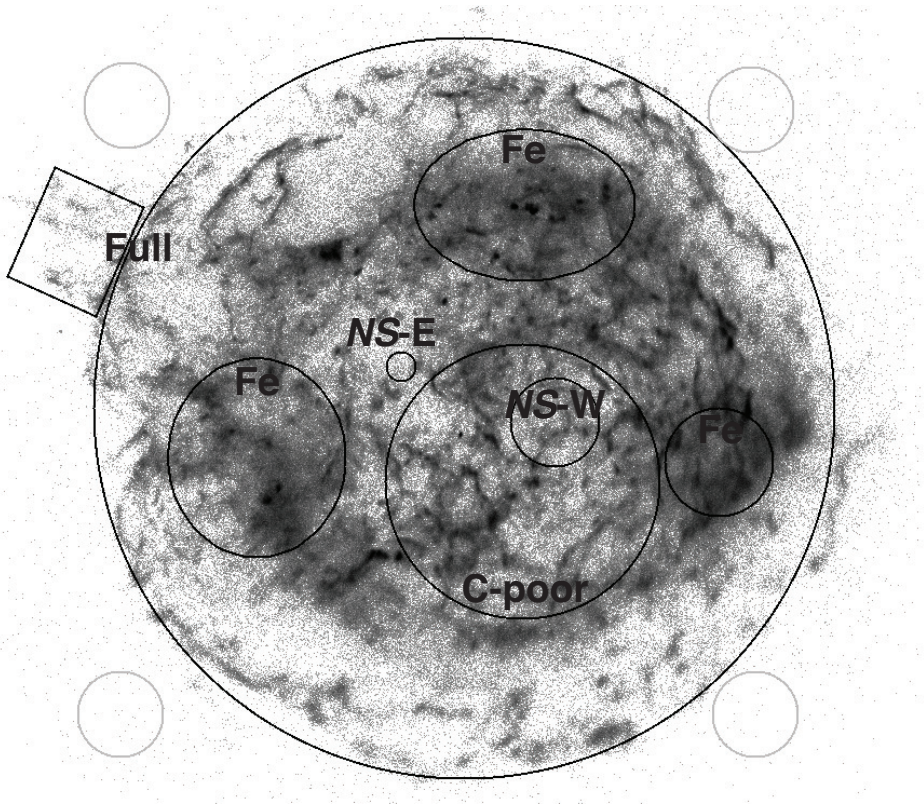

Figure 3-2: This 3.4–5.0 keV filtered *Chandra* image shows the search regions for ^{44}Sc. We revisit the large benchmark regions of Paper C: the whole remnant (Full), Fe-rich regions (Fe), and the continuum-poor interior (C-poor). The peak ^{44}Ti flux in the *NuSTAR* image coincide with the *NuSTAR*-East (*NS*-E) and *NuSTAR*-West (*NS*-W) regions. The grey circles outside the remnant constitute the background.

- Exposure time. *Chandra* has observed Cas A for over 1 Ms now, affording greater signal to noise for any tentative line detection.

- Velocity information. High spectral resolution infrared measurements of the unshocked ejecta are now available. If we believe the ^{44}Sc remains in the inner layers of ejecta, this data could inform its line shapes.

- Spatial distribution. With the *NuSTAR* data, the spatial distribution of the radioactive ejecta across the remnant could be better constrained to narrow the search.

- More rigorous statistics. Bayesian methods are now accessible to infer upper limits.

ObsID	Exposure (ks)
4634	149
4635	135
4636	143
4637	164
4638	165

Table 3.1: Cas A ACIS observations. All data were taken in April or May of 2004.

The 4.09 keV ^{44}Ti line has been observed in a different SNR, however, since the results of Paper C. Borkowski et al. [8] claim over 99.9% confidence for ^{44}Sc line emission from G1.9+0.3, the newly discovered — and youngest — Galactic remnant. A flux of 1.2 f_6 was reported from 237 ks of *Chandra* data. Broad spatial spectroscopy revealed nonthermal ^{44}Sc emission in the dim center and the brighter rim. The low ^{44}Sc flux puts the expected γ ray emission below the COMPTEL search limit of ~ 20 f_6 [26], consistent with non-detection.

Borkowski et al. utilize Bayesian statistics to quantify their claim of a line. They place educated (i.e., not uninformative) priors on the model parameters, motivated by past observations, and explore the posterior distribution via Markov Chain Monte Carlo sampling. The analysis employs a Bayesian likelihood ratio test [110] to compare between models[3].

My initial efforts to advance the ^{44}Sc X-ray search in Cas A are detailed in the rest of this chapter.

3.2 Methods: Creating an Upper Limit

As the first improvement on Paper C, we include more data, from the megasecond observation of Cas A (Hwang et al. [64]). The long exposure proves useful for the dimmest regions in the interior, where signal competes with statistical fluctuation. For the targeted search we use the five longest megasecond pointings, for a total of 755 ks, while for the broad remnant-wide

[3]The 'posterior predictive p-value' computes the traditional statistic many times via sampling to compare against a null model [92]. In this case, the null model, which gets rejected with over 99.9% significance, is that no line at 4.09 keV is present. Posterior predictive statements are indeed Bayesian, but in an unfamiliar way: inference is drawn from the posterior P(data replicates|data). Bayesian formulation of this classical test incorporates dependencies on both the data and nuisance parameters, acting to marginalize the nuisance parameters out of the null hypothesis by producing a 'mean' p-value.

Region	95% upper limit, this work (bin width as uncertainty)	2σ upper limit, Paper C
Whole remnant	24 (2.1)	38
Fe-rich area	19 (1.6)	15
Continuum-poor	9.8 (1.0)	23
NuSTAR-East	0.13 (0.014)	-
NuSTAR-West	1.9 (0.25)	-

Table 3.2: New Bayesian upper limits compared with Paper C. Fluxes in units of $f_6 = 10^{-6}$ photons/s/cm^2.

search, we use only one 164 ks pointing (Table 3.1).

We model the spectra with a simple absorbed power law plus two nuisance lines – Ar and Ca near 3.7 and 3.9 keV, respectively – and the ^{44}Sc line. The ^{44}Sc line is a delta function (modified by the ACIS response); we investigate the effect of other line shapes in Section 3.4. As with Paper C and Borkowski et al. [8], we refrain from trying a more physical model like `vnei` because we are mainly concerned about the shape of the spectrum. We performed standard data reduction with CIAO, and we fit the data products in the energy range of 3.4 – 5.0 keV. Backgrounds were chosen from outside the remnant's X-ray shell.

We first revisit the original search regions ('benchmark regions') of Paper C, to compare results. Next, we look closely at the small, dimmest parts of Cas A's interior that coincide with the *NuSTAR* detection. Even in these X-ray dim regions, only an upper limit can be placed. So, as a final stage of analysis, we ask what resolution and exposure would be required of an X-ray telescope for a detection. High resolution spectroscopy could enable the line to rise above the continuum emission.

We employ a Bayesian framework to marginalize nuisance parameters and provide a more natural interpretation of an upper limit. A maximum likelihood point estimate of, in this case, ^{44}Ti flux is influenced by the point estimates of parameters like the norm, column density, flux of nearby lines, and power law slope. It is proper, then, to 'average out' the correlations with these nuisance parameters, $\vec{x}$, via marginalization to obtain a posterior

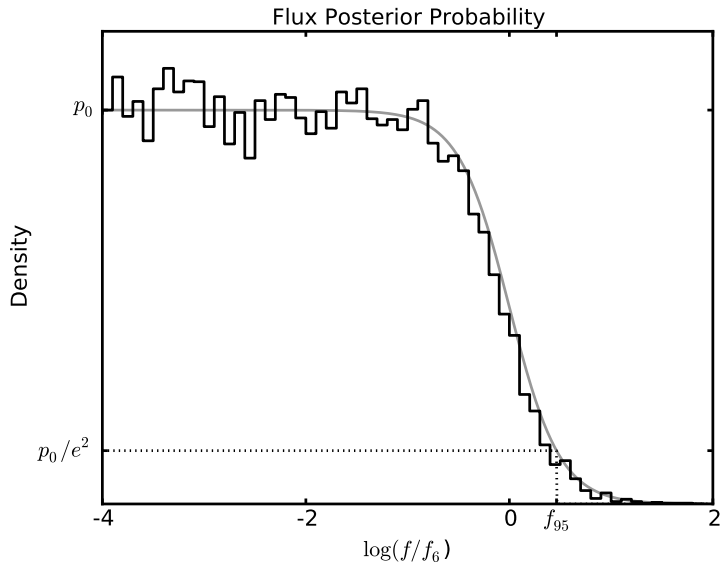

Figure 3-3: A schematic posterior for the ^{44}Sc flux shows how the CHMC algorithm modifies the flat prior down to a characteristic logistic ('S curve') shape. Our 95% upper limit corresponds to the height of a Gaussian function at 2σ. The bin width at the upper limit serves as a proxy for the uncertainty.

distribution of likely values given the data, D:

$$\begin{aligned} P(f_{Sc}|D) &= \int P(f_{Sc}, \vec{x}|D) d^N\vec{x} \\ &\propto \int P(D|f_{Sc}, \vec{x}) P(f_{Sc}, \vec{x}) d^N\vec{x}. \end{aligned}$$

The application of Bayes' rule in the second line allows us to compute the posterior with a manageable likelihood function, such as Poisson or normal (χ^2). We use uniform priors throughout the analysis. With the (normalized) posterior $P(f_{Sc}|D)$ in hand, the 95% upper limit is the value f_{95} such that:

$$\int_{-\infty}^{f_{95}} P(f_{Sc}|D) df_{Sc} = 0.95.$$

However, without a detection, this integral depends on the lower bound, set by the uniform prior. To fix this dependence, we redefine the definition of the 2σ upper limit to

align with that of a normal distribution, where the probability drops to $1/e^2$ of the peak probability. Figure 3-3 shows the characteristic logistic curve shape of the Sc flux posterior, modified from the uniform prior.

The integrals above can be numerically approximated with samples from a Markov Chain Monte Carlo. We employ the Constrained Hamiltonian Monte Carlo (CHMC) variant, which efficiently explores the parameter parameter space, using the local slope of the likelihood [95]. Parameters that vary over large scales, such as the norm or line fluxes, are transformed to logarithmic variables. The algorithm, including raw data import and background subtraction, was developed in Python[4]. The code reproduces the same real and synthetic spectra as ISIS [60], given identical data products and model parameters. The algorithm recovers known marginal distributions of a multidimensional test likelihood.

Eight chains run the CHMC algorithm in parallel on the same data[5]. Progress toward convergence to the target distribution is measured with Gelman-Rubin convergence statistic, $\hat{R}$ [42], a measure of how similar the between-chain variance is to the individual within-chain variances. Once convergence is detected ($\hat{R} < 1.10$ for each parameter) the chains continue for at least another 4000 samples each to produce the joint posterior. The upper limit is calculated from the marginalized posterior of the ^{44}Sc flux, as explained above, and we report the bin width at the upper limit as the 'error bar'. (The $\hat{R}$ metric quantifies how similar the within-chain and between-chain variances are. It is meant to supplant diagnosing convergence 'by eye', which is a common practice as well.) For some runs, we tuned the stepsize along certain parameter dimensions to achieve convergence.

For the faint *NuSTAR*-East and *NuSTAR*-West regions, likelihoods for 5 observations totaling 755 ks are assumed independent and added. The data themselves are not binned together in one spectrum; rather each observation is fit using its own Ancillary Response File (ARF) and Redistribution Matrix File (RMF). This method is preferred since different pointings result in different ARFs, and detector degradation and blocking filter contamination over time will change the RMF. However, this approach is most likely overly cautious

[4] `https://github.com/jmrv/scandium`

[5] The analysis computer has eight cores.

with these five datasets, taken within two weeks of each other at the same pointing.

3.3 Results

Our analysis yields upper limits for the three benchmark regions generally lower than Paper C. Table 3.2 shows these results, which in the case of the faintest region halves the previous upper limit.

The discrepancy with Paper C is large in likelihood space: the maximum likelihood drops nearly 10 decades with the '2-σ' ^{44}Sc fluxes of Paper C substituted. Indeed, synthetic spectra with Paper C's upper limit fluxes display a noticeable bump at 4.09 keV, a feature not seen in the actual spectra.

We believe the Paper C limits are too conservative, but revisit where our methods diverge. The three differences between this analysis and Paper C are the methods of inference, amount of data, and the use of backgrounds. As mentioned above, the CHMC algorithm has been verified with known probability distributions. Second, instead of two datasets totaling 100 ks, we employ one with 65% more counts. However, the results persist when the longer observation is pruned to the equivalent of 100 ks. Moreover, with thousands of counts per energy bin, upper limits should already be robust and not scale with the amount of data. The remaining analysis difference is the inclusion of backgrounds, which become comparable with the ^{44}Sc upper limit flux in the continuum-poor region. However, if we neglect backgrounds in our analysis, the upper limits for the benchmark regions do not increase, to within the uncertainty set by the posterior bin size. In this case, the power law norm compensates, disallowing higher ^{44}Sc flux. We believe that background modeling is critical for a faint line search, so we keep the feature in the model.

We now turn our attention to the inner regions of Cas A which show the brightest flux in the *NuSTAR* observation, and, fortuitously, the lowest flux in *Chandra*'s band. Without access to the raw data presented in Paper B, we estimated the size and location of the two brightest emission regions by comparing the contour map from the presentation to a *Chandra* image. This method admits a fair amount of systematic error, but suffices for a well-informed

preliminary investigation. The two brightest central regions from Paper B were extracted from the *Chandra* data using standard CIAO routines; we call them *NuSTAR*-East and *NuSTAR*-West (Figure 3-2). Judging by the public image, the regions may comprise at least 10% of the total ^{44}Ti flux for Cas A.

No line is visible at 4.09 keV in the combined 755 ks spectrum for either *NuSTAR*-East or *NuSTAR*-West, so we apply the same CHMC upper limit analysis as with the benchmark regions. This 755 ks exposure of *NuSTAR*-East provides the best local upper limit of the 4.09 keV ^{44}Sc signal: 1.6×10^{-7} ph/cm^2/s. *NuSTAR*-West yields an upper limit some $N\times$ higher (Table 3.2). As expected, increasing data constrains the upper limit more: 1/5 the exposure produces an upper limit of only 0.9 f_6 for *NuSTAR*-East, though with 3/5 the data, the limit already approaches the full-data limit.

3.4 Discussion

3.4.1 Concordance with Branching Ratios

Coincident γ-ray and X-ray observations of the same decay process will confirm laboratory-derived values for fluorescence yields. As mentioned above, these numbers may change in astrophysical plasmas due to excited or ionized states of the isotope.

Prior to *NuSTAR*, *INTEGRAL*/IBIS achieved the most precise measurement of the nuclear de-excitation lines: $25 \pm 3 f_6$, integrated across all of Cas A. This value translates to $\sim 4 f_6$ in the 4.09 keV Kα line for the 'Full' benchmark region, well below our limit (and that of Paper C). However, if the flux is concentrated in *NuSTAR*-East and *NuSTAR*-West, the upper limits may help confirm the reported branching ratio.

If we estimate that 10% of the extended ^{44}Sc nuclear line emission is concentrated in the small *NuSTAR*-East and *NuSTAR*-West regions, then we would expect $0.4 f_6$ split between them. Our derived upper limits sum to about $2 f_6$, though, implying that half of the emission would need to be concentrated in *NuSTAR*-East and *NuSTAR*-West for our limits to constrain fluorescence yields. However, the *NuSTAR*-East upper limit of 0.13 f_6 represents

a substantial fraction of this expected flux, suggesting either that the extracted region is off target from the true *NuSTAR* signal or that this region may actually constrain branching ratios.

The fluorescing ^{44}Sc atoms could be emitting a higher fraction of Kβ photons, so we apply the same analysis, including a line at 4.46 keV. The data can only constrain the Kβ flux to just below 1.5 f_6.

3.4.2 Line Shape Dependence

The ^{44}Ti ejecta is certainly moving, so we investigate the effect of Doppler shifting and broadening on the upper limits. Using the high resolution *Spitzer* Infrared Spectrograph, Isensee et al. [68] found Si II and O IV emission near the center of Cas A broadened and double-peaked, suggesting velocities of $\sim$ 2000 km/s. If the low-lying ^{44}Ti follows this unshocked material, we would expect similar ejecta velocities.

Our upper limits do not appear to depend on the assumed ^{44}Sc line shape. CHMC runs on the various regions with extreme profiles – either a delta function or emission from a thin 2000 km/s shell – yielded nearly identical results, though, with no significant differences in the inferred flux values for different line shapes. *Chandra*'s medium resolution ($\sim$ 100 eV at 4.1 keV) is far broader than any physical line broadening, so effectively smears the signals. The height ratio of the delta function to the shell, broadened with a 100 eV FWHM response, is about 1.2. Presumably we do not see 20% more flux because the data are sensitive to the whole line profile, not just the peak height.

3.4.3 Future Observations

Chandra evidently cannot see the ^{44}Sc signal, but can any other telescope? We attempt to resolve this question by simulating observations of *NuSTAR*-East for different detector resolutions and exposure times. Sharpening the spectral resolution of a telescope will individuate lines and let faint ones stick out above the continuum.

The simple simulation creates synthetic spectra at varying resolution, and calculates

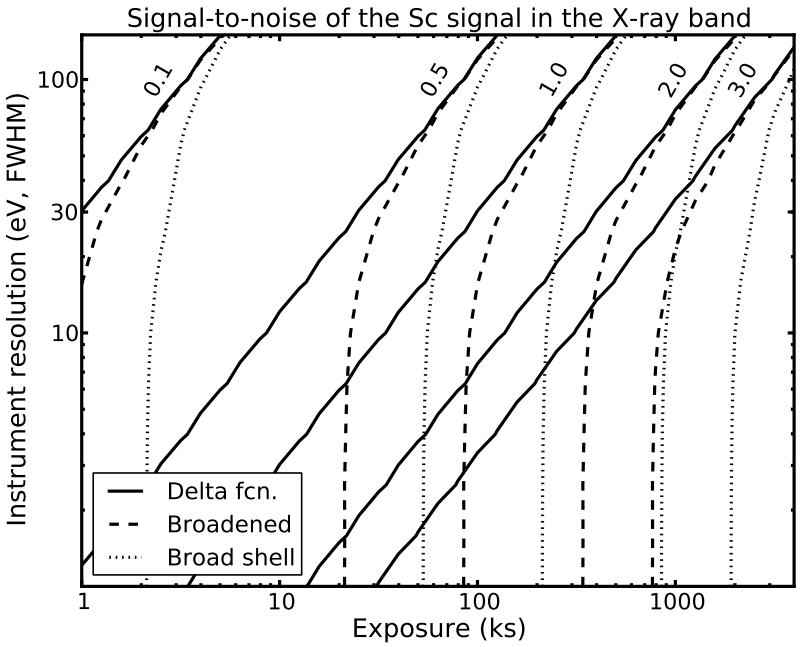

Figure 3-4: Simulations of S/N for various spectral resolutions show that a ^{44}Sc detection will require very long exposures, even for high resolution instruments. The signal has been fixed to the ^{44}Sc flux upper limit for *NuSTAR*-East, and a pristine spectral environment on top of the powerlaw continuum has been assumed. Adding nearby 'confusing' lines increases exposure times by $\sim 15\%$.

signal-to-noise (S/N) values by summing both signal and noise components across suitably defined selected energy bins. We define the signal region to enclose 68% of the ^{44}Sc flux, so that the noise, being a Poisson random variable, is the square root of all non-^{44}Sc components in the signal region. (This is a slightly sanguine region definition, as it assumes we roughly know the signal line shape. Summing only counts within one FWHM of 4.09 keV, though, misses Doppler-split data for higher resolutions.)

Spectral parameters are fixed to a nominal well fitting set for the *NuSTAR*-East data, and the ^{44}Sc flux is set to our derived upper limit, $0.13 f_6$, with the optimistic assumption that *Chandra* just nearly cannot see the signal. S/N values are generated for logarithmically spaced FWHM resolutions, R, at a fixed exposure, T, then scaled by $\sqrt{T}$ to fill out the plane. Only the mean Poisson count rate is calculated for each energy bin (width set to FWHM/7), as we are not interested in the realized random variations. The simulations employ the

Chandra effective area, so the results for other telescopes scale linearly with ARF values, ~ 390 cm^2 at 4.1 keV.

Figure 3-4 shows curves of constant S/N for three different line profiles: the ideal and unphysical delta function, a line broadened by 10 eV (FWHM), and a broadened 2000 km/s shell, meant to replicate the line profile from Isensee et al. [68]. The ideal S/N contours follow the expected $1\sqrt{T/R}$ relationship, dictated by the Gaussian signal which grows as $1/\sqrt{\sigma}$. The broadened signals display two salient characteristics. First, the signals require longer integration times for the same S/N as the delta function because their broadened signal encompasses more noisy spectral bins. Second, the S/N plateaus once the signal width and resolution become comparable, after which point the spectrum 'over-samples' the signal.

Astro-H, with $\sim$ 270 cm^2 of effective area and $\sim$ 7 eV resolution in the signal band would require over 1 Ms of data to reach a S/N of 3. This disappointing result underscores the need for both high resolution and high throughput X-ray instruments in the future. Furthermore, this does not take into consideration the larger pixel size, which will blend the higher continuum from brighter neighboring regions of Cas A into the signal pixel.

We have assumed a pristine spectral environment around the ^{44}Sc line, but small satellite lines may exist. Semi-physical models like `vnei` can inform that assumption, by taking into account the temperature and ionization age of the Ar and Ca species. Fits to the large benchmark regions exhibit no blended lines near near 4.09 keV, due to low ionization age ($\sim 7 \times 10^{10}$ s/cm^3) and temperature ($\sim$ 2 keV). Both Ar and Ca pollute the signal region with lines as τ and kT increase, though, at strengths comparable to the ^{44}Sc signal. While we doubt the material coherent with the ^{44}Sc will exhibit high τ or kT, as it appears to be in the unshocked interior, foreground or background material may emit such lines.

To get a sense for the effect of unknown nearby nuisance lines, we add three nuisance lines randomly within 20 eV of the ^{44}Sc line, at comparable strengths. Rerunning the above simulations – now including these lines in the noise components – decreases the S/N, as may be expected. We find that this nearby nuisance line confusion can require 15% longer exposures. This factor changes depending on the number and strength of the lines, but

illustrates the size of the problem.

Future observations that approach the limits of detectability will require sophisticated statistical techniques. The CHMC algorithm from this work is suitable, though not perfect; a posterior distribution which modifies the prior to produce both lower and upper limits would constitute a detection. A different technique, model comparison, can discern whether the more complicated with-line model is more likely than a without-line model. The nested sampling algorithm can compute this ratio numerically, and can employ CHMC as the sampling engine (Appendix B).

3.5 Summary and Path Forward

We have revisited the radioactive ^{44}Sc signal in Cas A in the *Chandra* band and improved on the limits established in Theiling and Leising [129]. The recent spatially resolved hard X-ray signal from *NuSTAR* allows us to place strict upper limits in the dimmest parts of the remnant. Despite improved knowledge of the whereabouts of ^{44}Ti in Cas A, the X-ray fluorescent signal remains undetected; only an upper limit on the flux can be placed with 750 ks of archival *Chandra* data. This result will likely stand when the *NuSTAR* results are released.

Estimates suggest that the large amount of archival *Chandra* data comes close to constraining fluorescence yields in the *NuSTAR*-East and *NuSTAR*-West regions. Better spectral resolution would greatly improve the ability to detect the line. Our simulations show that, for a focal plane with 10× better resolution, the same signal-to-noise could be achieved in 1/20 the time. We estimate that *Astro-H*, even with outstanding spectral resolution, would achieve a S/N of 3 only after 1 Ms of observation.

Chapter 4

The *Micro-X* Telescope: Design and Commissioning

The *Micro-X Imaging X-Ray Spectrometer* (*Micro-X* hereafter) will image SNRs with an array of microcalorimeters, obtaining fine-grained spectra at each pixel. The first *Micro-X* launch will signal a new generation of X-ray telescopes. High spectral resolution imaging spectroscopy — where each pixel on the focal plane functions as a high resolution spectrometer — will allow for precise understanding of the complicated emission from extended sources in the sky.

The previous chapters demonstrated what can be learned by pushing the best current generation X-ray telescope, *Chandra*, to its design limits. The Cas A knots analysis used two instruments on the observatory – one for high spectral resolution, one for high spatial resolution – and long exposure times to try to uncover faint signals. Next generation instruments, such as *Micro-X*'s microcalorimeter array, will combine both imaging and spectroscopy into a single high-throughput camera.

Micro-X employs the maturing technology of transition edge sensors (TESs) to achieve high spectral resolution at every pixel. The scientific rewards of using TESs come with the costs of operating them; TESs work best at extremely low temperatures, in the absence of any stray magnetic field. These stringent requirements constrain the design of many

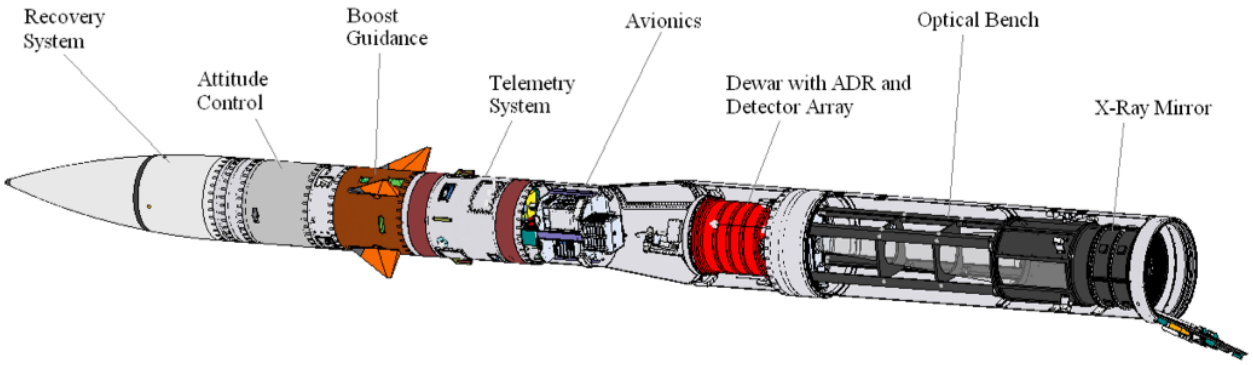

Figure 4-1: The MIT team is responsible for the construction and integration of the Science Instrument. The science instrument (the 3.66 m section from the avionics to the mirror) communicates with NSROC-provided telemetry for timing commands, data transmission, and power.

subsystems. One additional requirement drives the design: *Micro-X* must sustain a violent launch on a rocket. Fine-tuning a design which satisfies the delicacy of the focal plane and the robustness of the surrounding structure has proven to be the most challenging aspect of *Micro-X*'s development.

Micro-X will be borne on a sounding rocket, a launch method that affords a brief five minute glimpse above the X-ray-absorbing atmosphere. Sounding rocket payloads typically cost 2–3 orders of magnitude less than orbital missions, commensurate with the risk of new technologies like *Micro-X*[1].

Standard, NASA-provided support instruments will shepherd *Micro-X* upward: boost guidance (rough steering), attitude control (fine steering), telemetry (data downlink), flight termination (suicide button), and recovery (parachute). These payload sections, seen in Figure 4-1, sit atop the telescope section, which rests on the motor and booster. The motors burn out 45 s after launch and the payload to coasts up to 280 km in a parabolic arc, affording an estimated 330 s of observation time above 160 km before tumbling back through the atmosphere for a parachute landing, some 80 km uprange from the launch site.

[1]Additionally, at writing, NASA favors small sounding rocket missions for funding, partially in response to the 'lean startup' model of Space-X (Private presentation, Mason Peck, NASA Chief Technologist, Dec. 2012).

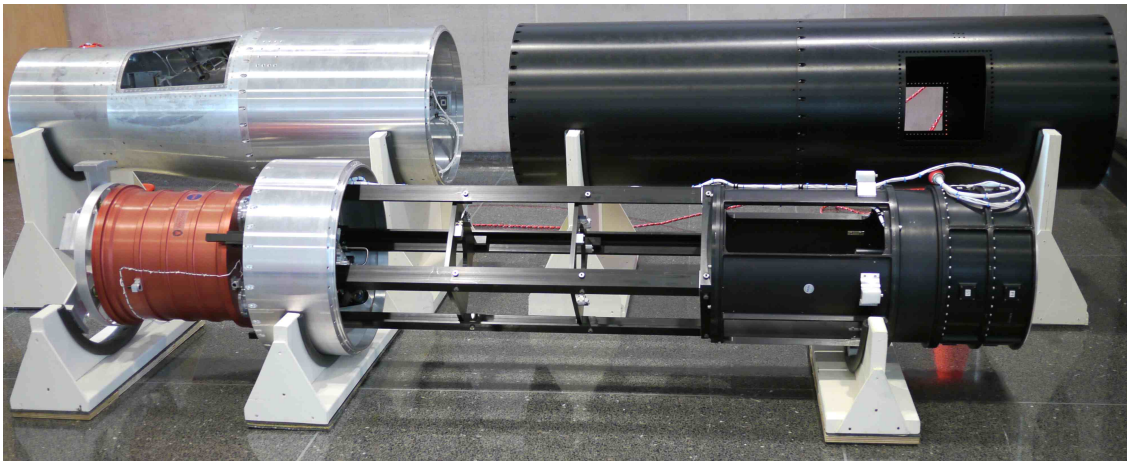

Figure 4-2: The *Micro-X* payload is being commissioned at MIT. The telescope is shown here removed from the skin. X-rays enter from the right at the grazing-incidence mirror and are focused over 2.1 m to the focal plane, housed in the cryostat (red). An optical bench (black) rigidly connects the mirror to the bulkhead (short skin section) and cryostat. The readout electronics are housed in the top of the transition skin piece (background, left).

Before delving into each of the subsystems, I will briefly describe how they fit within the *Micro-X* instrument. Figure 4-2 shows an exploded view of the hardware. X-rays come in from the right, grazing off the mirror foils and focusing down over 2.1 m through the optics bench to the focal plane. The cryostat and optical bench are joined mechanically at the bulkhead (short skin section). Two homegrown computers at room temperature read signals from the focal plane, which pass through two amplification stages at 50 mK and 2 K, and send them to the rocket's telemetry system. A compact, removable refrigerator section, called the Insert, houses the refrigerating magnet, amplifying electronics, magnetic and thermal shielding, and the focal plane, suspended by Kevlar strings for thermal isolation.

Micro-X was conceptualized by Enectali Figueroa-Feliciano, who has overseen its development as Principal Investigator since 2006. The vast majority of the design was engineered by Patrick Wikus, and drawn heavily from the sounding rocket experiment *X-ray Quantum Calorimeter* (*XQC*) [108]. The payload synthesizes expertise of microcalorimeters from NASA Goddard Space Flight Center (GSFC), sensitive cryogenic electronics readout from the National Institute of Standards and Technology (NIST), sounding rocket payload de-

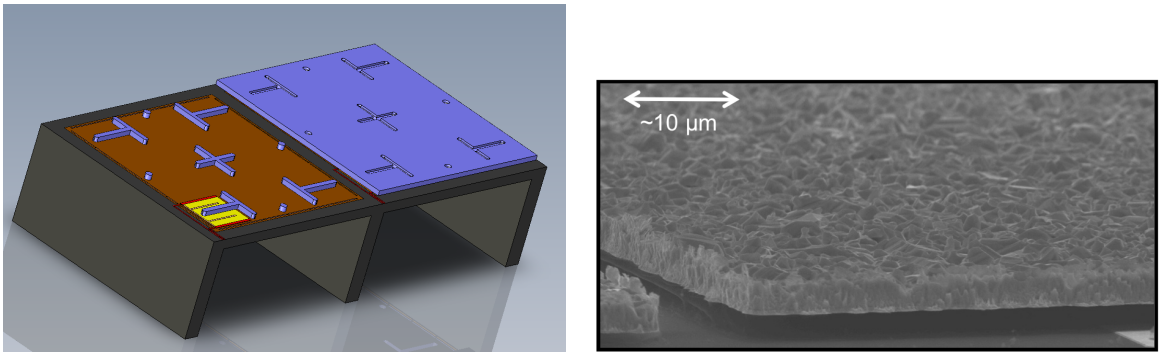

(a) (b)

Figure 4-3: The *Micro-X* pixels are microcalorimeters, which can measure the energy of individual photons absorbed into the pixel. At left, a CAD model of two scale model pixels shows one with the absorber removed to show the underlying TES (yellow). At right, the dual requirements for the absorber – low contact area and large fill fraction – produce overhanging corners, shown in this scanning electron microscope image, and made possible by advanced manufacturing techniques.

sign from the University of Wisconsin at Madison, and flagship X-ray experiments from the Massachusetts Institute of Technology (MIT), among other institutions.

The rest of this chapter will detail the individual subsystems which our team at MIT, the *Micro-X* PI institution, has integrated and tested. Components relating to the dewar will be given the most attention, commensurate to the time spent working with the cryostat over 6 years. 'Cryostat' and 'dewar' are used interchangeably for the red cylinder in Figure 4-2, as well as its diminutive backronym 'LMO', for the Lightweight MilliKelvin Observatory.

The payload development has been documented in several publications [141, 140, 144, 143, 34, 142, 27, 33].

4.1 A Novel Focal Plane

Micro-X employs microcalorimeters to precisely measure each incoming photon's energy. As implied by their name, microcalorimeters can detect tiny amounts of heat – in this case, the amount deposited by a single X-ray photon.

Microcalorimeters exhibit inherently higher energy resolution than CCDs: less than 2 eV

FWHM at 6 keV ($R = 3000$), compared to 150 eV for the ACIS instrument on *Chandra* ($R = 40$) [2]. Whereas microcalorimeters measure temperature increase, a continuous variable, CCDs count discrete electrons, which are subject to Poisson uncertainties. An X-ray incident on the Si of a CCD pixel produces a shower of photoelectrons, each of which requires 3.7 eV of energy to be liberated. This random process is governed by Poisson statistics. (The simple picture is modified by considering the amount of energy which disappears into the phonon system of the Si, known as the Fano factor.)

Microcalorimeters function with four main components: a stable low temperature reservoir, or 'cold bath'; an absorber, as the target for the photons; a weak thermal link between the absorber and cold bath; and a thermometer to read off the heat deposited in the absorber. Figure 4-3(a) shows a schematic of the *Micro-X* pixel-microcalorimeter design. When a photon hits the absorber, its energy quickly converts to thermal energy, raising the temperature of the absorber. This influx of energy then leaks out to the cold bath, producing a characteristic pulse-shaped trace in the thermometer signal (Figure 4-5). We can then back out the energy of the photon via the pulse height.

The concept of cryogenic microcalorimetry is not new. Niinikoski [96] proposed the use of small calorimeters at low temperatures to detect particles in 1974. The first suggestion to apply this technology to X-ray spectrometry came a decade later [93]. Since then, the technology has matured to the point where whole arrays of microcalorimeters have been commissioned for space flight: notably, the *XQC* sounding rocket [91] and the *Astro-E* and *Suzaku* XRS instruments [77, 76].

These previous instruments utilized semiconductor thermistors to sense the small temperature changes. Such devices are high impedance sensors, and consequently exhibit less sensitivity than superconducting devices. (This concept will be quantified below with the logarithmic sensitivity parameter α.)

We use transition edge sensors (TESs) as thermometers, which demonstrate remarkable energy resolution, significantly better than CCDs and even semiconductor thermistors. A

[2] `http://cxc.harvard.edu/cal/Acis/index.html`, `http://cxc.harvard.edu/proposer/POG/html/chap6.html#tth_chAp6`

TES is a superconductor held at its transition temperature where its electrical resistance vanishes[3]. At this point, any small change in temperature results in a large change in resistance, making a very sensitive thermometer. Since TESs are so essential to the design, the terms 'TES' and 'microcalorimeter' are used interchangeably in the field, and in the remainder of this thesis.

The 128 pixel focal plane array was designed by Megan Eckart at GSFC [141]. The design was governed by several constraints: (1) the bandpass of the mirror, (2) the speed of the multiplexing electronics, and (3) the expected performance of the refrigerator.

TESs are easily tuned by choosing appropriate materials for the absorber and thermal link, as the next section will show. In the case of *Micro-X*, we can achieve 99% quantum efficiency over the 0.2-3 keV bandpass by making the absorbers of layered Au and Bi, with thicknesses roughly twice the attenuation lengths (consideration 1). The pulses can be slowed down to address consideration 2 by adjusting the strength of the thermal link. Finally, the transition temperature of the superconducting TES can be tuned with the thickness of the TES layers, so multiple arrays can be manufactured around the expected base temperature of the refrigerator (constraint 3).

The array fabrication process is fickle, but produced several worthy arrays for flight. Results presented in this thesis derive from both a 170 mK transition temperature array, as well as a higher resolution 120 mK array.

I will lay out the basic theory of TESs, following the treatment of [35]; further references can be found therein.

4.1.1 Microcalorimeter Physics

TESs are simple devices. As calorimeters, these devices measure the deposition of heat into a system via a temperature increase. We are measuring individual photons, so the heat arrives in quanta: $\Delta T = E_\gamma / C$. For a known specific heat C, the photon energy E_γ can be

[3]This is the first of seven distinct uses of superconductors in the *Micro-X* payload: TESs, SQUIDS, wire to the focal plane, magnet, magnetic shielding, high temperature strips to route magnet leads across temperature gradients, and level gauge.

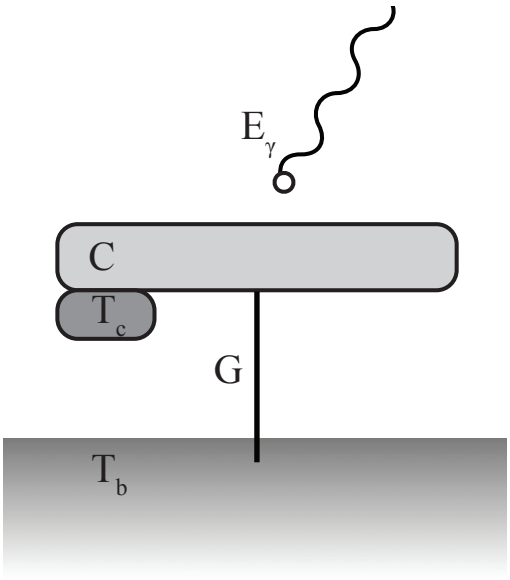

Figure 4-4: A microcalorimeter, in our setup, is simply an absorber with heat capacity C, connected to a superconducting thermometer (TES) with transition temperature T_c, heat sunk to a thermal bath at T_b through a weak thermal link of conductivity G. It can measure the energy, E_γ of an incident X-ray photon.

inferred.

In order to reset the detector back to the operating temperature, a weak thermal link is added between the device and a cold reservoir, or *bath*, as shown in Figure 4-4. The strength of the link is characterized by the thermal conductance G, which has units of power over temperature. Thus, with this energy drain, an incident photon produces a *pulse*, decaying down to the bath temperature after the initial rise in temperature. (The link is weak to allow time to measure the signal.)

The pulse height is proportional to photon energy, while decay time is a property of the system. To see this, we must solve the differential equation for this simple thermal system:

$$C\frac{dT(t)}{dt} = P - P_\ell(T(t), T_b) + E_\gamma \delta(t - t_\gamma), \tag{4.1}$$

$$P_\ell = G(T(t) - T_b) \tag{4.2}$$

On the right hand side of Equation 4.1, P is a constant power applied to the absorber (a

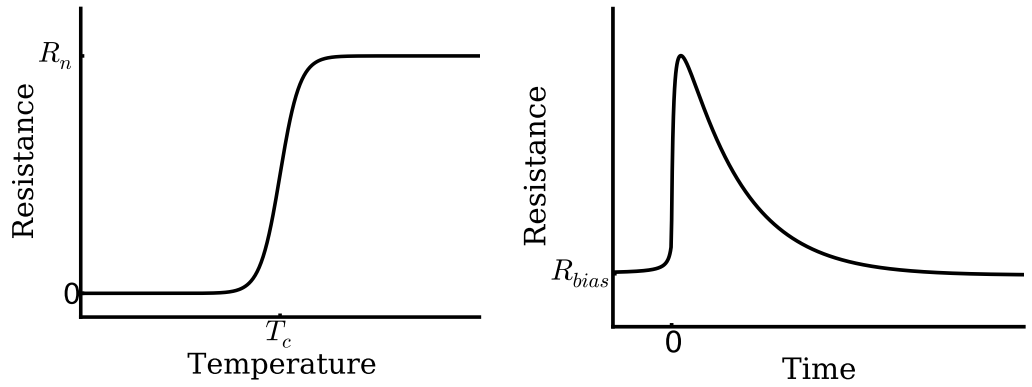

Figure 4-5: The resistance of a TES (a) changes sharply from $R_n \approx 10$ mΩ to 0 Ω over about 1 mK at $T_c \approx 100$ mK. When an X-ray warms up the TES, its resistance rises quickly before the heat drains through the thermal link, creating a pulse shape (b).

good approximation for now of the operating conditions[4]); the link power P_ℓ is negative as it flows out into the bath at temperature T_b; and the photon hits at time t_γ. This system admits the solution for $t_\gamma = 0$ of

$$T(t) = \frac{E_\gamma}{C} e^{-t/\tau_0} + \frac{P}{G} + T_b, \quad t \geq 0, \tag{4.3}$$

where $\tau_0 = C/G$. Thus, the pulse height is proportional to E_γ, while the decay time depends on G.

X-rays thermalize in an absorber via several processes. The above derivation assumes immediate, thorough thermalization, while in reality, the process to convert a photon's energy into heat in the absorber takes time. For the materials typically used in X-ray quantum microcalorimetry, the *photoelectric effect* is dominant (over Rayleigh scattering and Compton scattering). The incident X-ray absorbs into a single atom of the calorimeter, ejecting an electron[5]. This photoelectron rattles around with a high cross section, forming a local ball of energetic electrons. As the ball grows, the mean electron energy decreases and phonon creation becomes more favorable, further increasing the ball size, as the phonons have longer

[4]The definition can be generalized to $G \equiv dP_\ell/dT$, and the link is often described by $P_\ell \propto (T(t)^n - T_b^n)$.
[5]For X-rays of a few keV incident on bismuth or gold, only the M shell or higher is accessible.

mean free paths. The mean energies of all the particles in the absorber soon equilibrate to a higher value, signaling thermalization.

TESs self-regulate with electrothermal feedback (ETF), an elegant exploitation of the devices' simple properties. One way of maintaining a TES right in the superconducting-to-normal transition is with a voltage bias, whereby a constant, tuned voltage applies Joule heating power to the TES of $P(T) = V^2/R(T)$. In this configuration, the temperature will find an equilibrium value as the Joule heating power equilibrates with the leakage through the weak link. Both the power source and sink depend on temperature – as $\propto 1/T$ and $\propto T$, respectively – so perturbations can be schematized by the following relations:

$$T\nearrow \Rightarrow \quad R\nearrow \Rightarrow \quad P\searrow \Rightarrow \quad T\searrow \tag{4.4}$$

$$T\searrow \Rightarrow \quad R\searrow \Rightarrow \quad P\nearrow \Rightarrow \quad T\nearrow \tag{4.5}$$

This control loop affects the time constant of the system. Intuitively, the circuit creates a virtual heat link by reducing the amount of heat which would have been applied. This picture can be verified formally by replacing P in Equation 4.1 with the Joule power, generalizing G as dP/dT, and linearizing the resulting differential equation. The resulting effective time constant is

$$\tau_{eff} = \frac{\tau_0}{1 + \frac{\alpha P}{TG}} = \frac{1}{1/\tau_0 + 1/\tau_{ETF}}, \tag{4.6}$$

where

$$\alpha \equiv \frac{T}{R(T)} \frac{\partial R(T)}{\partial T} \tag{4.7}$$

is a measure of the sharpness of the superconducting transition, and

$$\tau_{ETF} \equiv \frac{\alpha P}{TG} \tau_0. \tag{4.8}$$

Equation 4.6 shows that ETF speeds up pulses. The *Micro-X* pulse decay constant is about 2 ms.

Engineering the decay time is important because *pileup* can be an issue for slower systems.

Pileup results when a photon hits a detector that is still recovering from the temperature spike of a previous photon. While the double-pulse can be modeled to recover both photons' energies, the exercise is rarely undertaken and the uncertainty will be larger, regardless. The resolution of a microcalorimeter depends on the available bandwidth[6], which pileup events reduce by shortening the usable portion of the pulse trace. Current methods to create high-resolution spectra segregate events based on their quality (uncertainty); thus, pileup events are usually thrown out. The SNR targets *Micro-X* will observe generally are not bright enough to cause concern for pileup, but the onboard calibration source (Section 4.6) must be throttled so as not to create too much pileup.

The resolution of a device is limited by the noise from phonons and the resistor's Johnson noise. Phonon noise arises from the thermal link to the cold bath: inherent power fluctuations between the two temperatures exhibit a nearly flat spectrum, but the temperature response of the detector rolls off above $1/2\pi\tau_0$. Johnson noise results from the random walk of electrons across the TES resistor, and is flat across the frequency space. It can be shown that, for a TES with these two noise sources, in the limit $\tau_{ETF}/\tau_0 \gg 1$, the energy resolution is

$$\Delta E_{fwhm} = 2.4\sqrt{4k_B T^2 C \frac{1}{\alpha}\sqrt{\frac{n}{2}}}. \tag{4.9}$$

Here n is the exponent from the parametrization of P_ℓ. While the resolution improves with higher α TESs, it does not depend on the amount of ETF[7].

4.1.2 The *Micro-X* Array

The *Micro-X* focal plane is a 12×12 array of square absorbers, each attached to their own TES. The microcalorimeters were specially designed within the bounds of the readout capabilities and science requirements. The design, summarized here, was published by Wikus et al. [141], and is detailed internally[8].

[6] We define bandwidth as the frequency range over which the S/N is greater than 1.

[7] This second result does not derive from Equation 4.9, but rather a more thorough analysis of the S/N, which ends up to be independent of the feedback loop circuit.

[8] `https://wikis.mit.edu/confluence/display/MICROX/TES+Detector+Array`

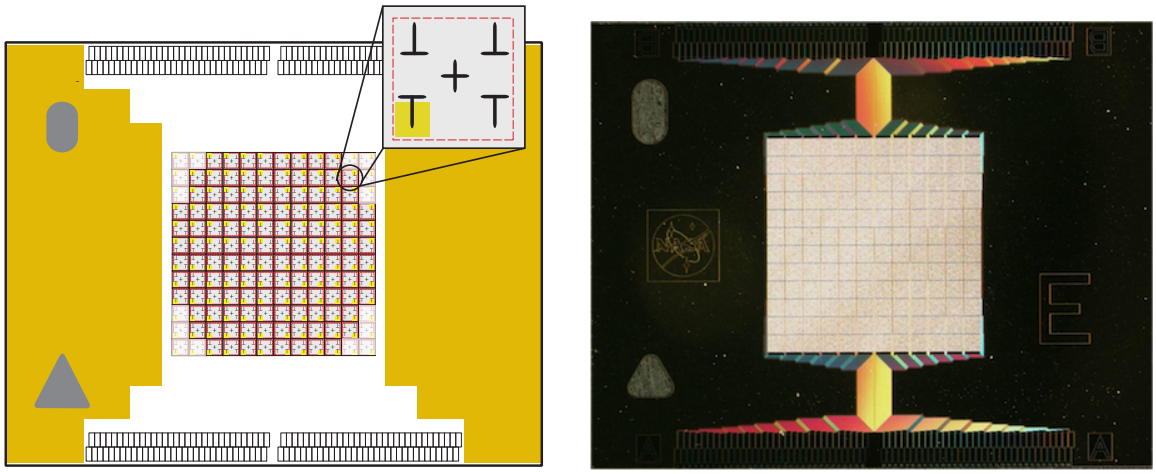

(a) (b)

Figure 4-6: The *Micro-X* array was designed (left, shown with a pixel and its TES) and manufactured (right, complete with circuit leads) at NASA GSFC.

Silicon microfabrication techniques, which rely on precise deposition and etching of microscopic layers on silicon wafers, and which have driven the boom of the semiconductor industry, are utilized to manufacture the *Micro-X* array. The GSFC foundry has become expert at producing specialized micro-structures and circuits, including microcalorimeter detector arrays.

The pixels are macroscopic, unlike a CCD: 590 μm on a side, which translates to about $1'$ on the sky. With 10 μm left between neighboring absorbers, the fill fraction is 97%. The focal plane layout and the actual array are shown in Figure 4-6. Due to the multiplexing readout design, only 128 pixels are read out; the dummy pixels are grayed in the schematic.

The absorbers are made of 3.4 μm of bismuth deposited onto 0.6 μm of gold. The Au helps the large absorber thermalize quickly with its low C, while the Bi provides structural stability and increases the quantum efficiency (photon stopping power), while not adding significantly to the heat capacity. Several absorber designs were tested, and this configuration of metals displayed the best thermal properties of the structurally stable design. The attenuation lengths of the materials change across the *Micro-X* band of 0.3–2.5 keV, so thicknesses were optimized to provide over two attenuation lengths of thickness for absorption across most of the band, while being constrained by the desire for thin, low-C absorbers. Consequently,

the quantum efficiency exceeds 98% through 4 keV, after which it drops, as higher energy X-rays easily penetrate the thin absorbers. The large absorbers are supported by T structures ('stems'): Figure 4-3(b) shows the Au/Bi absorber cantilevered over the Si floor and wires. One T-stem connects to the TES.

The TESs represent the culmination of many years of research and development at NASA GSFC. Each 140 μm square is a bilayer of molybdenum and gold. The superconductivity of the bilayer relies on the proximity effect for superconductors — also known as the Holm-Meissner effect — which 'tempers' the transition of the superconductor (Mo, 920 mK) with a normal metal (Au, 0 mK). (The theoretical is similar to the theory of *pn* junctions in semiconductors.) The T_c can be tuned by adjusting the thicknesses of the two layers; this is often done empirically, bracketing the desired T_c with a range of thicknesses on different chips from the same silicon wafer run. Noise is mitigated inside the TES by $\sim$ 350 nm zebra stripes of Au, which partially obstruct the flow of current from one side of the TES to the other[9].

Thermalization of the pixel proceeds from the absorber to the refrigerant salt as follows. The TESs and T-stems are connected to a 1 μm thick silicon nitride (SiN) membrane, which is supported by a grid of Si that follows the pixel borders (the grey table legs in Figure 4-3(a)). The Si grid is made from the larger monolithic chip on which the array and circuit leads rest. This chip is mounted on ball bearings inside the focal plane housing – the Front End Assembly, or FEA – and connected to the floor of the gold-plated housing with gold wire heat sinking bonds. The FEA is bolted to the cold finger, which runs to the salt pill. Control electronics ensure the FEA maintains a stable temperature, to provide a predictable cold bath for the TESs.

The final element required for a microcalorimeter is the weak thermal link. The link is created and tuned by perforating a square boundary in the SiN membrane, seen in the pixel inset of Figure 4-6[10]. G scales with the boundary of the perimeter, as this is the available

[9]This feature, like many aspects of TES design, was discovered by blind trial-and-error, and the quieting effect has not been fully explained theoretically.

[10]The area of the T-stems can also be adjusted to tune G.

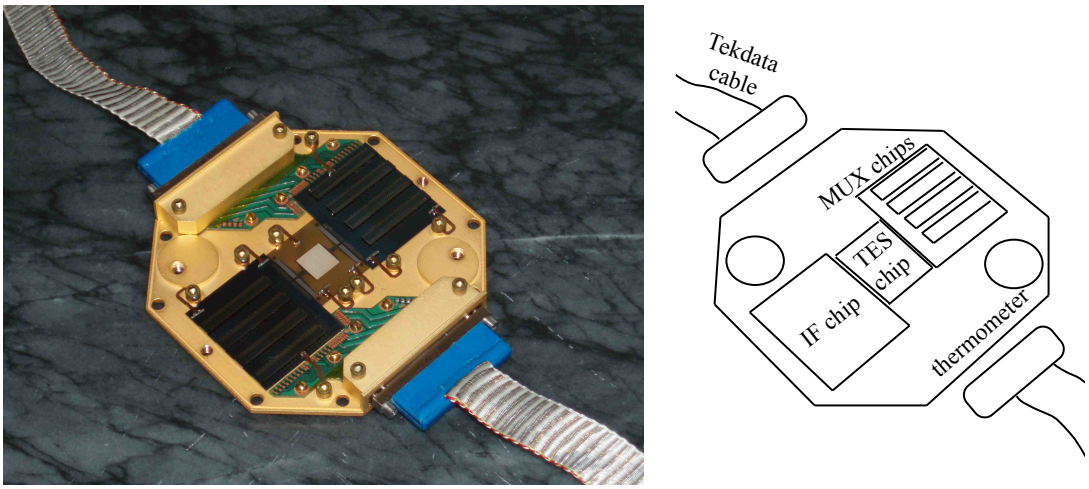

(a) (b)

Figure 4-7: The *Micro-X* FEA (Front End Assembly) houses the detector array within a gold-plated Mg box, in addition to the cold SQUID chip electronics and the control thermometers. The chips are kinematically mounted, to prevent over-constraint. The FEA lid, with its accompanying X-ray filter, is not shown.

boundary for phonon transmission, which thermalizes the device. The holes, therefore, slow the process of thermalization; their size and density can be bracketed over a range to produce the desired G value. Thus, the membrane, TES, T-stems, and absorbers form a single unit within the perforation boundary, warming together after a photon strikes the absorber before cooling through the perforations.

The G value was adjusted to produce pixels with 2 ms time constants. This timescale is slow for most TES designs; usually pileup concerns favor fast designs, and laboratory electronics can handle high readout frequencies. (The rule of thumb is to sample at least 10 times on the way up the fast X-ray pulse rise.) However, the rocket telemetry system is slower than laboratory readout systems, requiring slower pulses, even if only a fraction of pixels are recorded. To convert between timescales and thermal conductance, the telemetry constrains sampling to roughly 30 kHz per pixel, enforcing $G < 100$ pW/K.

4.1.3 The Front End Assembly

The Front End Assembly (FEA[11]) houses the focal plane array in a light-tight, thermally stable enclosure. The *Micro-X* FEA is made of lightweight magnesium, and gold-plated for better thermalization. The bottom portion, which holds the array, multiplexing readout chips, control thermometers, and breakout boards for the circuit, is shown assembled in Figure 4-7. A lid fits on top, with the 50 mK filters mounted in the center. The Airborn NK-Series 85 pin connectors for each science chain have been custom potted into a Tekdata weave cable to read out the TES signals. The entire FEA mounts to a cold finger which connects to the salt pill.

Each chip rests on three tungsten carbide balls, kinematically constrained by spring clips at six points of contact. Aluminum wire bonds connect signals across the various chips once the chips have been installed. I assembled the 120 mK flight array at GSFC and oversaw the wirebonding, taking advantage of the flight-quality wirebonding facilities and expertise. The bonds were tested by pulling to non-destructive forces, giving confidence they will survive launch.

The cold electronics in the FEA are described in Section 4.4.

4.2 The Cryostat

We must operate the focal plane at very low temperatures to achieve the high resolution possible with TESs. While the superconducting material is necessary for the operation of a TES, it is not sufficient. Indeed, superconductors exist at temperatures well above 100 K now, which can be cooled with cheap liquid nitrogen. However, these materials would perform poorly as X-ray detectors, since the photons carry only tiny amounts of energy. Recall Equation 4.9 from the previous section: $\Delta E_{fwhm} \propto T$, where T is the operating temperature. Thus the energy resolution only degrades at higher temperature. (The assumptions that validate Equation 4.9 no longer hold for high temperature superconductors, but the intuition

[11]The term has been adopted, inexplicably, from the 'nose' structure of a car, which holds the engine and front suspension.

is illustrative: the tiny heating signal must dwarf the random thermal noise in the system.)

Several refrigeration methods routinely reach temperatures in the mK range [106]:

1. **Evaporative cooling** of liquid helium (LHe) relies on the latent heat of evaporation to extract heat from a system by pumping vacuum on the vapor above the liquid. H. Kamerlingh Onnes, the discoverer of LHe superfluidity, which transitions at 2.17 K, set the lowest temperature record in his time via LHe evaporative cooling with large battery of pumps in 1922: 830 mK. Modern efforts have improved little on this mark, especially as novel refrigeration methods have been discovered.

2. **Adiabatic demagnetization**, explained in detail below and proposed in 1926, was the first novel cooling method after Onnes' roadblock with LHe evaporation. The method employs the magnetic disorder entropy of special materials' magnetic moments, cooling the sample to the mK range by absorbing heat into the slowly randomizing spin system, controlled by a changing magnetic field. Adiabatic demagnetization refrigerators (ADRs) offer only a single shot of cooling.

3. **Dilution refrigerators** rely on the spontaneous phase separation of ^{3}He and ^{4}He. The sample cools as it provides the energy necessary to move pure ^{3}He, supplied by a continuous cycle, across the phase boundary, 'diluting' the ^{4}He-rich phase. The continuous cooling method is the favored cold stage in most laboratories. Prevailing designs are not usually rocket-friendly, being large in volume, mass, and power usage to accomodate the continuous-cycle pumps and tiered radiation shields. However, the Planck space mission successfully employed a lighter, more compact, open-cycle (the ^{3}He was vented into space) dilution refrigerator which operated at 100 mK continuously for over years.

4. **Adiabatic nuclear demagnetization**, like (electronic) adiabatic demagnetization above, exploits the disorder entropy of *nuclear* magnetic moments. This method uses a dilution refrigerator as a first stage, and can cool to the microKelvin regime.

5. **Pomeranchuk coolers** exploit phase properties of ^{3}He to effect cooling via solidification of the liquid phase. Such technology has been superseded by dilution and nuclear refrigerators; like an ADR, the cooling method is single-shot.

ADRs have emerged as the favored cooling technology for space-based low temperature missions. *XQC*, *Astro-H*, and the proposed but unfunded *International X-Ray Observatory* utilize ADRs, primarily for their compactness.

4.2.1 Adiabatic Demagnetization

ADRs exploit paramagnetic salts as the refrigerant. The ion spins in such salts align in the presence of a magnetic field, but are otherwise randomized in a state of high disorder. The entropy of these materials depends on temperature and magnetic field, creating easy avenues to lower temperatures via higher fields.

The mathematical description of paramagnetic salts is simple, at least under the assumption of noninteracting magnetic dipoles. The entropy as a function of magnetic field and temperature is [106]:

$$S(B,T)/R = \bar{\mu}B/T(\coth(\bar{\mu}B/T) - \bar{J}\coth(\bar{\mu}\bar{J}B/T)) + \ln\left(\frac{\sinh(\bar{\mu}\bar{J}B/T)}{\sinh(\bar{\mu}B/T)}\right). \tag{4.10}$$

The quantity has been scaled to the ideal gas constant, $R = 8.31$ J mol^{-1}K^{-1}. The barred constants are:

$$\bar{\mu} = \frac{\mu_B g}{2k_B} \tag{4.11}$$

$$\bar{J} = 2J + 1, \tag{4.12}$$

where the Bohr magneton, $\mu_B = 9.27 \times 10{-}24$ J T^{-1}, and the Boltzmann constant, k_B, are physical constants, while the Landé factor, g, and the angular momentum, J, are properties of the salt.

Equation 4.10 reveals key characteristics of the system, which will enable its utility as a refrigerator.

1. If T can be kept constant while increasing B, then S will decrease. ADRs use a heat switch connection from the salt to a cold reservoir to dump the heat generated during magnetization.

2. If S is held constant while B decreases, then T decreases. This adiabatic step makes the refrigerator work.

3. In order to control T to a constant value, the system must compensate by reducing B as S increases. Heat leaks to the cold stage require the magnet to control the temperature like so.

These steps in sequence constitute a cycle of the ADR. Unlike continuous refrigeration methods, ADRs only afford a single shot of cooling, which lasts a few hours in the case of *Micro-X*.

Being an adiabatic process, the entropy from Equation 4.10 stays constant, ensuring the relationship between the initial (i) and final (f) states:

$$S(B_i, T_i) = S(B_i/T_i) = S(B_f/T_f) = S(B_f, T_f), \tag{4.13}$$

which implies

$$B_i/T_i = B_f/T_f. \tag{4.14}$$

However, this relationship in Equation 4.14 breaks down as $B_f \to 0$, since the internal interactions between the magnetic moments become important, invalidating the original assumption of non-interaction.

The available cooling power of the salt (per mole) at the end of a demagnetization cycle is simply

$$Q(B_f) = \int_{T_f}^{\infty} T(\partial S/\partial T)_{B_f} dT. \tag{4.15}$$

Thus, a larger final field (i.e., more current left in the magnet at base temperature) provides a larger cooling power.

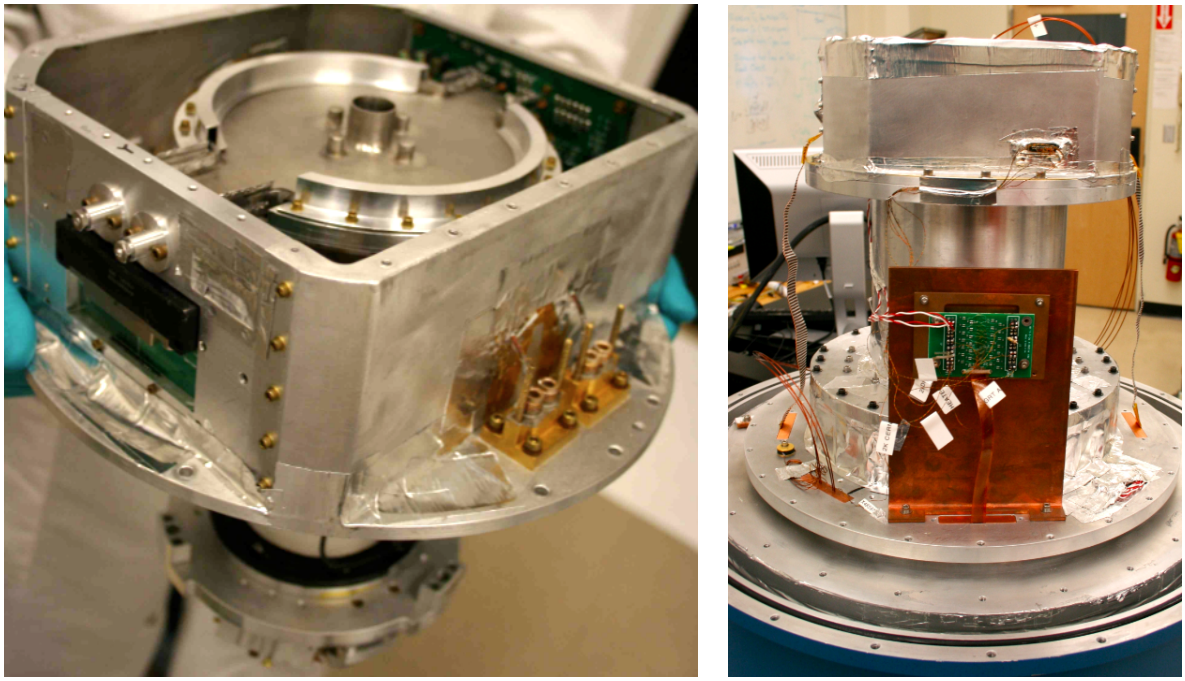

(a) (b)

Figure 4-8: The Insert (a) is the removable heart of the *Micro-X* ADR. The magnet (white, bottom) hangs off a supporting plate, where high-current connectors (gold, right) mount to supply it with current. The focal plane signals feed through a light-tight connection (black, left) in the 2 K thermal shield (silver box). A magnetic shield (top) encases the focal plane; X-rays enter from the top of the image. The Insert can be accommodated (b) in a larger laboratory dewar, Gonzo, which has a liquid nitrogen tank to provide longer hold times. The Insert is mounted on an adapter, and housekeeping wires are routed through a breakout board (rectangle, center).

4.2.2 The *Micro-X* ADR

A compact, modular, and lightweight ADR serves as *Micro-X*'s cooling system to maintain a cold bath for the TES array. The ADR design is drawn largely from the *XQC* rocket [108], which has launched successfully multiple times. The *Micro-X* ADR required new design features from the *XQC* template, including magnetic shielding, a different Kevlar suspension for the salt pill, and a more capacious interior to accommodate different electronics.

The ADR is composed of a superconducting magnet, a salt pill connected to the cold stage, and a heat switch, all of which is mounted together on a removable stage called the Insert. Layers of thermal shielding and resistance paths surround the 'heart' of the ADR to

keep it at an operating temperature below 2 K. The Insert (Figure 4-8(a)), which also holds magnetic shielding and control electronics, mounts within the central bore of a cylindrical liquid helium (LHe) tank. G10 tubes (Section 4.2.8) hold the LHe tank suspended within a vacuum vessel. Thin metal thermal shields surround the helium can and G10 cylinders, cooled by the LHe vapor.

Being removable, the Insert can be swapped into other dewars. The *Micro-X* ADR was tested through initial iterations in a large lab dewar, named Gonzo for its blue exterior, which allowed much longer cold runs due to the larger volume of cryogens. Figure 4-8(b) shows the ADR installed in Gonzo. (In addition to a larger volume of LHe, the dewar features a liquid nitrogen tank to function as an intermediate heat sink, taking advantage of the cryogen's larger heat capacity.) The Insert can also be mostly assembled and altered outside of LMO, a significant convenience.

As with most electromagnet designs, the *Micro-X* superconducting electromagnet is a wire coil wound around a cylinder. Being superconducting, the wire dissipates no heat at LHe temperatures, allowing higher current and, consequently, higher magnetic field. The magnet is a custom multifilament NbTi/Cu double solenoid, commissioned from Cryomagnetics, Inc. The primary coil surrounds the salt pill, while the shorter secondary 'bucking' coil is positioned just aft-ward, wound in the opposite direction to partially null the high magnetic field near the FEA. This configuration was modeled and optimized with the electromagnetic finite element model (FEM) package in COMSOL.

The superconducting electromagnet can apply fields up to 4 T, requiring about 9.5 A of current at maximum field, which is considered low-current for most ADRs. The low-current design eases the requirements on the power supply and the current leads, and results in an inductance of 22 H.

A primary design concern for superconducting magnets is quench protection. Quenches occur when a point of the magnet rises above the superconductor's critical temperature and turns resistive ('normal'), creating a runaway process as the point's resistive heating warms surrounding areas. Such local heating can arise from eddy currents or Joule heating of nearby

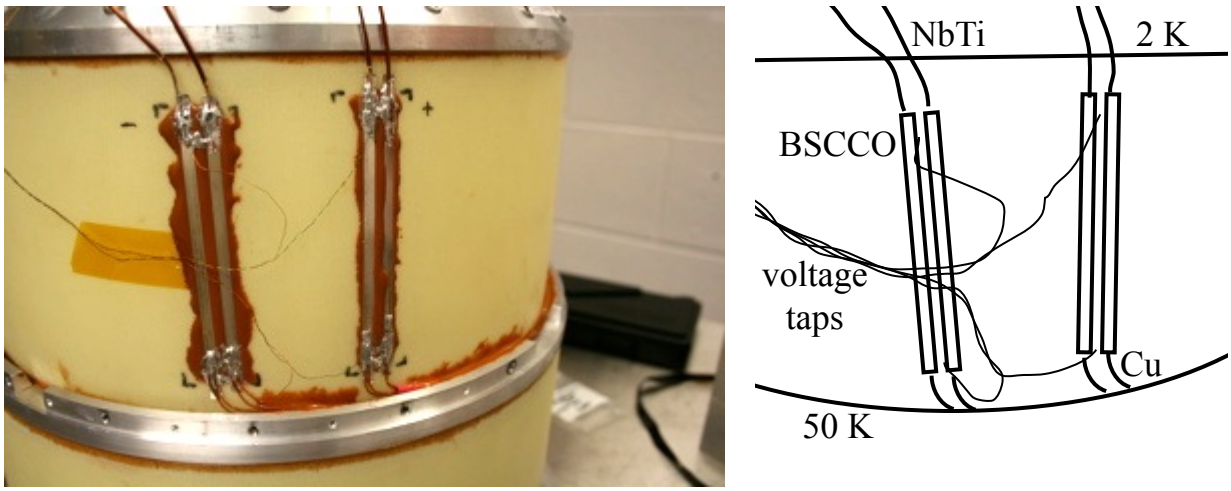

(a) (b)

Figure 4-9: High-temperature superconducting strips (BSCCO) conduct the magnet current between the 50 K shield (lower ring) and the LHe tank (top ring). The strips are epoxied onto the G10 with Armstrong A12. Two-wire voltage taps are placed at each solder joint. The copper wires coming from 300 K can be seen epoxied into the 50 K ring for heat-sinking.

resistive elements in the circuit. A quench dumps the stored magnetic energy into heat as the current drops to zero, resulting in large voltage spikes in addition to cryogen boil-off. Diodes in the magnet housing were included to help dissipate the energy released in the case of a quench. Several voltage taps, located at several points along the magnet circuit, monitor the induced voltage across the magnet coil as the current slowly ramps up (a few mA/s), to monitor for a quench. As the magnet is located in the vacuum space of LMO, not immersed in the LHe as in most applications, conduction is the only cooling method after a quench. The ADR has sustained several quenches, due both to initial 'settling' of the coils during the magnet commissioning, as well as design flaws (poor heat sinks, e.g.) which were later corrected.

The current path from the hermetic connector down to the magnet is broken into several stages. Eight copper wires (four redundant leads for supply, four for return) run from the top lid connector along the thermal shields, heat-sunk with aluminum tape. They are further heat sunk with epoxy at the final thermal stage before the LHe tank. At this point, the copper

leads hand over the current to strips of high-temperature bismuth-strontium-calcium-copper-oxide (BSCCO) superconductor[12]. BSCCO has a transition temperature above the liquid nitrogen (LN) boiling point. A solder joint is made with the low-melting-temperature solder alloy Cerrolow 117, which is similar to Wood's metal (Figure 4-9). Four BSCCO tape strips are epoxied along the G10 section, again for redundancy. Moving colder, NbTi wires carry the current along the LHe tank lid and down its side, terminating in plugs which mate to high-current connectors (HCCs)[13] on the Insert. The HCCs connect, finally, to the magnet, completing the circuit. The HCCs were chosen for their low voltage drop: $\sim 300\ \mu$V at a current of 4.25 A and a temperature of approximately 1.7 K (two HCCs serve each magnet wire, thus each see half of the full current).

During flight, the current is supplied by the ADR Electronics subsystem (Section 4.2.12, and powered by the 28 V rocket batteries. In lab tests, we employ a Cryomagnetics CS-4 Bipolar Superconducting Magnet Power Supply. The CS-4 can be controlled by software via a GPIB interface, and proved invaluable while the legacy ADR Electronics were debugged.

4.2.3 Salt Pill

Micro-X uses the refrigerant salt Ferric Ammonium Alum (FAA), chosen for its high magnetic moment, $J = 5/2$ and ready availability. Salt pills for ADRs are 'grown' layer by layer with the intent to produce a few number of large crystals, to aid in thermal conduction. Several salt pills have been grown at MIT, drawing on experience from GSFC. Sarah Heine oversees the salt pill growth and construction.

The current salt pill performs well, achieving base temperatures less than 45 mK (the lowest calibrated value of the GRT). However, it exhibits an abnormally low thermal conductance (G) of 0.3 μW/mK between the heat bus and the paramagnetic salt. This value was measured at 60 mK; the G value increases to 0.6 μW/mK at 100 mK. Typical values of similarly designed salt pills are 5 times higher, around 1.5 μW/mK at 60 mK [140]. We

[12]Bruker HTS GmbH provided free samples of the BSCCO. The outer metal cladding, which can conduct heat between the two stages, was not removed, as no increase in heat load was observed. The *XQC* ADR uses exposed high-temperature superconductor strips, etched with acid.

[13]SP4N and BL4N from Multi-Contact AG.

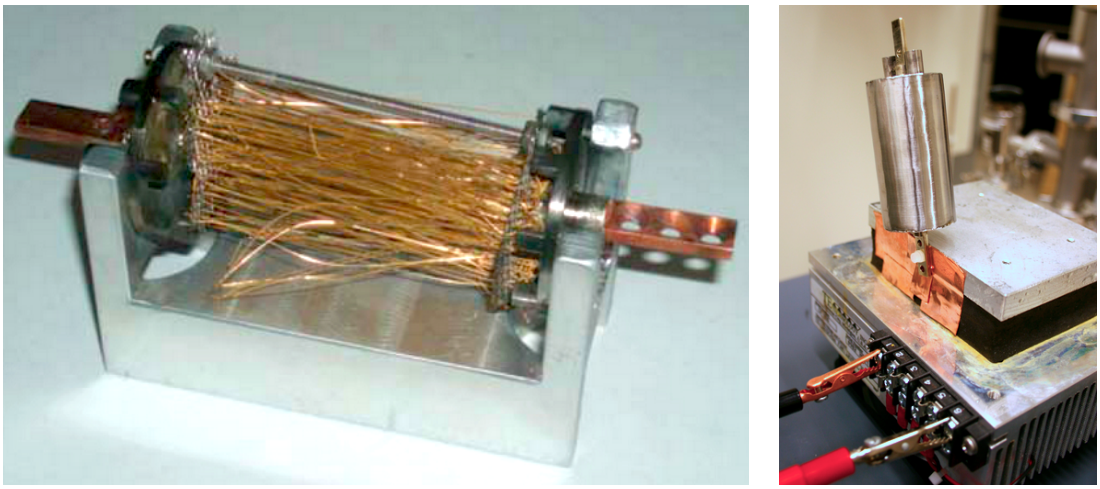

(a) (b)

Figure 4-10: The salt pill for the ADR was grown at MIT. Left, the gold wire cage, around which the salt will grow, has been brazed to the two endcaps and heat buses. Right, the FAA salt crystallizes within the stainless steel housing, connected to a cold plate.

have attributed this deficiency to substandard gold plating or brazing.

The G value can be measured by recording the temperature jump when a known heating power is applied. (The jump will be lower for higher thermal conductance, and the lower thermal resistance allows better equilibration between the two points.) The heating slope reveals the size of the external heat load; this procedure is described in Section 4.2.11 below.

Another version of the salt pill housing with the wire cage has been carefully treated at GSFC according to their in-house brazing practices, after being machined and assembled at MIT. The RRR of the original MIT salt pill cage was 100, while GSFC has produced components with $2-5\times$ higher values. After the growth of the salt in August 2013, the subsequent performance tests will determine whether the G value has improved. The current salt pill is suitable only for TES arrays with transitions greater than about 150 mK, due to the low G value.

The construction of the salt pill proceeds as follows, in its current form[14]. The salt is grown around 150 annealed gold wires (0.2 mm diameter), fixed in a cylindrical shape by a

[14]Detailed notes can be found at the internal site `http://space.mit.edu/figueroagroup/internal/micro-x/Subsystems/Salt\%20Pill\%20Files/SaltProcedure.html`, without the backslashes.

mesh cage. After the cage is built and cleaned, end-caps and copper heat buses are brazed to the wire bundles at either end with non-superconducting silver solder. Figure 4-10 shows the brazed inner cage. To protect the braze joints from the corrosive salt, the entire cage is gold plated and the joints painted with Stycast 1250 FT. A thin SS tube can now be welded to the cage, enclosing the volume for salt growth.

The salt is grown inside the pill housing by dissolving FAA in dilute sulfuric acid, pouring the solution into the housing via spouts on the top end-cap, and cooling the assembly to aid crystallization. The last step is shown in Figure 4-10. The salt solution must be kept between 37° C and 39° C, which requires an immersion circulator. Syringe pieces, including filters to keep out small undissolved crystals, are kept in the same bath to avoid premature crystallization before transfer to the pill. The salt pill can be filled with new solution every few hours. After draining the previous batch of solution from the salt pill (not all of the liquid solution will turn to salt), the fresh warm solution can be added. The housing is then attached to a cold plate, to facilitate crystal growth. 75 grams total of FAA are crystallized; the process takes 2–3 weeks.

The FAA is extremely corrosive. It can eat through certain alloys of SS, even, so care must be taken with material choices (316 SS) and skin contact. In addition, the salt corrodes above 39° C, so is susceptible to damage just from body heat. When not in use, the salt should be stored in a cool environment.

The role of the gold wires is to communicate between the thermal systems of the salt and the heat buses. The buses can then exchange heat from external systems, such as the FEA. *Micro-X* uses a somewhat unconventional two-bus system due to space constraints of the instrument. One bus is available for the heat switch to clamp, with the other forms a dedicated link to the FEA.

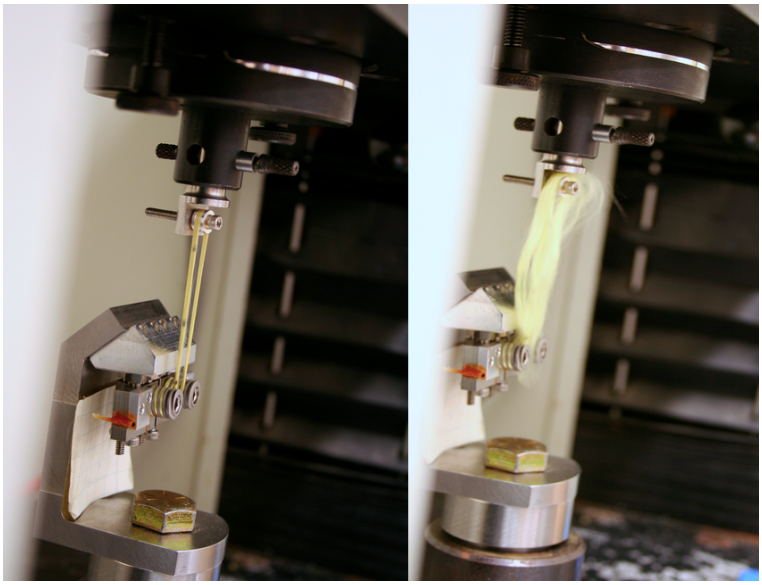

Figure 4-11: The Young's modulus of a Kevlar suspension piece is tested to failure. The screw and roller (top) would mount to the cold stage. The Kevlar passes around corkscrews before being affixed with epoxy (bottom left).

4.2.4 Kevlar Suspension

The entire cold stage — including the salt pill, the thermal buses, and the FEA — is suspended from the 2 K heat sink of the Insert by seven Kevlar[15] loops. Each string has a thickness — or more precisely, a linear mass density — of 4560 denier[16]. The suspended mass weighs approximately 260 g, kept lightweight by machining many parts out of Mg. Materials with low heat capacities were used wherever possible, as well as gold plating to increase thermal conductivity.

Six of the loops kinematically constrain the six degrees of freedom of the suspended mass, and one applies a constant pre-tensioning force through springs (seen variously in Figures 4-11, 4-12, and 4-23). The six non-tensioning The suspension was designed to withstand accelerations in excess of 200 g, and to have a high resonant frequency of a few hundred Hz, to minimize heating from rocket motor vibrations. As Section 4.3 will detail, the built suspension exhibited lower-than-expected resonance frequencies.

[15]The Kevlar-49 was donated by Dupont, and is a registered trademark.
[16]1 denier = 1 g/9000 m. The silk industry can be blamed for the bizarre unit.

In addition to vibration isolation, the Kevlar provides thermal isolation, as the fibers are a poor conductor of heat. A study of an alternative fibrous material, Zylon, performed in the laboratory dilution refrigerator, revealed Kevlar to be superior in this aspect [139]. The total heat load from the Insert through the Kevlar suspension system was measured at 0.9 μW, during commissioning in Gonzo. For comparison, roughly the same power leaks through the science chain connector weave, while stray radiation loading contributes at most 0.3 μW.

The Kevlar suspension mounts are a novel design, so required several iterations to perfect. Figure 4-11 shows one of the mounts up close. Kevlar winds around corkscrews on either end of the loop, with the intent to reduce the force on the glue joints which secures the string. The epoxy, used liberally elsewhere in LMO, is Armstrong A12.

To test the stiffness of the suspension, strings on individual mounts were pulled to breaking strength. (The Kevlar breaks before the metal parts or glue joints.) In initial designs, the Young's modulus of the system was determined to be 40% of its theoretical maximum, dictated by the characteristics of Kevlar. The suspension mounts were identified as the source of missing stiffness, as videos of the pull tests showed the corkscrews to move. With key parts of the suspension machined out of titanium, the elastic modulus increased to 65% of the theoretical maximum. The loop design may not let all the Kevlar fibers participate in load-bearing, resulting in this lowered stiffness. Additionally, we found that the suspension stiffens after pre-tensioning, presumably once the individual fibers have settled and the string has been allowed to slip along the corkscrews.

Sarah Heine performed the battery of suspension stiffness tests.

4.2.5 Heat Switch

The heat switch connects the salt to the ~ 2 K heat sink during magnetization, then releases it for demagnetization once the salt has thermalized. Its role is simple but crucial: to provide a high quality thermal connection to the salt pill bus. The current design, seen in Figure 4-12, is mounted to the Insert. A new version, mounted to the LHe tank with stronger heatsinking,

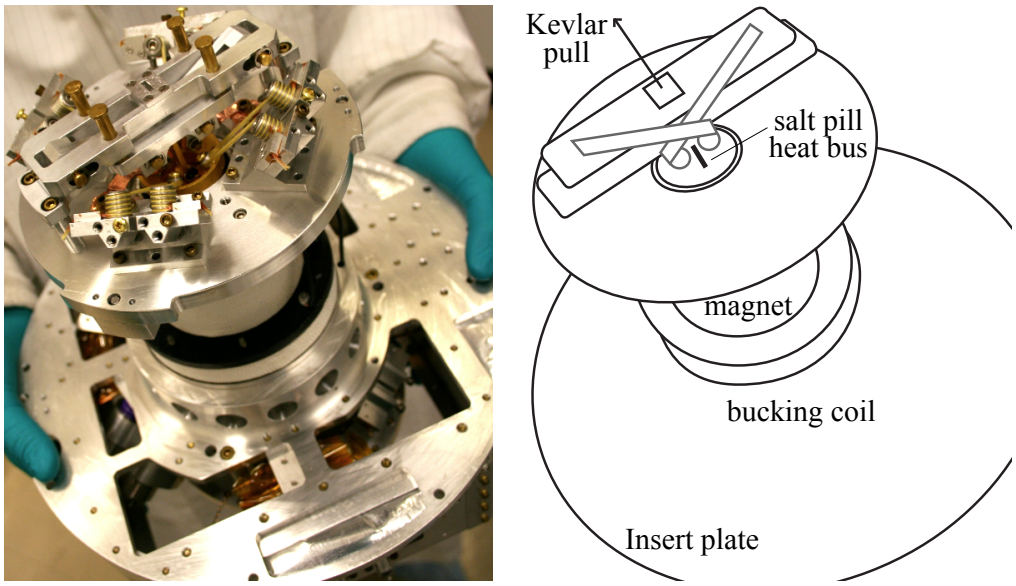

(a) (b)

Figure 4-12: A fore view of the Insert shows the heat switch, which operates by pinching the salt pill bus via a scissor mechanism (enhanced in the wire diagram). The magnet (including the bucking coil) is visible, as well as Kevlar suspension pieces which hold the salt pill in the bore of the magnet.

is being commissioned to correct some flaws of the current heat switch.

The simple clamp design engages when pulled by a Kevlar string, pinching the forward heat bus with copper rollers. At cryogenic temperatures, the force of a connection between two surfaces is often more important than the pressure; thus, a large contact area is not required to equilibrate the salt with the heat switch. The rollers are heat-sunk to the Insert with copper foils, and the Insert is bolted to the LHe tank with about a dozen bolts, creating the heat flow path.

The actuator mounts to the top vacuum lid of LMO, requiring low-conductivity Kevlar to make the connection from room temperature to 2 K. In addition, bellows must provide light tightness along the Kevlar's path through thermal shielding.

As the heat switch mounts to the Insert, and the Insert is removable from LMO, the component detaches from the Kevlar string/bellows assembly. The connection between the two halves of the heat switch assembly is a T piece which locks into a slot by 90 degree

rotation.

The best conductance of the heat switch was measured to be 0.6 μW/mK at 1.6 K, comparable to other designs [106]. The connection strength allows for sufficiently rapid thermalization of the salt during a magnet cycle.

However, the heat switch performance is often compromised by flaws. First, the jaws can become jammed in the four poles meant to guide the mechanism in a vertical motion. Second, the design is highly dependent on a precise length Kevlar string, which is difficult to achieve when knotting and gluing by hand. Lastly, the turnkey locking mechanism makes inserting the Insert quite difficult if the alignment is not perfect.

The forthcoming redesign will address each of these in turn, and will be commissioned in September 2013.

4.2.6 Liquid Helium Tank

A tank of LHe provides a low temperature starting point for the ADR cycle. The annular cylinder of welded aluminum holds 6.5 L. The LHe is pumped to below 2 K (typically 1.6 K), leaving about 3 L of superfluid LHe after the initial pumpdown. Just under 24 hours of hold time below 2 K are available with this design[17], due to a $\sim$ 85 mW heat load from the rest of the dewar components, cooled to intermediate temperatures. Being a compact space instrument with only modest hold time requirements, a LN tank is not incorporated into the design to intercept the heat load on the LHe.

The LHe tank has only two openings: one for a level gauge, the other for the fill tube. Both are sealed to the tank with indium. The level gauge is a strip of superconductor with $T_c > 4.2$ K; its measured resistance depends on the length submerged in LHe. The fill tube serves the dual purpose of liquid cryogen input and boil-off vapor exhaust. The tube is flexible, being made up of individual welded bellows, which allow for thermal expansion/contraction mismatches between thermal stages. The bellows' corrugation creates

[17]The hold time can be improved by more efficient heat transfer between the cooling He vapor and the dewar, with copper clamps, for instance. Such improvements are not exigent for a 5 minute observation, but would significantly improve the quality of life for graduate students, who would no longer need to fill LHe twice a day.

turbulence in the exhaust, improving the heat transfer between the metal and vapor.

To prevent LHe sloshing up the fill tube during deceleration in flight, a sintered stainless steel disk[18] with 1μm pore size caps the cold end of the fill tube. He vapor can pass through the disk, but liquid cannot: the design was tested by flipping LMO upside-down with a full LHe tank. This porous plug has a hole in the center for easier LHe filling and initial pumping, but is stoppered with a rod for flight.

As with any LHe dewar, we cool LMO first by filling the LHe tank with LN, as it is cheaper and has a higher heat capacity than LHe. The initial cooldown to 77 K primarily uses the N vapor, to keep the cooling rate of the Insert near 1 K/minute. This comfortable, empirical rate has prevented thermal gradients from stressing connections in the dewar with mismatched thermal expansion properties. After thermalization, the LN can be removed. LHe can then be filled at the same 1 K/minute rate until the Insert cools to 4 K and liquid can collect in the LHe tank without evaporating immediately.

4.2.7 Pumping Valve

As mentioned above, the LHe bath will be pumped to achieve a colder starting temperature for the salt demagnetization. In flight configuration, a flexible bellows connects the LHe fill tube to a skin-mounted assembly with a hermetic valve, a pressure relief valve, and a pressure gauge. This assembly — shown in Figure 4-14 — is the pumping valve. A flexible, ribbed vacuum hose completes the vapor path along the length of the launch rail to a pump[19], connected to the pumping valve via a slide-in press seal[20].

Just before launch, the pumping valve is closed. Exhaust pressure builds above the LHe bath until the payload passes 100 km, roughly 70 s into flight. At that point, the residual atmosphere of Earth is low enough to act as a mechanical pump, and the pumping valve opens, releasing vapor exhaust. The valve remains open until the payload falls below the

[18]Supplied by Mott Corporation.

[19]We use one Edwards XDS10 dry scroll pump, with displacement ~ 7 ft^3/min. This could be augmented with a second pump if the pump line length proves too long to achieve a satisfactory base temperature.

[20]Such a seal is termed 'flyaway', but is designed to disengage easily, so the pump doesn't take a ride.

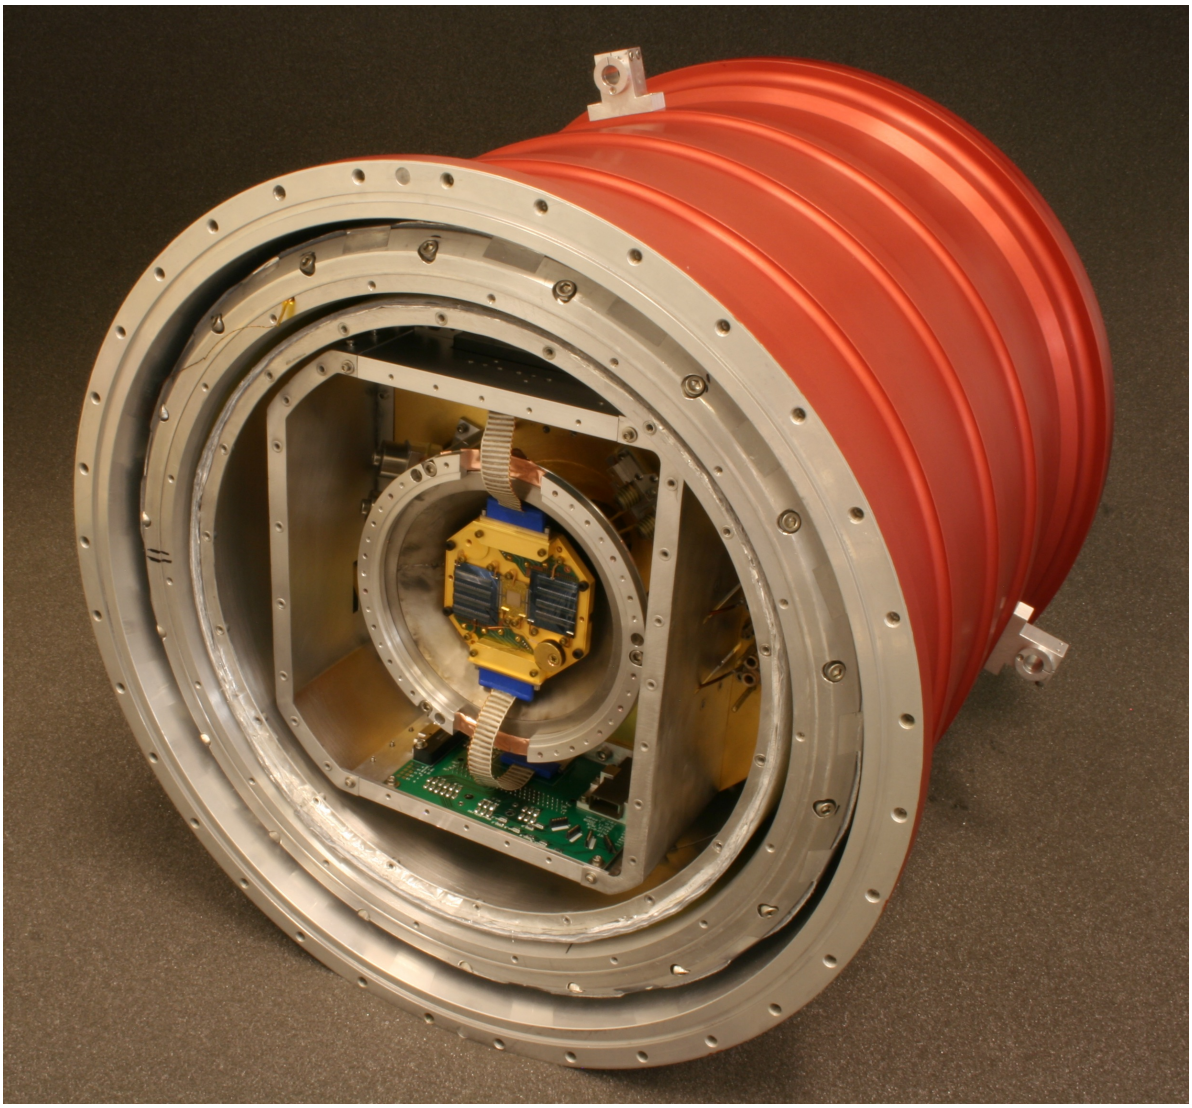

Figure 4-13: Layers of thermal shielding protect the TES array to maintain a $\sim$ 100 mK operating temperature. Moving outward: FEA housing (gold octagon), magnetic shield, 2K radiation shield (rounded box), 50 K shield, 150 K shield, and finally the vacuum jacket (red) at 300 K. All lids were removed for this image. The dewar is approximately 40 cm tall and wide.

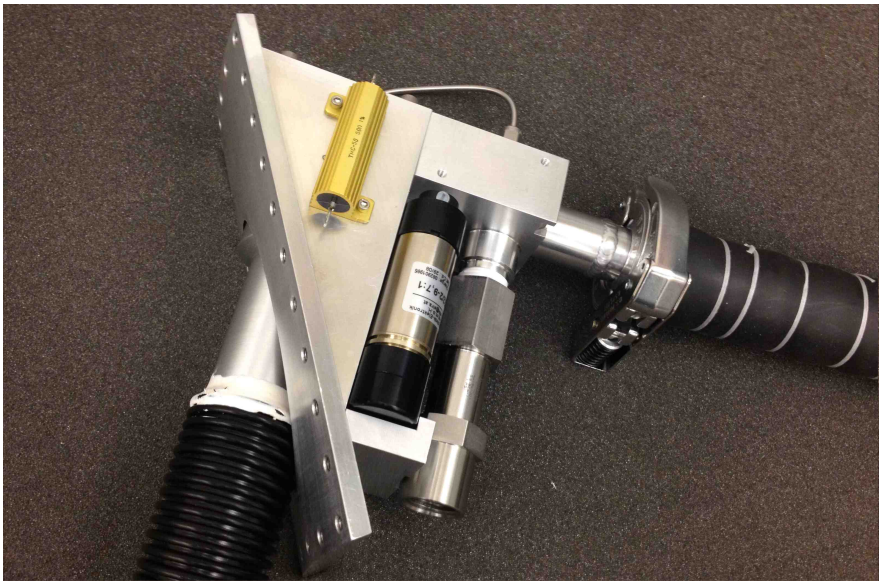

Figure 4-14: The pumping valve for the LHe bath mounts to the skin. From left to right: the flyaway vacuum tube which connects to the pump; the skin mount; heater (gold); the compact gate valve motor; over-pressure relief valve; vacuum hose to LHe fill tube, with damping rubber layer.

same altitude on the descent, maintaining a low LHe bath temperature. The valve remains closed on the ground after impact, naturally backfilling until recovery.

4.2.8 G10 Structure

The LHe tank and Insert are suspended from a series of shells, made of laminated glass and epoxy. This off-the-shelf material comes in several grades, and the *Micro-X* team calls it colloquially by the grade used in the payload: G10. The yellow material can be seen in Figure 4-9.

The G10 has three functions: thermal isolation, structural support, and vibration isolation. The practicability of the novel *XQC* payload, from which *Micro-X* adopted the nested shell support, relies heavily on the advantageous material characteristics of G10 that allow it to take on these three roles.

G10 has favorable thermal properties. The thermal conductivity along its length decreases from 0.85 W/m·K at 300 K to 0.3 W/m·K at 50 K [146]. (The thermal conductivity

of Al is over 200 times greater.)

G10 also has excellent longitudinal strength, while being fairly weak transversely, when pushed against its surface. For this reason we call the G10 cylinders *thrust tubes*, as their strength is critical in the thrust (z) axis. The stiffness translates into a natural spring constant, discussed in Section 4.3 as the second stage of vibration isolation between the focal plane and the rocket skin.

Each tube wall is 0.5 mm thick, glued into grooved aluminum rings[21], again with Armstrong A12. Pull tests demonstrated the success of the gluing procedure: the G10 broke before the glue joint.

Figure 4-16 shows the nested 'Russian doll' configuration of the G10 segments and the thermal shields. The outer G10 assembly bolts to the vacuum jacket at 300 K, hanging off the forward end. The inner G10 assembly sits on the aft ring of the outer assembly. The LHe tank hangs from the inner G10 assembly. This serpentine heat conduction path through thermally resistive G10 helps isolate the LHe from hot surfaces at room temperature.

4.2.9 Vacuum System

With few exceptions, a cryogenic system necessarily requires a vacuum environment to prevent rapid boil-off of the cryogens. Gas between two surfaces of different temperatures will create a thermal short, and will help to equilibrate the temperatures, an undesirable outcome for a focal plane at 100 mK.

The *Micro-X* vacuum vessel is a ribbed outer cylinder[22], sealed on the aft and forward ends by two O-ring connections to plates. The red anodized structure can be seen is Figure 4-13.

Several ports are cut into the plates, sealed by hermetic O-rings. Two 'military-style' circular connectors carry the signals for each science chain through the forward plate, as well as another military connector for housekeeping signals. A hole for the heat switch Kevlar

[21] Tubes with non-overlapping seams were custom ordered from Stevens Products

[22] Wallops machined the aluminum cylinder, as the facilities are specialized to cut and modify skin-like cylindrical forms.

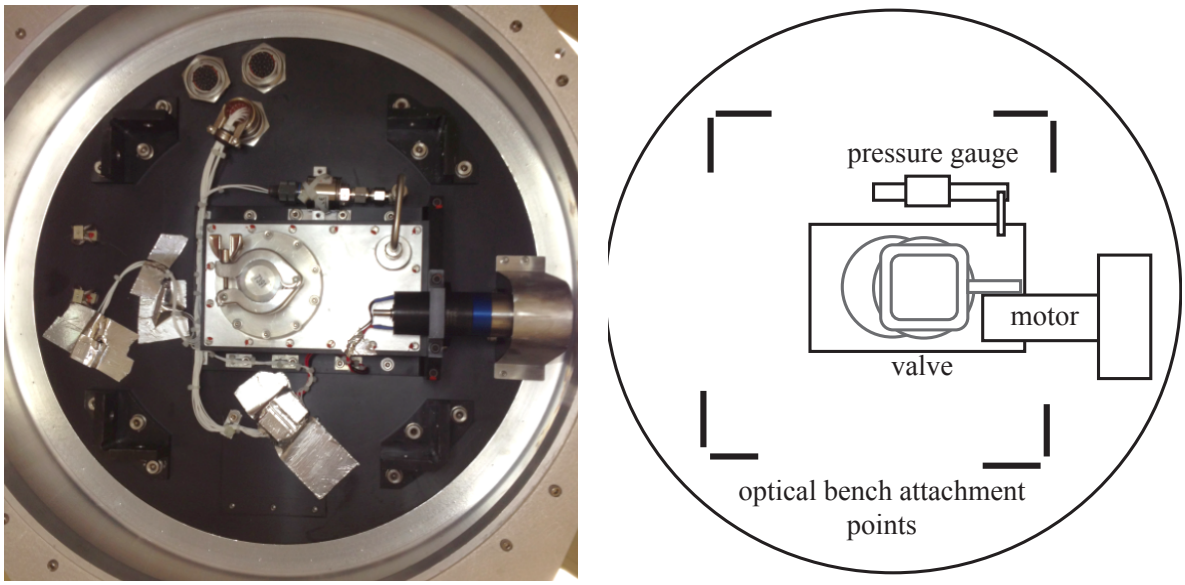

(a) (b)

Figure 4-15: The gate valve seals the cryostat from the optics section; it is opened at a suitable altitude to the vacuum of space. The grey components (vacuum port, simplified gate valve) in the wire diagram are hidden within the housing in the image.

sits in the middle of the forward plate. The fill tube, too, is accommodated on the forward plate. Vacuum is pulled through a standoff port on the forward plate, a 25 mm opening, sufficiently large for the small volume of the vacuum vessel. Finally, a single port on the aft plate lets X-rays through to the focal plane along the optical path. This opening on LMO connects to the bulkhead via bellows, which accommodate the movement of the dewar on the skin dampers.

The optical path is sealed by the *gate valve* on the aft side of the bulkhead, in the evacuated mirror section (Figure 4-15). The commercial valve is accommodated by custom housing with pressure sensors, and driven by an external motor.

Before cooldowns, LMO is evacuated[23] to approximately 10^{-6} bar before the application of cryogens. With the thin X-ray filters installed, a slow pump rate is required (see Section 4.8). A leak test must be performed before cooldown with a He leak checker.

After the dewar components have thermalized from a cooldown, the vacuum pump is

[23]We use a Pfeiffer HiCube80 turbo station, which is a turbo pump backed with a dry diaphragm pump, which has a displacement of ~ 0.3 ft^3/min.

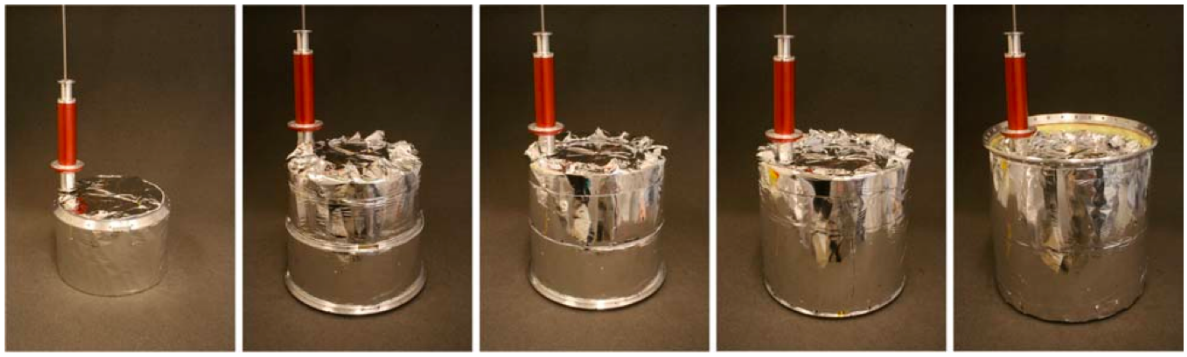

Figure 4-16: The layers of thermal shielding are built up like a Russian doll. From left to right: the LHe tank with MLI and the fill tube; the inner G10 assembly, with the LHe tank hanging from the top; a 50 K thermal shield is added; a 150 K shield is added; the assembly is dropped into the outer G10 assembly.

no longer needed, as *cryopumping* is stronger at such temperatures. Cryopumping is the phenomenon where molecules adsorb to cold surfaces, naturally pulling a vacuum. For LMO, cryopumping provides a vacuum of $\sim 10^{-9}$ bar.

4.2.10 Thermal Shielding

Layers of thermal shielding ensure that the helium stage and focal plane stay cold as long as possible. Figures 4-16 and 4-13 show the nested configuration of the various thermal layers. The *Micro-X* cryostat, like any cryostat, was designed to maximize the time between liquid helium fills. Liquid helium is expensive, and calibration runs require long periods of undisturbed data-taking.

In addition to physical thermal 'shorts', we must be wary of radiative infrared heat. Therefore, each thermal stage – the 100 mK focal plane, the 2K stage, the inner and outer G10 assemblies – is enclosed in a (infrared) light-tight box. (Pieces of metal tape are useful to completely seal seams and holes.) The sensitive components within each shield live in a radiation bath the temperature of the shield, with no direct line of sight to hotter shields. Only X-rays can penetrate the thin aluminum windows at each thermal shield (Section 4.8).

Two thin (0.5 mm) welded Al shells encase the G10 assemblies, heat-sunk to the G10's

metal support rings. These shields, topped by lids, intercept the cooling power of the He vapor by clamping to the fill tubes. The shields operate at about 55 K and 155 K, and are monitored by diode thermometers on the aft and forward lids for each shield. A box mounted on the Insert encases the FEA, suspension, and cold electronics in a 2 K environment. And the FEA, of course, does the same for the focal plane array, at 100 mK.

The outer shields are wrapped in multi-layer insulation (MLI). The crinkled, reflective layers of MLI are meant to trap stray infrared photons. Each shield is wrapped in two 20-sheet CoolCat 2NW MLI blankets from RUAG Space GmbH, and affixed by low emissivity Austrian Aerospace tape.

4.2.11 Heat Loads

When commissioning a new dewar, the first thing one does, after leak testing, is measure and diagnose heat leaking into the system. (Then, usually, the thermal design is iterated until the heat leaks are suppressed to an acceptable level.)

Heat loads on the FEA are measured by demagnetizing the salt fully and measuring the locally linear heating slope[24]. The parasitic heating load travels through the Kevlar suspension, the science chain signal weave, and via infrared radiation from the surrounding 2 K Insert. A calibrated power (in the few μW range) can then be applied, resulting in a steeper slope. If the measurements are performed in quick succession over small temperature ranges, the heat capacity of the cold stage cancels in the slope ratio, admitting direct measurement of the parasitic power. Radiative heat loads can be disentangled from this measurement. The suspended mass sits at a higher temperature than the Insert (by a few mK) when isolated from the heat switch at zero field, so the radiative loading can be determined given a value for G.

The heat load from the Insert on the FEA is 2.1 μW, an acceptable level for the hold times required for long calibration runs of the array. Radiative loading contributes 0.3 μW, while the remainder is split nearly evenly between the Kevlar suspension and science chain cables.

[24]The heat load on the LHe tank can be calculated using the hold time and latent heat of vaporization.

The two ribbon cables carrying focal plane signals are constructed of aramid string, woven around 43 twisted pairs of 50 μm NbTi wire[25]. The wire is superconducting, so conducts no heat from 2 K; however, the aramid weave, the CuNi cladding of the wires, and the Kevlar trussing do.

The thermal conductance, G, can be obtained with a similar setup – often as part of the heat load measurement. Again, a known heating power P is applied, but in this case, the ΔT above the base temperature is measured. The conductance is simply $G = P/\Delta T$. This measurement technique can be used at many points in the instrument: the link between the FEA and salt, across the heat switch, between the Insert and LHe tank, and even for the TES pixels.

4.2.12 Temperature Control

The temperature of the focal plane is controlled by a PID loop. The cold bath of the FEA must provide a stable temperature platform to better than 1 μK for high resolution spectra. Additionally, the ADR electronics must quickly return the focal plane to operating temperature after the heat-up incurred from launch vibrations. Figure 4-17 shows the performance of the *XQC* control electronics over the course of the flight.

Micro-X uses spare legacy electronics from *XQC*, which have been recommissioned at MIT. The system controls the focal plane temperature by controlling the voltage across the superconducting magnet, which is operated with a few hundred mA of residual current. The current drains slowly over time as the ADR converts the stored magnetic energy to thermal cooling power[26]. Controlling the magnet voltage — instead of the current — exploits the magnet's inductance to filter any power supply noise. The system provides housekeeping data as well, such as the LHe level and the Insert temperature. It also can ramp the magnet up to and down from full field.

The control circuit, shown in Figure 4-18, operates as a lock-in amplifier to read the temperature from a GRT. A 5 V sine wave (a) is converted to a small oscillating drive

[25]The cables were custom-made by Tekdata.

[26]Parasitic resistance also causes the current to drain.

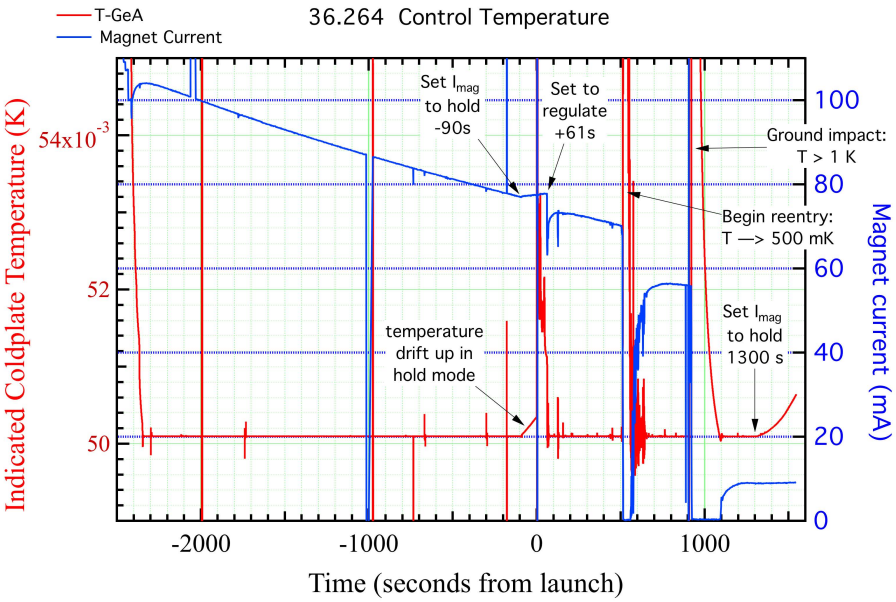

Figure 4-17: The time profile of the temperature control during launch will mimic *XQC*'s, shown here. Ahead of launch, the magnet current bleeds to keep the temperature constant. The focal plane warms up during launch, but recovers once the ADR electronics are commanded to 'regulate' mode.

current by a capacitor (b). The voltage drop across the GRT is amplified (c) and summed with a nulling signal based on a set-point resistance (d); the difference is further amplified (e, f). The set-point resistor determines the operating temperature of the FEA. A PID controller (g) operates off of this error signal, becoming the nulling signal at the amplifier input. A second loop is needed to compensate for the phase shift of the excitation signal — due to stray wire capacitance — which would otherwise skew the phased resistance measurement and produce ripple in the error signal. This 'quadrature loop' operates at a phase of $\pi/2$ to the control loop (i, j, k).

The PID controller outputs a control voltage across the magnet leads. The voltage error is integrated (l) to control two redundant MOSFET power supplies (m, n). These MOSFETs must have closely matching characteristics, otherwise the system will rail. A 'too cold' error signal ramps the magnetic field up with a positive voltage, while a 'too hot' signal ramps

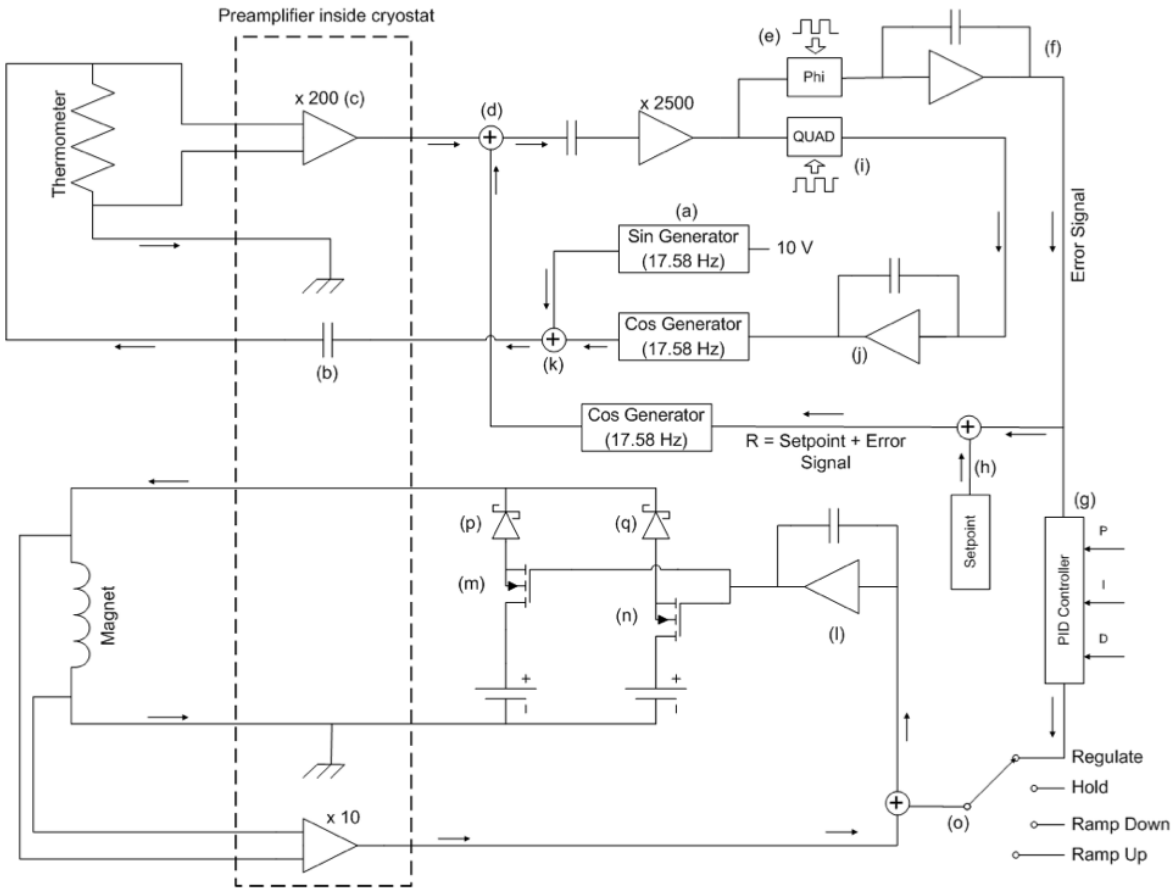

Figure 4-18: The ADR electronics control the focal plane temperature with a PID loop, adjusting the current through the magnet based on the error signal of a lock-in amplifier for the resistive thermometer. The components are explained in the text.

the field down. (For a superconducting magnet, $dI/dt \propto V$.) The PID loop hands off to constant voltages (o) to hold or ramp to full field. Finally, the current to the magnet passes through two Schottky diodes (p, q) which help to evenly distribute the load between the two supplies at high current, and isolate the two supplies in case one fails shorted.

This ADR control circuit has demonstrated 100 nK (rms) stability over a 1 Hz bandwidth in *XQC*. Commissioning and characterization of the ADR electronics in the flight dewar of *Micro-X* will take place in late 2013.

4.3 Vibration Isolation

The short rocket flight affords no time to re-cycle the ADR, so the focal plane must be launched close to the detectors' operating temperature at a fraction of 1 K. Thus, significant research and development has been directed towards vibration isolation of the focal plane.

It is no small task to keep a sub-Kelvin focal plane cold during a violent rocket launch. The motors output a formidable 12.7 g integrated over 20 – 2000 Hz for 10 s [126]. Without proper vibration isolation, a mere fraction of the mechanical energy from the shaking rocket motors will dissipate as heat, warming the focal plane above operable temperatures. Despite being derivative of the successful *XQC* isolation system, the initial *Micro-X* design could not sustain its vibration test without the detectors warming up above the helium bath temperature. The 2012/2013 launch window was spent diagnosing this problem and testing designs, efforts which I will detail here.

Micro-X employs three levels of vibration isolation with the intent to mitigate the heating to a couple microwatts. The resonant frequencies for each level are designed to be staggered, ensuring that the excitement of one stage does not affect the others much.

4.3.1 Low Frequencies: Rubber Dampers

As a first line of defense, LMO is suspended from the rocket skin by rubber ('elastomeric') dampers. (Other metal-to-metal connections from LMO to the skin, required for vacuum

Figure 4-19: Two low-frequency damper options are being considered to isolate LMO from the rocket skin vibrations. The Barry rubber dampers (left) have a higher natural frequency than the WRI's (right).

tightness – the pump line for the LHe exhaust, and the optical path to the gate valve – are made with bellows to take up the movement.) Two Barry Controls NC1010 dampers are attached in series at each of six attachment points to the dewar, three on the aft side and three on the forward side.

Laboratory tests revealed the resonance frequency of the rubber damper system with a fully loaded dewar to be just under 40 Hz.

Modeling efforts (see Section 4.3.4 below) indicate that an even lower resonance frequency for this first isolation stage could help reduce heating. Thus, we have investigated other damper options, primarily the wire rope isolators (WRIs) from Isotech (Figure 4-19). WRIs, being composed of flexible, thick gauge wire, feature much lower resonant frequencies, but are also more susceptible to shear. For this reason, WRIs are often mounted together at 45 degrees to the load, as this configuration reduces shear and provides a more predictable response.

We prototyped a WRI system to replace the rubber dampers. Initial tests found the natural frequency of LMO suspended by WRIs to be as low as 20 Hz, depending on the kind of damper.

The lower spring constant of the WRIs will result in larger displacements, though, than the rubber dampers under the same acceleration. This consequence presents the alarming possibility of collision between the dewar and surrounding components, especially during

reentry, as somewhat tight clearances already exist between the skin-mounted components and the suspended dewar. If we adopt the WRI suspension, we may have to remove the clamping mechanisms for this reason. Without clamping mechanisms to perfectly align the array with the mirror's focal point, the alignment lasers will be required to reconstruct the pointing over the flight.

I modeled the response of various WRI dampers over the course of the rocket's trajectory, in order to find the maximum displacement. The aim of this full-fledged simulation was to find the most extreme fluctuations, which are not easily predicted with analytical formulae, given the non-white noise spectrum and variable acceleration of the rocket motors. The simulation numerically integrates the harmonic oscillator equation of motion, given the acceleration power spectrum of the motors, the time profile of the rocket's acceleration, the dampers' spring constant as a function of displacement, for various Q-factors. The Q-factor (for 'quality') is a measure of the peakiness of the frequency response on resonance, and can be related directly to the damping coefficient in the damped harmonic oscillator equation of motion.

The simulation suite matches data. The position monitor lasers (see Section 4.9 below) tracked the position of LMO during the Wallops sine sweep (Section 4.3.5), with the rubber Barry dampers installed. A simple test case of the code, with a driving force at the resonant frequency, reproduced the observed displacement amplitude of the dewar as the sweep passed the dampers' resonance, to within 20%.

Figure 4-20 shows one run of simulated z-axis displacement: the two motors thrust the dewar downward — the Terrier Mk70 booster from $t = +0$ s to $t = +6$ s, and the Black Brant IX from $t = +12$ s to $t = +42$ s — while coasting in between, while the natural ringing of the dampers plus the vibration spectrum of the motors add displacement on top. The total displacement is generally dominated by the large swing induced by the burnout of the first stage. The likelihood of extreme displacements can then be calculated with many simulations. I will not detail the full results here, but a comparison between the original rubber dampers and the current favored WRI model is illustrative: the '5σ-high' maximum

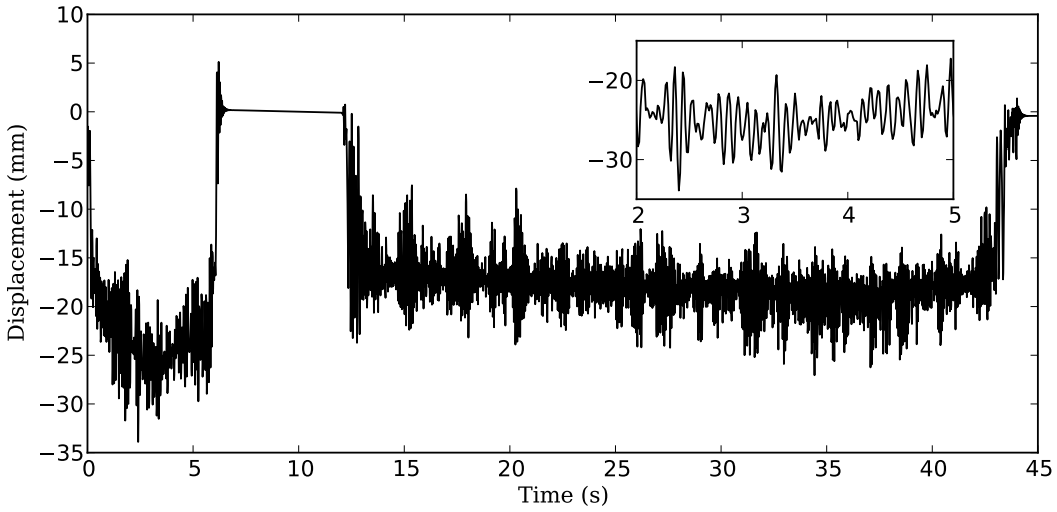

Figure 4-20: The displacement of LMO suspended on WRI dampers can be simulated over the launch, given the WRI characteristics and the acceleration profile of the rocket motors. These simulations are useful to determine the maximum displacement of the dewar and design for sufficient clearance. The graph shows the initial thrust from the Terrier booster, followed by coasting and then thrust from the Black Brant burn.

displacement are 15 mm and 58 mm, respectively. The mounts for the WRI were designed with this displacement in mind, though the reentrant impulse force (e.g., from hitting a sand dune) could result in much larger displacements. For this reason, if the *Micro-X* design employs the WRI's, rubber padding may be installed to prevent such a damaging shock.

The decision whether to adopt the WRI design will be made after tests at NASA Wallops in the August 2013.

4.3.2 Medium Frequencies: G10 Shells

The G10 structure, described in Section 4.2.8, functions as the second level of isolation. G10 has excellent strength in compression and tension, providing good 'reentrant' support and a stiffer spring constant than the skin dampers. Stress tests performed at MIT demonstrated a breaking strength of roughly 1300 N on average, and an elasticity modulus around 10 GPa.

The G10 assembly exhibited a resonant frequency of 200 Hz in laboratory tests, higher than the naive calculation from the pull-test values. The reason for this stiffer spring constant

(it is also stiffer than the *XQC* system) is currently unknown. The team has focused on the other two stages for redesign, though, postponing any further research and development into the G10.

4.3.3 High Frequencies: Kevlar

The Kevlar suspension serves the dual purposes of thermal and vibrational isolation of the salt pill and detectors from the 2 K stage. Kevlar's favorable thermal properties were discussed in Section 4.2. Its stiffness is also favorable, as its natural frequency under tension can be tuned higher than 200 Hz, away from the other isolation stages.

The Kevlar suspension consists of seven pairs of strings (Figure 4-23); six to constrain all the degrees of freedom, and one to provide tension, taking up slack with springs. The tensioning springs provide a force of $\sim$ 900 N, providing forces ranging between 250 and 400 N on the individual strings, which depends on the geometry.

The *Micro-X* team (primarily Sarah Heine) has invested significant research and development to stiffen the suspension closer to its theoretical maximum. See Section 4.2.4 for more details on the pull tests.

The current suspension design (i.e., individual suspension pieces including the Kevlar, the attachment roller to the cold stage, and the corkscrew mounting piece) achieves $\sim$ 65% of the theoretical Young's modulus of Kevlar alone. The hidden softness can most likely be attributed to slippage along the corkscrews, and internal stresses of the metal parts (many of which are now made of stiffer Ti).

Resonant modes of the Kevlar suspension are governed primarily by the relative geometry of the seven strings, though, not the strength along individual strings. That is, the softest, lowest frequency modes act perpendicular to the strings. Such modes must be diagnosed by modeling and prototyping (Sections 4.3.5 and 4.3.6, below).

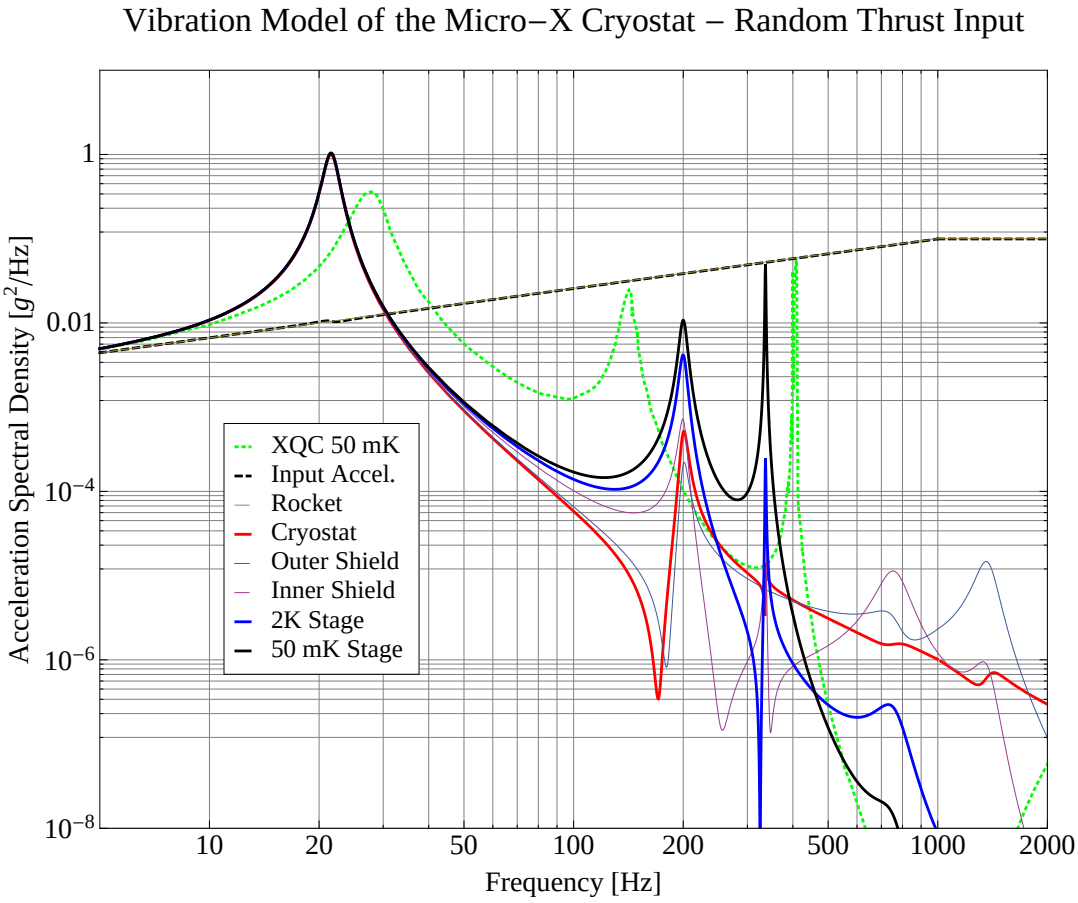

Figure 4-21: The analytic model for six chained damped oscillators, representing the main vibration elements of the payload, models the focal plane heating of *XQC* well, and can set goal resonance frequencies for the various systems to keep the heating below 2μW. Notice how the LHe tank (200 Hz) and Kevlar (320 Hz) resonances are closer than in *XQC*; *Micro-X* tries to compensate for this overlap with a lower frequency damper response at 20 Hz.

4.3.4 Simple Analytic Model

Even though the vibration isolation stages operate along all three axes, a simple 1D model is useful to build intuition and predict vibrational heating for different configurations. Tali Figueroa built the model in Mathematica, which is described thoroughly in internal documentation[27].

The lumped-mass model connects six masses with damped oscillators: the rocket skin, LMO, the outer G10 assembly, the inner G10 assembly, the LHe tank, and the FEA. The equations of motion for each stage couple between each other, resulting in a system of second-order differential equations:

$$\mathbb{I}\ddot{X} + \mathbf{C}\dot{X} + \mathbf{K}X = \mathbb{I}F(t). \tag{4.16}$$

The entries of the damping coefficient matrix $\mathbf{C}$ determine the dissipation, and thus the heating, which is a function of the *relative* velocity between the FEA and Insert.

The individual k's between stages (in the $\mathbf{K}$ matrix) are taken from measured material data, such as pull tests, but are tweaked to match vibration results (below). For instance, the measured Young's modulus of G10 predicts a lower resonant frequency than observed. (We have not been able to determine what aspect of the cryostat design stiffens the G10 assembly.)

Figure 4-21 shows the predicted responses of the various levels in response to the rocket motors. The simulation correctly models the heating of *XQC* (0.9 μW) and predicts that the increased masses of *Micro-X* will result in similar heating (given the same k's as the *XQC* payload). Data from the Wallops vibration (below) revealed that the G10 stage and Kevlar suspension had resonances too close; this model helped establish goal frequencies for all three damping systems to push the heating below 2 μW.

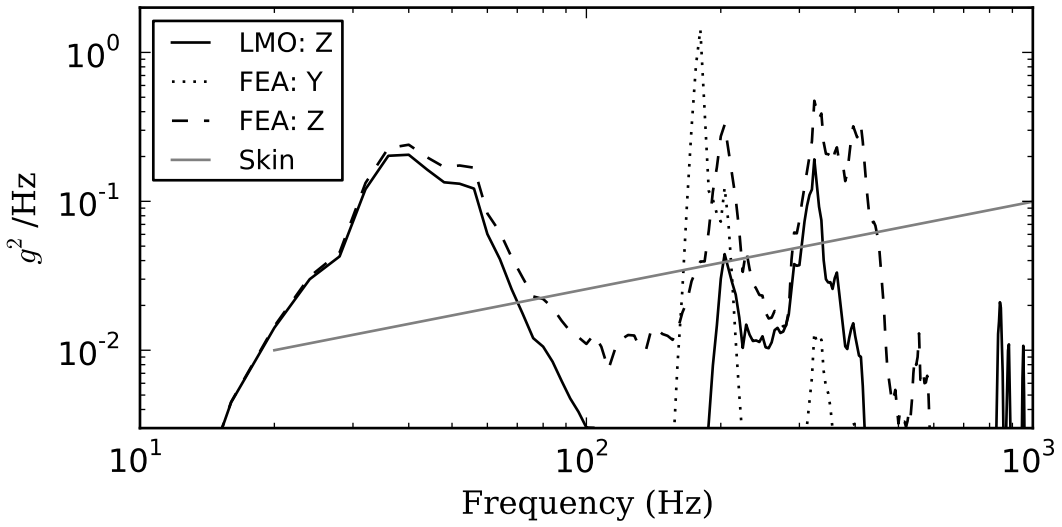

Figure 4-22: The accelerometer responses from the Wallops vibration test revealed an overlap of resonances around 200 Hz, especially along the y-axis of the FEA. This test was a flight-level random input along the z-axis, controlled on the skin to the grey line. The Kevlar suspension redesign effort aimed to lessen the $y - z$ coupling, push the Kevlar resonances higher, above 300 Hz. Softer skin dampers are also being investigated, to push the first broad 40 Hz peak lower.

4.3.5 Wallops Vibration Test: Results and Redesign

The team put the *Micro-X* payload through launch-level vibration loads at the NASA Wallops Flight Facility in June 2012. The Wallops 'shake' facility can apply specified accelerations, controlled via a feedback loop to account for the response of any components mounted to the rocket skin. The sophisticated shake tables provide either broadband random noise or sine sweeps, along a desired axis.

The payload passed the room temperature mechanical worthiness tests, with all components surviving acceleration and no indication of loosening or deformation.

Acceleration response data were taken at several locations in the payload, including on the cryostat lid (after one level of isolation) and at the location of the array (on a fake FEA, three levels of isolation deep). The responses at those locations from a sine sweep are shown

[27] `https://wikis.mit.edu/confluence/download/attachments/83631555/micro-x-vibration-modeling.pdf`

in Figure 4-22. The accelerometer on the top of the cryostat shows the broad response from the rubber dampers near 40 Hz, as well as some back-reaction from internal components around 200 Hz and 320 Hz, which rise up to the input level. The FEA experienced similar accelerations along the thrust axis (z-axis), as well as an extremely strong response along the y-axis at 180–200 Hz.

The design failed the 'cold' test, though. The focal plane, cooled to ~ 100 mK, warmed to inoperable temperatures when subjected to launch-level broadband noise. No subsequent tests, such as a sine sweep to identify specific coupling modes, were performed due to time constraints and, frankly, oversight.

The primary suspect to explain the heating was vibrational mode coupling. A mechanical touch, causing a thermal short, between the suspended cold stage and the surrounding 2K stage was quickly ruled out; no electrical connection between the two stages was observed during the vibration, and no screws or other components came loose to bridge the gap. Thus, the postmortem examination fell upon computer modeling to diagnose what vibrational modes were unanticipatedly coupling between isolation stages.

It should be mentioned that such a dramatic failure of the vibration isolation system was not expected. The *XQC* focal plane, suspended with a much simpler Kevlar design, routinely survives launches on Black Brants, needing only a few milliAmps of current to recover to base temperature. The *XQC* design was largely ad hoc, and was never modeled with computer simulations to predict heating. So, it was assumed the *Micro-X* design would not need modeling either, as long as the suspension was tensioned to a suitable force.

Detailed finite element method (FEM) modeling followed the shake test to diagnose the heating in the *Micro-X* design. Initial simulations suggested that the angles of the Kevlar strings conspired to create a lower-frequency resonant mode of the Kevlar suspension, one that would couple z-axis vibration with the perpendicular axes. The coupling primarily affects the y-axis, the direction along which the tensioning spring pulls. The mode, visualized best as the FEA teetering on the tip of a rocking bowling pin, displayed a resonant frequency just below 200 Hz, very close to the observed peak in the y-axis response spectrum. The

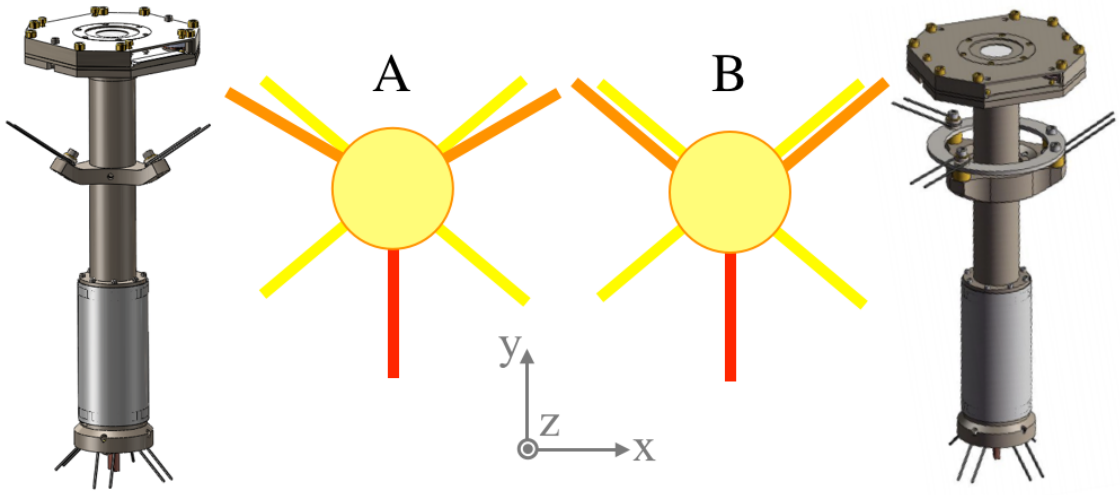

Figure 4-23: The iterations of Kevlar suspension designs reflects growing understanding of its resonant modes. The initial design (A), angled the two aft strings (orange in top-down cartoon; upper strings in CAD model) for mechanical convenience. To correct for low-frequency coupling, design B angles those strings perpendicular to the salt, and decreases the angle between them. Design C (not shown) further modifies the cold-side connection with another stiffening ring and thicker Kevlar. The red string is connected to the tensioning spring.

overlap of this mode with the principal frequency of the G10 shells around 200 Hz was inferred to cause the heating. (A 310 Hz response in the x direction also fits the data nicely. The model predicts the z-axis to be quite stiff, with a mode above 600 Hz.)

The FEM modeling produced some descriptive success, being able to explain the large y-axis coupling at the correct frequency, but had not yet predicted resonances of the suspension. The team applied the FEM model to a variety of new Kevlar suspension configurations to determine an optimal redesign. Figure 4-23 shows the results of the redesign effort, which are intended not only to push the resonant modes well above 300 Hz, but also to be easily machined and accommodated by the current Insert design.

Design B improves on the original design A in three main ways, while only changing the aft suspension (the three strings closest to the FEA, including the one attached to the tensioning spring):

- The two non-tensioning Kevlar strings are thickened. The increased heat load to the

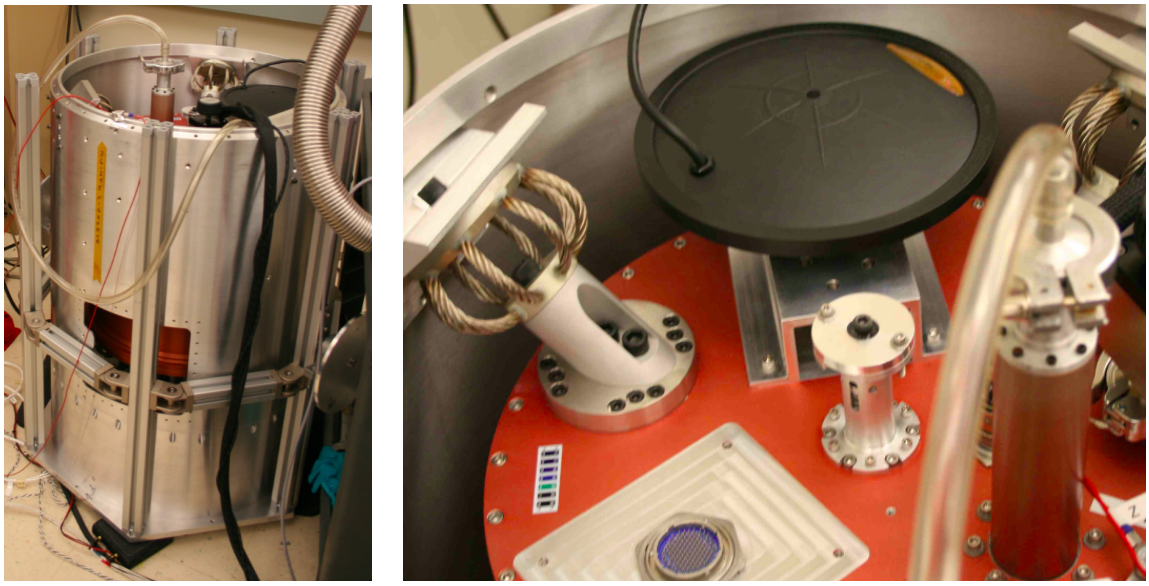

(a) (b)

Figure 4-24: Two audio transducers serve as cheap, in-lab replacements for a full vibration table, helping diagnose resonances of redesigned components. Bigfoot (left) outputs more power at low frequencies, so is able to shake the entire skin section to test the skin dampers. The UFO (right) has better response at high frequencies, and we mount it inside the first level of damping, directly on LMO, to probe LHe tank and Kevlar suspension responses.

FEA was deemed acceptable.

- The same strings are flattened to be parallel to the Insert surface, making it harder to couple vibrations from the z-axis into x and y modes.

- In order to stiffen the y mode, the angle between the same strings was decreased to 100 degrees – as close to 90 degrees as the Insert design would allow.

Design C stiffens the *connection points* with a second ring and shoulder bolts.

The FEM modeling predicted resonances at 320 Hz, 360 Hz, and above 500 Hz for Design C. Designs B and C were fabricated for testing.

4.3.6 Laboratory Vibration Testbed: Results and Redesign

In order to verify the resonances of new suspension designs, while at the same time avoiding extended and costly trips to the Wallops facility, we commissioned a cheap in-laboratory

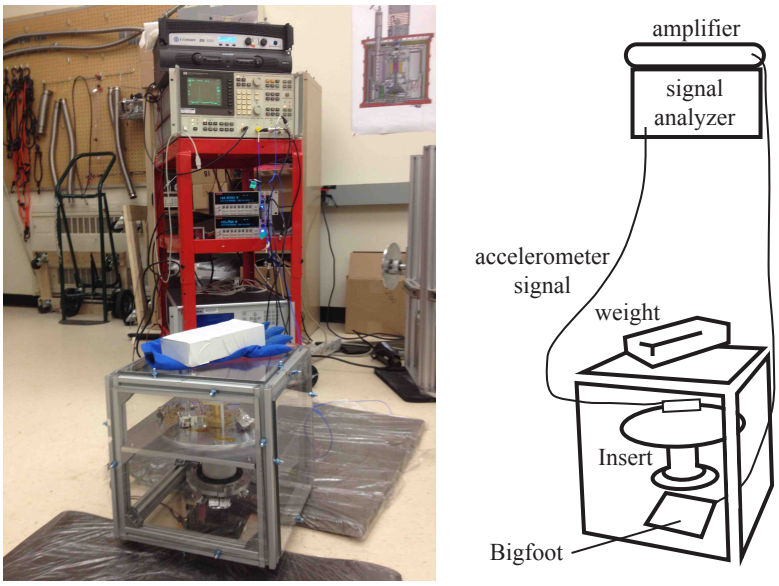

(a) (b)

Figure 4-25: The Insert-standalone configuration can measure responses of the Kevlar suspension alone. Bigfoot sine-sweeps the Insert, which is encased in a polycarbonate/aluminum holder and weighted down by a lead brick. The accelerometer signal on the suspended mass is read in by a dynamic signal analyzer.

vibration testbed. The laboratory shake facility consists of home theater components, amplifiers, and signal generators. We employ two *audio tactile transducers*[28] with different frequency responses as the vibration elements to perform sine sweeps.

The two transducers are shown in Figure 4-24: the Clark Synthesis TST329 Gold Transducer ('UFO') has better response at high frequencies, while the Crowson Technology EDH-100 ('Bigfoot') outputs greater power at low frequencies. The Bigfoot is thus our replacement for the Wallops shake table, as we can place it underneath the skin with the dewar mounted inside. We use the weaker UFO mounted directly on the top of LMO, one vibration stage in. Both can be used to excite the Insert directly when disassembled from the rest of the dewar and held in a cube (Figure 4-25), which I'll call the Insert-standalone setup.

I should stress that this homegrown setup cannot replicate the highly controlled Wallops setup perfectly. Rather, we use the transducers to find resonant peaks, while keeping in

[28]The common term for this component among home theater enthusiasts is a 'butt-kicker', as it gets attached to the underside of a seat to enhance loud explosions.

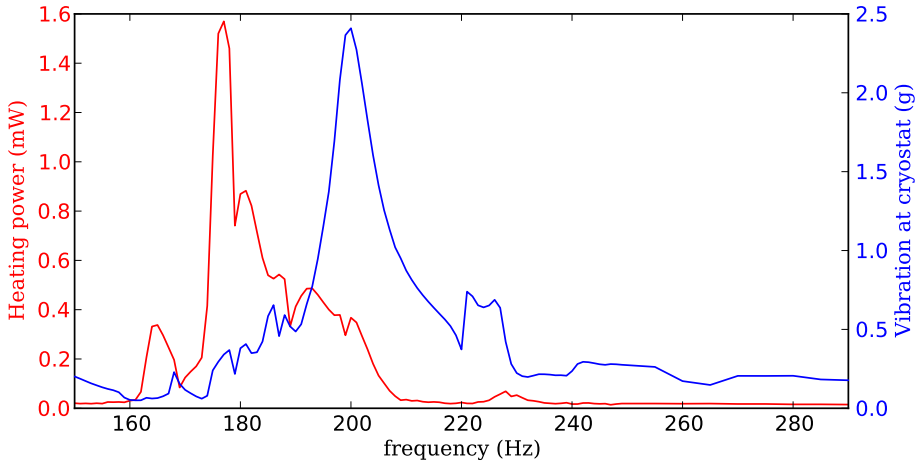

Figure 4-26: The laboratory vibration testbed can identify heating responses of the Kevlar suspension at 4 K with a Bigfoot sine sweep. Design A exhibits strong heating at 180 Hz (left axis, red), close to but distinct from the LHe tank and G10 response (right axis, blue). The 180 Hz mode was observed with the Insert isolated from the dewar, and also predicted in FEM modeling.

mind that their strength can easily be dependent on the internal, uncontrolled response of the transducers.

As a first step with the lab testbed, we subjected Design A to sine sweeps in several configurations to reproduce the results from Wallops. In the Insert-standalone configuration, we found strong responses in the suspension at 180 Hz, 200 Hz, and 350 Hz, replicating the Wallops responses. (After experimenting with different setups, including the cube suspended by bungee cords, we found the most faithful setup to be the cube sandwiched between the Bigfoot and a stack of weights.)

We can also shake the dewar at LHe temperatures, slowly sweeping the frequency of the transducer while measuring heating responses. Figure 4-26 shows the heating response of the FEA (only at 4 K, not 150 mK as at Wallops) for Design A. For this measurement, Bigfoot provided the acceleration to the rocket skin, while the simultaneous acceleration response at the dewar was measured. A strong heating response is evident at 180 Hz, very near to the 200 Hz back-reaction response of the LHe tank suspended from the G10 assembly. (The heating slope was converted to a calibrated power with a known heater.) This measurement,

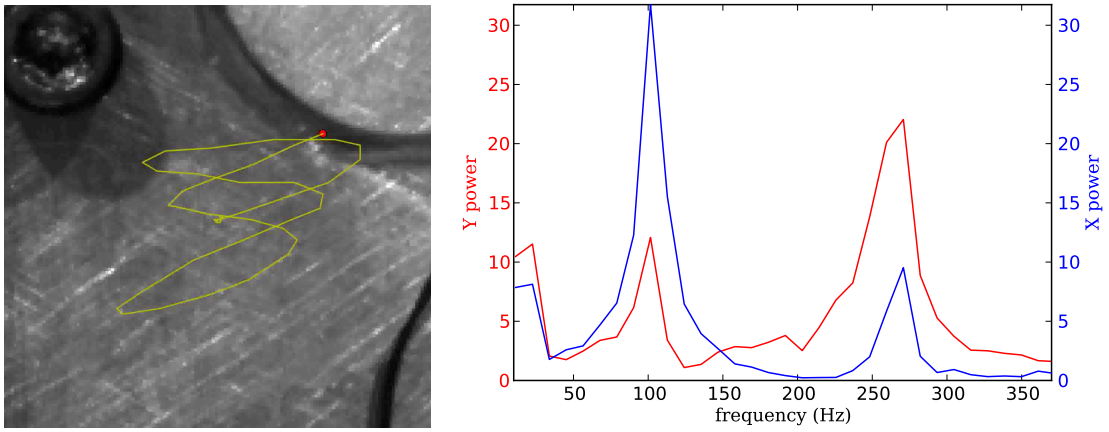

(a) (b)

Figure 4-27: The response of the suspended mass to an impulse force, captured with high-speed video, can be analyzed to sub-pixel resolutions to determine the natural resonant frequencies. At left, a frame from the tracking video, with the motion tracker overlaid, at 50× magnification. The linear motion is FFT'ed to reveal 100 Hz and 280 Hz natural frequencies along both the x- and y-axes; the former is a rotational mode, while the latter is the 'bowling pin' mode described in the text.

which was not performed at Wallops, confirmed that this low-frequency mode – observed in the Insert-only tests, and modeled by FEM – was the heating culprit.

The home theater testbed was able to replicate – and further explain – results from Wallops, providing confidence that it would be able to quickly appraise the success of redesign efforts. Designs B and C were subjected to room temperature vibration tests in the Insert-standalone configuration, and indeed the x and y modes were pushed higher, above 300 Hz, but not above 400 Hz as predicted by the model. Furthermore, weak evidence for a z-axis mode – also unpredicted – was observed.

The FEM model was further refined to better describe attachment points, but still could not accurately predict every resonant mode. Thus, our efforts to refine Design B to Design C were largely heuristic, as prototyping was deemed quicker and more expedient than further complicated modeling. Various Kevlar thicknesses, looping designs (crossed, parallel, and spiral attachment points), screw types, and stiffening rings were tested until the x and y modes were pushed tens of Hz higher, around 330 Hz.

High speed video[29] was taken to determine the nature of the $\sim$ 300 Hz mode, whether rotational or the 'bowling pin' motion predicted by the FEM. For Design B, video taken at 2900 frames per second revealed that a 280 Hz driving force along the z-axis coupled into both the x- and y-axes, producing elliptical motion of the FEA, a bowling pin mode. As a further confirmation of this fundamental mode, an impulse force was applied to the FEA[30], and the video of the ringing response was analyzed. I developed a method to track the motion of the FEA frame-by-frame at resolutions below one pixel[31]. (The scale of $\sim$ 0.15 mm/pixel was set by the focal length of the camera.) The power spectrum of the ringdown (Figure 4-27) showed strong natural frequencies at 100 Hz (rotational) and 280 Hz (x/y translational). These tests confirmed the nature of the bowling pin mode, but also provided striking visual evidence of its lower-than-predicted frequency.

Sine sweeps at LHe temperature confirmed heating at this higher mode (280 Hz for Design B, 320 Hz for Design C), but also revealed heating again at 200 Hz. Thus, the problem inherent in Design A has not been solved. Neither the FEM model nor the room temperature tests predicted a mode at 200 Hz. Design C, with stiffer attachment points, shows decreased 200 Hz heating but more high frequency heating. This trend may imply that the large response of the Insert at 200 Hz affects the suspended FEA as well, producing heating not via out-of-phase rubbing, but from the slight deformation of the attachment points.

4.3.7 Design Outlook

The FEM modeling remains incomplete as a predictive physical model, though it provides valuable intuition for how the suspended mass behaves under acceleration. Thus, the most recent incremental redesigns have been made by trial-and-error. The home theater testbed, too, while informative, cannot answer the most pressing question for *Micro-X*: will the focal plane overheat on launch? For these reasons, the current Design C is set to be tested at

[29]The MIT Edgerton Center, named for the father of modern high speed photography, Harold Edgerton, loaned the specialized video camera and lighting equipment for this task.

[30]A mass model, not the actual focal plane with delicate components, was installed.

[31]`https://github.com/jmrv/micro-x-vibe/blob/master/fitMotion.py`

Wallops in August 2013, even though we have not completely eliminated resonant modes near 200 Hz. The tests will better quantify our progress in vibration isolation. The WRI suspension will also be tested to confirm better low-frequency response than the rubber dampers.

One further fix under consideration is the use of a tuned mass damper (TMD). A TMD is simply a mass on a spring, tuned to a target frequency, whose job is to take up vibrational energy at that frequency, reducing the available energy for other components attached to the system. Thus, a TMD tuned near a heating resonance, mounted on the exterior of LMO, could reduce the energy injected into the stages inside the dewar. Initial tests indicate the TMD may work for modes along its axis (the z-axis in our case), but exacerbate heating in perpendicular modes. The Wallops tests will also investigate the effects of a TMD.

If all else fails, the *XQC* suspension design can be adopted to the *Micro-X* payload. This was not done initially because the *XQC* suspension is harder to assemble, and does not accommodate a magnetic shield.

4.4 Detector Readout

The TESs on the focal plane are addressed and read out by a complicated set of electronics, which traverse several temperature stages. To reduce the number of wires on the focal plane, the 128 pixels are addressed by a *multiplexer* (MUX), which communicates in a pulsed, round-robin style to subsets of pixels that share a bias line. As a whole, this collection of signals and controls is called the *science chain*, to distinguish it from housekeeping signals.

Let us loosely follow a bundle of science chain signals before getting lost in specifics. Address signals and constant amplifier biases are generated in the *MUX stack*, which resides at the top of the science instrument, on the avionics deck. They pass through the *Interface (IF) stack* where they are, at this point in their journey, merely aggregated into one connector at LMO's top lid. Most of these signals pass directly down to the focal plane after routing through another breakout, the *2K board*, built into the wall of the 2 K shield on the Insert. The TESs and first two stages of amplification on the cold stage are turned on (biased).

Now, any signal from the TES can work its way back out, by translating changing resistance into flux, which couples to the first two readout stages. Back at the 2K board, the signal is amplified by cold electronics before being sent up to the IF stack, for further amplification. Once returned to the MUX stack, the signal is both recorded and used as the error signal in a control loop, which adjusts the output signals for another trip through the science chain.

With this broad overview, I will now refine the description, starting at the focal plane.

4.4.1 SQUIDs

The tiny TES signals are amplified by SQUIDs (Superconducting QUantum Interference Devices). This now-common readout method came into maturity in the 1990s [5], with the first three-stage amplification system (which *Micro-X* and collaborators adopt) published in 1993 [138]. SQUIDs are useful at cryogenic temperatures, where other low noise amplifiers cease to operate.

SQUIDs — originally studied for their ability to expel flux in quanta — make useful circuit elements when they are biased above a critical current, becoming resistive and developing a voltage drop. In this resistive mode, the voltage across the device oscillates with increasing flux. The sinusoid-like signal is nearly linear between the peaks, in the middle of the oscillation, so signals are locked at this midpoint with a PID loop.

Each TES couples to a SQUID (SQ1 in Figure 4-28) via an inductor, turning the current through the TES — from a fixed voltage — into flux. Any change in the TES resistance couples into SQ1, then is collected by the summing coil (SC) and sent to the control electronics via another coupling to SQ2. The flux feedback offset on SQ1, sent by the MUX stack, adjusts based on this signal, keeping SQ1 in the linear regime. (The feedback adjusts the current in the TES, too, due to mutual inductance.)

Individual SQUIDS can be ganged together into phase-coupled arrays, which function as signal amplifiers by way of combined inductive coupling. The SQUID arrays in *Micro-X* join 100 together. SQUID arrays are biased similar to individual SQUIDs, providing enough current to develop a voltage drop, then adding a suitable offset flux to keep the system in

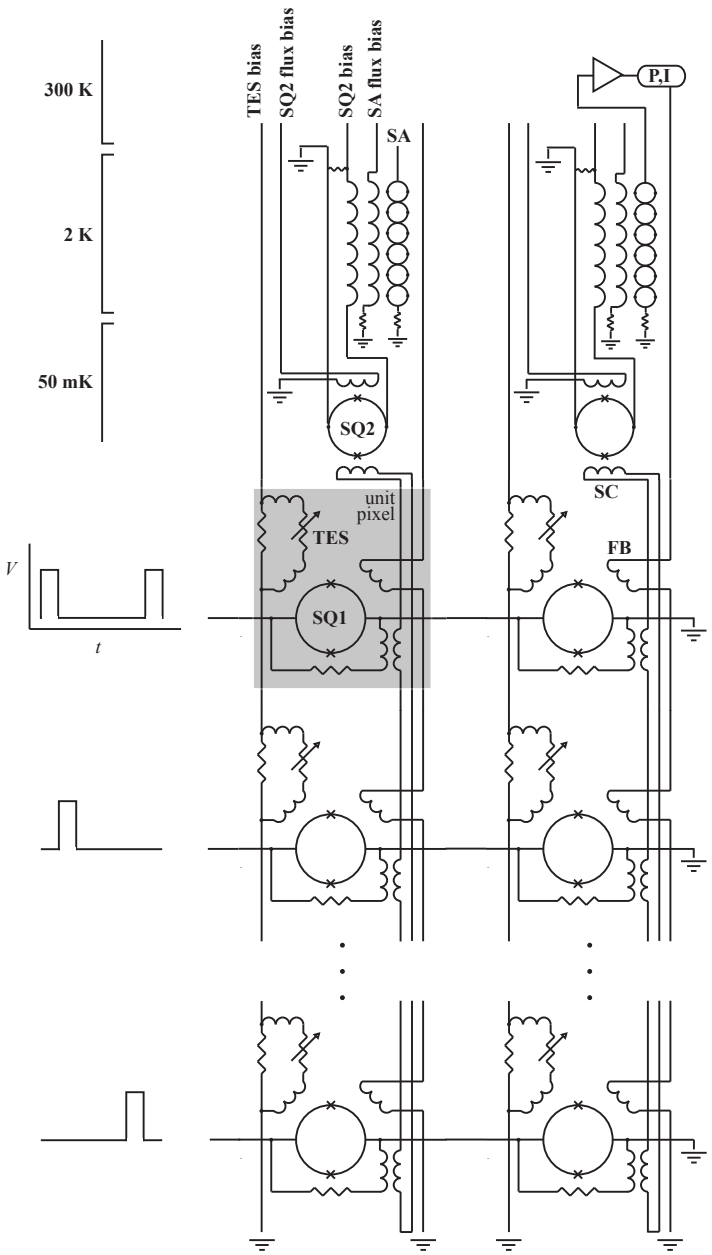

Figure 4-28: The multiplexing readout of the TESs uses SQUIDS at low temperatures to amplify the small signal. The changing resistance of the TES induces a flux through a biased SQUID (SQ1), adjusting the voltage in response. The voltage is translated again to flux and brought out to the second stage SQUID (SQ2) by the summing coil (SC). SQ2, in turn, transfers the flux to the SQUID array (SA), which amplifies the signal for a room temperature amplifier. The MUX stack outputs a feedback (FB) voltage based on the error signal from the amplifier, completing a PID loop, as the FB signal works to null the induced signal from the TES. The SQ1s are biased sequentially by the MUX stack. *Micro-X* employs 4 columns of 16 rows for each half of the focal plane array, instead of the 2 columns and N rows schematized here.

the array's linear regime.

The signal from SQ2 is amplified by one such SQUID series array. The series array sends the signal upward to the IF stack, where a room-temperature amplifier (100×) magnifies the deviation of the signal from the setpoint. This error signal is processed by the MUX stack, and compensated for by adjusting the flux feedback on the lowest-stage SQUID, SQ1.

The SQUID stages must be tuned before operation. Two knobs are available for each stage: the bias voltage, which supplies the current to drive the SQUID into the resistive mode, and a flux offset voltage, which adjusts the phase of the oscillatory response. Using these knobs, the three SQUID stages can be 'aligned' across the midpoints, so that a linear response in the SQ1 signal begets a linear response in the SQ2, which begets a linear response in the SQUID array.

The various circuit elements are laid out on different physical chips. At the center of the FEA, the detector array chip contains only the TES resistors. The SQ1s and their feedback coils for one column are etched onto a long *MUX chip*, developed by NIST. (The chips actually contain 32 SQ1s, but only 16 are wired per column. The surplus of SQUIDs allows us to choose the set of 16 with the most uniform characteristics.) A single SQ2 presides over a whole MUX chip. Four MUX chips are epoxied to an *interface chip*, which routes wires between the components. The two interface chips flank the detector chip. All connections between the chips are made with wire bonds, and route to the 85-pin FEA connectors via a breakout board. The SQUID arrays live on the 2K electronics board, shielded in a Nb box.

4.4.2 Time-Division Multiplexing

Multiplexing the addressing and feedback signals simultaneously helps solve two problems by reducing the number of wires: heat load on the focal plane and space constraints (thinning the forest of wirebonds between chips). All 16 SQ1's in a column economically share one SQ2, one summing coil, one SQ1 feedback coil, and one TES bias line. While the TES bias line is set to a constant voltage, the SQ1's and feedback loop must be addressed in a round-robin fashion to ensure that only one pixel contributes to the shared signal line at a

time.

Moreover, the row address signals (turning on SQ1) are shared across columns, further reducing the required wire number.

This method is known as time-division multiplexing. Other multiplexing schemes exist to sample a signal with less power, including frequency division multiplexing (the signals are chopped in frequency space) and code division multiplexing (the signals are not restricted in time or frequency, but are encoded with something similar to a one-time pad).

In a single column — defined by a common amplifying chain and summing coil — the rows of SQ1's get addressed with a step function on/off signal. With zero bias current, a SQUID acts as an electrical short (being superconducting), and does not transmit any changing TES signal. (The summing coil, then, only functions to 'sum' many null signals with the only on-signal.) The *Micro-X* columns are 16 rows long, so the row 0 SQ1 waits 15 clock cycles to turn on again, with the adjusted bias computed by the MUX stack. With the three SQUID stages initially tuned in the middle of their linear regimes, the feedback signal on SQ1 functions to keep a highly nonlinear system poised on a small linear span. This feedback method is known as *flux-locking*.

NIST has been a lead developer in SQUID MUX technology, demonstrating systems able to out 1000-pixel arrays with TES time constants as low as 0.05 ms [25].

4.4.3 The Data Stream

The *Micro-X* science chain is composed of two independent 4-column-16-row SQUID MUX systems, just enough signals to read the 128-pixel array. The design is failure-tolerant, as each independent chain is powered by its own power card, controlled by its own MUX stack, and encoded and broadcast by its own telemetry system.

The full 66 Mb/s (per science chain) data stream is recorded by the Master Control Board of the MUX stack to flash memory. The same board doles out a decimated portion of the array to telemetry (the encoders can only handle 20 Mb/s).

The digitizing MUX stack is fairly configurable, allowing for experimentation in lab

with P and I values for the PID control, as well as settling times after which to read the multiplexed signal. The parameters, though — easily adjusted in software with a control suite and accompanying GUI — will be fixed for launch, once calibration determines optimal operating conditions.

4.5 Magnetic Shielding

Both TESs and SQUIDs are sensitive to magnetic fields. Thus, we need to shield the FEA from both static and varying magnetic flux. The most common way to create a low-field environment for these components is to encase them in a superconducting box[32].

An additional property of superconductors, besides zero resistance below T_c, is the ability to expel magnetic field below the critical temperature, known as the Meissner effect. Surface currents arise as the metal cools in the presence of a magnetic field, and persist below T_c, cancelling the field in the bulk of the material (as long as the strength does not exceed the critical field). Niobium, a Type II superconductor[33], is a popular material choice for its high superconducting temperature (9.2 K) and workable material properties (it can be rolled, welded, and milled). *Micro-X* uses it, additionally, for its low mass, compared to lead. (Many laboratory ADRs encase the superconducting magnet in a heavy iron alloy cylinder.)

The TESs themselves can tolerate a static residual field of approximately 1 μT. However, we must consider the noise induced in the TES bias current through the inductor. If the magnetic field is time-varying, it is estimated the readout cannot sustain powers greater than ~ 3 pT/$\sqrt{\text{Hz}}$ for the *Micro-X* design, much lower than the static limit. The magnetic shielding was designed to these constraints.

The first stage of magnetic shielding is a 'bucking coil' between the primary magnet and FEA, which winds the magnet current in the opposite direction. The coil, in series with

[32]Thus, superconductors are at once the antidote and the poison in this design: superconducting TESs offer the promise of high resolution, but strong fields are required for ADRs, effected by a superconducting magnet, requiring (superconducting) shielding to protect the delicate TESs. It is difficult to discern whether this design is elegant or just reactive.

[33]Type I superconductors go normal above a critical field. Type II superconductors first develop magnetic vortices above a critical field, then go normal at a higher field.

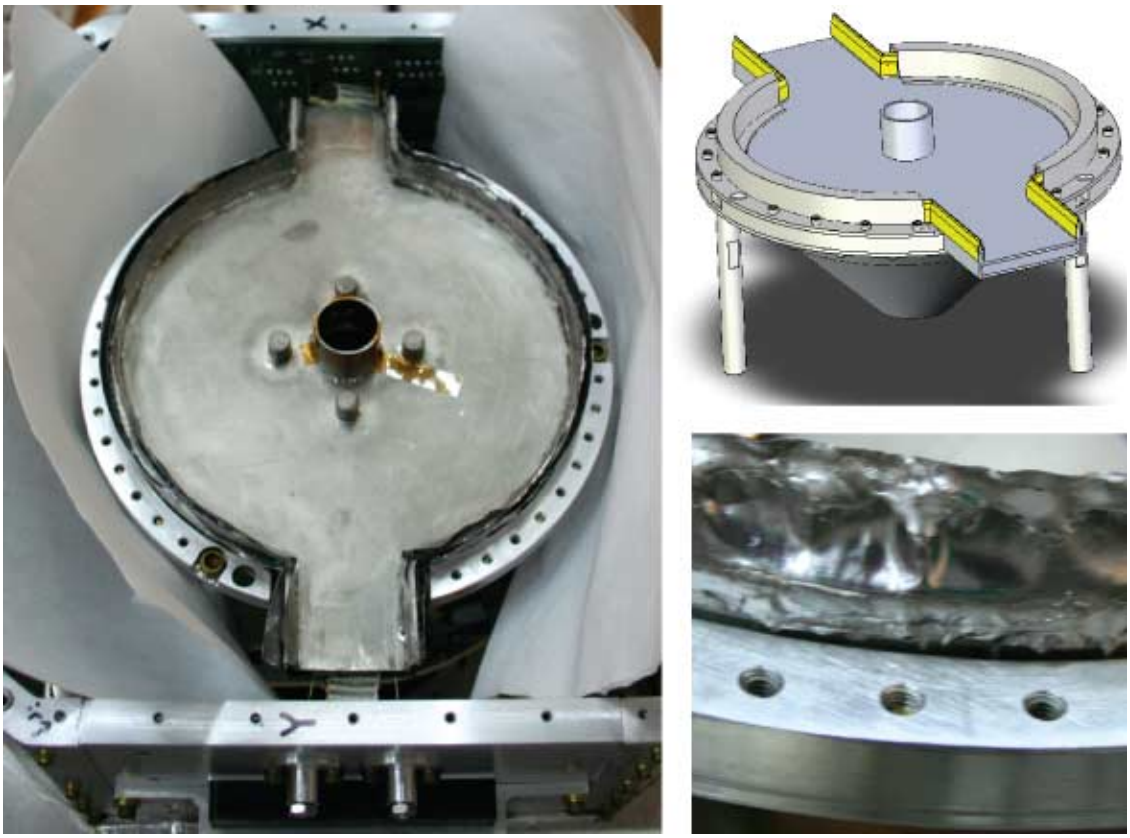

Figure 4-29: The superconducting Nb magnetic shield is designed to create a field-free environment for the TESs and SQUIDs, protecting both against the large quasi-static field of the magnet, as well as any external fluctuations. The current design connects the lower conical piece to the lid with a soldered lead zipper (lower right inset; highlighted in upper right inset). The four pegs surrounding the X-ray snout (left) are screw mounts to hold the calibration source inside.

the primary magnet, partially cancels the field in the detector region without diluting its strength at the salt pill.

The current design of the superconducting magnetic shield is shown in Figure 4-29. It consists of a welded Nb cone (so shaped to accommodate the Kevlar suspension), a fitted lid, and a lead tape 'zipper'. Each of these components evolved from modeling and performance tests.

The first design failed, leaking too much field. The leakage suppressed the T_c of the array by tens of mK, and introduced kinks in the response curves of the SQUIDs.

The initial shield converged by way of FEM simulations. Different enclosures were tested in an electromagnetic model of the *Micro-X* dewar, with appropriate geometric relations, current densities in the magnet coils, and permeability of materials, using the COMSOL 2D E&M package with axial symmetry. For ease of assembly and manufacturing, the lid was offset from the Nb cone, creating a $\sim$ 1 cm gap for the focal plane wires. The axial symmetry condition — favored for its shorter solution times and simplicity — suggested this gapped design, which turned out to be susceptible to transverse fields and misalignment. In this original design, we used magnetometers to measure the residual field of the magnet at the FEA, within the Nb shield. The attenuation ratio perpendicular to the optical axis was just above 10^{-5}, three orders of magnitude higher than predicted by the COMSOL model. Once the FEM model was extended to 3D, methodical analyses revealed the design to be very susceptible to small deviations from axial symmetry.

This iteration also included a high-permeability shield, mounted on the Insert surface midway between the Nb shield and magnet. The thin shield was made of ultra-low carbon steel, meant to 'soak up' a portion of the field, reducing the strain on the Nb shield. While high permeability materials are often used in cryogenic designs, we found this component became permanently magnetized, completely defeating its purpose. It was removed, consequently.

We redesigned the Nb shield with 'snouts' for the science chain cables, but otherwise aggressively sealed the rest of the circumference. Conversations with the team in charge of

the IXO shielding influenced us to connect the two Nb pieces with a lead tape zipper. Nb is not an easy substrate to solder on, so the procedure to make the lead zipper is difficult and delicate. The procedure, adapted from Garfield [40], requires Nb to be roughened (removing oxide) and wetted with indium. Thin lead[34] tape, stripped of its adhesive, can then be soldered to the Nb/In surface with Roses alloy. Test joints were stress tested in LN to determine the optimal soldering temperature.

The zippered shield exhibits much less susceptibility to changing magnetic fields outside the dewar. The benchmark test involves spinning a permanent magnet on a nearby chair, looking for the effect on the locked SQ1 signal. Without soldering a zipper shut, the $\sim 1\mu$T field swings would cause a very measurable signal of a few mV, a few percent of the pulse signal. Moreover, the T_c of the array has risen back to the actual T_c value in the zippered configuration.

The Nb pieces are clamped with Al fixtures, and the assembly attaches to the insert via standoff posts. The Nb follows the Insert temperature closely, confirmed with a thermometer GE-varnished[35] to the Nb, so we do not need to worry about it approaching its critical temperature during regular operation.

As a final precaution, we have investigated lining the Nb shield with high permeability sheets of Metglas, inspired by the use in the SCUBA-2 experiment [59]. Initial tests suggest the Metglas helps screen static fields, achieving the correct T_c, but cannot shield against changing fields. The solder zipper may render the Metglas component redundant, with respect to static field shielding, so the working configuration for future integration runs will not include it.

4.6 Calibration Source

Extensive in-laboratory calibration will be performed ahead of launch. Fluorescence lines from a radioactive ^{55}Fe source outside LMO illuminate the TES array, calibrating the pulse

[34]Lead is a Type I superconductor, with $T_c = 7.2$ K.
[35]A common non-permanent adhesive for cryogenic applications.

height gain to the known energy of the photons. Initial calibration results using this method are detailed in Section 5.1. This data will inform the optimal device settings, individual pixel gains, crosstalk between pixels, and noise characteristics.

On-board calibration during the flight is also crucial, however. The rapidly changing environment – microphonics from vibration, changing magnetic field, temperature swings – could influence the response of individual pixels. Thus, the *Micro-X* design incorporates a persistent source of known fluorescent lines mounted on the lid of the magnetic shield. Any drift, for instance, of the pulse height of such lines over the course of the observation can be calibrated out in post-processing.

The opportunity also exists to add a calibration source on the gate valve, as *XQC* does, in order to calibrate through the entire filter stack. Such a configuration could monitor ice growth on the filters before and after the observation. (The gate valve retracts during the observation, though, to allow science photons through, prohibiting mid-flight calibration through the filters.) Low energy lines are required below the 540 eV O K-edge to diagnose icing, so the ^{55}Fe source is unsuitable alone, but can be used to fluoresce a target material such as Teflon, which would yield C, N, and O lines.

Micro-X cannot accommodate a gate valve source, though, and relies only on the always-on FEA source for in-flight calibration. The gate valve is situated farther away from the focal plane than in the *XQC* payload, requiring a prohibitively strong radioactive source to produce an appreciable calibration flux. The gate valve was situated farther from the dewar originally to avoid microphonics from the motor, as well as potential O-ring cooling, which could result in gate valve failure. (Moreover, the Cm source in *XQC* is not commercially available now, and the desired strength of the Po source is not available to non-government entities.)

The calibration source is composed of an annulus of ^{55}Fe which illuminates a co-axial annulus of KCl (Figure 4-30). The bombardment of KCl results in fluorescent lines outside of the science band: 2.62 keV (Cl Kα), 2.82 keV (Cl Kβ), 3.31 keV (K Kα), and 3.59 keV (K Kβ) [130]. Lead walls absorb stray X-rays from ^{55}Fe. Inspiration for the annular design was

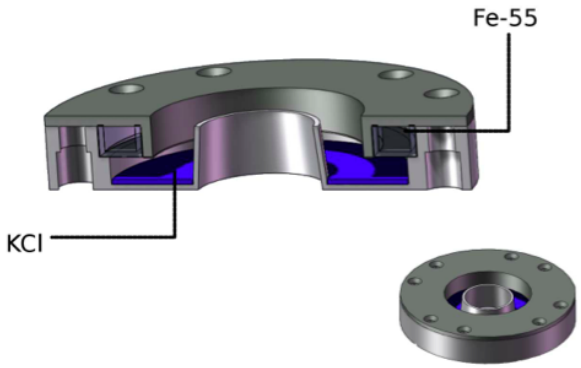

Figure 4-30: The annular calibration source design, depicted in this CAD model, fluoresces X-rays from ^{55}Fe off of a KCl target, producing a constant flux of lines outside the science band.

drawn from Yap and Saligan [147]. The unique source shape was made by mixing ^{55}Fe with epoxy and pouring it into an annular acrylic (Polymethyl methacrylate, or PMMA) holder to cure.

The geometry of the source, lead shield, and fluorescence target was tuned (within constraints) to produce a fairly uniform exposure across the focal plane. The uniformity and strength was characterized with CCD images from the MIT X-ray Calibration Facility. The measurements, taken in July of 2011, revealed a strength of 8.2 photons/pixel/s on average. This rate can be throttled with a suitable thickness of Al (e.g., aluminized Mylar flim) to produce a rate commensurate with the strength of the science target, to minimize pileup while producing enough calibration events.

4.7 Optics

As mentioned in Chapter 1, imaging is just as important in X-ray astronomy as in the visible or radio bands. Sources like SNRs are extended; many sources are dim, requiring a mirror to function as a collector; concentrating the signal on the detector typically increases signal-to-noise; point sources can be dissociated from the background, and even be dispersed through gratings to produce high resolution spectra. The optical principles for such high energy

photons, though, results in a different configuration for X-ray mirrors: grazing incidence.

X-ray instruments, being rocket-borne, must additionally take the weight of a mirror assembly into account[36]. The *Chandra* High Resolution Mirror Assembly weighs in just shy of 1000 kg, due to the large quantity of glass. The mirror achieved 0.5″ resolution with thick (2 cm), precisely polished glass segments underneath every sub-micron reflecting surface. Other instruments, such as *Micro-X*, can only afford lightweight reflecting foils in the weight budget.

The *Micro-X* mirror, refurbished from a combination of previous missions, completed commissioning as one of the earliest subsystems. The mirror and optical bench are contained in their own payload section, between the bulkhead and a mechanized shutter door. (In this setting, the optical bench is not a table with controlled damping, but a blackened aluminum structure to hold the mirror aligned to the detectors.) This payload section is hermetic, and will be pumped to microbar levels ahead of launch, to help prevent moisture from freezing on the mirror foils or the filters. The shutter door will also protect the mirror from dust during reentry, to preserve it in better condition than the SXS launch left it.

4.7.1 Grazing-Incidence Mirrors

The familiar day-to-day characteristic of X-rays is to penetrate human bodies or luggage, implying they would pass right through a parabolic mirror, too. Indeed, the only way to reflect X-rays is to direct them at a metallic surface at a very low angle, increasing the effective number of electrons a photon coherently scatters with.

An X-ray photon incident on a material obeys Snell's law of refraction. It can be shown that the critical angle, θ_c, below which total reflection occurs, follows:

$$\theta_c \propto \sqrt{Z}/E.$$

[36]Large ground-based telescopes must consider weight as well, but only at much higher masses: the previous incarnation of the Green Bank Telescope 90 m radio dish collapsed, while the 16,000 kg Large Synaptic Sky Survey mirror will require an exoskeleton to minimize distortion from its own weight.

Here, E is the energy – the derivation assumes X-ray energies, away from absorption edges – and Z is the atomic number, which enters into the relation due to the dependence on free electron density. Higher energies, therefore, require lower grazing angles and longer focal lengths. High-Z materials are favorable for the mirror coating: gold and iridium are preferred, while beryllium and aluminum work for soft X-ray telescopes [2]. For a bandpass up to 10 keV, 1/2 degree grazing angles are required for typical mirror materials. The *amplitude* of the reflected signal is a function of angle and index of refraction, and polarization can generally be ignored (which makes X-ray polarization studies difficult).

For real materials, not equations, surfaces are not perfectly smooth, resulting in scattering which decreases the true reflected amplitude. Metrology of Angstrom-scale deviations is not feasible, so the effects of scattering must be determined empirically in laboratory tests, or in a statistical framework. Additionally, reflecting materials are deposited on a substrate like glass and typically require an adhesion layer (chromium is used for gold coatings). These layers must be taken into account for higher grazing angles and near absorption edges, where X-rays penetrate. (Mirrors are not made of pure metal to avoid deformation from thermal expansion, and to reduce weight.)

A mirror with a parabolic profile will focus all rays parallel to its axis – a good approximation for photons originating at cosmic distances – to a single point. Two reflections, though, are required to form an image over a finite field of view [145]. Wolter investigated several configurations which would create identical optical path lengths for all rays forming the image: two employ a parabola-then-hyperbola sequence, a third uses parabola-then-ellipsoid. Most astronomical X-ray mirrors use the Type I Wolter optic, since it allows the shortest focal length for a given aperture.

To fill out the collecting area, the parabolic and hyperbolic segments are revolved around the axis. The *effective area* of a paraboloid mirror is $2\pi r\ell\alpha$, with α the angle between a ray and the shell of length ℓ at radius r. For larger effective area, shells are nested at various r; *Chandra*, for instance has four shells. The *actual* polished area of $2\pi r\ell$ can be many times larger, with α typically below 20 milliradians.

Wolter also demonstrated that such a double-grazing incidence setup acts like a thin lens, allowing undistorted imaging close to the focal axis. While the imaging surface is actually curved like a shallow dish, a flat detector plane suffices for most telescopes, as long as the imaging plane can be adjusted to provide the best response for a point source like a quasar[37]. The effective area suffers for sources off-axis, as photons photons get vignetted by the nested shells or come in at higher (less efficient) angles. The point spread function also worsens off axis. Scattering increases with energy (actually as E^2), resulting in energy-dependent effective areas and point spread functions.

Lighter, cheaper thin-foil mirrors were developed for missions which did not require sub-arcsecond imaging [104]. With thinner substrates, many more mirror segments can be nested within one another to achieve similarly high throughput. The conception and advancement of the thin-foil technology is due almost exclusively to Peter Serlemitsos at GSFC [104, 120].

The paraboloid and hyperboloid segments are approximated by more easily manufactured conical frusta. Under this approximation, foils can be mass produced as different sections of the same flat annulus, saving money. The mirror is built up in quadrant sections due to this geometrical feature. Additionally, a cheap substrate like aluminum can be used if it is smoothed first with a lacquer coating. The reflecting layer (gold, typically) can be deposited on top of the lacquer. The conical approximation blurs the psf by arcseconds, typically, exceeding any off-axis aberrations.

Many major X-ray observatories, including *ASCA*, *Suzaku*, and the now defunded *IXO*, have employed thin-foil mirrors. The initial technology continues to improve, with smoother epoxy deposition, multilayer foils, more accurate foil shaping on mandrels, and stiffer housings.

4.7.2 The *Micro-X* Mirror

The maiden flight of *Micro-X* will reuse the mirror from the Supernova X-Ray Spectrometer (SXS), launched from Woomera, Australia in February 1988 [90, 120]. The mirror assembly

[37] *Chandra* adjusted its image plane to within 50 μm, on-orbit.

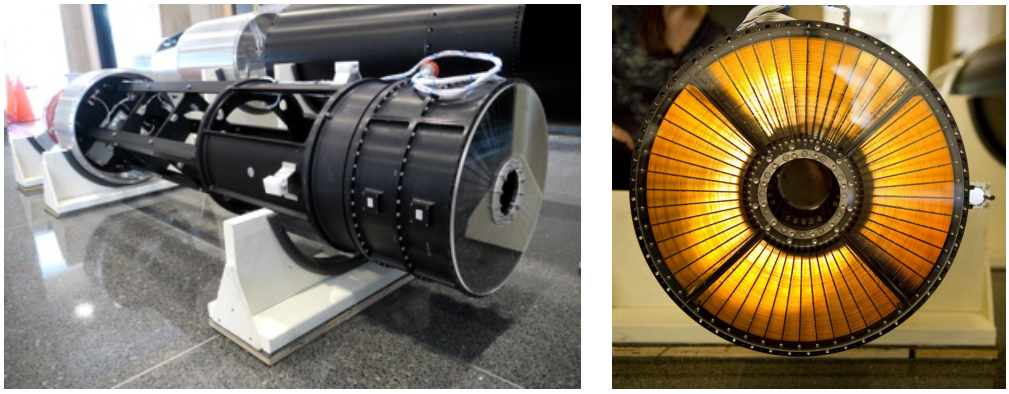

(a) (b)

Figure 4-31: The *Micro-X* mirror is made from nested thin gold-plated foils, a design which allows high throughput without heavy, expensive glass substrates. The assembly – refurbished from the first successfully launched mirror of such a design – achieves 2.4′ HPD resolution, adequately coupled to the 1′ size of the pixels on the sky.

was developed at NASA GSFC in 1987 ahead of the larger Broad Band X-ray Telescope mission, and is quite historical: it was the first conical X-ray mirror to fly in space. The goal of the instrument was to observe young SN 1987A, but failed to do so because of a pointing error[38].

The mirror consists of 68 concentric foils, ranging from 80 mm in diameter to 200 mm, and has a focal length of 2.10 m (Figure 4-31). The foils are held fixed in individual quadrant housing by precisely machined combs. As mentioned above, these foils were formed from lacquer-dipped aluminum ($\sim$ 0.1 mm thick), and evaporatively gold plated, in a similar process to the ASCA mirrors. The entire mirror was coated with a fine layer of dust from impact, but otherwise survived intact.

The *Micro-X* team's initial efforts to clean the mirror foils actually resulted in degraded reflectivity. Facing the possibility of a complicated and costly refurbishment effort, the opportunity to reuse spare foils from Astro E-2 (Suzaku) arose. These newer generation foils, which use epoxy instead of lacquer, feature much better smoothness over millimeter scales.

[38]The telescope did observe X-ray emission from LMC X-1 in a *different* part of the Large Magellanic Cloud, though.

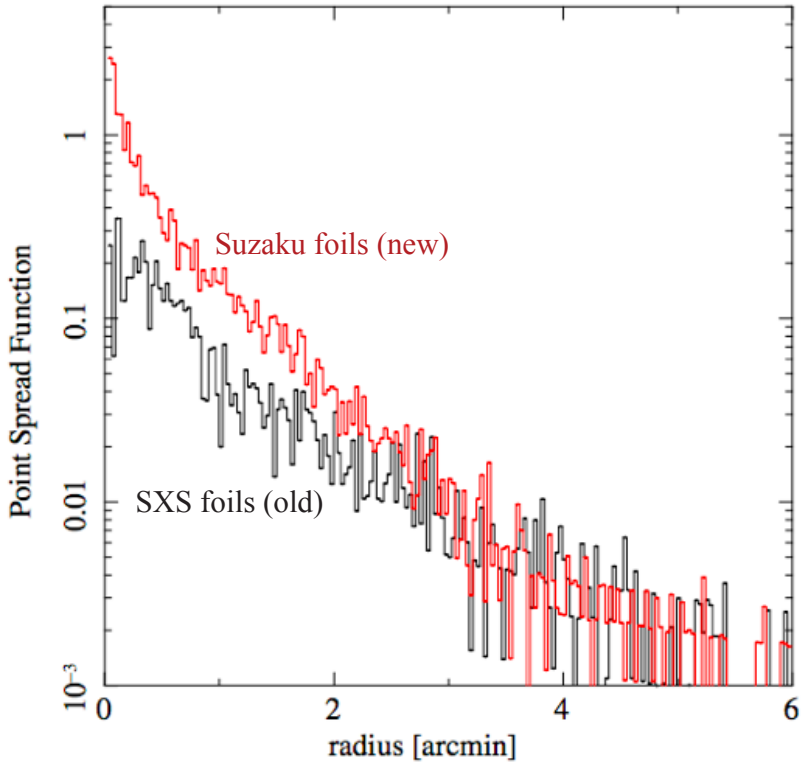

Figure 4-32: The HPD of the X-ray mirror's PSF improved from 3.6′ to 2.4′ by retrofitting spare Suzaku foils for the SXS mirror housing.

The focal length of Suzaku (4.75 m) is over twice that of *Micro-X*, so the foils were re-cut to modify the individual curvatures. Thanks to the existing infrastructure devoted to thin-foil mirrors at GSFC, the trimming and re-installation of the spare foils was straightforward.

The focus of each mirror quadrant was tuned on an optical parallel beam at the GSFC X-ray Optics Lab. The foils were then staked with epoxy before undergoing a vibration test.

The mirror quadrants were calibrated individually on an X-ray beamline after initial optical alignment. The new foils outperformed even the clean spare SXS foils, improving the half-power-diameter (HPD) from 3.6′ to 2.4′ (Figure 4-32). The effective area also increased 35%. Both the reduction of 'orange peel' roughness and the improved substrate manufacturing contributed to these improvements.

Thus, though old, the mirror adequately matches the on-sky scale of the pixels, which is $\sim 1'$.

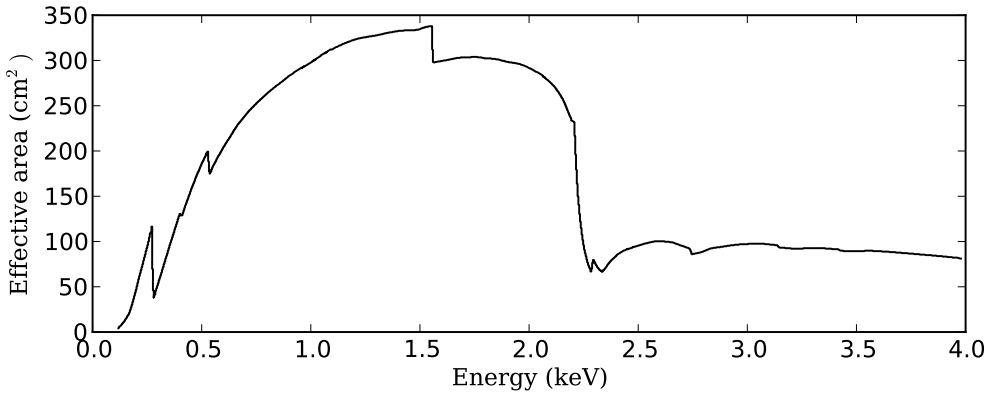

Figure 4-33: The nested shell design achieves high effective area. The curve accounts for the filter transmissions.

After individual calibrations, the quadrants were assembled together to align the focal points and axes. The mechanical housing was tuned to converge the four focal points to within the same 0.1 mm square area. The focal length was determined to be 2100 ± 5 mm and the optical axes were aligned within $0.5'$. No image degradation occurred after the rocket vibration test.

The optical bench was also reused, though modified slightly to accommodate a different payload geometry at the bulkhead. A 'magnetic broom' was added to sweep electrons out of the focal path. An undergraduate researcher designed the configuration of eight NdFeB magnets to deflect electrons with energies as high as 100 keV.

Future upgrades to the *Micro-X* payload will include a higher resolution mirror provided by Marshall Space Flight Center. The full shells (not segmented foils) will use the same electroformed nickel replicated technique used to build the XMM-Newton mirrors. The five rhodium-coated shells have a (modest) target $30''$ resolution and a 2.5 m focal length, though will have just over half the effective area.

4.7.3 Alignment

As detailed above, the focus of the mirror was determined with an optical parallel beam source. To record this focal point, a picture was taken with a digital camera mounted in

the bore of the mirror. The camera was kinematically mounted (fully constrained but not overdefined) to the mirror housing with three springs compressed by screws, and the zoom lens was fixed thereafter with epoxy.

To align the mirror to the TES array, the entire science instrument must be assembled with an open path between the detectors and the mirror, mounted to the optical bench, which is in turn mounted to the bulkhead. The reference camera will then be re-installed, and live images of the focal plane will be compared to the reference photo taken at GSFC. Shims, placed underneath the feet of the optical bench, will adjust the angle between the mirror and array until the focal point lies in the center of the array[39].

With the mirror aligned to the detectors, the mirror may now be aligned to the *star tracker*. The star tracker is an optical camera, which, as the name suggests, tracks the bright stars in the field of view to provide fine guidance information for the rocket avionics. *Micro-X* will use the ST 5000 model, which was developed at the University of Wisconsin with the goals of low cost, high cadence (10 Hz) on low bandwidth downlinks, large field of view (5.4×7.4 degrees), and sub-arcsecond resolution [103]. The ST 5000 tracks the eight brightest stars in the field to determine the pointing.

The star tracker replaces the reference camera before flight in the bore mount, and is aligned by NSROC. A permanent flat mirror, attached to the side of the X-ray mirror housing, was aligned perpendicular to the mirror's axis during calibration at GSFC. NSROC will use this flat to perform the alignment, a digital task which will not require physical adjustment of the star tracker.

4.8 X-ray Filters

While the *Micro-X* design goes to great lengths to stop infrared light from reaching the focal plane, it still, obviously, must allow X-ray light in. What we need, then, is a set of filters in the focal path, not unlike $H\alpha$, UBV, or $u'g'r'i'z'$ filters in ground-based astronomy.

X-ray filters, like X-ray mirrors, require very different geometries than their optical coun-

[39]Full alignment procedures are detailed in the internal document Mirror Alignment Procedure RevB.pdf.

terparts. X-rays attenuate quickly in most materials, so we use an extremely thin film of aluminum – $\sim 1000\times$ thinner than a human hair – to reflect optical and infrared light but transmit X-rays. Such filters are sometimes called X-ray filters, optical blocking filters, or infrared filters; for succinctness I will refer to them as just filters. (These filters bear no resemblance to the earliest versions on Giacconi's Aerobee payload, which were made from thin mica coated with lampblack [47].)

As with many components, the design of the filters relies heavily on *XQC*'s. Luxel, an outside contractor, manufactured the thin films with a highly specialized process. I oversaw the manufacturing, then performed calibrations at the National Synchrotron Light Source (NSLS) of Brookhaven National Laboratory, after initial tests at the MIT CCD Laboratory.

4.8.1 The Filter Stack

Five filters sit in the optical path: one on each of the outer thermal shields, one on the 2K shield, and two on the focal plane housing. The 150 K, 50 K, and 2 K filters are angled with respect to one another to decrease the number of times a trapped infrared photon can bounce between two reflective layers, decreasing the probability of transmission. Only 10 mm separates these warmer filters from each other.

The filters are made via a proprietary process, whereby a thin aluminum layer ($\sim$ 250 Å) is deposited on a substrate, which in our case is $\sim$ 500 Å of polyimide. Polyimide is a strong polymer with good thermal properties, and it is beloved by most cryogenic engineers in the form of Kapton. The film is epoxied to the filter body.

Needless to say, the filters are extremely delicate, as stipulated by the transmission requirements, but robust enough to be installed by hand in the dewar and survive some pressure differential while pumping down the vacuum space. Preventing light leaks in a thermal shield also seals large air leaks, so pressure differentials do arise across the shield boundaries. Channels are etched on the filter mounts to ensure proper pressure equalization before bursting the film. As destructive tests are expensive, though, we only have rules of thumb: the larger, mesh-supported filters can sustain a 40 Torr differential, while the free-standing small filters

can sustain 1 Torr. The initial pumpdown of the vacuum space is done slowly with a needle valve, with $P/(dP/dt) > 15$ minutes.

The aperture of the two filters at 50 mK is 15 mm, to accommodate the field of view of the 12 mm × 12 mm array. The film hangs across a Mg annulus without any support, which is a large area compared to most free-standing Al filter films in other X-ray experiments. The filters on the thermal shields all have a diameter of 28 mm, so must be supported to prevent breakage. We use a thin Si mesh, manufactured by University of Wisconsin at Madison, to prop up the filter film.

Any residual moisture around the filters will freeze on the cold surface, lowering the transmission. The extra O from H_2O gouges the spectrum below the oxygen K edge at 540 eV[40]. Both heaters and baffles are used to combat this problem. Surface-mount resistors act as heaters for the meshed filters, wired together by gold traces on the Si frame. Tens of volts are applied across the heaters in parallel, producing nearly 1 W of heating power to sublimate the water ice. Aluminum baffles are also used to create a maze-like path to the cold filter surface, increasing the probability of adsorption of a H_2O molecule before reaching the filter.

The *XQC* payload has seen thin layers of ice accumulate on the filters during cooldowns, on the scale of nanometers. However, up to two microns of ice have formed during launch, by some yet unexplained mechanism. (Vibration testing alone does not replicate the thick icing.) While devastating to the low-energy science goals of *XQC*, this problem will not affect the *Micro-X* as much, because most of the expected signal of interest lies above 500 Å.

The thicknesses of the filters were determined by modeling the allowable phonon noise in the pixels[41]. The acceptable infrared 'loading' per pixel was set at $< 1\%$ of the X-ray power per pixel, or ~ 10 fW of IR leakage. Starting with a blackbody spectrum at 300 K, the energy-dependent power spectrum was calculated at each stage, including geometrical effects and the shields' blackbody temperatures, all the way to the focal plane, where the spectrum

[40]For this reason, the filters are stored in a 'dry box' with nitrogen gas, to hinder the growth of oxide

[41]`http://space.mit.edu/figueroagroup/internal/micro-x/adr/20090625_filtersims/filter_sims.html`

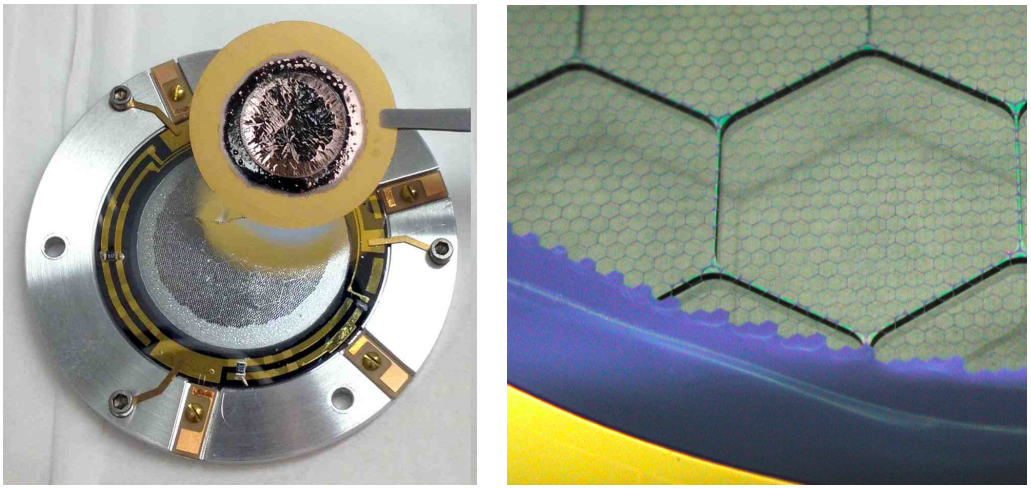

(a) (b)

Figure 4-34: The X-ray filters are made of extremely thin Al, which need to be supported with Si meshes for the larger apertures. A large-aperture filter and FEA-mounted filter (a) are shown together for comparison. The Si mesh (b) is made of hexagons, which fills 3% of the aperture.

was integrated to yield the total IR load. Polyimide was assumed to be transparent in the IR waveband [148], Al transmission values were taken from Henke et al. [57], and Al_2O_3 formation was taken into account.

The simulations explored both the number and relative angles of the filters, and strongly favored two 50 mK filters, angled with respect to the 2 K filter. With realistic shield temperatures and filter angles, the final predicted value is $\sim$ 5 fW leaked power per pixel, an acceptable value.

4.8.2 Supporting Large Apertures with Silicon Mesh

As mentioned above, thin, large diameter aluminum/polyimide films require additional support. A mesh of Si works well because it can be easily made with modern lithographic techniques, and it thermalizes well without conducting electricity. Large diameter filters are important to *XQC*, with its large field of view, so the team has invested a lot of effort into the manufacture of Si meshes. Kari Lynn Kripps, under the supervision of Dan McCammon, has headed the development.

The current generation of Si meshes requires several fabrication steps. The initial 200 μm thick Si wafer has an implanted insulating SiO_2 layer 25 μm deep, which prevents 'lateral leakage' through the Si. Fine and coarse hexagonal unit cells (Figure 4-34) are made by deep reactive-ion etching to the oxide layer from the top and bottom, respectively, after defining the patterns with photo resist. A thermal oxide is grown on top for two reasons: first, the layer will insulate the gold leads from the Si, and second, it rounds the edges of the mesh to prevent pinprick holes from developing in the thin film, which will rest on top. Gold leads are applied around the circumference of the filter with an electron beam evaporator. Finally, on the back side, holes for a kinematic mount (like the one for the detector chip) are made, and the surface is aluminized to reflect IR photons into the blackened baffle.

Si is opaque to X-rays in our waveband, so the transmission decrement is set by the filling fraction of the hexagonal mesh, only about 3%.

4.8.3 Calibration considerations

As the filters are made primarily of just two simple materials with well known transmission properties, only a few measurements of transmission at different energies suffice to infer a filter's response curve over the whole science band. However, several effects become important enough to warrant more detailed calibration: several photoabsorption edges display 'structure' beyond the predicted sheer drop-off, the filter adhesive can effect the transmission, and, of course, more data yields much lower uncertainties during analysis. Thus, I undertook a calibration campaign at NSLS to measure transmissions as a function of energy at spacings of only 2-20 eV, instead of much larger separations with Manson and radioactive sources at MIT.

The characteristic notched shape of a thin film transmission curve derives from X-ray photons ejecting electrons. As a photon's energy increases, the atoms in the filter put up less of a fight letting the photon through. Thus, generally, we see the transmission value asymptote toward 1 with increasing energy. However, once a photon becomes energetic enough to eject an atom's inner electrons, the cross section to that atom increases sharply.

The jump in cross section cuts an 'edge' into the filter's transmission, as more photons are absorbed. At the energies we deal with, innermost shell (K) edges for C, N, O, and Al – the constituent elements for polyimide and the filter film – form the most prominent features at 284, 410, 543, and 1560 eV. The depth of these edges tells us about the thickness of the filter, which we can compare to Luxel's quoted values.

The NSLS facility operates an extreme UV ring and an X-ray ring, with many beamlines fanning out from it. Three national labs – Lawrence Livermore, Los Alamos, and Sandia – jointly operate the two beamlines I worked on. The UV beam can supply photons from 50 eV to 1 keV, while the X-ray beam operates from 0.8 keV to 6 keV[42]. The gratings for each beamline have resolutions dE/E of 3.3×10^{-3} and 0.8×10^{-3}, respectively. Filters and multiple reflections off finely tuned coated mirrors reduce higher order scattered light contamination to below 1%.

Using the infrastructure of the *Astro-H* calibration team, I calibrated 5 flight filters and 3 spares at NSLS. The experiment chamber at the end of the (by that point monochromatic) beam is simple, consisting of a vertical mounting stage, one (movable) photodiode in the beam in front of the stage, and one in back. The stage can be moved in three directions, allowing us to aim the 1.6 mm beam in the middle of a large mesh hexagon. During data taking, the filter remains still on the stage to maintain its precise position, while the front photodiode moves in and out, alternately blocking and revealing the beam to the second diode. The two sensors were calibrated against each other over the passband, and thus their calibrated ratio produces a transmission measurement. The measurements take into account the changing electron beam current in the synchrotron, which decays over time between injections.

The energy scale of the UV gratings is systematically low by ~ 10 eV, but the discrepancy is accounted for by aligning the absorption edges of C, N, and O to their measured values.

The *Astro-H* calibration team has made exquisite measurements of their meshed filters. They discovered that the epoxy makes a healthy second order (energy dependent) contribu-

[42] `http://www.bnl.gov/u3cx8a/`

tion to the transmission, as it wicks up to the mesh. Decreasing the amount of epoxy gained the system an extra 2–3 % of transmission, to beat the target values. Finely mapping the K edges also 'recovered' a couple percent of transmission. *Micro-X* cannot, at the moment, invest the time and money to replicate the *Astro-H* efforts, but they are worth noting when the time comes to translate *Micro-X* technology to an orbital mission.

The analysis methods and results of the calibration run are discussed in Section 5.3.

4.9 Position Monitors

As I have described, LMO is mounted on dampers to attenuate some of the vibrational energy transmitted by the motors through the skin. The dewar will deflect several mm in Earth's gravitational field, depending on its orientation. Thus, the optical axis in flight – with the payload pointed roughly horizontal at Puppis A, depending its position during the time of year – will most likely be different than the axis calibrated in lab. So, being a focusing imager, the *Micro-X* telescope must correct for this offset, or prevent it.

The original vibration isolation design with rubber dampers included both: prevention and correction.

To *prevent* misalignment, clamping jaws were designed to constrain the dewar to the bulkhead after the motors burned out and the vibrations stopped. Individual 28 V motors move metal clamping fingers into precisely positioned divots, in tabs extending off the aft base plate of LMO. With the CMs actuated, LMO is again rigidly aligned to the mirror, as during optical axis calibration. The CM system was successfully commissioned before the Wallops vibration test.

If the new WRI design is adopted, the CM system will most likely not be used. The throw (largest travel distance) of the CMs is limited by the constrained geometry between LMO and the bulkhead. As the WRIs will allow larger displacements under the same acceleration profile, the tabs will collide with the PM.

The alignment system has a redundant *correction* component, though: laser position monitors (PMs). The PM system is composed of four lasers mounted on LMO which shine

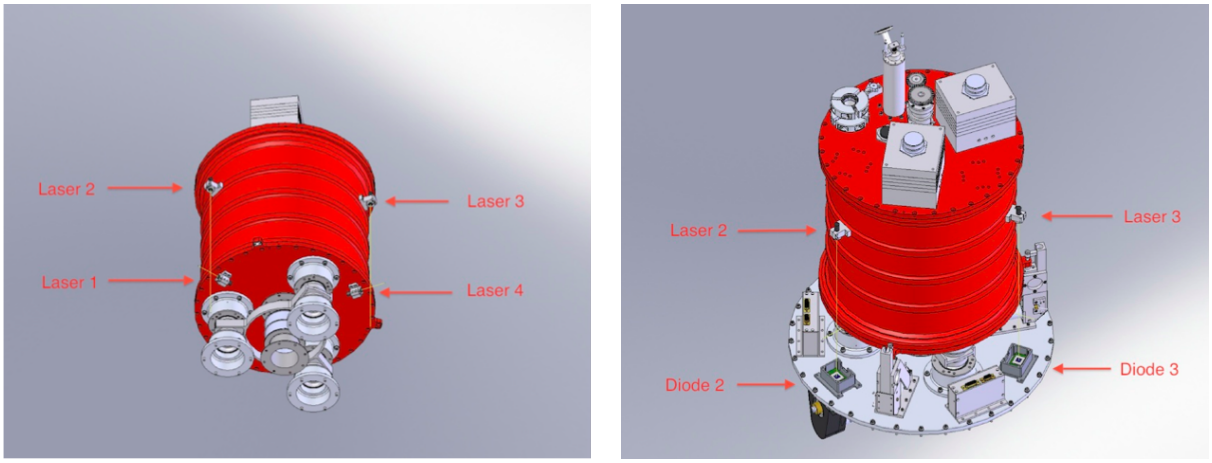

(a) (b)

Figure 4-35: Four pairs of lasers and photodiodes monitor the position of LMO relative to the rocket skin. Diodes 1 and 4 are mounted on the skin, and not shown in these views. These CAD images show the rubber damper configuration, including the clamping mechanisms.

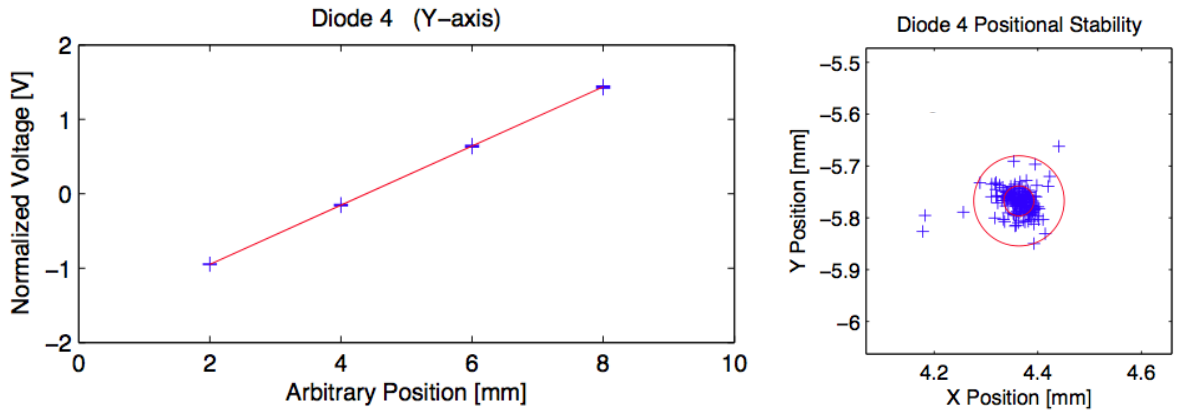

(a) (b)

Figure 4-36: The laser PM diodes exhibit a linear response to the position of an incident laser. The calibration of 0.398 V/mm for the y-axis of Diode 4 (a) enables an estimate of 29 μm standard error induced by output voltage fluctuations (b). The box of Figure b is the size of one *Micro-X* pixel.

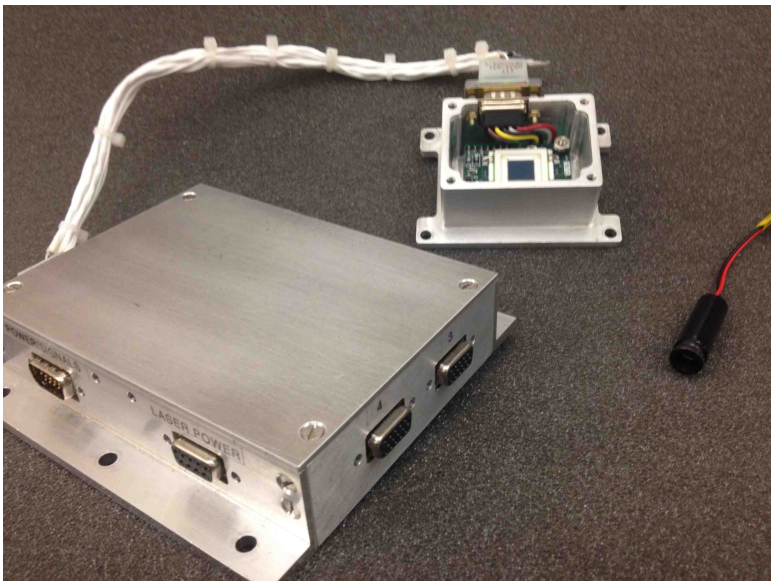

Figure 4-37: The position monitor subsystem, shown here dismounted from the payload and partially disassembled, determines the 3D position of LMO, mounted on soft dampers. A laser (right) shines on a photodiode (center, within white box). The signal is read out by custom electronics (left); the entire subsystem uses four laser/diode pairs.

on four photodiodes, which are mounted on the fixed reference system of the rocket housing – see Figure 4-35 for a CAD rendering. The photodiodes produce x- and y-position-dependent voltages, over their 10 mm $\times$ 10 mm detector chips. Recall, the *Micro-X* detector plane is 7.2 mm $\times$ 7.2 mm.

With the original rubber damper design, the PMs were conceived as a check on the repeatability of the CMs. However, if the CMs are eliminated in favor of the WRI design, the PMs will serve as the primary in-flight alignment correction. The image coordinates will then need to be reconstructed post-flight by correcting the star tracker's pointing with the solid angle and roll defined by the two pairs of lasers.

All of the components are powered and read with a custom board made by University of Florida collaborators (Figure 4-37). The system utilizes RS-422 standards to report positions (x and y voltages) to the telemetry section. The x voltages are negative, while the y voltages are positive, to distinguish the two defined axes; $+x$ points towards the connector of the custom diode housing, and the coordinate system is right-handed as seen by the laser beam.

The individual diodes exhibit slightly different responses, so it is important not to mix them up. Otherwise, the calculated position of LMO will be incorrect.

Phil Oakley performed calibration measurements of the individual diodes, the details of which can be found in internal documentation[43]. The diodes were moved on a fine resolution x/y stage under a fixed flight laser to calibrate Volts/mm linear conversion factors. Each axis of each diode has its own conversion, around 0.400 V/mm. The output voltages for a fixed laser position were stable to a couple percent, translating to $4 - 29$ μm precision, depending on the diode (Figure 4-36). Such uncertainty is well within the 600 μm pitch of each *Micro-X* pixel. The maximum standard error introduced by brightness variations of the lasers over their 3 – 6 V operating range was 50 μm.

4.10 Ground Support Equipment

The ground support equipment (GSE) commands and powers the instrument before flight. During flight, the various subsystems — the MUX, the gate valve, the pumping valve, and so on — are operated on timers and supplied by batteries from the TM. On the ground, for tests in the laboratory and on the range (i.e., at WSMR), the GSE takes on these responsibilities. Additionally, we have added functions to the GSE to aid with other laboratory tasks, such as temperature control (while the ADR electronics are being commissioned) and single-pixel characterization with modular batteries. The NSROC TM team also has their own GSE, which the *Micro-X* team connects through to communicate with the instrument. Thus, the term 'GSE' is neither single-valued nor specific; however, I will use it to describe the laboratory system which I have built to control the dewar and mechanical subsystems.

The GSE was designed to be compact, modular, and mobile, so it can be moved between different integration environments. (At WSMR, for instance, the GSE must follow the instrument as it moves to different test tables.) Figure 4-38 shows how these aims were achieved, by housing a control computer, power supply, and auxiliary instruments within a small rolling server rack. Power and signals are distributed to the subsystems through a

[43] `Laser_Position_System_Summary.pdf`

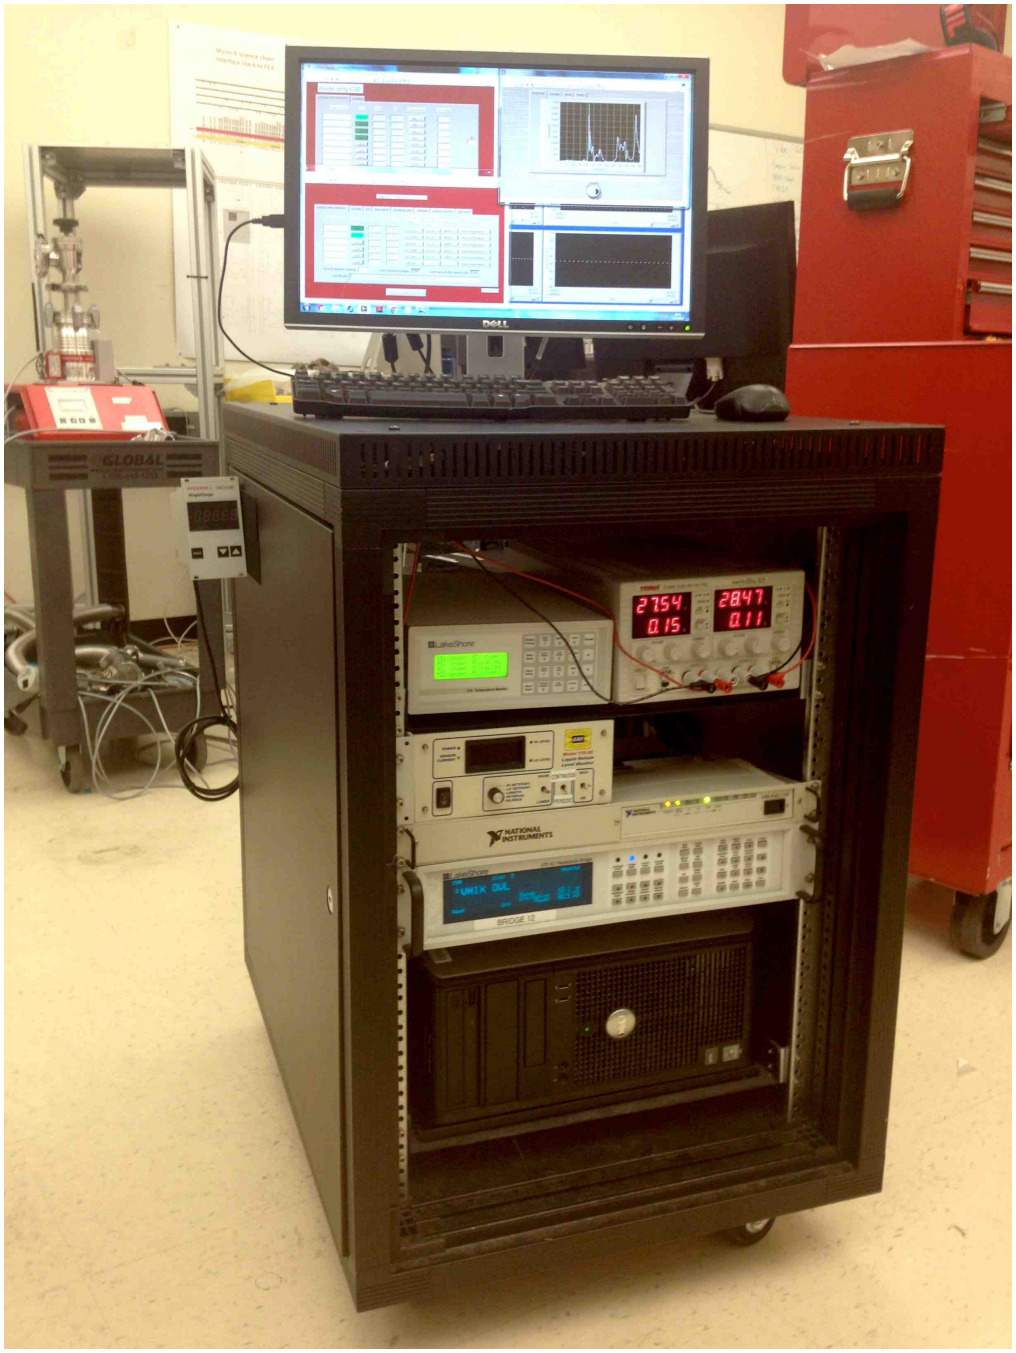

Figure 4-38: A wheeled miniature server rack houses the GSE, providing a portable control and readout platform for the instrument. The current laboratory configuration is shown, from bottom moving upward: the control computer with data acquisition and synthesis card; resistive thermometer readout; LHe level gauge; diode thermometer readout; 28 V dual power; monitor with various running GUIs.

breakout deck (Figure 4-39) hidden underneath the work surface, to a bank of connectors. These connectors have the same definitions of those at the interface to the TM, on the avionics deck, allowing for easy debugging.

In the on-range configuration, the GSE receives housekeeping signals from subsystems and commands them through digital signals. These systems are:

- ADR electronics: housekeeping data (Insert temperature, voltage taps, etc.) are communicated through RS-422, and commands ('ramp up', 'control temperature', etc.) can be sent, via the same protocol.

- Position monitor: the positions of the four lasers are passively read out.

- Gate valve: a motor on the bulkhead opens a hermetic valve, revealing the focal plane to X-rays. A logic board in the payload ensures operation only when the exterior pressure is lower than the dewar pressure (but it can be overridden).

- Heat switch: the heat switch must be operated remotely — as opposed to manually — on the launch rail due to safety considerations. Thus, a stepper motor engages the Kevlar-and-clamp assembly a couple hours before launch.

- Pumping valve: the helium bath is pumped on the launch rail. A valve on the skin seals this pathway just before launch, and reopens on an altitude timer, closing again on descent. Similar to the gate valve, a logic board connected to a pressure gauge ensures proper operation. As soon as the valve closes, the GSE sends a signal to backfill the long pumping hose with gas, helping the flyaway press seal on the pumping valve detach easier upon launch. The backfill is operated by a relay-controlled valve, backed by a high pressure nitrogen gas canister.

- Clamping mechanisms: send signals to the motors to secure the dewar.

The GSE, TM GSE, or batteries provide power for these systems, depending on the configuration.

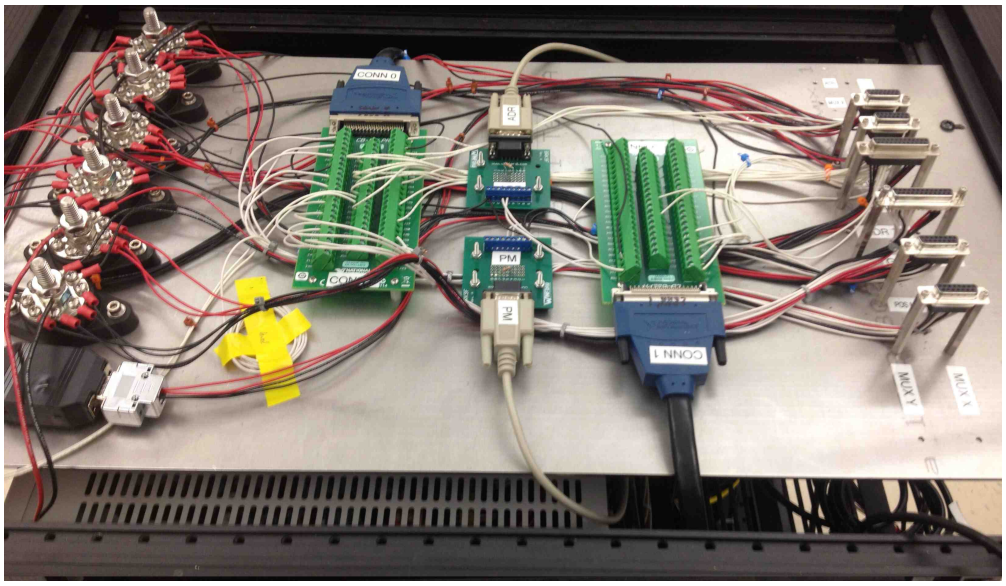

Figure 4-39: The GSE signal breakouts are hidden underneath the top of the rack. From left to right: power distribution blocks, digital I/O (NI PXI 6224), RS-422 signals (position monitors and ADR electronics), second digital I/O breakout, connectors to interface to instrument or NSROC TM.

I developed a suite of software to decode, monitor, and send these signals. The back-end communication and front-end GUI were both written in LabVIEW. LabVIEW was chosen for its strong instrument communication libraries, extremely easy GUI customization, and intuitiveness for non-programmers[44]. 5 V on/off signals ('gate valve close', 'heat switch open', limit switch indicators, etc.) are read and created by a NI PXI-6224 board. The LabVIEW software interfaces easily with the 6224 ports. RS-422 signals (ADR electronics, position monitors) are read into the computer with a USB converter box.

While the computer and power supply persist across GSE configurations, other components in the rack change to provide support for specific commisioning tests. All such electronics can be controlled remotely, and I wrote software to manipulate them via GPIB or RS-422 protocols, incorporating them into the main GSE GUI. The various components which have been part of the GSE include:

- Resistance thermometry readout. The Lakeshore resistance bridge is a mainstay of

[44]The benefits of distributed version control, algorithmic simplification, and readability of a proper programming language, though, would make a strong case against LabVIEW if starting from scratch.

cryogenics, and was copied for the ADR electronics.

- Diode thermometry readout. Many non-essential temperatures are read out with cheap two-wire diodes, such as the thermal shields.

- Low noise voltage supplies to bias individual TESs. Automated voltage sweeps determine the level to bias the TESs at.

- Superconducting magnet power supply. Combined with voltage tap monitoring, the focal plane temperature can be controlled with a PID loop.

- Pressure monitor of the vacuum space.

- LHe level gauge.

- Dynamic signal analyzer: accelerator responses were measured in the lab vibration testbed.

Most of the development has focused on integrating these components into the GSE, to aid in commissioning.

As a final note, the science chain readout is technically also part of the GSE. The current working version consists of four components. First, the data from the MUX stack passes through a TM simulator, which downsamples and reformats the data according to the specifications of the 20 Mb/s encoder. (The actual TM GSE will replace this component on range.) Second, the data stream passes to a dedicated server blade (Windows OS). A client computer, third, displays the multiplexed data for selected pixels with a configurable GUI (Linux OS). Finally, the parameters of the MUX stack, such as individual pixel biases, are managed on a separate computer with a lightweight Mac OS program. These components (and duplicates, for the second science chain) will be housed in their own racks for mobility on range, but currently sit in a desktop configuration.

4.11 Rocket Systems

The *Micro-X* instrument will be supported onboard the rocket vehicle with standard NSROC-supplied subsystems. Most of the information presented here was drawn from the NASA Sounding Rocket Program Handbook, a comprehensive reference for any variety of payload [126]. These systems are rather 'plug-and-play', offered as part of a configurable launch vehicle, so the *Micro-X* team need only worry about the interfaces (connections to telemetry, pointing coordinates), leaving the testing and operation to NSROC teams.

4.11.1 Motors

The Terrier-Black Brant IX family of vehicles were developed specifically for scientific payloads which required higher altitudes. A Terrier MK 70 solid state booster provides the first stage of acceleration, for 6 s. The stage then separates from the sections above it, allowing the payload to coast for a few seconds before the Black Brant performs its 27 s burn. While the payload sections are bolted together, the motors are stacked, and separate by springs during coasting. The Black Brant provides a thrust of $\sim 17,000$ lbs with 2200 lbs of propellant before also being jettisoned, to allow a clear line of sight out the aft section for the X-ray mirror.

The motor manufacturing process produces enough variation in the strength (and yield) of the boosters that 3σ-low and -high values are reported to the science team, often translated to flight times, once the payload mass is known.

4.11.2 Telemetry

Moving upward along the launch rail, the science payload of *Micro-X* sits on top of the Black Brant, and the telemetry (TM) section connects to the top of it.

Micro-X will store the entire timestream of data on flash memory cards, in the MUX stacks for each science chain. The option to telemeter the data down via a persistent radio link to a base station is valuable, though, in the event of payload loss or memory failure

after a successful observation.

Each science chain produces data at over 60 Mbit/s, well beyond the rate the current generation of encoders can attain. Thus, the telemetered data will be decimated. Two 20 Mbit/s WFF-93 Pulse Code Modulation encoders will be used for each of the science chains, as well as a 1.5 Mbit/s encoder for housekeeping signals. The data are sent down to a WSMR base station via three independent S-band links (2.2 GHz). The 20 Mbit/s models are new technology; the highest rate encoders available to sounding rockets previously was 10 Mbit/s. At the time of writing, no science payload has used the 20 Mbit/s version.

No uplinks for commands will be used for *Micro-X*, as all events (start of temperature control, gate valve opening, etc.) are on timers, fixed to the launch time. Future launches could employ a command link to steer the payload to point at separate targets, or to correct mis-pointings.

The TM section, consisting of the three antennas and encoders, as well as timers, batteries, and power management electronics, sits just above the science instrument. A bank of connectors at the top of the *Micro-X* electronics deck defines the interface between the two sections (Figure 4-1 (b)). The connector and pin definitions were agreed upon early in the development of the payload, to allow the *Micro-X* science team and NSROC TM team to develop the sections independently.

4.11.3 Boost Guidance

As a first order method of stabilization, the rocket is spun up to ~ 4 Hz, much like a bullet. Counterweights are deployed after the motor burnout to de-spin the payload, removing angular momentum.

The *Micro-X* payload also employs active stabilization. The S-19 Boost Guidance System (BGS) helps control the trajectory of the rocket during the early portion of flight. The S-19 is known as a constant attitude control system, as it maintains a programmed pitch and yaw throughout the first 18 s of launch. By using fins to correct gyroscope-detected deviations from the set angle, the BGS improves the trajectory accuracy by as much as a factor of 10,

and allows launches in windier conditions. The fins, known as canards, are controlled via an analog feedback circuit which deflects the two pairs of axles with pneumatic servo motors.

4.11.4 Attitude Control System

Once the payload has begun its free-fall parabolic trajectory after the motor burnout, and the atmosphere becomes thinner, fine pointing of the entire payload becomes feasible with small bursts of gas. Fine pointing of the payload will allow a long exposure of the target field of view, making the Attitude Control System (ACS) a critical element of the flight vehicle. The control loop incorporates the electronic signal from the star tracker (Section 4.7.3) to point at a fixed point on the celestial sphere. The microprocessor of the ACS calculates the gas outputs of pneumatic valves from two reserves of Argon, typically, at different pressures.

4.11.5 Nosecone

The Ogive[45] Recovery System Assembly (ORSA) deploys the parachute for reentry from within the nosecone. The nosecone is jettisoned on the upward trajectory, revealing the parachute, which is opened on the descent via an altitude timer. Post-flight calibration data can be taken on the relatively smooth descent.

4.12 Launch Logistics and Countdown Timeline

Micro-X will launch from White Sands Missile Range (WSMR) from a Navy-supported launch rail[46]. *XQC* has launched from the same rail, so its infrastructure is well known to the collaboration.

Electronic signals are routed out of the payload via umbilical cables, which terminate

[45]The archaic word 'ogive' describes the profile of a rounded, pointed arch, like the ribs of a Gothic church. Modern rocketry has adopted the word for the 3D shape produced by revolving such a profile about its axis of symmetry.

[46]Sounding rockets typically launch hanging off an extensible rail at a steep angle, in contrast to the common notion of a launch pad for larger vehicles. The conspicuous US Navy facilities on the Army base are known as 'the desert ship'.

in a blast-proof 'hutch' at the base of the rail. This hutch also houses the pumps for the LHe bath. The signals connect to a block house some 100 m away through an underground fiber optic link. The telemetry team feeds the link into their GSE, through which our GSE communicates to the instrument.

Integration of the payload takes place for 2–3 weeks before launch. The time allows the NASA Sounding Rocket Office Contractor (NSROC) teams — telemetry, guidance, etc. — and the *Micro-X* team time to ensure their subsystems are working, and to interface the various sections together. All subsystems except the motors integrate together in the capacious Vehicle Assembly Building (VAB), a mile or so from the launch rail. (The newly renovated VAB even includes clean tents for contaminant-sensitive missions.) The Navy and NSROC usually overlap their support windows for multiple payloads.

In the final days before launch, end-to-end tests, known as 'horizontals', proceed to check that all systems function individually and together. Timers, telemetry data quality, gas ballast, and even pyrotechnic deployments are tested through the whole launch timeline. The final run-through test — 'range vertical' — occurs on the launch rail, extended to the launch angle.

Table 4.1 outlines the events which proceed on timers during the launch.

Launch will be scheduled for solar midnight, with the possibility of delay due to wind conditions. The science team accompanies the payload recovery team on the helicopter ride down range, to ensure neither the LHe nor the calibration source presents a safety issue.

Time (s)	Altitude (km)	Event
T – 9600	1.2	Ramp ADR up
T – 6300	1.2	HS open; Ramp ADR down
T – 5100	1.2	ADR hold
T – 2400	1.2	ADR regulate; calibration data
T – 150	1.2	PV close
T – 90	1.2	ADR hold
T	1.2	Ignition
T + 0.6	1.2	Rail release
T + 6	3	Terrier burnout
T + 12	6	Black Brant ignition
T + 18	9	S19 canards decouple
T + 45	45	Black Brant burnout
T + 60	77	De-spin
T + 68	92	Nose tip eject
T + 68	92	PV open
T + 70	96	Open shutter door
T + 70	96	Clamping mechanisms engage
T + 77	106	ADR regulate
T + 80	114	Start MUX
T + 111	160	Open GV
T + 111	160	Start observation

T + 273	283	Apogee
T + 438	160	Close GV
T + 438	160	End observation
T + 467	112	Close PV
T + 470	107	Release clamping mechanisms
T + 478	92	Close shutter door
T + 619	5	Parachute deploys
T + 921	1.2	Impact

Chapter 5

The *Micro-X* Telescope: Data

The original launch window for *Micro-X* was slated for the first months of 2011, an adequate cushion from the time the payload development was funded in 2007 (proposed in 2006). As has been mentioned many times in the previous chapter, the *Micro-X* design is largely based off the successful *XQC* payload, a fact which encouraged the brisk timeline. However, a number of setbacks have pushed the launch past the original target date. At writing, the maiden launch date cannot be set until vibrational coupling issues are solved.

Most of the hurdles are directly related to the differences between the two payloads:

- The SQUID MUX readout, necessary for a large array of TESs, has proved very susceptible to noise. The design transfer of laboratory 'crates' and 'towers' to compact flight electronics altered the team's intuition of grounding, resulting in long iteration timescales while unfamiliar signals were interpreted. In addition, low yields plagued the custom *Micro-X* Interface Chip, which were subject to inter-layer electrical shorts.

- It was difficult to make the magnetic shielding for the TESs and SQUIDs impervious to stray fields. Some shielding intuitions — gleaned from modeling and from the rules-of-thumb in the cryogenic community — proved to be outright wrong. Again, being a cryogenic experiment, the iteration timescales were long to remedy these errors and converge on a reliable design.

- Vibrational coupling between isolation stages has plagued the experiment of late. The Kevlar suspension was altered from the *XQC* design to accommodate the magnetic shield, thereby introducing softer resonant modes.

The delays caused by these issues have prohibited not only flight data, but also rigorous calibration of the focal plane TES array. Results presented in this chapter represent the latest understanding of the science instrument, which is scheduled to improve dramatically with an integration run at GSFC in the fall of 2013.

5.1 Array Calibration

Characterization of the focal plane array has taken place in three dewars, with different electronics configurations, ranging from a laboratory setup to flight-like layout.

- GSFC lab readout. The reliable testbed features thorough magnetic shielding and 'legacy' readout electronics.
- Insert in Gonzo. The flight-configuration cold cabling mated to legacy readout through a breakout box.
- LMO flight configuration, with end-to-end flight electronics and cabling. (Some signals were picked off the flight cables for diagnostics and isolation.)

The data presented here was obtained with the 170 mK (T_c) array. The array has become the workhorse for debugging the science chain due to the poor performance of the salt pill. A 120 mK array will be characterized during the fall 2013 GSFC integration run.

5.1.1 Pulse Data Reduction

With the GSFC microcalorimeter group's partnership in *Micro-X*, the MIT team can avail ourselves not only of highly valuable institutional knowledge, but also a mature, automated suite of analysis software.

The software takes in a timestream of pulses (usually snapshotted with triggers to reduce file sizes), and outputs spectra with noise estimates. Some manual input is required to define 'good time intervals' and cuts for the noise level and pulse heights, but otherwise the pulse fitting is automated. A GUI provides intuitive data visualizations to perform the manual tweaks. (The process differs significantly from event-file-based data reduction in *Chandra*, for instance.)

The engine of the software is the *optimal filter*. An optimal[1] filter assumes that a slice of data containing a pulse can be modeled as a pulse template times a scaling factor, plus noise. The pulse template and model noise are created by averaging many pulses and many noise traces. It can be shown analytically that the maximum likelihood[2] pulse height depends only on these templates and a constant scale factor [35]. The pulse and noise templates need only be calculated once, then inverse-Fourier transformed into the time domain. This feature ensures a speedy algorithm.

With the software, individual device characterization is also easy. For instance, we can compare pulse heights or decay times for different TES bias points. Or, different TESs can be compared to diagnose variations across the array, due to stray frozen-in magnetic fields or changing heat loads.

5.1.2 Pixel Characterization

Megan Eckart performed the initial array characterization at GSFC. A representative sample of pixels were subjected to a standard battery of tests.

Once the SQUID amplification stages are adjusted, there are generally two tunable parameters which will determine the array performance under operation. The first is the bath temperature, while the second is the TES bias voltage, which will determine how far into

[1]The conditions for optimality are, in fact not met. First, the detector response is not linear, since the TES will saturate for high energy X-rays, changing the pulse shape. Second, the noise is not stationary (independent of time), since the TES Johnson noise changes as the TES changes resistance. Moreover, the L/R response of the control circuit also changes, affecting the roll-off response frequency. The method is still used, though, due to convenience, speed, and accuracy.

[2]The use of maximum likelihood relies on the assumption of random normal errors in frequency space of the estimated noise over the average noise, allowing the use of χ^2.

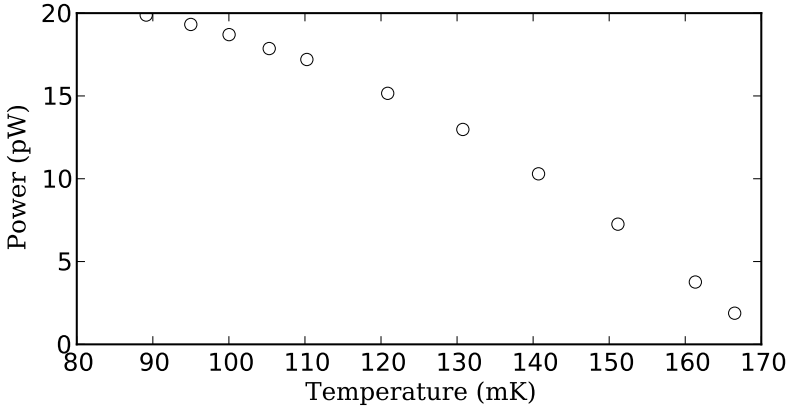

Figure 5-1: The G value of individual pixels can be measured by finding what power (i.e., bias voltage) is necessary to bring the TES exactly halfway into the transition, at various bath temperatures. This GSFC measurement found $G = 410$ pW/K at T_c (governed by the slope of the fitted curve), within the designed range.

the superconducting transition the device sits while quiescent.

The strength of the thermal link, G, can be determined by varying the bath temperature and seeing what bias voltage is required to set the TES halfway up the transition. (Biasing to 50% of the transition ensures that $P(T_c) = 0$.) This measurement describes a curve of input power as a function of temperature, from which the G value can be read off as a slope. Figure 5-1 shows the measurement for a pixel of the array, which confirms that the perforated membrane produced a G value in the desired range, approximately 410 pW/K at the transition.

With a measurement of $G(T_b)$, the bath temperature T_b is set to produce the desired time constant for the X-ray pulses.

The characteristics of the TESs change with bias. Pulse traces at different biases are analyzed for their heights, time constants, and NEP. (The biases are reported as percent of the TES normal resistance, but recorded in lab as voltages across the TES, typically ~ 1 V.)

Figure 5-2 shows how the pulse height — for a monochromatic X-ray source — can change by a factor of 3 depending on the TES bias point. Lower bias points utilize a higher dynamic range to produce larger pulses; the TES has more room to run.

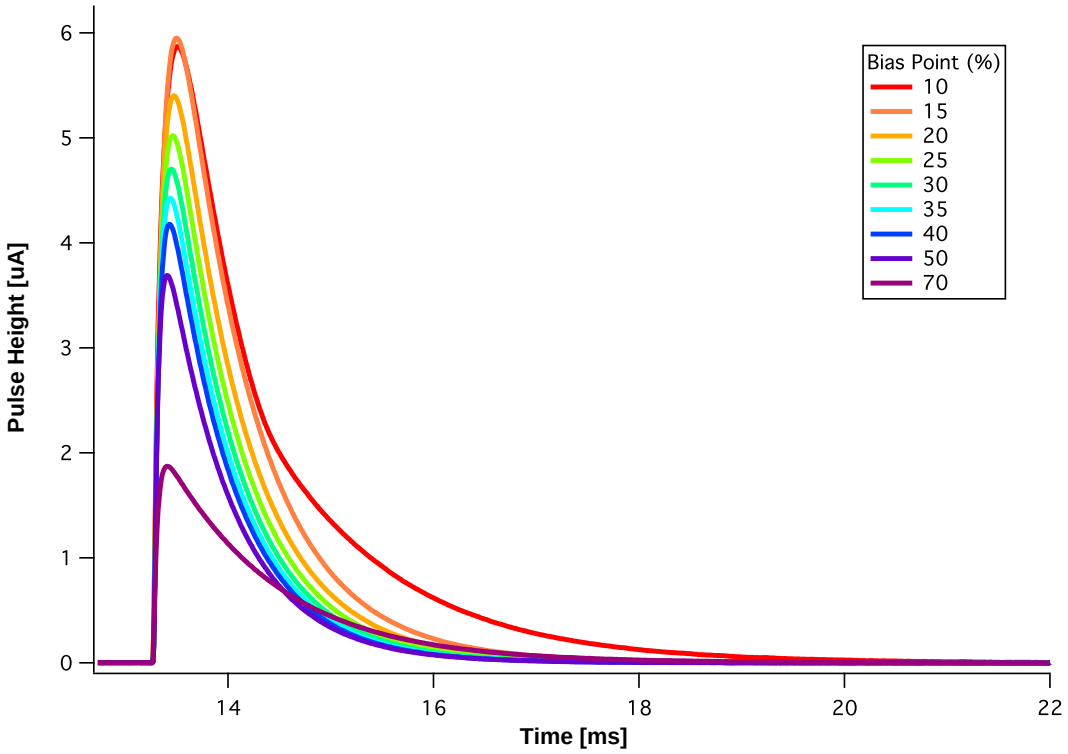

Figure 5-2: The pulse height changes with bias point: at lower biases, the TES has more dynamic range, allowing for higher pulse heights (and, thus, higher S/N). This measurement, taken at GSFC, was replicated at MIT. The legend indicates the color for different percentages of the normal resistance.

The decay time, τ, will also change with bias point. By eye, it is evident in Figure 5-2, that the average pulses at 10% and 70% bias are slower. This observation is quantified in Figure 5-3. The NEP, too, varies with bias, and the minimum of the curve usually determines what bias voltage to set for the pixel during operation.

These results were repeated at MIT, but not replicated, due to magnetic field leakage and readout noise. As detailed in Section 4.5, an impermeable design for the superconducting shield was achieved only after many iterations. Some components of the dewar (including the magnetic shielding itself) were found to be permanently magnetized, resulting in inhomogeneous fields which varied from cooldown to cooldown. This stray field suppressed the T_c of the array by as much as 70 mK, and slowed the pulses by factors of 2 or more. The power-versus-temperature curve (Figure 5-1), for instance, was pushed down and flattened,

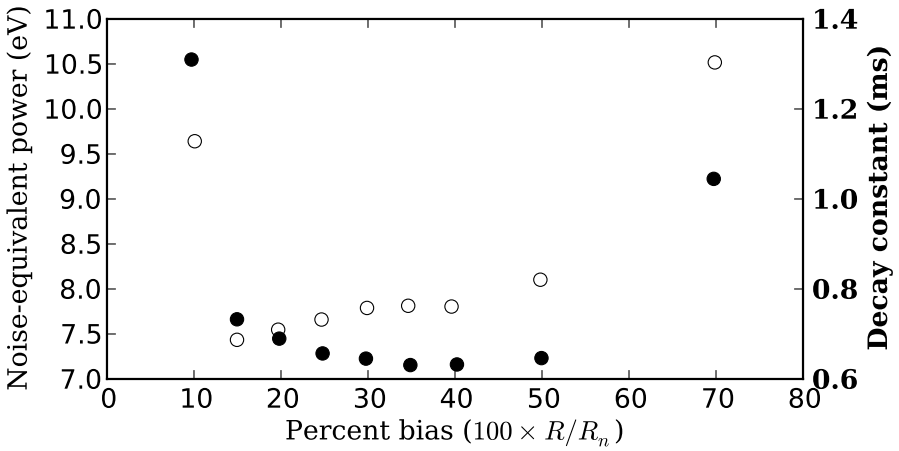

Figure 5-3: The noise-equivalent power (NEP), in eV, varies depending on the bias point of the TES (left axis; hollow points). Thus, an optimal bias point can be chosen to minimize the noise. This measurement, taken at GSFC, indicates that the TESs should be biased 15% of the way up the transition. The pulse decay times (right axis; black points) also vary with bias. This plot quantifies the different pulse shapes at the extreme biases seen in Figure 5-2.

behavior which was mimicked at GSFC with an applied field.

With the operating temperature and voltages chosen, calibration with an X-ray source can proceed under temperature control.

5.1.3 Spectra

The standard calibration source is radioactive ^{55}Fe, which decays via electron capture, yielding Mn K fluorescent lines near 6 keV. The relevant lines are Mn Kα_2 at 5.888 keV, Mn Kα_1 11 eV higher, and Mn K_β at 6.490 keV. The pair of Kα lines, being closely spaced, is useful to demonstrate high resolution spectra.

Figure 5-4 shows the pulse heights of ^{55}Fe decay X-rays over time in the MIT lab dewar. The plot shows temperature-controlled data to start, then drifts once the control is released. (The data were taken to determine if the lab GSE temperature controller added noise; in typical operation, the temperature control operates full time.) This plot represents the heads-up information available to the analyst after a first pass through the data.

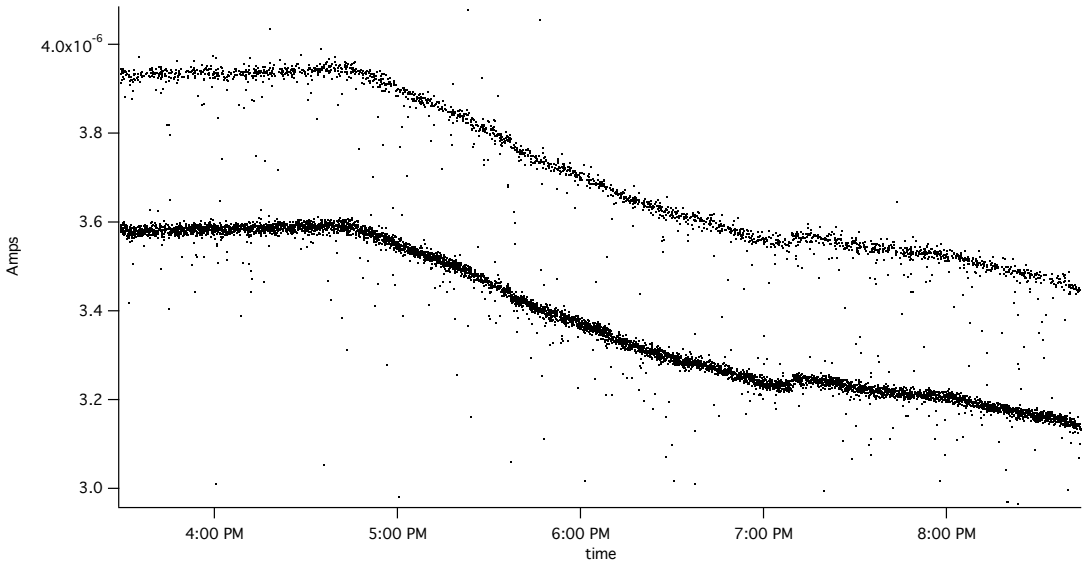

Figure 5-4: The heights of individual pulses are easily tracked with the analysis software GUI. The higher yield Mn Kα line blend is the lower, thicker stripe, with the Mn Kβ line above it, at $\sim$ 600 eV higher energy . For the first hour of this (illustrative) run in the MIT lab dewar, the temperature was controlled with the lab GSE. Then, the array was allowed to drift up in temperature, resulting in higher bias points, and thus lower pulse heights. This test confirmed that the temperature control added no further noise to the system. The SQUID amplification chain loses lock towards the end of the run, resulting in a glitch as the system finds another equilibrium point.

The pulse heights, determined via an optimal filter, are binned into a histogram to produce a spectrum. Spectra of the Mn Kα pair — one from GSFC, one in the best conditions in Gonzo — are compared in Figure 5-5. The 9.5 eV resolution (FWHM) achieved at GSFC approaches the lowest achievable resolution for this specific device configuration, determined by theory. It is significantly higher than the routinely fabricated $\sim$ 2 eV devices.

The 25 eV resolution achieved at MIT — in a noisy, magnetically impure environment — is poor, but still over 5 times better than the ACIS CCDs at the same energy. Again, the resolution is expected to improve below 10 eV with a new array (120 mK T_c), heavily filtered flight power supplies, improved magnetic shielding, and a more consistent grounding scheme for the flight electronics. Integration during the fall of 2013 will test this improved setup in the flight dewar.

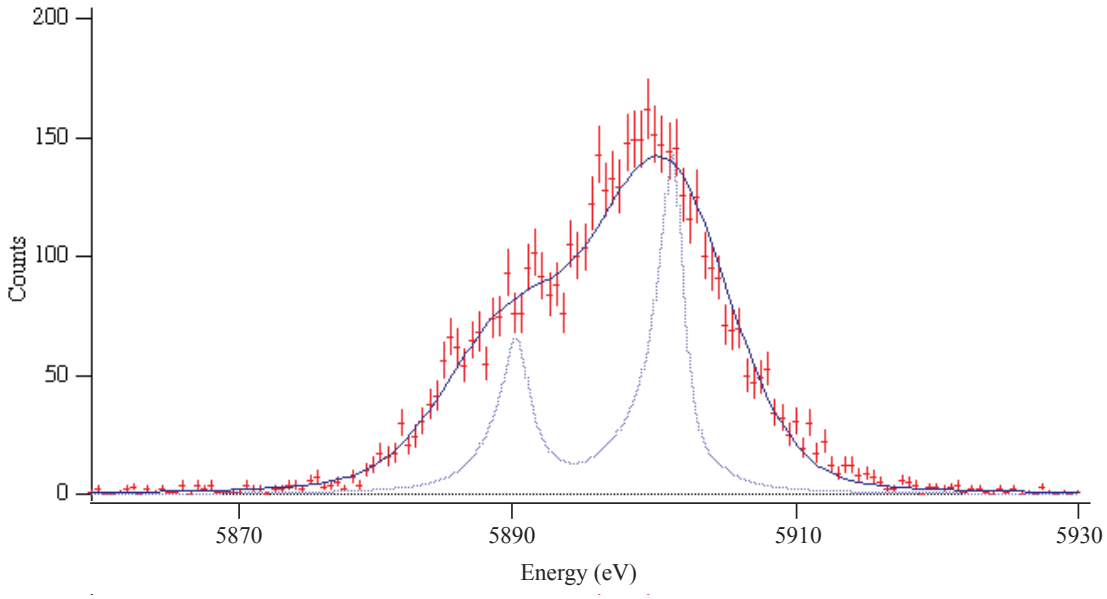

(a)

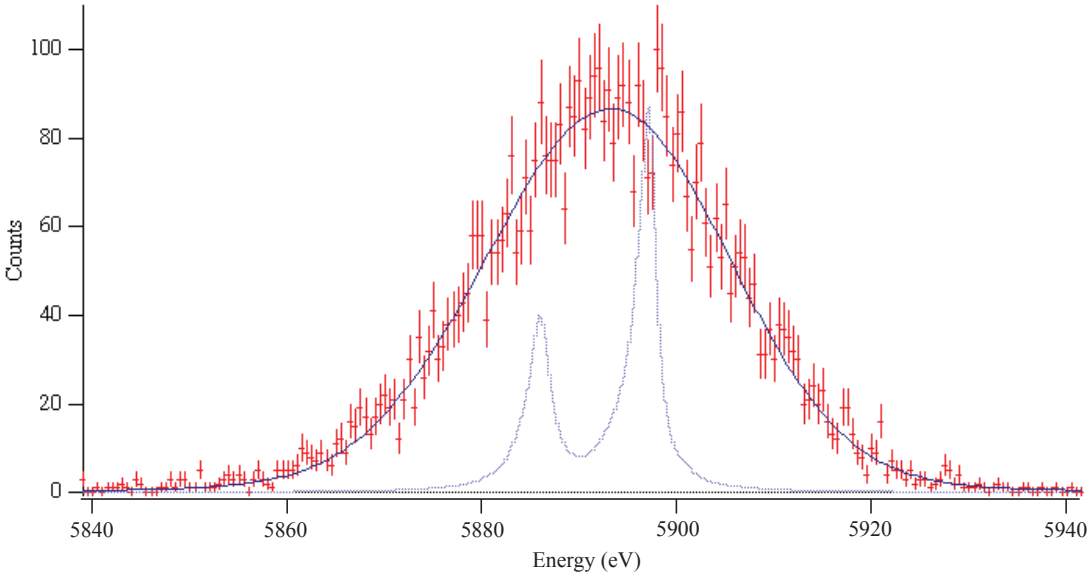

(b)

Figure 5-5: (a) The 170 mK test array was calibrated at GSFC, exhibiting an energy resolution of 9.5 eV (FWHM), consistent with its theoretical baseline resolution. (b) The best performance of the same array at MIT, in the lab dewar with legacy electronics, was nearly 3× worse, at 25 eV. The setup suffered, at the time, of RF pickup and penetrable magnetic shielding.

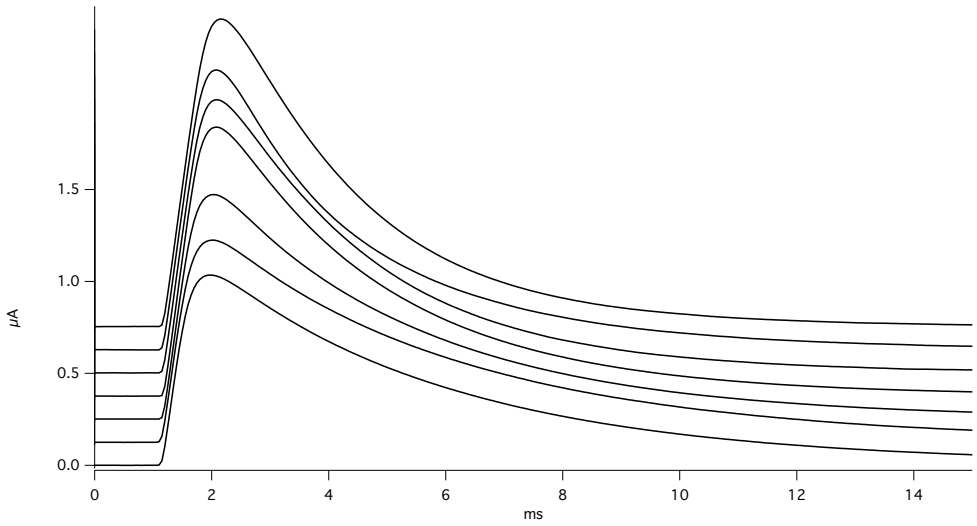

Figure 5-6: The average pulses on seven *Micro-X* pixels demonstrates that multiplexing works in the flight dewar with the flight electronics. The 5.9 keV Mn Kα pulses have been offset on the y-axis for clarity. Note the undesirable different pulse shapes, which may be due to local magnetic field differences.

The above results at MIT were furthered with a muliplexing demonstration on several pixels in LMO, the flight dewar, with the flight electronics. The average pulses for the 7 pixels[3] are stacked in Figure 5-6. The long and variable decay times evident in Figure 5-6 reveal the need for better magnetic shielding, which was not finalized at the time.

X-ray pulses were only achieved in the flight configuration after grounding return lines to the dewar's vacuum jacket. This configuration is different than most laboratory systems of TESs, which float the signal lines as they go in and out of the dewar, tying them to a ground further upstream. This simple difference — realized after months of incomprehensible data, and requiring cables to be remade — illustrates just one of many pitfalls of translating laboratory systems to space hardware.

Further array characterization in a flight configuration has been on hold while diagnosing vibration coupling and building a quiet, isolated power converter for the electronics.

[3]A combination of broken pixels and software limitations led to this number.

5.2 Noise Floor

In an ideal microcalorimeter readout system, the noise from the TES dominates all other sources, such as amplifiers. Johnson noise (of the TES resistor) and phonon noise (of the thermal link) — both inescapable, natural noise sources — will dominate in this regime. The former is white across our band of interest ($\sim$ Hz to MHz). The latter arises from power fluctuations, which translate to temperature fluctuations in the TES. The temperature fluctuations are damped, though, by the thermal link, rolling off above $1/(2\pi\tau)$, where again $\tau = C/G$.

Adding electrothermal feedback (ETF) suppresses the Johnson noise at low frequencies, as the monitor circuit can correct for the fluctuations. However, the circuit itself adds noise in this regime, cancelling the benefit. It is beyond the scope of this work to show that ETF does not improve (or degrade) the resolution of microcalorimeters [35].

However, so far, noise from external sources — introduced along the TES bias line, for instance — has dominated the inherent noise in the TES itself. We estimate the current best setup at MIT is $10\times$ noisier than the legacy system at GSFC. This added noise translates, for example, to distinct fuzziness on pulse signals in the timeseries.

The system noise can be analyzed in real time with SQUID/MUX GUI. The software Fourier transforms incoming (noise, not pulse) traces, updating a noise spectrum in real time. This feature has been the primary debugging method, allowing the team to iteratively change the grounding, filtering, or external environment, then view the effect.

We can also isolate individual signal lines and discover their susceptibility to noise, a measurement known as a *transfer function*. The method proceeds by applying white noise across a line — a feedback loop or bias, for instance — and recording the response at the TES (the SQUID 1 feedback loop). The response tells us, as a function of frequency, whether the input signal is suppressed or enhanced, and will reveal characteristic roll-offs for certain signals. If the input noise is amplified, we know to take special care to provide a clean voltage signal for that line.

Radio frequency (RF) coupling has proven a serious problem for the laboratory and flight

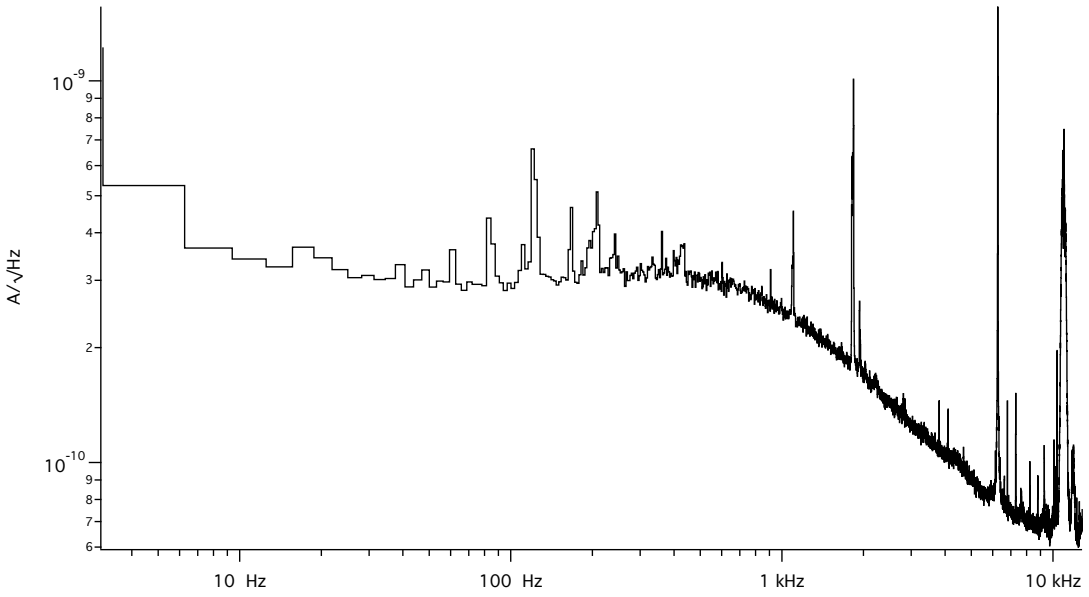

Figure 5-7: The best average noise of the readout system is 5× greater than achieved at GSFC. The roll-off is due to the electronics L/R, while RF pickup reveals itself in the lines above 1 kHz. Aggressive grounding and low-pass filtering will improve future iterations.

dewars. The building wireless signal polluted the bias signals, adding square pulse 'announce' signals to the constant voltages. The wires from 300 K proceed to the focal plane in twisted pairs, but still picked up the signals, functioning like antennae. Dissipative filters and ground planes were added to the Interface Stack to address these problems. The dissipative filters were easily added as flexible connector inserts, which were custom made to short the ground pins to the dewar chassis and pass the signal pins through a low pass filter, via miniature RC circuits.

In an effort to understand the largest noise sources, the entire noise budget of the circuit was derived analytically. The derivation refers all noise sources at all amplification stages to the noise at the TES. Such values are known as *input-referred* levels, as every source gets translated to a common denominator, which can be compared directly to the TES Johnson noise. The derivation and full expression for the input-referred noise can be found in internal documentation[4]. It considers noise from power sources, Johnson noises of resistors at various

[4]`https://wikis.mit.edu/confluence/download/attachments/81035444/ScienceChainNoise.pdf`

temperatures, intrinsic noise of SQUID elements, and the mutual inductances between stages.

The analytical model yielded limited predictive success. The feedback loop on the second stage of SQUID amplifiers, for instance, is a much more influential noise component than the model suggests, as revealed by measurements.

Thus, redesign decisions have been governed by empirical findings. Integrating the MUX and IF stacks together revealed the need for higher capacity ground lines. (The two systems were supplied by different teams, so this need was only found during integration.) The susceptibility to noise — especially on the TES bias lines and the second stage SQUID — prompted the decision to isolate the power unit, outside the MUX stack, to provide the cleanest voltages to the electronics.

The spectrum in Figure 5-5(b) was taken in the best achievable noise environment, but with readout noise well above the Johnson noise of the TESs. Figure 5-7 shows this noise environment. While the expected TES Johnson noise was 50 pA/$\sqrt{\text{Hz}}$, the measured noise during operation was 300 pA/$\sqrt{\text{Hz}}$. (The pulse signals are typically a few μA.) This noise is much higher when biased, compared to zero-bias superconducting or normal states, implying excess noise coming down the TES bias line.

Noise reduction is a critical aim of the next end-to-end integration of LMO and the electronics.

5.3 Filter Calibration

5.3.1 Calibration Considerations

Before describing the analysis procedure, I will revisit the experimental setup in more detail.

Two photodiodes are required to perform a transmission measurement where the delicate, carefully aligned filter remains fixed – one in front of and one behind the filter. After the user centers the beam in the middle of a large mesh hexagon, the beamline software takes over to perform a standard set of measurements at the user-prescribed energies. At each discrete energy (typically separated by 2 – 20 eV), the photodiodes take 20 measurements.

The so-called 'monitor push-pull method' proceeds as follows. Five measurements on the 'detector' diode are made with the 'monitor' diode out of the beam, allowing simultaneous monitor background measurements. Ten measurements follow with the monitor in the beam, while the detector takes background readings. The sequence finishes with five measurements with the monitor diode retracted again.

An upstream monochrometer ensures a beam of photons with a single energy, while a set of filters downstream cleans the beam of higher order harmonics and scattered light. At the time of measurement, the absolute calibration of the energy scale was biased by approximately 10 eV across the band, and the estimated spread in the photon energies for the beam was estimated to be 1 eV. The former systematic source of error can be accounted for by floating the energy scale during the fitting procedure. The latter source of error is ignored, because most measurements are taken with 20 eV spacing.

Several aspects of the data-taking process inform the calibration analysis. First, the energy calibration of the system is off. Second, the beam flux changes at each energy, due to the efficiencies of the various monochrometer configurations. Third, the efficiency of the diodes changes with energy. Fourth, the synchrotron beam decays over the course of the slow energy scans.

The first three concerns are related to energy-dependent responses. As mentioned above, the energy shift is easily accounted for, as the absorption edges are measured to high precision in the literature. The second point highlights the importance of background measurements, as the background level becomes comparable to the lowest signal measurements, while for other energies, the signal is orders of magnitude above the background. A simple calibration run without the filter in place solves concern number three. With this diode calibration data, the individual responses can be normalized out to obtain the correct transmission value.

The decay of the beam, concern number four, turns out not to noticeably affect the 20 transmission measurements at a single energy. The slow drift, therefore, does not need to be subtracted out. This fact is important to the fitting method I describe below.

5.3.2 Extrapolation

The ideal calibration would utilize well trusted photon sources and detectors to produce an *empirical* response across the whole science band, to whatever energy precision necessary. While the National Synchrotron Light Source represents a good approximation to such an ideal instrument, resource constraints (time, mostly) prevented me from covering the whole science band of *Micro-X*, up to 2.5 keV. Only two of the eight filters could be calibrated on the X-ray beam beyond the Al K edge at 1.56 keV.

However, the calibration curves are fit well by analytical models with only a few free parameters, namely the thicknesses of the filter materials. The fits to the low-energy data, therefore, can be extrapolated to higher energies to provide a good estimate of the unmeasured values. Moreover, each data set can tightly constrain the calibration of filters from the same manufacturing run, and inform the transmission values of filters with similar thickness parameters.

Thus, for future analyses with astrophysical data, both the measured and extrapolated/interpolated values can be be used, to suit not only the specific energies but also the desired precision.

5.3.3 Analysis Pipeline

To be explicit, the data of interest are the ten measurements of the two diodes and their background levels at each energy for each filter. A standard analysis proceeds by averaging the measurements, subtracting the background means, and taking the ratio at each energy bin. A least-squares algorithm then produces a best-fit transmission curve for the averaged data across all energies. This step is simple and takes only a few milliseconds to compute.

The transmission of light through a material is governed by its (energy-dependent) attenuation length:

$$T(E) = e^{-t/\tau(E)} \ .$$

Here, T is the transmission, t the material thickness, and τ the attenuation length. Trans-

mission is multiplicative, so for two materials,

$$T(E) = e^{-t_1/\tau_1(E) - t_2/\tau_2(E)} \ .$$

Attenuation lengths were taken from the online Center for X-ray Optics database[5] [57]. Thus, the only parameters for the transmission curve are the thicknesses of Al, polyimide, and O (to account for Al_2O_3 formation), and a nuisance energy shift. Si mesh thicknesses and fill factors are set to 25 μm and 3%, but could be floated for a more careful analysis.

A best fit is only a point estimate, though; it provides no information about the answer's robustness or certainty. Here I present a simple numerical technique to infer the standard error of best fit parameters: bootstrapping. Note, the definition of 'standard error' coincides exactly with the concept of '1 sigma' for physicists[6].

Bootstrapping, though drolly named, is a rigorous statistical method that relies on resampling the given data to produce sampling distributions for parameters of interest (see Appendix B). Synthetic datasets are created by resampling the original data *with replacement*: 10 new values of the monitor signal are drawn from the 10 original, and likewise for the detectors and associated backgrounds. Repeated values are likely, therefore, but the assumption is that the observed data well approximates the underlying distribution from which they are drawn. The algorithm repeats the fitting procedure outlined above many times on the synthetic datasets, yielding a collection of best fit values for statistics of interest, like Al thickness[7]. For this work, 10,000 synthetic datasets were created. Standard errors can then be read off from the resulting distributions.

As mentioned above, the limitations of the calibration measurement require extrapolation to higher energies. For those energies, the bootstrapped transmission curves produce the standard error of the mean transmission value. For measured energies, the bootstrapped mean of the data suffices. Future analyses with astrophysical or laboratory data should

[5] `http://henke.lbl.gov/optical_constants/`

[6] As mentioned elsewhere in this thesis, 1 sigma confidence intervals are not the same as the middle 68% of a posterior probability distribution function. See Appendix B.2. Bootstrapping is a frequentist method, so the reported errors are confidence intervals.

[7] `https://github.com/jmrv/filters`

Filter name	Aluminum thickness (Å)		Polyimide thickness (Å)	
	Luxel	Fitted	Luxel	Fitted
FM3	257	227 – 230	484	723 – 728
FM4	257	233 – 236	484	729 – 735
FM5	275	300 – 302	451	679 – 684
FM6	275	279 – 285	451	596 – 611
FM7	255	196 – 198	492	734 – 737
FM8	255	175 – 181	492	663 – 675
FO1	208	199 – 202	492	914 – 918
FO2	208	212 – 214	492	888 – 894

Table 5.1: Quoted and fitted filter film thicknesses.

compute an array of standard errors across the energy band, using either the bootstrapped extrapolated or measured values as necessary.

Caveats for this method are outlined in Appendix B, but they are few. The main assumption posits the samples are identically distributed, as well as independent. I assert the former by ignoring the slow time dependence of the signal, and the latter by assuming the refractory period for a photodiode current measurement is shorter than the sampling time. The simplicity of bootstrapping contrasts favorably with the rich, informative histogram which results after only a few seconds of computation. Any lopsidedness or multimodality becomes immediately evident, as do abnormalities in the fitting algorithm.

The above analysis can easily be replaced with Bayesian numerical methods, such as MCMC, if the analyst wishes.

5.3.4 Calibration results

The first conspicuous result is that the quoted thicknesses from the manufacturer differ significantly from the fitted values. Each filter is manufactured to specified thicknesses of Al and Polyimide by Luxel Corporation. The measured (i.e., inferred) thickness can then be compared to the claimed manufactured thicknesses. Table 5.1 shows the discrepancies between fitted and reported values, which can exceed 20% for Al thickness and 80% for polyimide thickness. Calibration is necessary, otherwise large biases are introduced by relying

on the reported thicknesses.

Second, the analytic model can be too simplistic. While the broad curves between the edges are well behaved, the level of intricacy in the edge structure is surprising. The structure arises from more complicated electron transitions than simple K-shell ejections. While of interest to atomic physics, for our purposes, measuring the structure is most useful to recover a few percent of effective area at these energies. These intricate edges are beyond the lower border of the Fe L forest of lines, which is of primary interest to *Micro-X*, but the edge structure would greatly inform results from a soft X-ray mission like *XQC*.

Figure 5-8 shows two examples of broad-band fits and uncertainty envelopes, derived from the bootstrapping. The filter with less calibration data exhibits higher uncertainty in the transmission model. The transmission of an entire stack can be calculated as a product of these individual bootstrap samples for each filter.

5.4 Simulated Data

Micro-X will target SNRs for its observations. The first flight will image a region of Puppis A known as the Si knot. Hwang et al. [65] found traces of enhanced Si emission over a broad, but coherent patch of the remnant with *Suzaku* data. They interpreted this region to be an old, evolved, but still intact ejecta knot (perhaps not dissimilar from the ones studied in Chapter 2), presenting a unique follow-up opportunity for *Micro-X*. The blended *Suzaku* spectra required the use of equivalent widths to quantify the Si abundance, while *Micro-X* would be able to see the lines directly, and study the dynamics of the evolved plasma knot as well.

With a sub-arcminute optics upgrade, *Micro-X* will also be able to target Cas A, seen in Figure 5-9. While the instrument cannot match *Chandra*'s imaging, our simulations show that the pixel that contains the brightest knot (highlighted in the rightmost panel of Figure 5-9) will be able to distinguish velocity shifts in the Si lines.

The Puppis A Si Knot and Cas A, the first two targets for *Micro-X*, have already been observed by other instruments. The best-fit `vnei` model of the CCD spectra can then be

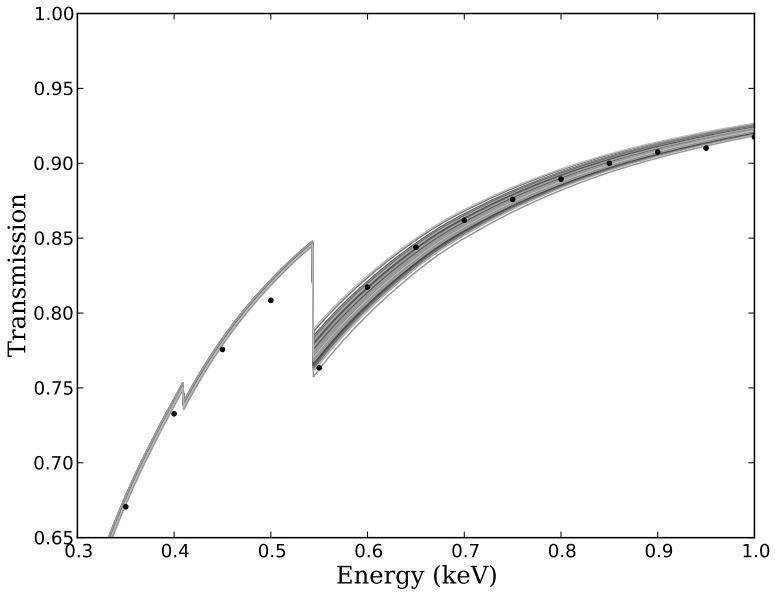

(a)

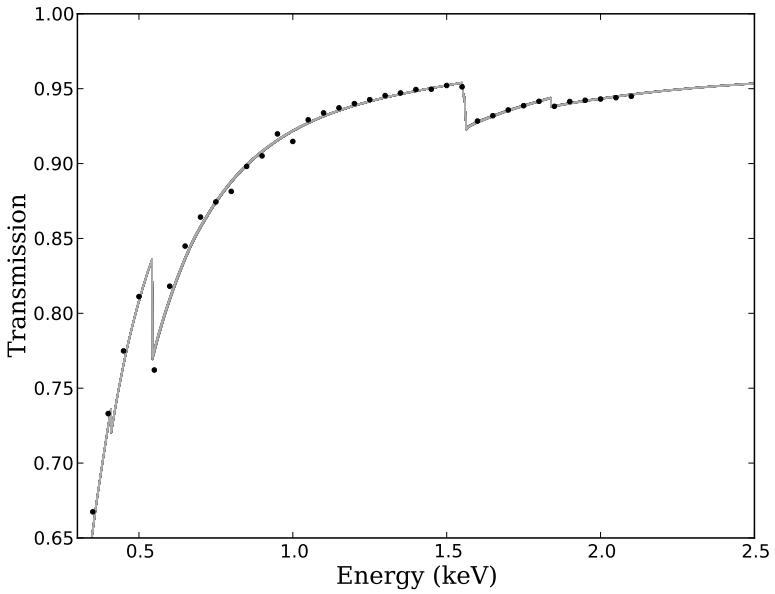

(b)

Figure 5-8: 100 individual, overlaid bootstrap samples (grey lines) for two different filters show the uncertainty envelopes of the fits to the transmission data (black dots). The inferred transmission curve for FM6 (a) has greater uncertainty because the filter was calibrated only to 1 keV. The bootstrap samples for FM5 (b) display a much tighter spread, due to calibration out to 2 keV.

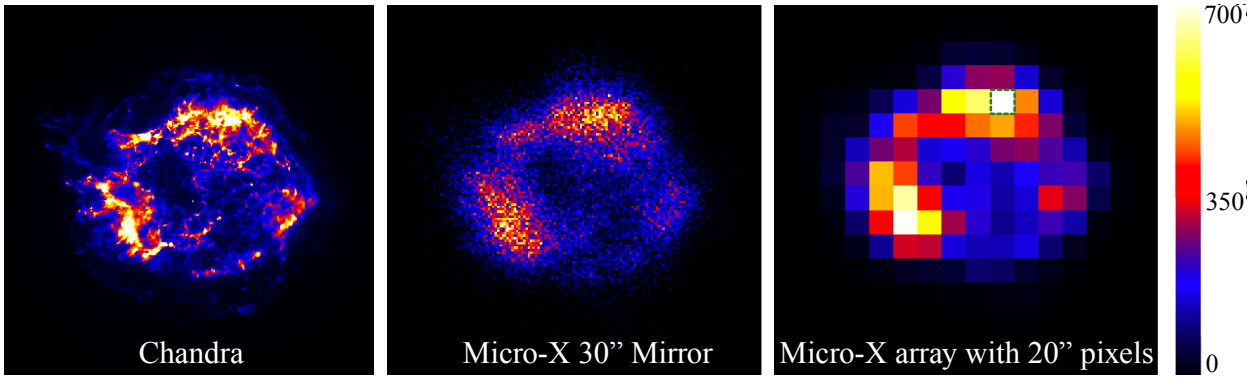

Figure 5-9: A planned upgrade to the *Micro-X* mirror would allow observation of Cas A. The simulated image (with smaller pixels) is shown at right, in counts. Velocity shifts over 1000 km/s within the brightest pixels would be discernible.

passed through the sharper *Micro-X* spectral response, revealing the underlying forest of lines (Figure 5-10). The actual spectrum will most likely differ from the `vnei` model, but it is difficult to predict how.

We simulate additional complexity beyond the `vnei` model by adding velocity shifts and turbulence. Figure 5-10 shows how the O He-like triplet is visibly split by 1500 km/s Doppler shifts with the *Micro-X* resolution.

The telemetered data will made available to the team live, inviting the opportunity to dither the telescope if desired. For instance, without the clamping mechanisms to fix LMO to the star tracker optical axis, the dewar may end up misaligned. While a command uplink will not be used for the first flight, future flights could employ re-pointing.

To explore the possibility of re-pointing the payload during flight, I simulated a heads-up display of incoming pulse counts[8]. The simulation deliberately points the incoming image off-target, then uses a library of pre-calculated templates to find the best guess for the pointing offset. The simulation showed that the pointing offset can confidently be determined within 10 seconds for Cas A, validating the procedure and the feasibility of re-pointing early in the observation. (A broader, more diffuse source would take longer to determine the pointing.)

Image templates — energy information is discarded — are made ahead of time to dras-

[8] `https://github.com/jmrv/pointing`

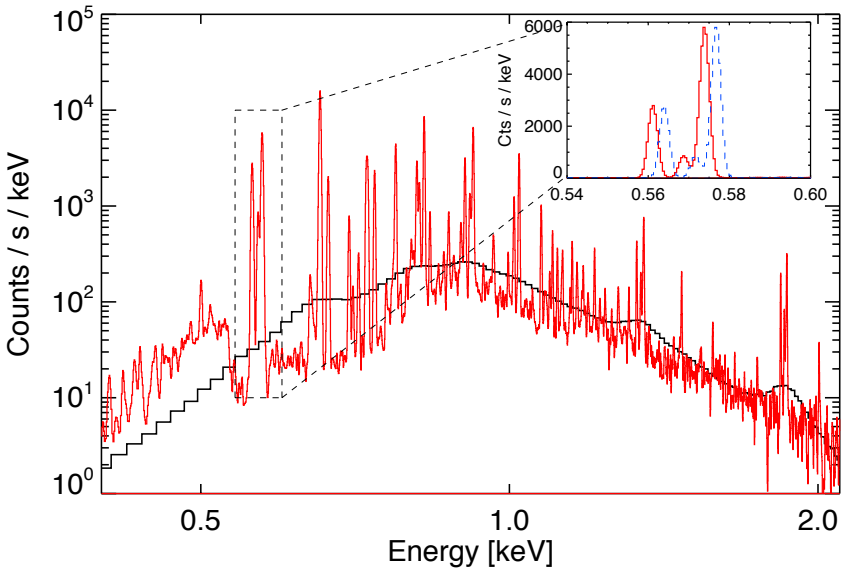

Figure 5-10: A simulated *Micro-X* spectrum of the Si knot of Puppis A reveals an information-dense forest of lines, as compared to the ACIS spectrum (black). The inset of the O He-like triplet demonstrates the instrument's ability to distinguish Doppler shifts (1500 km/s shown).

tically reduce computation time while live data comes in, allowing for quicker re-pointing decision-making. The Harvard tool `simx`[9] produced the template images at an array of RA and DEC pointings, using *Micro-X* telescope parameters (pixel scale size, mirror response, etc.). A sample heads-up snapshot, shown in Figure 5-11, shows a (tentative) mask pattern for the telemetered pixels. The Poisson likelihood of the incoming data, given a model templates, is calculated for each template to determine where the array is pointed.

The *Micro-X* team has also built a more faithful and complicated image simulator than `simx`. Kosuke Sato developed a simulation suite for the *Micro-X* optical path, to produce accurate simulated images[10]. The `uxcsim` software employs ray-tracing from a source image (from a higher spatial resolution instrument), through the specific geometry of the *Micro-X* mirror and filter stack, down to the focal plane. Opaque mirror housing, off-axis vignetting, multiple bounces, and the instrument coordinate system are all taken into account. The

[9]'Potential photons from possible satellites': `https://hea-www.harvard.edu/simx/`

[10]`https://wikis.mit.edu/confluence/download/attachments/36996611/caldb_uxcsim_20110321.pdf`

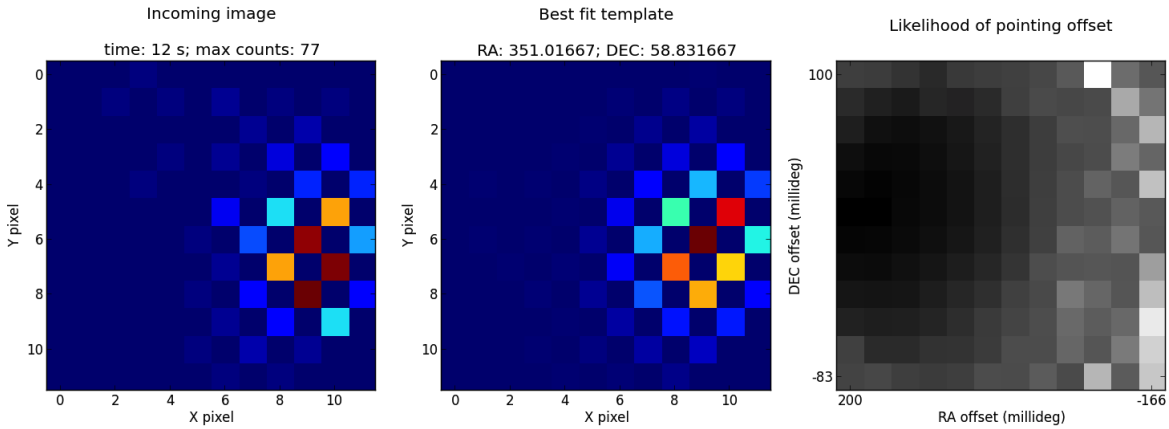

Figure 5-11: The heads-up display simulator takes in the decimated, telemetered count data (left) and compares it to a pre-made library of image templates pointing at different coordinates. The likelihood of each pointing coordinate is quickly calculated using this cached data (right) to determine the most likely pointing offset (center). The display updates every second with fresh data and re-calculated likelihoods.

software builds off previous work from the *Suzaku* mirror group [69].

5.5 Outlook

Micro-X has come a long way from conception to construction. The payload development has produced valuable intuition and data for magnetic shielding, large-absorber microcalorimeters, compact control electronics, and Kevlar isolation. This thesis provides a snapshot of the payload as-built ahead of its imminent launch.

As I have mentioned at various points prior, the date of *Micro-X*'s maiden launch remains uncertain until the vibration isolation performs to specification. Upcoming tests at the Wallops shake facility will accurately asses the progress of the Kevlar suspension redesign and the WRI skin dampers. The team takes some reluctant solace in the vibration coupling issues which have delayed the *Astro-H* payload. High performance space-based cryogenics is difficult, even for well funded observatory missions.

Multiplexed TES microcalorimeters are the inevitable future of X-ray astronomy. The development of *Micro-X* has already advanced this next generation of technology, and its

first five-minute exposure will demonstrate the worth of its complexity with high quality spatio-spectral data, as yet unattained by other instruments.

Micro-X's TES array and SQUID multiplexer will be the first such systems in space, and its planned upgrades — higher resolution optics, on-board pulse triggering software — will keep the instrument pertinent even as *Astro-H* launches, filling the gap left by the death of the microcalorimeter instrument on *Suzaku*.

In the young field of X-ray astrophysics, telescopes have been the primary driver of new scientific understanding. I am confident *Micro-X* will continue the trend and inspire instrument design for a new generation to come.

Appendix A

Common Abbreviations and Symbols

ACS	Attitude Control System
ADR	Adiabatic Demagnetization Refrigerator
BEK	Bright Eastern Knot
CAD	Computer Aided Design
CIE	Collisional Ionization Equilibrium
COG	Center Of Gravity
CM	Clamping Mechanism
DFB	Digital Feedback
ETF	Electro-thermal Feedback
FAA	Ferric Ammonium Alum
FEA	Front End Assembly
FEM	Finite Element Method
FOV	Field Of View
G	Thermal Conductance
GSE	Ground Support Equipment
GRT	Germanium Resistance Thermometer

GSFC	Goddard Space Flight Center
GUI	Graphical User Interface
GV	Gate Valve
HCC	High Current Connectors
HMC	Hamiltonian Monte Carlo
I/F	Interface
I/O	Input/Output
IR	Infrared
k	Spring constant
k_B	Boltzmann Constant
LHe	Liquid Helium
LMO	Lightweight Millikelvin Observatory (*Micro-X* flight dewar)
LN	Liquid Nitrogen
MCMC	Markov Chain Monte Carlo
MIC	Mission Initiation Conference
MLI	Multi-layer Insulation
MIT	Massachusetts Institute of Technology
MUX	Multiplexing
NASA	National Aeronautics and Space Administration
NEI	Non-equilibrium Ionization
NIST	National Institute of Standards and Technology
NSLS	National Synchrotron Light Source
NSROC	NASA Sounding Rocket Office Contractor
PI	Principal Investigator
PM	Position Monitor
PV	Pumping Valve
RA	Row Address
RF	Radio Frequency

RRR	Residual Resistance Ratio
S/R	Signal to Noise Ratio
SNR	Supernova Remnant
SQUID	Superconducting QUantum Interference Device
SS	Stainless Steel
SXS	Supernova X-ray spectrometer
TES	Transition Edge Sensor
TM	Telemetry System
TMD	Tuned Mass Damper
VAB	Vehicle Assembly Building
`vnei`	Variable (Abundance) Non-equilibrium Ionization
WFF	Wallops Flight Facility
WSMR	White Sands Missile Range
XQC	X-Ray Quantum Calorimeter

Appendix B

Statistical Inference and Numerical Methods

The common and important task of inference in science draws its power from the rigor of statistics. In this framework, data can be compared to a model and prior knowledge to produce a quantitative, reproducible conclusion. While scientists are obligated to design good experiments and take clean measurements, basic inference allows us to push experiments just beyond the technological limits.

With data in hand, an experimenter usually wants to know which values of model parameters are the most likely, or which theoretical model is the most likely. The natural starting point is Bayes' Rule, the simple rearrangement of the product rule for probabilities:

$$\begin{aligned} P(\theta|\mathcal{D},\mathcal{I}) &= \frac{P(\mathcal{D}|\theta,\mathcal{I}) \times P(\theta|\mathcal{I})}{P(\mathcal{D}|\mathcal{I})} \\ \text{posterior} &= \frac{\text{likelihood} \times \text{prior}}{\text{evidence}}. \end{aligned} \tag{B.1}$$

In this general formula, I've labeled model parameters as θ, the collected data as $\mathcal{D}$, and any useful previous data – including the implicitly assumed model or even an experimenter's hunches – as additional information, $\mathcal{I}$. This nearly tautological statement is extremely powerful: it flips the probability of some unknown parameters ($P(\theta|\mathcal{D},\mathcal{I})$) on its head in a

more tractable form: the product of an easily calculated likelihood of *data* being true *given* model parameters ($P(\mathcal{D}|\theta,\mathcal{I})$) modified by some prior knowledge ($P(\mathcal{D}|\mathcal{I})$), normalized by the 'evidence', which will feature prominently in model comparison. Here, $P(x|y)$ is the conditional probability of x being true, given that we have observed some information or fixed some variable y.

Despite this simple starting point, two distinct philosophies exist in statistics: the Bayesian viewpoint, and the Frequentist. Bayesians hold the data sacred, and assign uncertainty to the inference of the parameters. They are also take much more leeway with the prior, using intuition or previous experiments to assign informative priors. Frequentists, on the other hand, assume that one true parameter set exists, so the uncertainty arises from the notion that the data set is just one of many possible instances. Thus Frequentists perform, in essence, a thought experiment of re-taking the data many times until probability can be interpreted as a frequency of occurrence.

This thesis exhibits one example of simple Bayesian inference:

- Inferring the ^{44}Sc flux upper limit by marginalizing over nuisance parameters;

as well as two Frequentist methods:

- Determining confidence intervals of plasma model parameters with the χ^2 distribution;
- Generating standard errors with the bootstrap method.

Each philosophy and method has its own strengths.

B.1 Bayesian Parameter Inference

Inferring the probability distribution of a parameter θ is ostensibly simple with Equation B.1. The likelihood is motivated by the type of measurement (e.g., Poisson is appropriate for counting photons), and the prior is specified by the analyst. The denominator — the evidence — is usually not calculated directly; the posterior is normalized so that the integral across its support is 1.

In nearly every application, the model has multiple parameters, so $\boldsymbol{\theta}$ is a vector. Some will be nuisance parameters, though, which need to be *marginalized* out. In a two-parameter model, with θ the one of interest and θ_n a nuisance, marginalization proceeds as:

$$P(\theta|\mathcal{D},\mathcal{I}) = \int d\theta_n P(\theta,\theta_n|\mathcal{D},\mathcal{I}). \tag{B.2}$$

Again, the result is simple to write down, but the integral is difficult to perform except in very few textbook analytic cases.

B.2 Frequentist Confidence Intervals

As stated above, if asked for an estimate of a parameter θ, the Bayesian analyst provides a marginalized posterior distribution, as with Equation B.2. If pressed further to specify an interval covering all probable values of θ except for a fraction a, the analyst would define an interval $\mathcal{J}(d)$ for a given dataset $\mathcal{D} = d$:

$$P(\theta \in \mathcal{J}(d)|\mathcal{D} = d) = 1 - a. \tag{B.3}$$

The Frequentist constructs a similar interval differently. The data — not the model parameters — are taken to be the random variables, thus making the interval $\mathcal{J}(\mathcal{D})$ a random variable. Thus, the interval must provide the guarantee

$$P(\theta^* \in \mathcal{J}(\mathcal{D})|\theta = \theta^*) = 1 - a, \tag{B.4}$$

where θ^* is the true parameter.

Cash [16] wrote a seminal paper on parameter estimation in X-ray astronomy in 1976, soon after the launch of *SAS-3*, and just before the *HEAO* satellites would come online. His general prescription for generating multi-dimensional confidence intervals for photon count data persists to this day in modern spectral fitting packages.

His result applies to the germane statistic

$$S = \sum_i \frac{(d_i - m_i)^2}{m_i}, \tag{B.5}$$

where d_i is the number of data counts in bin i, while m_i is the count predicted by the model. The S statistic is commonly called χ^2 by physicists, but I will reserve that name for the distribution S follows. Implicit in the use of S as a faithful statistic is that the number of counts is large. If $\mathcal{L}$ is the calculated likelihood of n counts, the quantity $-2\ln\mathcal{L}$ is distributed as χ^2, modulo a term of order $1/\sqrt{n}$.

As above, some of the parameters in θ will be nuisance parameters. Of the p parameters, assume that only the first q are interesting. First, find the maximum likelihood corresponding to S_{min}. Set the desired significance, α, and calculate the target value $\chi^2_q(\alpha)$. Fix a set of q interesting parameters, varying the $p - q$ nuisance parameters to find a minimum of S, which we denote S^*_{min}. The quantity $S^*_{min} - S_{min}$ will be distributed as χ^2 with q degrees of freedom. Thus, the surface of

$$S^*_{min} = S_{min} + \chi^2_q(\alpha) \tag{B.6}$$

defines the region within which the true q parameters exist, with probability $1 - \alpha$.

The result is remarkably general, not requiring linear models or even a Normal likelihood (though, again, large n is assumed).

B.3 Model Comparison

The selection between two models, to determine which offers more predictive power for a set of data, is a compelling technique. I did not get the chance to use model comparison in this work, but mention it briefly here because I think many problems in X-ray astrophysics could benefit from the method. As more complex models are proposed to account for observational nuances, the analyst must be careful not to overfit.

Bayesian model comparison computes the ratio of evidences, $P(\mathcal{D}|\mathcal{I},\mathcal{M}_1)/P(\mathcal{D}|\mathcal{I},\mathcal{M}_2)$,

where I've added the model $\mathcal{M}$ which was previously implicit. I'll mention three methods to compute this term:

- Nested sampling [122]. This MCMC sampling method sums the evidence via a clever change of variables. Additionally, the posterior can be generated from the evidence samples, with appropriate area weighting.

- Posterior Predictive p-values. This method, popularized by Protassov et al. [110] in astrophysics, applies Bayesian statistics within a Frequentist framework. True model comparison is not performed, though, as only 'null' models can be rejected; the rejection offers no alternative model.

- Bayesian Information Criterion (BIC). This correction to the likelihood ratio test penalizes the more complicated model with [number of free parameters] $\times$ ln([number of data points]), approximating the fully Bayesian method under certain assumptions [54].

B.4 Numerical Methods

Numerical statistical methods – as opposed to traditional analytic derivations meant to minimize calculation – have become popular as computing power has become cheaper and more readily available. The central limit theorem, t-tests, and closed form solutions for linear regression are useful tools, but represent the shortcuts of an age where computation was laborious.

Three primary reasons stand out, to me, for this shift. Primary among them is, simply, the outstanding performance of such methods. Computers explore large parameter spaces more efficiently than human intuition alone, leading to often counter-intuitive inferences. Moreover, instead of just maximum likelihood values and error bars, the whole posterior distribution can be recovered, making any skewness, heavy tails, or multimodality plainly evident.

Secondly, related to the first aspect, the methods I describe below are *non-parametric*, meaning they do not rely on the assumption that the data follow some well known, parametrized distribution. Most critiques of 'orthodox' statistics stem from the implicit assumptions (such as normality) an analyst must make to use the tidy formulae. An analyst can be more comfortable, therefore, relaxing the assumption that the data is normally distributed and letting the algorithm result speak for itself.

Lastly, I believe the skills necessary for computer programming make statistical algorithms more intuitive. Learning rules of error propagation by rote does not lend to intuition, but algorithms may. Today's new physicists learn to code out of necessity, but are still taught statistics as a disparate collection of rule-based tools. Algorithmic statistics, therefore, become an easier point of entry.

B.4.1 Markov Chain Monte Carlo

Markov Chain Monte Carlo (MCMC) is a popular technique to sample from a complicated posterior distribution, thereby providing a numerical approximation to the marginal distribution (Equation B.2). As opposed to evaluating the posterior on fine grid, MCMC is less prone to miss small areas of parameter space with high probability, and is more efficient for high dimensional problems. This is useful for likelihoods which are easily calculated for a given point, but cannot be specified over the whole parameter range without computational difficulty.

The simplest idea of a MCMC algorithm is a particle roaming around the parameter space in discrete steps ('links'), taking a random walk with a preference for higher likelihoods. Each link of the chain depends on the previous link, and links may be repeated if the tentative step is not accepted. The 'preference' is expressed in a transition probability (also known as acceptance probability), which also allows for transitions to lower probability states. The chain grows until it is deemed converged to the target probability distribution, after which it 'burns in' for a suitable number of links.

The burned-in links, then, form the discrete approximation to the continuous posterior.

The marginal posterior distribution for a parameter is formed, simply, by binning the samples into a histogram, ignoring the nuisance parameter values. I leave the formal treatment of Markov chains and their convergence to MacKay [89], Sivia and Skilling [122], and Brooks et al. [12].

The requirements on the transition probability and proposal link-generating algorithm are few, but convergence to the actual probability distribution is guaranteed. However, there is no guarantee for *when* the chain converges. Intuitive heuristics — such as comparing the behavior of multiple parallel chains [12] — have been developed, though.

The freedom to choose the transition probability and random walk algorithm has yielded diverse algorithms, each with their own strengths. This work employs Hamiltonian Monte Carlo (HMC), as described in MacKay [89]. As its name suggests, the particle is made to obey Hamiltonian mechanics, and is endowed with momentum to roll along the likelihood contours (the potential energy), sensing local slopes as it proceeds. To propose a new link, a particle is kicked with a random force vector and allowed to roll for a specified length of time. The extra information from the slope allows quicker exploration than a naive random walk.

I implemented this algorithm in Python[1] to infer line strengths in *Chandra* spectra[2]. Several free parameters can be tuned — the stepsize, the number of timesteps to find a new link, the masses of the particle (which can be different along each parameter axis) — but most of these were fixed for simplicity. To achieve convergence, some masses were tuned, and prior ranges were restricted to create similar dynamic ranges for each parameter. Care must be taken not to let the kinetic energy dominate the potential energy (likelihood), otherwise errors in the discretized trajectory will add to create spurious samples.

[1] Other languages are faster, but require more overhead for development. Python allows quick prototyping, on timescales which dominate computational durations.

[2] `https://github.com/jmrv/scandium`

B.4.2 Bootstrapping

Bootstrapping is a simple but underused numerical method for estimating the variance, or *standard error*, of fit parameters [28]. I will note up front, though, that bootstrapping is a Frequentist method, and produces confidence intervals for parameters, not posterior probability density functions.

The method relies on the notion of resampling the original dataset to create many synthetic datasets. From those samples, a distribution of derived values is built up for the value of interest, for instance the mean or a fit parameter, either of which is termed a *statistic*. This synthetic distribution approaches the *sampling distribution* for the parameter with more data and resamples. A dataset can thereby pull itself up by its own bootstraps, producing not only the one estimate of a statistic, but a whole distribution. The sampling distribution can be thought of as a shortcut around the full joint posterior distribution directly to the empirical distribution of the statistic of interest, given the data points at hand[3]. The standard deviation of the sampling distribution is the standard error, the quantity which physicists refer to as '1 sigma' error bars on a parameter.

Formally, we can consider n random samples from an unknown distribution F:

$$X_i \sim F, \;\; i = 1, 2, ...n.$$

The random variable we care about (mean, fit parameter, etc.), R, may depend on $\mathbf{X}$ and F, so we write it $R(\mathbf{X}, F)$. The sample probability distribution, $\hat{F}$, is easily constructed from the data by assigning equal probability to each of the observed values of $\mathbf{X}$. Now we can perform the resampling to create a bootstrap sample $\mathbf{X}^*$:

$$X_i^* \sim \hat{F}, \;\; i = 1, 2, ...n.$$

[3]Despite its foreign name, the idea of a sampling distribution is natural to physicists. For instance, one can show that with two normal distributions, $\mathcal{N}(\mu_1, \sigma_1^2)$, $\mathcal{N}(\mu_2, \sigma_2^2)$, the difference between the sample means (the *statistic* in question) is distributed as $\bar{X}_1 - \bar{X}_2 \sim \mathcal{N}(\mu_1 - \mu_2, \sigma_1^2/n_1 + \sigma_2^2/n_2)$. The sample distribution corresponds to directly with basic error propagation.

Note that the sampling is done *with replacement*, so duplicates will arise. Then, the bootstrap distribution of $R^* = R(\mathbf{X}^*, \hat{F})$ approximates the desired sampling distribution of R. This distribution can be built up by performing many random samples of $\mathbf{X}^*$ from $\hat{F}$ and binning the resulting R^* values into a histogram. The number of required resamples depends on R, so there is no general formula, but usually a few thousand are sufficient. However, as computing power has increased, the rule of thumb now is as many resamples as possible. Like MCMC, the bootstrap method belongs to the class of *non-parametric* estimation techniques, as it does not assume the data are drawn from a particular distribution (e.g., normal).

The generality and robustness of bootstrapping is remarkable, especially in light of the minimal computation effort. To use an example from this thesis, performing a least-squares fit on filter transmission data takes only a few milliseconds on a laptop. Satisfied with the look of this maximum likelihood fit, but unsure of how to put error bars on the thickness of Al, I need only perform this fit a few thousand more times on resampled data to build up a distribution of fitted Al thicknesses. After only a minute, the sampling distribution materializes, producing error bars, along with other information about the skewness, long tails, or multi-modal nature of the distribution. I need not make assumptions of normality to propagate uncertainties, nor do I need to fine tune a MCMC algorithm to converge to the target distribution. In addition, I have found that bootstrapping is useful to check the robustness of fitting algorithms, to diagnose over-reliance on initial guesses, and to cross-check with back-of-the-envelope error propagation.

Some general caveats apply to bootstrapping. First, like many numerical statistical techniques, bootstrapping theory offers only asymptotic consistency, not a cutoff point (for number of resamples, or datapoints) after which the answer is correct to some precision. Second, the construction of $\hat{F}$ relies on the assumption that the X_i are independently and identically distributed from F. This assumption is extremely commonplace in statistics, though. Third, since the method cannot 'comprehend' values more extreme than those observed, the standard error estimates may be slightly optimistic. This pitfall, of course, afflicts nearly any analysis with limited data. Fourth, bootstrapping resides within a decidedly frequentist

framework: the resulting distribution says nothing about the likelihood of the parameter, only the frequency of the statistic given the data. The output is a confidence interval, not a posterior distribution (see Section B.2, following). Some analysts may find this aspect unsavory and unintuitive, so a Bayesian bootstrap has been developed by Rubin [115] to produce a bootstrapped posterior distribution. The method has its own drawbacks, though, including a very specific form of the prior and complicated handling of fit parameters.

The ease of bootstrapping will often outstrip most of the concerns; it should be added to more physicists' toolboxes.

B.4.3 Confidence Contours

Confidence intervals were developed, in part, like many Frequentist methods, to obviate complicated Bayesian integrals. As models have become more complicated with larger and larger parameter spaces, though, the computational effort to compute confidence contours has risen. As an example, the confidence contours from Chapter 2 took a few hours to run, for each knot.

From where does this computational intensity stem? The locus of points defined by Equation B.6 requires that S^*_{min} be calculated across a grid of points in the q-dimensional subspace of interesting parameters. For each point in the grid, the algorithm — typically a gradient ascent of some variety — must vary the $p - q$ nuisance parameters to find the minimum value of the S statistic. This minimization step alone is computationally expensive, for large models, and the computer time is compounded by repeating the process over a grid. The grid, too, must not be too coarse, in order to discern the boundary shape.

B.5 Calibration Data Products: Considerations

Two calibration products have become standard to include with public X-ray telescope data. The first is the Ancillary Response File (ARF), the effective area of the optics as a function of energy. Typically, the ARF comprises only the response of the telescope's mirror, which is

calibrated separate from the detectors and filters. (It would make sense to include the filters and mirror in one file, but the physical layout of the instrument usually dictates otherwise; calibration of the detectors must take place with filters installed, usually. Moreover, the detector and mirror often belong to distinct teams for large missions.)

The second product, the Redistribution Matrix File (RMF), contains a 2D matrix that describes how impinging photons of one energy are 'smeared', probabilistically, to a broader range of detected energies. This smearing is the energy resolution of the instrument. The RMF for a perfect detector is the identity matrix, while for *Chandra* the spread is roughly 100 eV, and grows with increasing energy. Pileup – when two or more photons hit a pixel in close succession, mimicking one higher energy X-ray – is not included in the RMF, but should be accounted for separately. Both the ARF and RMF are encoded in the FITS format.

Modern X-ray analysis rarely takes the instrumental calibration uncertainties into account. Lee et al. [83] have unequivocally demonstrated the perils of accepting the mean calibration as the true value. In their work, they offer a solution to account for the systematic uncertainties by incorporating principal components of the ARF into a MCMC inference engine.

This thesis, unfortunately, did not account for the calibration uncertainties in the two *Chandra* analyses. *Micro-X*, which has not fully defined its ARF and RMF yet, could amend its calibration products with uncertainties, to provide fellow astronomers a better understanding of the instrument's eventual public data.

Bibliography

[1] E. Anders and N. Grevesse. Abundances of the elements - Meteoritic and solar. *Geochim. Cosmochim. Acta*, 53:197–214, January 1989. doi: 10.1016/0016-7037(89)90286-X.

[2] K. Arnaud, R. Smith, and A. Siemiginowska. *Handbook of X-ray Astronomy*. Cambridge Observing Handbooks for Research Astronomers. Cambridge University Press, 2011. ISBN 9780521883733.

[3] B. Aschenbach, A. F. Iyudin, and V. Schönfelder. Constraints of age, distance and progenitor of the supernova remnant RX J0852.0-4622/GRO J0852-4642. *A&A*, 350: 997–1006, October 1999.

[4] M. Balucinska-Church and D. McCammon. Photoelectric absorption cross sections with variable abundances. *ApJ*, 400:699, December 1992. doi: 10.1086/172032.

[5] A. Barone. *Principles and Applications of Superconducting Quantum Interference Devices*. World Scientific, 1992. ISBN 9789810209117.

[6] J. M. Blondin and D. C. Ellison. Rayleigh-Taylor Instabilities in Young Supernova Remnants Undergoing Efficient Particle Acceleration. *ApJ*, 560:244–253, October 2001. doi: 10.1086/322499.

[7] J. M. Blondin, K. J. Borkowski, and S. P. Reynolds. Dynamics of Fe Bubbles in Young Supernova Remnants. *ApJ*, 557:782–791, August 2001. doi: 10.1086/321674.

[8] K. J. Borkowski, S. P. Reynolds, D. A. Green, U. Hwang, R. Petre, K. Krishnamurthy, and R. Willett. Radioactive Scandium in the Youngest Galactic Supernova Remnant G1.9+0.3. *ApJ*, 724:L161–L165, December 2010. doi: 10.1088/2041-8205/724/2/L161.

[9] Kazimierz J. Borkowski, William J. Lyerly, and Stephen P. Reynolds. Supernova remnants in the sedov expansion phase: Thermal x-ray emission. *ApJ*, 548(2):820, 2001.

[10] H. Bradt, W. Mayer, S. Naranan, S. Rappaport, and G. Spada. Evidence for X-Radiation from the Radio Galaxy M87. *ApJ*, 150:L199, December 1967. doi: 10.1086/180125.

[11] H. Bradt, S. Naranan, S. Rappaport, and G. Spada. Celestial Positions of X-Ray Sources in Sagittarius. *ApJ*, 152:1005, June 1968. doi: 10.1086/149613.

[12] S. Brooks, A. Gelman, G. Jones, and X.L. Meng. *Handbook of Markov Chain Monte Carlo*. Chapman & Hall/CRC Handbooks of Modern Statistical Methods. Taylor & Francis, 2011. ISBN 9781420079425.

[13] B. E. Burke, R. W. Mountain, P. J. Daniels, M. J. Cooper, and V. S. Dolat. CCD soft X-ray imaging spectrometer for the ASCA satellite. *IEEE Transactions on Nuclear Science*, 41:375–385, February 1994. doi: 10.1109/23.281527.

[14] C. R. Canizares, K. A. Flanagan, D. S. Davis, D. Dewey, J. C. Houck, and M. L. Schattenburg. High resolution spectroscopy of two oxygen-rich SNRs with the Chandra HETG. In S. S. Holt and U. Hwang, editors, *Young Supernova Remnants*, volume 565 of *American Institute of Physics Conference Series*, pages 213–221, May 2001. doi: 10.1063/1.1377096.

[15] C. R. Canizares, J. E. Davis, D. Dewey, K. A. Flanagan, E. B. Galton, D. P. Huenemoerder, K. Ishibashi, T. H. Markert, H. L. Marshall, M. McGuirk, M. L. Schattenburg, N. S. Schulz, H. I. Smith, and M. Wise. The Chandra High-Energy Transmission Grating: Design, Fabrication, Ground Calibration, and 5 Years in Flight. *PASP*, 117: 1144–1171, October 2005. doi: 10.1086/432898.

[16] W Cash. Generation of Confidence Intervals for Model Parameters in X-ray Astronomy. *A&A*, 52:307, October 1976.

[17] J. E. Davis. Event Pileup in Charge-coupled Devices. *ApJ*, 562:575–582, November 2001. doi: 10.1086/323488.

[18] Tracey DeLaney, Lawrence Rudnick, MD Stage, JD Smith, Karl Isensee, Jeonghee Rho, Glenn E Allen, Haley Gomez, Takashi Kozasa, William T Reach, et al. The Three-Dimensional Structure of Cassiopeia A. *ApJ*, 725(2):2038, 2010.

[19] K. P. Dere, E. Landi, P. R. Young, G. Del Zanna, M. Landini, and H. E. Mason. CHIANTI - an atomic database for emission lines. IX. Ionization rates, recombination rates, ionization equilibria for the elements hydrogen through zinc and updated atomic data. *A&A*, 498:915–929, May 2009. doi: 10.1051/0004-6361/200911712.

[20] D. Dewey. Extended Source Analysis for Grating Spectrometers. In G. Branduardi-Raymont, editor, *High Resolution X-ray Spectroscopy with XMM-Newton and Chandra*, December 2002.

[21] D. Dewey and M. S. Noble. Simulation and Fitting of Multi-Dimensional X-ray Data. In D. A. Bohlender, D. Durand, & P. Dowler, editor, *Astronomical Data Analysis Software and Systems XVIII*, volume 411 of *Astronomical Society of the Pacific Conference Series*, pages 234–+, September 2009.

[22] D. Dewey, T. Delaney, and J. S. Lazendic. Cas A: The Bright X-ray Knots and Oxygen Emission. In *Revista Mexicana de Astronomia y Astrofisica Conference Series*, volume 30 of *Revista Mexicana de Astronomia y Astrofisica Conference Series*, pages 84–89, August 2007.

[23] D. Dewey, S. A. Zhekov, R. McCray, and C. R. Canizares. Chandra HETG Spectra of SN 1987A at 20 Years. *ApJ*, 676:L131–L134, April 2008. doi: 10.1086/587549.

[24] R. Diehl and F. X. Timmes. Gamma-Ray Line Emission from Radioactive Isotopes in Stars and Galaxies. *PASP*, 110:637–659, June 1998. doi: 10.1086/316169.

[25] WB Doriese, JA Beall, WD Duncan, L Ferreira, GC Hilton, KD Irwin, CD Reintsema, JN Ullom, LR Vale, and Y Xu. Progress toward kilopixel arrays: 3.8 eV microcalorimeter resolution in 8-channel SQUID multiplexer. *Nucl. Phys. A*, 559(2):808–810, 2006.

[26] C. Dupraz, H. Bloemen, K. Bennett, R. Diehl, W. Hermsen, A. F. Iyudin, J. Ryan, and V. Schoenfelder. COMPTEL three-year search for galactic sources of ˆ44ˆTi gamma-ray line emission at 1.157MeV. *A&A*, 324:683–689, August 1997.

[27] M. E. Eckart, J. S. Adams, S. R. Bandler, R. P. Brekosky, A.-D. Brown, J. A. Chervenak, A. J. Ewin, F. M. Finkbeiner, R. L. Kelley, C. A. Kilbourne, F. S. Porter, J. E. Sadleir, S. J. Smith, E. Figueroa-Feliciano, and P. Wikus. Large-Absorber TES X-ray Microcalorimeters and the Micro-X Detector Array. In B. Young, B. Cabrera, and A. Miller, editors, *American Institute of Physics Conference Series*, volume 1185 of *American Institute of Physics Conference Series*, pages 699–702, December 2009. doi: 10.1063/1.3292437.

[28] Bradley Efron. Bootstrap methods: another look at the jackknife. *The Annals of Statistics*, pages 1–26, 1979.

[29] C. I. Ellinger, P. A. Young, C. L. Fryer, and G. Rockefeller. A Case Study of Small-scale Structure Formation in Three-dimensional Supernova Simulations. *ApJ*, 755:160, August 2012. doi: 10.1088/0004-637X/755/2/160.

[30] Ryoji Enomoto, T Tanimori, T Naito, T Yoshida, S Yanagita, M Mori, Philip Gregory Edwards, A Asahara, Geoffrey Vincent Bicknell, S Gunji, et al. The acceleration of cosmic-ray protons in the supernova remnant RX J1713. 7–3946. *Nature*, 416(6883): 823–826, 2002.

[31] A.C. Fabian, K.A. Pounds, and R.D. Blandford. *Frontiers of X-Ray Astronomy*. Cambridge Planetary Science. Cambridge University Press, 2004. ISBN 9780521534871.

[32] R. A. Fesen, J. A. Zastrow, M. C. Hammell, J. M. Shull, and D. W. Silvia. Ejecta Knot Flickering, Mass Ablation, and Fragmentation in Cassiopeia A. *ApJ*, 736:109–+, August 2011. doi: 10.1088/0004-637X/736/2/109.

[33] E. Figueroa-Feliciano, J. S. Adams, R. Baker, S. R. Bandler, T. A. Delaney, D. Dewey, W. B. Doriese, M. E. Eckart, M. Galeazzi, R. Goeke, R. Hamersma, G. C. Hilton, U. Hwang, K. D. Irwin, R. L. Kelley, C. A. Kilbourne, S. W. Leman, D. McCammon, T. Okajima, F. S. Porter, C. D. Reintsema, J. M. Rutherford, T. Saab, P. Serlemitsos, Y. Soong, S. N. Trowbridge, and P. Wikus. New Science Case for the Micro-X High Energy Resolution Microcalorimeter X-ray Imaging Rocket. In B. Young, B. Cabrera, and A. Miller, editors, *American Institute of Physics Conference Series*, volume 1185 of *American Institute of Physics Conference Series*, pages 426–429, December 2009. doi: 10.1063/1.3292368.

[34] E. Figueroa-Feliciano, J. S. Adams, R. Baker, S. R. Bandler, D. Dewey, W. B. Doriese, M. E. Eckart, R. Hamersma, G. C. Hilton, U. Hwang, K. D. Irwin, R. L. Kelley, C. A. Kilbourne, S. E. Kissel, S. W. Leman, D. McCammon, P. H. H. Oakley, T. Okajima, R. H. O'Neal, R. Petre, F. S. Porter, C. D. Reintsema, J. M. Rutherford, T. Saab, P. Serlemitsos, Y. Soong, S. N. Trowbridge, and P. Wikus. Update on the Micro-X Sounding Rocket payload. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 8443 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, September 2012. doi: 10.1117/12.927167.

[35] Enectali Figueroa-Feliciano. *Theory and Development of Position-Sensitive Quantum Calorimeters*. PhD thesis, Stanford University, 2001. Published by Goddard Space Flight Center.

[36] W. Forman, C. Jones, L. Cominsky, P. Julien, S. Murray, G. Peters, H. Tananbaum, and R. Giacconi. The fourth Uhuru catalog of X-ray sources. *ApJS*, 38:357–412, December 1978. doi: 10.1086/190561.

[37] A. R. Foster, R. K. Smith, N. S. Brickhouse, T. R. Kallman, and M. C. Witthoeft. The Challenges of Plasma Modeling: Current Status and Future Plans. *Space Sci. Rev.*, 157:135–154, December 2010. doi: 10.1007/s11214-010-9732-1.

[38] A. H. Gabriel and C. Jordan. Interpretation of solar helium-like ion line intensities. *MNRAS*, 145:241–+, 1969.

[39] T. J. Gaetz, Y. M. Butt, R. J. Edgar, K. A. Eriksen, P. P. Plucinsky, E. M. Schlegel, and R. K. Smith. Chandra X-Ray Observatory Arcsecond Imaging of the Young, Oxygen-rich Supernova Remnant 1E 0102.2-7219. *ApJ*, 534:L47–L50, May 2000. doi: 10.1086/312640.

[40] James RB Garfield. Superconducting contacts for use in niobium thin film applications. *Review of Scientific Instruments*, 68(4):1906–1907, 1997.

[41] G. P. Garmire, M. W. Bautz, P. G. Ford, J. A. Nousek, and G. R. Ricker, Jr. Advanced CCD imaging spectrometer (ACIS) instrument on the Chandra X-ray Observatory. In J. E. Truemper and H. D. Tananbaum, editors, *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 4851 of *Society of Photo-Optical*

Instrumentation Engineers (SPIE) Conference Series, pages 28–44, March 2003. doi: 10.1117/12.461599.

[42] A. Gelman and D.B. Rubin. Inference from iterative simulation using multiple sequences. *Statistical Science*, 7:457–511, 1992.

[43] R. Giacconi, W. P. Reidy, T. Zehnpfennig, J. C. Lindsay, and W. S. Muney. Solar X-Ray Image Obtained Using Grazing-Incidence Optics. *ApJ*, 142:1274–1278, October 1965. doi: 10.1086/148404.

[44] Riccardo Giacconi. History of X-ray telescopes and astronomy. *Experimental Astronomy*, 25(1-3):143–156, January 2009.

[45] Riccardo Giacconi. The high energy X-ray universe. *PNAS*, 107(16):7202–7207, 2010. doi: 10.1073/pnas.0913079107.

[46] Riccardo Giacconi and Bruno Rossi. A 'Telescope' for Soft X-Ray Astronomy. *Journal of Geophysical Research*, 65:773, February 1960.

[47] Riccardo Giacconi, Herbert Gursky, Frank R. Paolini, and Bruno B. Rossi. Evidence for x Rays From Sources Outside the Solar System. *Phys. Rev. Lett.*, 9:439–443, Dec 1962. doi: 10.1103/PhysRevLett.9.439.

[48] E. V. Gotthelf, B. Koralesky, L. Rudnick, T. W. Jones, U. Hwang, and R. Petre. Chandra Detection of the Forward and Reverse Shocks in Cassiopeia A. *ApJ*, 552(1): L39, 2001.

[49] B. Grefenstette, H. An, S. E. Boggs, F. Christensen, W. W. Craig, C. Freyer, C. J. Hailey, F. Harrison, B. Humensky, S. Jakobsen, V. Kaspi, T. Kitaguchi, L. A. Lopez, K. Madsen, H. Miyasaka, K. Mori, M. Nynka, M. Pivovaroff, S. P. Reynolds, D. Stern, N. J. Westergaard, D. R. Wik, W. Zhang, A. Zoglauer, and NuSTAR Team. The NuSTAR Program for Supernova Remnants. In *American Astronomical Society Meeting Abstracts*, volume 221 of *American Astronomical Society Meeting Abstracts*, page 244.08, January 2013.

[50] H. Gursky, R. Giacconi, P. Gorenstein, J. R. Waters, M. Oda, H. Bradt, G. Garmire, and B. V. Sreekantan. A Measurement of the Location of the X-Ray Source SCO X-1. *ApJ*, 146:310–316, October 1966. doi: 10.1086/148891.

[51] N. J. Hammer, H.-Th. Janka, and E. Muller. Three-dimensional Simulations of Mixing Instabilities in Supernova Explosions. *ApJ*, 714(2):1371, 2010.

[52] F. A. Harrison, S. Boggs, F. Christensen, W. Craig, C. Hailey, D. Stern, W. Zhang, L. Angelini, H. An, V. Bhalerao, N. Brejnholt, L. Cominsky, W. R. Cook, M. Doll, P. Giommi, B. Grefenstette, A. Hornstrup, V. Kaspi, Y. Kim, T. Kitaguchi, J. Koglin, C. C. Liebe, G. Madejski, K. Kruse Madsen, P. Mao, D. Meier, H. Miyasaka, K. Mori,

M. Perri, M. Pivovaroff, S. Puccetti, V. Rana, and A. Zoglauer. The Nuclear Spectroscopic Telescope Array (NuSTAR). In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 7732 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, July 2010. doi: 10.1117/12.858065.

[53] F. A. Harrison, W. W. Craig, F. E. Christensen, C. J. Hailey, W. W. Zhang, S. E. Boggs, D. Stern, W. R. Cook, K. Forster, P. Giommi, B. W. Grefenstette, Y. Kim, T. Kitaguchi, J. E. Koglin, K. K. Madsen, P. H. Mao, H. Miyasaka, K. Mori, M. Perri, M. J. Pivovaroff, S. Puccetti, V. R. Rana, N. J. Westergaard, J. Willis, A. Zoglauer, H. An, M. Bachetti, N. M. Barrière, E. C. Bellm, V. Bhalerao, N. F. Brejnholt, F. Fuerst, C. C. Liebe, C. B. Markwardt, M. Nynka, J. K. Vogel, D. J. Walton, D. R. Wik, D. M. Alexander, L. R. Cominsky, A. E. Hornschemeier, A. Hornstrup, V. M. Kaspi, G. M. Madejski, G. Matt, S. Molendi, D. M. Smith, J. A. Tomsick, M. Ajello, D. R. Ballantyne, M. Baloković, D. Barret, F. E. Bauer, R. D. Blandford, W. Niel Brandt, L. W. Brenneman, J. Chiang, D. Chakrabarty, J. Chenevez, A. Comastri, F. Dufour, M. Elvis, A. C. Fabian, D. Farrah, C. L. Fryer, E. V. Gotthelf, J. E. Grindlay, D. J. Helfand, R. Krivonos, D. L. Meier, J. M. Miller, L. Natalucci, P. Ogle, E. O. Ofek, A. Ptak, S. P. Reynolds, J. R. Rigby, G. Tagliaferri, S. E. Thorsett, E. Treister, and C. M. Urry. The Nuclear Spectroscopic Telescope Array (NuSTAR) High-energy X-Ray Mission. *ApJ*, 770:103, June 2013. doi: 10.1088/0004-637X/770/2/103.

[54] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning*, volume 1. Springer New York, 2001.

[55] J.L. Heilbron. *The Oxford Guide to the History of Physics and Astronomy*. Number v. 10 in The Oxford guide to the history of physics and astronomy. Oxford University Press, USA, 2005. ISBN 9780195171983.

[56] E. A. Helder, J. Vink, A. M. Bykov, Y. Ohira, J. C. Raymond, and R. Terrier. Observational Signatures of Particle Acceleration in Supernova Remnants. *Space Sci. Rev.*, 173:369–431, November 2012. doi: 10.1007/s11214-012-9919-8.

[57] B.L. Henke, E.M. Gullikson, and J.C. Davis. X-ray interactions: Photoabsorption, scattering, transmission, and reflection at e = 50-30,000 ev, z = 1-92. *Atomic Data and Nuclear Data Tables*, 54(2):181 – 342, 1993. ISSN 0092-640X. doi: 10.1006/adnd.1993.1013.

[58] J. A. Hinton and W. Hofmann. Teraelectronvolt Astronomy. *ARA&A*, 47:523–565, September 2009. doi: 10.1146/annurev-astro-082708-101816.

[59] Matthew Hollister, Helen McGregor, Adam Woodcraft, Dan Bintley, Michael MacIntosh, and Wayne Holland. Cryogenic magnetic shielding for SCUBA-2. In *Astronomical Telescopes and Instrumentation: Synergies Between Ground and Space*, pages 702023–702023. International Society for Optics and Photonics, 2008.

[60] J. C. Houck and L. A. Denicola. ISIS: An Interactive Spectral Interpretation System for High Resolution X-Ray Spectroscopy. In N. Manset, C. Veillet, & D. Crabtree, editor, *Astronomical Data Analysis Software and Systems IX*, volume 216 of *Astronomical Society of the Pacific Conference Series*, pages 591–+, 2000.

[61] J. P. Hughes, C. E. Rakowski, D. N. Burrows, and P. O. Slane. Nucleosynthesis and Mixing in Cassiopeia A. *ApJ*, 528:L109–L113, January 2000. doi: 10.1086/312438.

[62] U. Hwang and J. M. Laming. Where Was the Iron Synthesized in Cassiopeia A? *ApJ*, 597:362–373, November 2003. doi: 10.1086/378269.

[63] U. Hwang and J. M. Laming. The Circumstellar Medium of Cassiopeia A Inferred from the Outer Ejecta Knot Properties. *ApJ*, 703:883–893, September 2009. doi: 10.1088/0004-637X/703/1/883.

[64] U. Hwang, J. M. Laming, C. Badenes, F. Berendse, J. Blondin, D. Cioffi, T. DeLaney, D. Dewey, R. Fesen, K. A. Flanagan, C. L. Fryer, P. Ghavamian, J. P. Hughes, J. A. Morse, P. P. Plucinsky, R. Petre, M. Pohl, L. Rudnick, R. Sankrit, P. O. Slane, R. K. Smith, J. Vink, and J. S. Warren. A Million Second Chandra View of Cassiopeia A. *ApJ*, 615:L117–L120, November 2004. doi: 10.1086/426186.

[65] U Hwang, R Petre, and K Flanagan. X-Ray-emitting Ejecta in Puppis A Observed with Suzaku. *ApJ*, 2008.

[66] Una Hwang and J Martin Laming. A Chandra X-ray Survey of Ejecta in the Cassiopeia A Supernova Remnant. *ApJ*, 746(2):130, February 2012.

[67] K. Isensee, G. Olmschenk, L. Rudnick, T. DeLaney, J. Rho, J. D. Smith, W. T. Reach, T. Kozasa, and H. Gomez. Nucleosynthetic Layers in the Shocked Ejecta of Cassiopeia A. *ApJ*, 757:126, October 2012. doi: 10.1088/0004-637X/757/2/126.

[68] Karl Isensee, Lawrence Rudnick, Tracey Delaney, J D Smith, Jeonghee Rho, William T Reach, Takashi Kozasa, and Haley Gomez. The Three-dimensional Structure of Interior Ejecta in Cassiopeia A at High Spectral Resolution. *ApJ*, 725:2059, December 2010.

[69] Y. Ishisaki, Y. Maeda, R. Fujimoto, M. Ozaki, K. Ebisawa, T. Takahashi, Y. Ueda, Y. Ogasaka, A. Ptak, K. Mukai, K. Hamaguchi, M. Hirayama, T. Kotani, H. Kubo, R. Shibata, M. Ebara, A. Furuzawa, R. Iizuka, H. Inoue, H. Mori, S. Okada, Y. Yokoyama, H. Matsumoto, H. Nakajima, H. Yamaguchi, N. Anabuki, N. Tawa, M. Nagai, S. Katsuda, K. Hayashida, A. Bamba, E. D. Miller, K. Sato, and N. Y. Yamasaki. Monte Carlo Simulator and Ancillary Response Generator of Suzaku XRT/XIS System for Spatially Extended Source Analysis. *PASJ*, 59:113–132, January 2007.

[70] H. Itoh. Theoretical Spectra of the Thermal X-Rays from Young Supernova Remnants. *PASJ*, 29:813–830, 1977.

[71] A. F. Iyudin, R. Diehl, H. Bloemen, W. Hermsen, G. G. Lichti, D. Morris, J. Ryan, V. Schoenfelder, H. Steinle, M. Varendorff, C. de Vries, and C. Winkler. COMPTEL observations of Ti-44 gamma-ray line emission from CAS A. *A&A*, 284:L1–L4, April 1994.

[72] A. F. Iyudin, B. Aschenbach, W. Becker, K. Dennerl, and F. Haberl. XMM-Newton observations of the supernova remnant RX J0852.0-4622/GRO J0852-4642. *A&A*, 429: 225–234, January 2005. doi: 10.1051/0004-6361:20041779.

[73] F. Jansen, D. Lumb, B. Altieri, J. Clavel, M. Ehle, C. Erd, C. Gabriel, M. Guainazzi, P. Gondoin, R. Much, R. Munoz, M. Santos, N. Schartel, D. Texier, and G. Vacanti. XMM-Newton observatory. I. The spacecraft and operations. *A&A*, 365:L1–L6, January 2001. doi: 10.1051/0004-6361:20000036.

[74] C. C. Joggerst, A. Almgren, and S. E. Woosley. Three-dimensional simulations of rayleigh-taylor mixing in core-collapse supernovae with castro. *ApJ*, 723(1):353, 2010.

[75] G. C. Jordan, S. S. Gupta, and B. S. Meyer. Nuclear reactions important in α -rich freeze-outs. *Phys. Rev. C*, 68(6):065801, December 2003. doi: 10.1103/PhysRevC.68.065801.

[76] Richard L Kelley, Kazuhisa Mitsuda, Christine A Allen, Petar Arsenovic, Michael D Audley, Thomas G Bialas, Kevin R Boyce, Robert F Boyle, Susan R Breon, Gregory V Brown, et al. The Suzaku High Resolution X-ray Spectrometer. *Publications of the Astronomical Society of Japan*, 59:S77–S112, 2007.

[77] RL Kelley, MD Audley, KR Boyce, R Fujimoto, KC Gendreau, Y Ishisaki, D McCammon, T Mihara, K Mitsuda, SH Moseley, et al. The Microcalorimeter Spectrometer on the ASTRO-E X-Ray Observatory. *Nucl. Phys. A*, 444(1):170–174, 2000.

[78] K. Koyama, R. Petre, E. V. Gotthelf, U. Hwang, M. Matsuura, M. Ozaki, and S. S. Holt. Evidence for shock acceleration of high-energy electrons in the supernova remnant SN1006. *Nature*, 378:255–258, November 1995. doi: 10.1038/378255a0.

[79] O. Krause, S. M. Birkmann, T. Usuda, T. Hattori, M. Goto, G. H. Rieke, and K. A. Misselt. The Cassiopeia A Supernova Was of Type IIb. *Science*, 320:1195–, May 2008. doi: 10.1126/science.1155788.

[80] S Kullback and R A Leibler. On Information and Sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, March 1951.

[81] J. M. Laming and U. Hwang. On the Determination of Ejecta Structure and Explosion Asymmetry from the X-Ray Knots of Cassiopeia A. *ApJ*, 597:347–361, November 2003. doi: 10.1086/378268.

[82] J. S. Lazendic, D. Dewey, N. S. Schulz, and C. R. Canizares. The Kinematic and Plasma Properties of X-Ray Knots in Cassiopeia A from the Chandra HETGS. *ApJ*, 651:250–267, November 2006. doi: 10.1086/507481.

[83] Hyunsook Lee, Vinay L Kashyap, David A van Dyk, Alanna Connors, Jeremy J Drake, Rima Izem, Xiao-Li Meng, Shandong Min, Taeyoung Park, Pete Ratzlaff, Aneta Siemiginowska, and Andreas Zezas. Accounting for Calibration Uncertainties in X-ray Analysis: Effective Areas in Spectral Fitting. *ApJ*, 731:126, April 2011.

[84] M. D. Leising. Prospects for X-Ray Studies of Supernova Electron Capture Radioactivity. *ApJ*, 563:185–190, December 2001. doi: 10.1086/323776.

[85] A. M. Levine. *The Ultrasoft Cosmic X-ray Background*. PhD thesis, Massachusetts Institute of Technology, 1976.

[86] W. H. G. Lewin, J. Doty, G. W. Clark, S. A. Rappaport, H. V. D. Bradt, R. Doxsey, D. R. Hearn, J. A. Hoffman, J. G. Jernigan, F. K. Li, W. Mayer, J. McClintock, F. Primini, and J. Richardson. The discovery of rapidly repetitive X-ray bursts from a new source in Scorpius. *ApJ*, 207:L95–L99, July 1976. doi: 10.1086/182188.

[87] L. A. Lopez, E. Ramirez-Ruiz, D. Huppenkothen, C. Badenes, and D. A. Pooley. Using the X-ray Morphology of Young Supernova Remnants to Constrain Explosion Type, Ejecta Distribution, and Chemical Mixing. *ApJ*, 732:114–+, May 2011. doi: 10.1088/0004-637X/732/2/114.

[88] Kenneth SK Lum, Claude R Canizares, George W Clark, Joan M Coyne, Thomas H Markert, Pablo J Saez, Mark L Schattenburg, and PF Winkler. The spectral archive of cosmic X-ray sources observed by the Einstein Observatory Focal Plane Crystal Spectrometer. *ApJS*, 78:423–503, 1992.

[89] D.J.C. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, 2003. ISBN 9780521642989.

[90] F. Marshall, E. Boldt, S. Holt, R. Kelley, R. Mushotzky, R. Petre, P. Serlemitsos, J. Swank, and A. Szymkowiak. X-ray Spectrometers for Observing SN1987a. In *Bulletin of the American Astronomical Society*, volume 19 of *Bulletin of the American Astronomical Society*, page 735, March 1987.

[91] D McCammon, R Almy, S Deiker, J Morgenthaler, RL Kelley, FJ Marshall, SH Moseley, CK Stahle, and AE Szymkowiak. A Sounding Rocket Payload for X-Ray Astronomy Employing High-Resolution Microcalorimeters. *Nucl. Phys. A*, 370(1):266–268, 1996.

[92] Xiao-Li Meng. Posterior Predictive p-Values. *The Annals of Statistics*, 22(3):1142–1160, September 1994. doi: 10.1214/aos/1176325622.

[93] SH Moseley, John C Mather, and Dan McCammon. Thermal Detectors as X-Ray Spectrometers. *Journal of Applied Physics*, 56(5):1257–1262, 1984.

[94] S. S. Murray, G. Fabbiano, A. Epstein, R. Giacconi, and A. C. Fabian. High-resolution X-ray observations of the Cassiopeia A supernova remnant with the Einstein Observatory. *ApJ*, 234:L69–L72, November 1979. doi: 10.1086/183111.

[95] Radford M. Neal. MCMC using Hamiltonian dynamics. In S. Brooks, A. Gelman, G. L. Jones, and X. Meng, editors, *Handbook of Markov Chain Monte Carlo*, pages 113–162. Chapman & Halll/CRC, Boca Raton, FL, USA, 2010.

[96] T. O. Niinikoski. *Liquid and Solid Helium*, chapter Cosmic-Ray Disturbances in Thermometry and Refrigeration, pages 145–147. Halsted, 1975.

[97] M. S. Noble and M. A. Nowak. Beyond XSPEC: Toward Highly Configurable Astrophysical Analysis. *PASP*, 120:821–837, July 2008. doi: 10.1086/590324.

[98] Minoru Oda. High-Resolution X-Ray Collimator with Broad Field of View for Astronomical Use. *Applied Optics*, 4(1):143–143, 1965.

[99] S. Orlando, F. Bocchino, M. Miceli, O. Petruk, and M. L. Pumo. Role of ejecta clumping and back-reaction of accelerated cosmic rays in the evolution of type ia supernova remnants. *ApJ*, 749(2):156, 2012.

[100] Tony Pan, Daniel J Patnaude, and Abraham Loeb. Super-luminous X-ray Emission from the Interaction of Supernova Ejecta with Dense Circumstellar Shells. *arXiv.org*, astro-ph.HE, March 2013.

[101] D. J. Patnaude and R. A. Fesen. Small-Scale X-Ray Variability in the Cassiopeia A Supernova Remnant. *AJ*, 133:147–153, January 2007. doi: 10.1086/509571.

[102] D. J. Patnaude, J. Vink, J. M. Laming, and R. A. Fesen. A Decline in the Nonthermal X-ray Emission from Cassiopeia A. *ApJ*, 729:L28+, March 2011. doi: 10.1088/2041-8205/729/2/L28.

[103] Jeffrey W. Percival, Kenneth H. Nordsieck, and Kurt P. Jaehnig. The ST5000: a high-precision star tracker and attitude determination system. *Proc. SPIE 7010*, 2008. doi: 10.1117/12.787917.

[104] R. Petre. Thin Shell, Segmented X-Ray Mirrors. *X-Ray Optics and Instrumentation*, 2010. doi: 10.1155/2010/412323.

[105] E. Pfeffermann, U. G. Briel, H. Hippmann, G. Kettenring, G. Metzner, P. Predehl, G. Reger, K.-H. Stephan, M. V. Zombeck, and J. Chappell. The focal plane instrumentation of the ROSAT telescope. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 733 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 519, February 1987.

[106] F. Pobell. *Matter and Methods at Low Temperatures*. Springer Berlin Heidelberg, 1996. ISBN 9783540585725.

[107] D. Porquet, R. Mewe, J. Dubau, A. J. J. Raassen, and J. S. Kaastra. Line ratios for helium-like ions: Applications to collision-dominated plasmas. *A&A*, 376:1113–1122, September 2001. doi: 10.1051/0004-6361:20010959.

[108] F. S. Porter, R. Almy, E. Apodaca, E. Figueroa-Feliciano, M. Galeazzi, R. Kelley, D. McCammon, C. K. Stahle, A. E. Szymkowiak, and W. T. Sanders. The XQC microcalorimeter sounding rocket: a stable LTD platform 30 seconds after rocket motor burnout. *Nucl. Phys. A*, 444:220–223, April 2000. doi: 10.1016/S0168-9002(99)01391-1.

[109] K. Pounds. *Frontiers of X-Ray Astronomy*, chapter Forty Years on from Aerobee 150. In *Cambridge Planetary Science* Fabian et al. [31], 2004. ISBN 9780521534871.

[110] R. Protassov, D. A. van Dyk, A. Connors, V. L. Kashyap, and A. Siemiginowska. Statistics, Handle with Care: Detecting Multiple Model Components with the Likelihood Ratio Test. *ApJ*, 571:545–559, May 2002. doi: 10.1086/339856.

[111] S. Rappaport, H. V. Bradt, S. Naranan, and G. Spada. Low-energy X-ray Spectra of Sco X-1 and Four Sagittarius Sources. *Nature*, 221:428–432, February 1969. doi: 10.1038/221428a0.

[112] T Rauscher, A Heger, R D Hoffman, and S E Woosley. Nucleosynthesis in Massive Stars With Improved Nuclear and Stellar Physics. *ApJ*, 576:323, September 2002.

[113] M. Renaud, J. Vink, A. Decourchelle, F. Lebrun, P. R. den Hartog, R. Terrier, C. Couvreur, J. Knödlseder, P. Martin, N. Prantzos, A. M. Bykov, and H. Bloemen. The Signature of ^{44}Ti in Cassiopeia A Revealed by IBIS/ISGRI on INTEGRAL. *ApJ*, 647:L41–L44, August 2006. doi: 10.1086/507300.

[114] S. P. Reynolds, B. M. Gaensler, and F. Bocchino. Magnetic Fields in Supernova Remnants and Pulsar-Wind Nebulae. *Space Sci. Rev.*, 166:231–261, May 2012. doi: 10.1007/s11214-011-9775-y.

[115] Donald B Rubin. The Bayesian Bootstrap. *The Annals of Statistics*, 9(1):130–134, 1981.

[116] John Rutherford, Daniel Dewey, Enectali Figueroa-Feliciano, Sarah N T Heine, Fabienne A Bastien, Kosuke Sato, and C R Canizares. A Decade-Baseline Study of the Plasma States of Ejecta Knots in Cassiopeia A. *ApJ*, 769(1):64, May 2013.

[117] A. Sandage, P. Osmer, R. Giacconi, P. Gorenstein, H. Gursky, J. Waters, H. Bradt, G. Garmire, B. V. Sreekantan, M. Oda, K. Osawa, and J. Jugaku. On the optical identification of SCO X-1. *ApJ*, 146:316, October 1966. doi: 10.1086/148892.

[118] E.M. Schlegel. *The Restless Universe: Understanding X-Ray Astronomy in the Age of Chandra and Newton*. Oxford University Press, USA, 2002. ISBN 9780198033844.

[119] E Schönfeld and H Janssen. Evaluation of atomic shell data. *Nucl. Phys. A*, 369(2): 527–533, 1996.

[120] P. J. Serlemitsos. Conical foil X-ray mirrors - Performance and projections. *Appl. Opt.*, 27:1447–1452, April 1988. doi: 10.1364/AO.27.001447.

[121] P. J. Serlemitsos, L. Jalota, Y. Soong, H. Kunieda, Y. Tawara, Y. Tsusaka, H. Suzuki, Y. Sakima, T. Yamazaki, H. Yoshioka, A. Furuzawa, K. Yamashita, H. Awaki, M. Itoh, Y. Ogasaka, H. Honda, and Y. Uchibori. The X-ray telescope on board ASCA. *PASJ*, 47:105–114, February 1995.

[122] D.S. Sivia and J. Skilling. *Data analysis: a Bayesian tutorial*. Oxford science publications. Oxford University Press, 2006. ISBN 9780198568315.

[123] R. K. Smith and J. P. Hughes. Ionization Equilibrium Timescales in Collisional Plasmas. *ApJ*, 718:583–585, July 2010. doi: 10.1088/0004-637X/718/1/583.

[124] R. K. Smith, G.-X. Chen, K. Kirby, and N. S. Brickhouse. A New Calculation of Ne IX Line Diagnostics. *ApJ*, 700:679–683, July 2009. doi: 10.1088/0004-637X/700/1/679.

[125] T. J. Spisak. SAS-3: An Informal History. Consisting of available photographs and monographs by the SAS-3 crew on the occasion of the last SAS Bash, May 18, 1979, 1979.

[126] Suborbital & Special Orbital Projects Directorate. NASA Sounding Rocket Program Handbook, 2005. URL `http://sites.wff.nasa.gov/code810/files/SRHB.pdf`.

[127] L.-S. The, M. D. Leising, J. D. Kurfess, W. N. Johnson, D. H. Hartmann, N. Gehrels, J. E. Grove, and W. R. Purcell. CGRO/OSSE observations of the Cassiopeia A SNR. *A&AS*, 120:C357, December 1996.

[128] L.-S. The, D. D. Clayton, R. Diehl, D. H. Hartmann, A. F. Iyudin, M. D. Leising, B. S. Meyer, Y. Motizuki, and V. Schönfelder. Are ^{44}Ti-producing supernovae exceptional? *A&A*, 450:1037–1050, May 2006. doi: 10.1051/0004-6361:20054626.

[129] M. F. Theiling and M. D. Leising. Investigation of ^{44}Ti decay X-ray lines in CasA. *NewAR*, 50:544–547, October 2006. doi: 10.1016/j.newar.2006.06.054.

[130] A. C. Thompson, editor. *X-ray Data Booklet*. Lawrence Berkeley National Laboratory, University of California, third edition, September 2009.

[131] Yasunobu Uchiyama, Felix A Aharonian, Takaaki Tanaka, Tadayuki Takahashi, and Yoshitomo Maeda. Extremely fast acceleration of cosmic rays in a supernova remnant. *Nature*, 449(7162):576–578, 2007.

[132] J. Vink. Supernova remnants: the X-ray perspective. *A&A Rev.*, 20:49, December 2012. doi: 10.1007/s00159-011-0049-1.

[133] J. Vink, J. S. Kaastra, and J. A. M. Bleeker. A new mass estimate and puzzling abundances of SNR Cassiopeia A. *A&A*, 307:L41–L44, March 1996.

[134] J. Vink, J. M. Laming, J. S. Kaastra, J. A. M. Bleeker, H. Bloemen, and U. Oberlack. Detection of the 67.9 and 78.4 keV Lines Associated with the Radioactive Decay of ^{44}Ti in Cassiopeia A. *ApJ*, 560:L79–L82, October 2001. doi: 10.1086/324172.

[135] W. Voges, B. Aschenbach, T. Boller, H. Bräuninger, U. Briel, W. Burkert, K. Dennerl, J. Englhauser, R. Gruber, F. Haberl, G. Hartner, G. Hasinger, M. Kürster, E. Pfeffermann, W. Pietsch, P. Predehl, C. Rosso, J. H. M. M. Schmitt, J. Trümper, and H. U. Zimmermann. The ROSAT all-sky survey bright source catalogue. *A&A*, 349:389–405, September 1999.

[136] Jessica S Warren, John P Hughes, Carles Badenes, Parviz Ghavamian, Christopher F McKee, David Moffett, Paul P Plucinsky, Cara Rakowski, Estela Reynoso, and Patrick Slane. Cosmic-ray acceleration at the forward shock in Tycho's supernova remnant: Evidence from Chandra X-ray observations. *ApJ*, 634(1):376, 2005.

[137] Martin C. Weisskopf. The making of the Chandra X-Ray Observatory: The project scientists perspective. *PNAS*, 107(16):7135–7140, 2010. doi: 10.1073/pnas.0913067107.

[138] R.P. Welty and J.M. Martinis. Two-stage integrated SQUID amplifier with series array output. *IEEE Transactions on Applied Superconductivity*, 3(1):2605–2608, 1993. ISSN 1051-8223. doi: 10.1109/77.233523.

[139] P. Wikus, E. Figueroa-Feliciano, S. A. Hertel, S. W. Leman, K. A. McCarthy, and J. M. Rutherford. The thermal conductivity of high modulus Zylon fibers between 400 mK and 4 K. *Cryogenics*, 48:515–517, November 2008.

[140] P. Wikus, J. M. Rutherford, S. N. Trowbridge, D. McCammon, J. S. Adams, S. R. Bandler, R. Das, W. B. Doriese, M. E. Eckart, E. Figueroa-Feliciano, R. L. Kelley, C. A. Kilbourne, S. W. Leman, F. S. Porter, and K. Sato. Flight Cryogenic System for the Micro-X Sounding Rocket. In *Proceedings of the International Cryocooler Conference*, International Cryocooler Conference, May 2008.

[141] P. Wikus, W. B. Doriese, M. E. Eckart, J. S. Adams, S. R. Bandler, R. P. Brekosky, J. A. Chervenak, A. J. Ewin, E. Figueroa-Feliciano, F. M. Finkbeiner, M. Galeazzi, G. Hilton, K. D. Irwin, R. L. Kelley, C. A. Kilbourne, S. W. Leman, D. McCammon, F. S. Porter, C. D. Reintsema, J. M. Rutherford, and S. N. Trowbridge. The Detector and Readout Systems of the Micro-X High Resolution Microcalorimeter X-Ray Imaging Rocket. In B. Young, B. Cabrera, and A. Miller, editors, *American Institute of Physics Conference Series*, volume 1185 of *American Institute of Physics Conference Series*, pages 434–437, December 2009. doi: 10.1063/1.3292371.

[142] P. Wikus, J. M. Rutherford, J. S. Adams, Y. Bagdasarova, S. R. Bandler, M. Bautz, K. Boyce, G. Brown, S. Deiker, W. B. Doriese, E. Figueroa-Feliciano, K. Flanagan, M. Galeazzi, G. C. Hilton, U. Hwang, K. D. Irwin, R. L. Kelley, T. Kallman, C. A. Kilbourne, S. Kissel, S. W. Leman, A. Levine, M. Loewenstein, D. Martinez-Galarce, R. Mushotzky, D. McCammon, D. Najjar, R. Petre, F. S. Porter, C. D. Reintsema, T. Saab, N. Schulz, P. Serlemitsos, R. Smith, J. N. Ullom, and K. Yoha. Micro-X, the TES X-ray Imaging Rocket: First Year Progress. *IEEE Transactions on Applied Superconductivity*, 19:553–556, June 2009. doi: 10.1109/TASC.2009.2019129.

[143] P. Wikus, J. S. Adams, Y. Bagdasarova, S. R. Bandler, W. B. Doriese, M. E. Eckart, E. Figueroa-Feliciano, R. L. Kelley, C. A. Kilbourne, S. W. Leman, D. McCammon, F. S. Porter, J. M. Rutherford, and S. N. Trowbridge. The Adiabatic Demagnetization Refrigerator for the Micro-X Sounding Rocket Telescope. In J. G. Weisend, editor, *American Institute of Physics Conference Series*, volume 1218 of *American Institute of Physics Conference Series*, pages 633–640, April 2010. doi: 10.1063/1.3422413.

[144] P. Wikus, J. S. Adams, R. Baker, S. R. Bandler, W. Brys, D. Dewey, W. B. Doriese, M. E. Eckart, E. Figueroa-Feliciano, R. Goeke, R. Hamersma, G. C. Hilton, U. Hwang, K. D. Irwin, R. L. Kelley, C. A. Kilbourne, S. W. Leman, D. McCammon, T. Okajima, R. H. O'Neal, Jr., F. S. Porter, C. D. Reintsema, J. M. Rutherford, P. Serlemitsos, T. Saab, K. Sato, Y. Soong, and S. N. Trowbridge. Progress on the Micro-X sounding rocket x-ray telescope: completion of flight hardware. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 7732 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, July 2010. doi: 10.1117/12.857034.

[145] H Wolter. Glancing incidence mirror systems as imaging optics for X-rays. *Ann. Physik*, 10:94, 1952.

[146] Adam L Woodcraft and Adam Gray. A Low Temperature Thermal Conductivity Database. In *AIP Conference Proceedings*, volume 1185, page 681, 2009.

[147] C.T. Yap and P.P. Saligan. A Monte Carlo approach to X-ray fluorescence analysis at low energies using an annular 55Fe source. *Nucl. Phys. A*, 251(1):140 – 143, 1986. ISSN 0168-9002. doi: http://dx.doi.org/10.1016/0168-9002(86)91160-5.

[148] Z.M. Zhang, G. Lefever-Button, and F.R. Powell. Infrared refractive index and extinction coefficient of polyimide films. *International Journal of Thermophysics*, 19: 905–916, 1998. ISSN 0195-928X. doi: 10.1023/A:1022655309574.