Novel X-Ray and Antinucleus Searches for Dark Matter

by

Brandon Michael Roach

B.S., University of Notre Dame (2017)

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

June 2023

Author . . . .
Department of Physics
May 18, 2023

Certified by . . . .
Kerstin M. Perez
Associate Professor of Physics
Thesis Supervisor

Accepted by . . . .
Lindley Winslow
Associate Department Head of Physics

Novel X-Ray and Antinucleus Searches for Dark Matter

by

Brandon Michael Roach

Submitted to the Department of Physics
on May 18, 2023, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Physics

Abstract

Over a century of cosmological observations suggest that only one-fifth of the matter density of the Universe resides in the familiar subatomic particles of the Standard Model. The remaining eighty percent is known as dark matter (DM), whose presence has so far only been inferred by its gravitational effects on SM particles at cosmic scales. This dissertation describes indirect searches for DM decaying or annihilating into Standard Model particles, particularly x-rays and antinuclei.

A variety of DM candidate particles are expected to decay or annihilate into x-ray photons, which can be detected by space-based telescopes. For example, keV-scale sterile neutrinos arise in many new-physics scenarios, and their decay would produce a distinctive x-ray line. I describe three searches for x-ray line emission from decaying sterile-neutrino DM using data from the NuSTAR x-ray observatory, thereby setting world-leading constraints on the decay rate of this DM candidate in the mass range 6–40 keV.

Low-energy cosmic antinuclei are also a powerful probe of DM. In particular, low-energy antideuterons are expected to be a nearly background-free channel for DM detection, owing to their suppressed production in cosmic-ray collisions. The General Antiparticle Spectrometer (GAPS) balloon experiment will employ a novel exotic-atom-based detection technique to achieve world-leading sensitivity to low-energy antinuclei. The GAPS experiment will contain a large-area tracker consisting of more than 1100 lithium-drifted silicon [Si(Li)] detectors, which serve as the antinucleus stopping target, x-ray spectrometer, and charged-particle tracker. I describe the x-ray testing procedure used to validate the performance of these detectors for flight. This testing also validates that thick, large-area Si(Li) detectors can be mass-produced and operated at temperatures as high as $-40\,^{\circ}$C, with potential applications throughout nuclear physics, particle physics, and astrophysics.

Thesis Supervisor: Kerstin M. Perez
Title: Associate Professor of Physics

Dedication

For Mom, Dad, and Nan, who first taught me how to be a good man, and
for Mary Beard, who first taught me how to be a researcher. Rest in peace, Mary.

Acknowledgments

A dissertation ultimately bears only one name on its cover, but it represents the contribution of dozens (or even hundreds!) of people over the course of many years.

First, I would like to thank my family, particularly my parents, Jim and Paula Roach. In a very literal way, I would not be here today were it not for the two of you. You both sacrificed so much to give me an opportunity to follow my passion and love of learning, and words as common as "thank you" seem to fail. I would also like to thank my Aunt Nanette, the aunt who has supported me so much that she basically became a "second mom." I know it has been difficult with me living so far away, but I look forward to being only a few hours' drive from you starting in a few months!

In addition to my family, there are many teachers and professors throughout the years who fostered my love of science. I would especially like to thank Mary Beard, Mark Caprio, Elizabeth Centers, Philippe Collon, Justin Crepp, Antonio Delgado, Jana Faltova, Chris Gismondi, Ben Herman, Colin Jessop, Kevin Lannon, Arielle Phillips, Daniel Robertson, and Anna Simon-Robertson.

Next, I would like to thank my thesis advisor, Kerstin Perez. You brought me into your group as someone who had no experience in astroparticle physics, and I cannot thank you enough for everything you have done for me these past six years. You were my advocate, my sounding board for crazy ideas. You brought together an amazing group of people, from whom I have learned so much. Tyler and Field, you were great officemates and remain my valued friends. Mengjiao, you are a role-model as a researcher and a postdoc, and you have taught me a great deal about electronics, vacuum/gas systems, data analysis, and so much else.

Finally, I want to thank Ian Bouche and Anika Katt for their tireless work on Si(Li) testing during 2021–2022. It was truly a pleasure working with such motivated and thoughtful student-researchers! Kelsey Yee, I now hand off the torch to you. I know you will do great things.

I am extremely grateful to the faculty and staff at MIT for their assistance with my studies and research. In particular, I would like to thank Joe Cucinotta and Jack McGlashing for their tireless work in Facilities at LNS. On behalf of the GAPS Collaboration, I am also grateful to the staff at MIT Bates (particularly Bob Abruzzio, Peter Binns, Jim Kelsey, Anne Maloney, Hamid Moazeni, and so many others), who kept us safe and on-schedule while we were working on construction. Next, I would like to thank the staff of the Academic Programs Office, particularly Sydney Miller (and Turbo!) and Cathy Modica. You do so much to keep the Department running smoothly, and we are so lucky to have you. Lastly, I would like to thank the other members of my thesis committee, Erin Kara and Salvatore Vitale, for their insightful questions and comments throughout my dissertation-writing process.

Next, I would like to thank my research colleagues. First, my NuSTAR colleagues: John Beacom, Brian Grefenstette, Shunsaku Horiuchi, Roman Krivonos, Kenny Ng, Steve Rossland, and Dan Wik. You trusted me, an incoming graduate student with no experience in high-energy astrophysics analysis, to lead much of the sterile-neutrino work. Thank you for always being there to answer my questions about the finer points of statistics and spectral modeling, and for putting your trust in me. Next, my GAPS colleagues. There are far too many to list, but I would especially like to thank my colleagues at Columbia University (Gabriel Bridges, Chuck Hailey, Derik Kraych, Mike Law, Evan Martinez, and Nate Saffold) and the University of Hawai'i at Mānoa (Cory Gerrity, Achim Stoessl, and Philip von Doetinchem) for their help with Si(Li) detector testing and GEANT simulations. I am also thankful for the hospitality of the UCLA TOF group (particularly Sydney Feldman, Takeru Hayashi, and Rene Ong) during my time in Los Angeles. Finally, I would like to thank Marco Cirelli, Nicolao Fornengo, Jordan Koechler, and Elena

Pinetti for bringing me into your group to work on MeV-to-GeV DM constraints! I look forward to seeing where this research might take us.

I am incredibly grateful to my friends, those I have met in Boston and those I have known since my days at Notre Dame (or even longer!). You have supported me and encouraged me throughout this journey, and I am a better person for having known each of you. In addition to those mentioned previously, I would especially like to thank Grace Agolia, Connor Bagwell, Sean Brudney, Nicolas Dixneuf, Bobby Dorigo-Jones, Lisa Drummond, Michael Foley, Kmolin Fu, Chris Gomez, Kevin Greenman, Robert Johnston, Nick Kamp, Jeff Krupa, Kevin Montes, Kaliroë Pappas, Chiara Salemi, Kaitlin Salyer, Alex Slavsky, Joe Smolsky, Jacqueline Tetrault, Adam Trebach, Loyd Waites, Colin Whitehead, Logan Whitehead, Payton Whitehead, Will Wolf, and Grace Yang. I would also like to thank the Morin and Reynolds families. You opened your doors to me when I was unable to return to Michigan during the holidays, and showed me a warmth and generosity of spirit that genuinely brings tears to my eyes. And Kylie, thank you for being one of my best friends. I will always be grateful for your friendship and compassion.

Finally, I would like to thank the faculty at the University of Chicago Kavli Institute for Cosmological Physics, particularly Abigail Vieregg and Alex Drlica-Wagner, for offering me a postdoctoral fellowship in your department and research group, respectively. I look forward to returning to the Midwest, and to continuing my research adventure at the interface of particle physics, silicon detectors, and cosmology.

Contents

List of Figures

List of Tables

List of Abbreviations

ADC — analog-to-digital converter

ACIS — Advanced CCD Imaging Spectrometer

ASI — Agenzia Spaziale Italiana

AMS — Alpha Magnetic Spectrometer

ARF — auxiliary response file

ASIC — application-specific integrated circuit

ATLAS — A Toroidal LHC Apparatus

BBN — Big-Bang nucleosynthesis

BESS — Balloon-borne Experiment with Superconducting Spectrometer

BGO — bismuth germanium oxide

BSM — beyond the Standard Model

BTFR — baryonic Tully-Fisher relation

CCD — charge-coupled device

CERN — Conseil européen pour la recherche nucléaire

CMB — cosmic microwave background

CMS — Compact Muon Solenoid

CP — charge-parity

CR — cosmic ray

CSBF — Columbia Scientific Balloon Facility

CU — Columbia University

CXB — cosmic x-ray background

CXO — Chandra X-ray Observatory

CZT — cadmium zinc telluride

DM — dark matter

ENC — equivalent noise charge

EPIC — European Photon Imaging Camera

EW — electroweak

FEB — front-end board

FOV — field of view

FPM — focal plane module

FSR — final-state radiation

FWHM — full width at half-maximum

GAPS — General Antiparticle Spectrometer

GBM — Gamma-ray Burst Monitor

GC — Galactic Center (Milky Way)

GFP — GAPS Functional Prototype

GR — guard ring

GRXE — Galactic ridge x-ray emission

GTI — good time interval

GW — gravitational wave

HEAO — High Energy Astronomy Observatory

HG — high-gain

HNL — heavy neutral lepton

HRC — High Resolution Camera

HPGe — high-purity germanium

HVPS — high-voltage power supply

HXD — Hard X-ray Detector

ICM — intracluster medium

IGM — intergalactic medium

IKI — Institut Kosmicheskikh Issledovaniy

INFN — Istituto Nazionale di Fisica Nucleare

INTEGRAL — International Gamma-Ray Astrophysics Laboratory

ISM — interstellar medium

ISR — initial-state radiation

JAXA — Japan Aerospace Exploration Agency

JFET — junction field-effect transistor

LDB — long-duration balloon

ΛCDM — cold dark matter with cosmological constant

LEP — Large Electron-Positron Collider

LG — low-gain

LHC — Large Hadron Collider

LIGO — Laser Interferometer Gravitational-Wave Observatory

LOS — line of sight

LTB — local trigger board

LVPS — low-voltage power supply

MOND — modified Newtonian dynamics

MOS — metal oxide semiconductor

MPE — Max Planck Institute for Extraterrestrial Physics

NASA — National Aeronautics and Space Administration

NFW — Navarro, Frenk, and White (DM profile)

NuSTAR — Nuclear Spectroscopic Telescope Array

νMSM — Neutrino Minimal Standard Model

OHP — oscillating heat pipe

OM — optics module

PBH — primordial black hole

QCD — quantum chromodynamics

RAT — Readout and Trigger box

RGS — Reflection Grating Spectrometer

RMF — redistribution matrix file

RMS — root mean square

SAA — South Atlantic Anomaly

SDD — silicon drift detector

SHM — statistical hadronization model

Si(Li) — lithium-drifted silicon

SiPM — silicon photomultiplier

SLC — Stanford Linear Collider

SM — Standard Model (of particle physics)

SPI — Spectrometer on INTEGRAL

SRG — Spectrum-Roentgen-Gamma

SSL — Space Sciences Laboratory

SUSY — supersymmetry

SXI — Soft X-ray Imager

SXS — Soft X-ray Spectrometer

TES — transition edge sensor

TOA — top of the atmosphere

TOF — time of flight

TOI — top of the instrument

UCB — University of California, Berkeley

UCLA — University of California, Los Angeles

UHM — University of Hawai'i at Mānoa

VEV — vacuum expectation value

WDM — warm dark matter

WIMP — weakly interacting massive particle

XIS — X-ray Imaging Spectrometer

XMM — X-ray Multi Mirror Mission

XRISM — X-Ray Imaging and Spectroscopy Mission

XRS — X-Ray Spectrometer

XRT — X-Ray Telescope

Statement of Contributions

In Chapter 3, I describe three NuSTAR x-ray searches for sterile neutrinos in which I had a leading role. For the first, an analysis of archival data from the M31 galaxy, I led the x-ray data analysis and wrote approximately half of the resulting journal article [1]. For the second (dedicated observations near the Galactic bulge [2]) and third (archival observations of the Galactic halo [3]), I led the data analysis and contributed the majority of text and figures in the resulting papers.

On the GAPS side, I led the Si(Li) detector testing campaign at MIT, supervised by our postdoctoral researcher Mengjiao Xiao (Chapter 6). I was responsible for the testing and analysis of several hundred Si(Li) detectors between 2019–2022, as well as training and supervising an undergraduate researcher (Anika Katt), a post-baccalaureate researcher (Ian Bouche), and a then-new graduate student (Kelsey Yee) who were integral in the testing campaign. In collaboration with the GAPS Si(Li) detector groups at MIT, the University of Hawai'i at Mānoa, and Columbia University, I led the analysis of x-ray testing results from more than 1100 flight detectors, and contributed most of the text and figures to an upcoming paper on the subject. This paper was submitted to IEEE Transactions on Nuclear Science at the same time I was completing this dissertation [4], and has been expanded upon here with the permission of my coauthors. I was also extensively involved in the mechanical, electrical, and thermal integration and testing of the Si(Li) detectors in the GAPS Functional Prototype and in the full flight instrument (at MIT Bates and at UC Berkeley). Finally, though this is outside the scope of my dissertation, I collaborated with colleagues from INFN and MIT to validate the performance of the flight ASIC in the laboratory.

1

Introduction

Nothing is too wonderful to be true.

—Sir Michael Faraday

At first glance, it might seem that the disciplines of experimental particle physics and astrophysics could not be more different. Particle-physics experiments often represent the pinnacle of controlled experimental conditions—as an example, the Large Hadron Collider (LHC) beam pipes have a vacuum colder and more rarefied than interplanetary space. Operators can choose the energies and particle types brought to collision, with the outcomes of the interactions limited only by the inherent randomness of quantum mechanics. In contrast, astrophysics experiments have almost no control over the signals (electromagnetic radiation, cosmic rays, and recently gravitational waves) that the Universe sends their way. Despite these challenges, particle-physics and astrophysics experiments are both concerned with the same fundamental questions: what are the constituents of the Universe, and how do these constituents interact to form the great variety of complex structures (from the subatomic to the intergalactic) that we observe?

Observations of the large-scale Universe suggest that only $\sim$5% of the cosmic energy density consists of familiar atomic matter. Approximately 25% appears to be in the form of invisible matter that does not interact with the known fundamental forces aside from gravity. This *dark matter* (DM) is the subject of my dissertation. (The remaining $\sim$70% of the Universe's energy density seems to reside in empty space itself, the even more enigmatic *dark energy* responsible for the accelerating expansion of the Universe.) In this chapter, I will briefly describe the problem of DM—in particular, why cosmologists believe it exists and how one might go about finding it.

1.1 The two pillars of modern physics

As Shakespeare's Hamlet reminds us, "There are more things in heaven and Earth, Horatio, / Than are dreamt of in your philosophy." For millennia, philosophers and scientists have wondered whether the Earth is made of the same matter—and follows the same physical laws—as the heavens. With the developments of mathematical theories of gravity (both Newtonian and relativistic), electromagnetism, and quantum mechanics, astrophysicists acquired powerful tools to study the distant Universe. The compositions of distant stars and galaxies could be probed with atomic spectral lines, and their velocities measured with the Doppler effect. With the discoveries of Hubble and others in the early twentieth century that many distant "nebulae" were galaxies composed of billions of stars—and that almost all of those galaxies were moving away from Earth—the field of physical cosmology truly began. These two pillars of modern physics—quantum mechanics (in the form of the Standard Model of particle physics) and general relativity—are both crucial to defining the DM problem, and so deserve further explanation.

1.1.1 The Standard Model of particle physics

The Standard Model of particle physics (hereafter simply "the Standard Model" or SM) is the culmination of humanity's centuries-long efforts to determine the fundamental nature of matter (for historical reviews, see, e.g., Refs. [5, 6]). In the SM, all of the complexity of the material Universe is reduced to a collection of quantum fields representing the elementary particles, of which there are two types: bosons with integer spin, and fermions with half-integer spin, as shown in Fig. 1-1. The fermions are the "matter" particles, and come in two broad categories: the quarks and the leptons. The quarks and leptons can both be split into three "generations" of increasing mass, with the two members of each generation differing in electric charge and weak isospin by one unit. The vector (spin-1) bosons, including the photon (γ), the gluon (g), and the $W^{\pm}/Z$ bosons, are responsible for mediating the electromagnetic, strong, and weak interactions, respectively. Finally, the recently-discovered Higgs boson (H) is a massive scalar (spin-0) responsible for breaking the high-energy symmetry between the electromagnetic and weak interactions, for granting mass to the $W^{\pm}$ and Z while keeping the photon massless, and for indirectly granting mass to the fermions.

Despite its successes, the SM has several glaring omissions. Perhaps the most notable is the absence of gravity. In the language of Einstein's general relativity, gravity is not so much a force as a response to the matter/energy content of space-time. Attempts to create a theory of quantum gravity (e.g., involving the exchange of spin-2 gravitons) are often plagued by divergences and non-renormalizability (see, e.g., Refs. [7–10] for reviews). Another issue is the observed matter-antimatter asymmetry in the Universe—if the SM interactions produce matter and antimatter in equal amounts, then where is all the antimatter? There are a number of other outstanding issues—neutrino masses[1], the strong CP problem[2], the hierarchy problem[3], etc—some or all of which may be related to a beyond-the-Standard-Model (BSM) theory. This dissertation will focus one one problem in particular: the SM only accounts for 5% of the matter/energy content of the Universe. To determine the composition of the large-scale Universe, of course, we require a change of scale.

---

[1] See Chapter 2.

[2] The apparent fine-tuning needed to ensure the lack of observed CP violation in QCD.

[3] The apparent fine-tuning of the Higgs boson mass with respect to large quantum corrections.

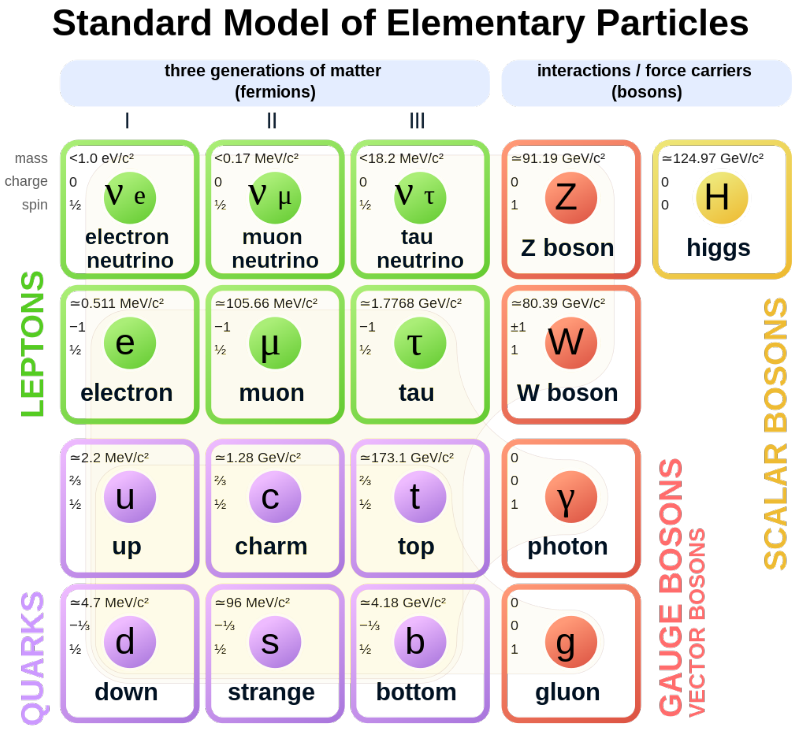

Figure 1-1: The Standard Model of particle physics following the discovery of the Higgs boson in 2012. The fermions include the leptons (green) and quarks (purple) and are grouped into three generations of increasing mass. The gauge (spin-1) bosons are shown in red, and the Higgs scalar (spin-0) is shown in yellow. © MissMJ, Cush, minor edit by Acrux13, CC-BY-3.0 via Wikimedia Commons.

1.1.2 The concordance model of cosmology

As the domain of very large length scales and very massive objects, cosmology is principally concerned with the effects of gravity. The expansion of the Universe—as well as its apparent homogeneity and isotropy at large scales—can easily be accommodated within general relativity. In relativistic cosmology described by the Friedmann-Lemaître-Robertson-Walker (FLRW) metric [11–14], the large-scale Universe is described by a small number of measurable parameters. Chief among these is the Hubble parameter[4] H_0 describing the expansion rate of the Universe, which in turn depends upon the densities ρ_i of its various constituents (matter, radiation, vacuum energy, etc). It is conventional to normalize the density of each component to the critical density[5] $\rho_{\text{crit}} \equiv 3H^2/(8\pi G_{\text{N}})$, defining the parameters $\Omega_i \equiv \rho_i/\rho_{\text{crit}}$. These parameters are also related to the overall curvature of the Universe:

$$\begin{cases} \rho > \rho_{\text{crit}} \rightarrow \text{closed (positive curvature, spherical)} \\ \rho = \rho_{\text{crit}} \rightarrow \text{flat (zero curvature)} \\ \rho < \rho_{\text{crit}} \rightarrow \text{open (negative curvature, hyperbolic)} \end{cases} \tag{1.1}$$

What are the constituents of the Universe? A wealth of cosmological data (some of which are described later in this chapter) are consistent with the following breakdown (see, e.g., Refs. [15, 16]). First, the overall curvature of the Universe is consistent with zero, i.e., $\rho = \rho_{\text{crit}}$ (equivalently, $\sum_i \Omega_i = 1$). The Universe is composed of $\sim$30% matter ($\Omega_m \approx 0.3$), with "matter" referring to any species that clumps gravitationally. Curiously, all of the baryonic[6] matter described by the Standard Model only accounts for $\sim$5% of the total energy density ($\Omega_b \sim 0.05$), with $\sim$25% apparently being comprised of *dark matter* (DM). As I will discuss in the next section, this DM cannot interact appreciably with the SM outside of gravita-

---

[4] In cosmology, a subscript 0 refers to the value measured today. It is also common to write cosmological parameters with factors of h, defined by $H \equiv h \times 100\,\text{km}\,\text{s}^{-1}\,\text{Mpc}^{-1}$.

[5] Named such because a Universe with $\rho = \rho_{\text{crit}}$ will expand forever, stopping only as $t \rightarrow \infty$.

[6] Cosmologists often use "baryonic" as a synonym for "atomic" matter (i.e., protons, neutrons, and electrons), whereas particle physicists define "baryons" as three-quark bound states (e.g., protons and neutrons). I will generally follow the cosmologists' definition unless it becomes confusing, in which case I will mention explicitly.

tional forces; furthermore, it must have a small velocity dispersion (i.e., be "cold") at the era of structure formation, hence the name *cold dark matter* (CDM). Finally, the majority of the Universe's energy density appears to reside in empty space, as the mysterious *dark energy*. Current data are consistent with a cosmological constant Λ ($\Omega_\Lambda \sim 0.7$) causing the recent accelerated expansion phase [17]. Thus, the concordance model of modern cosmology is often abbreviated ΛCDM (cold DM plus cosmological constant).

1.2 Cosmological evidence for dark matter

In this section, I will outline some of the cosmological evidence for the existence of (cold) DM, and for its large abundance compared to normal baryonic matter in the Universe. This is by no means an exhaustive account, so the interested reader is referred to the many review articles listed throughout.

1.2.1 Galaxy clusters

Galaxy clusters, containing dozens to hundreds of galaxies, are some of the largest gravitationally-bound objects in the Universe. Thus, their dynamics provide a probe of matter distribution and gravitation in the Universe at large scales.

One method of studying cluster dynamics is to measure the velocities of their constituent galaxies. For a cluster of galaxies in steady state, interacting solely under their mutual gravity, the (scalar) virial theorem states that $2T + V = 0$, where T is the (time-averaged) total kinetic energy and V is the (time-averaged) total potential energy [18]. If the individual galaxies have masses m_i and position vectors $\boldsymbol{r}_i$, one may define the total mass M, the average velocity dispersion $\langle v^2 \rangle$, and the effective gravitational radius R_G [19]:

$$\begin{cases} M = \sum_i m_i \\ \langle v^2 \rangle = \dfrac{1}{M} \sum_i m_i v_i^2 \\ R_\mathrm{G} = 2M^2 \left(\sum_{i \neq k} \dfrac{m_i m_k}{|\boldsymbol{r}_i - \boldsymbol{r}_k|} \right)^{-1} \end{cases} \tag{1.2}$$

In most cases, it is only possible to observe the projected (i.e., two-dimensional rather than three-dimensional) separations r_{ik}^proj and one component of the velocity (usually along the line of sight). With these restrictions, and assuming an isotropic distribution of cluster positions and velocities, one obtains

$$\begin{cases} \langle v^2 \rangle = 3\sigma_v^2 \\ R_\mathrm{G}^\mathrm{proj} = 2M^2 \left(\sum_{i \neq k} \dfrac{m_i m_k}{r_{ik}^\mathrm{proj}} \right)^{-1} \end{cases} \tag{1.3}$$

where σ_v^2 is the line-of-sight velocity dispersion and $R_\mathrm{G}^\mathrm{proj}$ is the projected cluster radius. With these results, the mass M of the cluster can be expressed in terms of observable quantities [19]

$$M \approx \frac{3\pi R_\mathrm{G}^\mathrm{proj} \sigma_v^2}{2G_\mathrm{N}} \approx 1.1 \times 10^{15}\, M_\odot \left(\frac{R_\mathrm{G}^\mathrm{proj}}{1\,\mathrm{Mpc}} \right) \left(\frac{\sigma_v}{10^3\,\mathrm{km\,s^{-1}}} \right)^2, \tag{1.4}$$

where G_N is the Newtonian gravitational constant.

In the early 1930s, Fritz Zwicky applied this technique to the Coma cluster, a rich cluster of $>10^3$ galaxies, and obtained $M \sim 4.5 \times 10^{13}\, M_\odot$. Converting from apparent brightness to luminosity by use of the Hubble constant, Zwicky found that the mass-to-light ratio of the Coma cluster was $\sim 500\, M_\odot\, L_\odot^{-1}$ [20]. Compared to the value of ~few $M_\odot\, L_\odot^{-1}$ expected from stars and luminous gas, this was a significant discrepancy. (We now know that Zwicky's mass-to-light ratio was a factor ~8 too high, owing to his use of $\sim 550\,\mathrm{km\,s^{-1}Mpc^{-1}}$ for the Hubble constant, but the discrepancy clearly remains.) In time, further clusters with a similar mass-to-light ratio were discovered, with several explanations proposed (see, e.g., Ref. [21] for a historical review). Some astronomers believed that galaxy clusters were dynamically unstable, existing for only a short time in cosmic history and thus violating the steady-state assumption of the virial theorem. Other astronomers believed that the mass was non-luminous baryonic matter, e.g., in the form of cold gas, compact objects, or dwarf galaxies. Of course, there was also the suspicion that the inferred mass was some new type of invisible, non-baryonic dark matter.

With the advent of sensitive x-ray telescopes in the latter half of the twentieth century, astronomers gained another tool for studying galaxy clusters. This subfield is far too vast to fully describe here—the reader is referred to, e.g., reviews in Refs. [22–28]. In brief, the x-ray emission from clusters includes both continuum (e.g., bremsstrahlung) and line-emission components from collisionally-excited gas in the intracluster medium (ICM) at temperatures $\sim$few $\times\ 10^7$ K (corresponding to x-ray energies $\sim$few keV). For a galaxy cluster in hydrostatic equilibrium, the total mass $M_{\rm tot}$ enclosed within a radius r is given by (e.g., Ref. [27])

$$M_{\rm tot}(r) = -\frac{krT_{\rm gas}(r)}{\mu m_{\rm u} G_{\rm N}} \left(\frac{\partial \ln T_{\rm gas}}{\partial \ln r} + \frac{\partial \ln n_{\rm gas}}{\partial \ln r} \right), \tag{1.5}$$

where k is the Boltzmann constant, $T_{\rm gas}$ is the gas temperature, $\mu \approx (2X + 0.75Y + 0.5Z)^{-1}$ is the mean molecular weight of the gas, X is the mass fraction of hydrogen in the gas, Y is the mass fraction in helium, Z is the mass fraction of all heavier elements, $m_{\rm u}$ is the atomic mass unit, and $n_{\rm gas}$ is the number density of gas atoms. By measuring the x-ray surface brightness profile of the cluster, as well as the relative contributions of continuum and line emission, it is possible to construct temperature, density, and abundance profiles of gas within the cluster, and to independently determine $M_{\rm gas}$ and $M_{\rm tot}$. The $M_{\rm tot}$ values inferred in this way are generally consistent with those obtained from gravitational lensing and from the Sunyaev-Zel′dovich effect[7] (see, e.g., Refs. [30–34] for recent work along these

---

[7] The upscattering of CMB photons by energetic electrons in the ICM (see, e.g., Ref. [29]).

lines).

To conclude the discussion on galaxy clusters, it is necessary to mention the (in)famous cluster 1E 0657–56, better known as the *Bullet cluster*. This object is in fact two galaxy clusters that passed through each other, causing much of the ICM to be stripped from the clusters (owing to the electromagnetic interactions between the charged particles in the ICM). Using gravitational lensing, it is possible to map the mass distribution of the clusters emerging on the other side of the collision. All told, the x-ray-emitting gas comprises $\sim$20% of the mass density of the Bullet cluster, with the (visible) stellar mass of the galaxies comprising another $\sim$10% [35] The vast majority of the cluster mass is spatially consistent with the visible galaxies (suggesting that, unlike the intracluster gas, it is collisionless), but is not detected electromagnetically. The Bullet cluster is often presented as some of the most direct evidence of DM (e.g., the title of Ref. [36]) and is a demanding test for any modified-gravity theory to meet [37], with the latter requiring as-yet-undetected "dark baryons" (cold molecular gas bound to cluster galaxies [38]), eV-scale sterile neutrinos [39], or other exotica.

1.2.2 Galactic rotation curves

The orbits of stars and other baryonic matter (e.g., hydrogen clouds) provide a valuable tracer of the mass distribution within galaxies. Consider a disk-shaped galaxy with mass profile $M(r)$ (i.e., there is mass M within a radius r). From Newtonian gravity, a star in circular orbit at a distance r from the galactic center would have an orbital speed [40]

$$v(r) = \sqrt{r\frac{\partial \Phi_{\rm N}}{\partial r}}, \tag{1.6}$$

where $\Phi_{\rm N}$ is the $1/r$ Newtonian gravitational potential. In galactic astronomy, $v(r)$ is known as a rotation curve. Assuming that $\Phi_{\rm N}$ eventually decreases with distance (i.e., that all of the mass of the galaxy is enclosed within some finite radius, perhaps comparable to the size of the visible disk), the rotation curve at large r would be expected to take the Keplerian form $v(r) \simeq \sqrt{G_{\rm N}M_{\rm tot}/r}$.

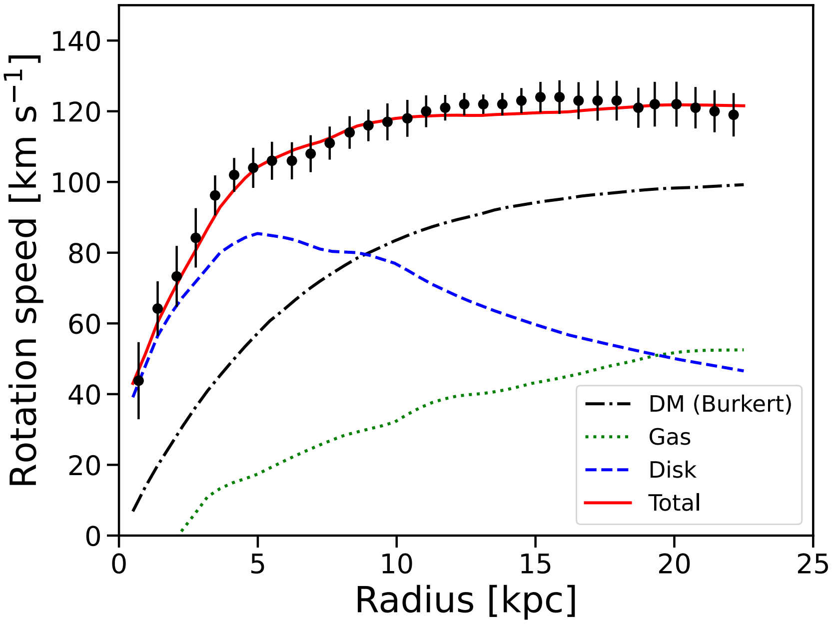

Figure 1-2: Rotation curve of spiral galaxy NGC 4559 (black points) from Ref. [41]. The best-fit mass model (solid red) consists of DM (Burkert profile, black dashed-dotted), disk (blue dashed), and gas (green dotted) added in quadrature [42, 43].

Beginning in the 1950s, a concerted program was launched to measure the rotation curves of galaxies in the nearby Universe. Most methods relied on the Doppler shift of atomic spectral lines of known wavelength, e.g., the 21-cm radio line from neutral atomic hydrogen (HI) and the 2.6-mm microwave line from $^{12}C^{16}O$ (see, e.g., Ref. [44]). The rotation-curve data compiled by Rubin and others for dozens (and eventually hundreds) of galaxies quickly began to tell a very different story than the expected Keplerian $r^{-1/2}$ scaling. As shown in Fig. 1-2 for a representative galaxy NGC 4559, at large distances from the centers of galaxies and far from any visible matter in the disk, many of the rotation curves remained flat. Taken at face value, this suggested that galaxies were embedded in diffuse "halos" of DM much larger than the visible extent of the galaxies (though see the section on modified gravity later in this chapter for a counter-proposal).

Galactic rotation curves are also a powerful probe of the DM density profiles[8]

---

[8] Strictly, the DM density is $\rho_\chi(\boldsymbol{r})$, but most observations are consistent with (and thus most analyses

$\rho_\chi(r)$ within galaxies. There are many different ways to classify DM density profiles, with a particularly relevant one being the behavior of ρ_χ at small radii. *Cuspy* profiles have densities which rise sharply at small radii, and include the famous Navarro-Frenk-White (NFW) profile [45]

$$\rho_{\rm NFW}(r) \propto \left(\frac{r}{r_s}\right)^{-1}\left(1+\frac{r}{r_s}\right)^{-2} \tag{1.7}$$

derived from DM-only N-body simulations, where r_s is the scale radius corresponding to $\mathrm{d}\ln\rho_{\rm NFW}/\mathrm{d}\ln r = -2$. In recent decades, rotation-curve measurements in a variety of galactic systems have suggested that the DM density is approximately flat near the centers of galaxies, producing *cored* profiles[9] (see, e.g., Ref. [47] for a recent review). Such cored profiles include the Burkert[10]profile [48]

$$\rho_{\rm Bur}(r) \propto \left(1+\frac{r}{r_s}\right)^{-1}\left(1+\frac{r^2}{r_s^2}\right)^{-2} \tag{1.8}$$

and the Einasto profile [49, 50]

$$\rho_{\rm Ein} \propto \exp\left[-\left(\frac{r}{r_s}\right)^{\frac{1}{n}}\right], \tag{1.9}$$

where n is some positive real number ($n < 1$ profiles being more cored, and $n > 1$ profiles being more cuspy). As an intermediate between cored and cuspy profiles, one may also use the shallow NFW (sNFW) profile

$$\rho_{\rm sNFW} \propto \left(\frac{r}{r_s}\right)^{-\gamma}\left(1+\frac{r}{r_s}\right)^{\gamma-3}, \tag{1.10}$$

with some recent analyses favoring $\gamma \approx 0.7$ [51, 52]. When studying the Milky

---

assume) that the DM halo is approximately spherically-symmetric, i.e., $\rho(\boldsymbol{r}) \simeq \rho(r)$.

[9] In the standard ΛCDM cosmology, cored profiles may arise from baryonic feedback "flattening" the DM density near galactic centers. Alternatively, cored profiles may be a consequence of self-interacting DM (see, e.g., Ref. [46]).

[10] Note that different DM profiles may define the scale radius r_s differently, and thus values for different profiles should be treated with care.

Way, most DM profiles are normalized to the local DM density ρ_0 in the vicinity of the Sun, with recent analyses giving ρ_0 between $\sim$0.3–0.5 GeV cm^{-3} (see, e.g., Refs. [53–56]).

1.2.3 Cosmic microwave background

One of the central predictions of the hot Big Bang model is the presence of a relic photon background. Following the end of the Big Bang nucleosynthesis epoch ($t \sim$ few minutes), the baryonic matter in the Universe consisted of an optically-thick[11] hydrogen-helium plasma in a sea of thermal photons. For several hundred thousand years, as the Universe expanded, the photon temperature T_γ and the plasma temperature evolved in lock-step. Approximately 3.5×10^5 years after the Big Bang, the temperature had sufficiently decreased for electrons to bind to protons for the first time, forming neutral hydrogen atoms. This dramatically decreased the opacity of the plasma, allowing photons to stream freely and thus decoupling T_γ from the temperature of the baryons. These relic photons would continue to stream through the Universe, their wavelengths redshifting with the cosmic expansion, until ultimately being detected at Earth as an almost-perfect black-body spectrum peaking at $\lambda \sim 2\,\mathrm{mm}$: the cosmic microwave background (CMB). Measurements by the FIRAS instrument on the Cosmic Background Explorer (COBE) satellite find an average temperature $T_0 = 2.728(2)\,\mathrm{K}$ [57].

As conclusively demonstrated by many experiments, though the CMB temperature is uniform[12] to better than one part in 10^4 across the sky, there are anisotropies. Defining $\delta T(\hat{\boldsymbol{n}}) \equiv T(\hat{\boldsymbol{n}}) - T_0$ to be the temperature fluctuations about the average T_0 in a direction $\hat{\boldsymbol{n}}(\theta, \phi)$, we define the temperature contrast $\Theta(\hat{\boldsymbol{n}}) \equiv \delta T(\hat{\boldsymbol{n}})/T_0$. After subtraction or modeling of Galactic and other astrophysical foregrounds, the CMB brightness temperature field over some region of sky $\mathcal{S}$ can be decomposed

---

[11] That is, the mean free path of the photons was much smaller than the horizon size.

[12] Not including the dipole induced by Earth's orbit around the Sun, the Sun's orbit around the Galaxy, etc.

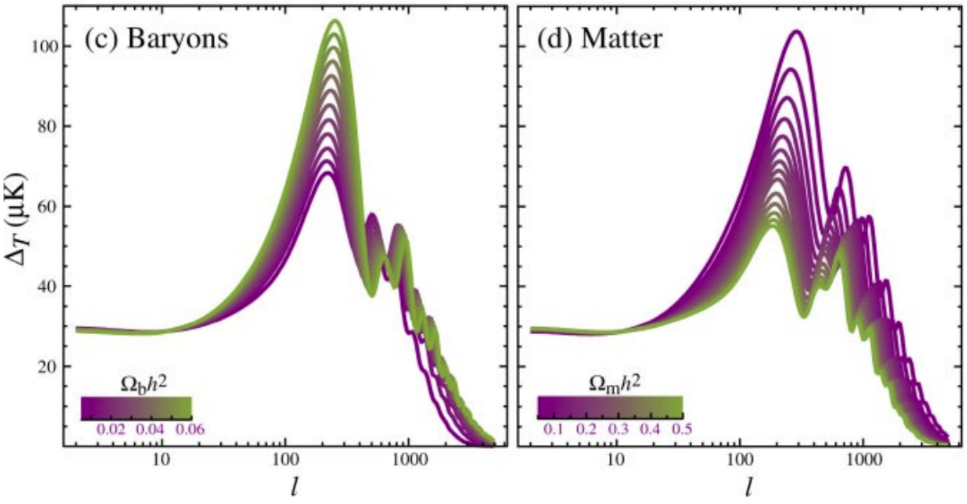

Figure 1-3: Sensitivity of the CMB temperature angular power spectrum—plotted as power per logarithmic increment in l, i.e., $\Delta_T^2 \equiv T_0^2 l(l+1)C_l^{\mathrm{TT}}/(2\pi)$—to changes in $\Omega_b h^2$ (**left**) and $\Omega_m h^2$ (**right**) about a fiducial cosmology with $\Omega_{\mathrm{tot}} = 1$, $\Omega_b h^2 = 0.02$, $\Omega_m h^2 = 0.147$, and $\Omega_\Lambda = 0.65$. Reprinted from Ref. [58] with permission. © 2002 by Annual Reviews.

into a sum of spherical harmonics $Y_{lm}(\hat{\boldsymbol{n}})$ with moments Θ_{lm} given by [58]

$$\Theta_{lm} = \int_{\mathcal{S}} Y_{lm}^{\star}(\hat{\boldsymbol{n}})\Theta(\hat{\boldsymbol{n}})\mathrm{d}\Omega, \tag{1.11}$$

where Ω is the solid angle and the star (*) denotes complex conjugation. These moments obey the orthonormality condition $\langle \Theta_{lm}^{\star}\Theta_{l'm'} \rangle = \delta_{ll'}\delta_{mm'}C_l^{\mathrm{TT}}$, where the "TT" refers to the temperature autocorrelation and the angle brackets refer to averaging over the solid angle. In most cases, it is some quantity proportional to C_l^{TT} to which cosmologists refer when they speak of "the CMB power spectrum." (Recently, however, there has also been substantial interest in polarization-sensitive CMB measurements, e.g., of C_l^{EE}, C_l^{TE}, and C_l^{BB}, with E and B referring to the two linear polarizations—see, e.g., Refs. [58, 59] for reviews.)

A plot of the CMB angular power spectrum is shown in Fig. 1-3. Perhaps the most striking features of the CMB power spectrum are the succession of *acoustic peaks* beginning at $l \approx 200$. The standard model of cosmology with inflation predicts an essentially scale-invariant spectrum of quantum fluctuations in the plasma density. As an extremely brief summary, these density fluctuations translate into sinusoidal waves in the baryon-photon plasma, propagating at the speed of sound $c_s \simeq 1/\sqrt{3}$ ($c = 1$ units). The acoustic peaks simply correspond to modes caught at their maxima or minima at the time of recombination (since power is proportional to amplitude squared). The first peak at $l \approx 200$ represents the mode with wavelength equal to the sound horizon scale at recombination, $s_{\rm rec} = c_s \eta_{\rm rec}$, where η is the conformal time. The higher-l "harmonics" correspond to modes that underwent multiple compression/rarefaction cycles (odd-numbered peaks correspond to compression in gravitational potential wells, and even-numbered peaks to rarefaction). As shown in Fig. 1-3, the amplitude of the odd-numbered peaks increases with $\Omega_b h^2$ (all other parameters held constant), since the increased baryon loading allows for higher compression within the potential wells; however, the additional baryons also damp the oscillations more rapidly. Thus, the second and third peaks are particularly sensitive to the amounts of dark versus baryonic matter in the Universe (see, e.g., Ref. [15]).

1.2.4 The matter power spectrum and structure formation

The density perturbations in the early Universe, as revealed by the CMB, provided the "seeds" around which larger structures (galaxies and galaxy clusters) eventually grew. Defining the density contrast $\Delta(\boldsymbol{x}) \equiv [\rho(\boldsymbol{x}) - \rho_0]/\rho_0$, where ρ_0 is the average density, the matter power spectrum in wavenumber space $\mathcal{P}_m(k)$ is defined as the autocorrelation of the density contrast field (see, e.g., Ref. [19])

$$\langle \widetilde{\Delta}(\boldsymbol{k}) \widetilde{\Delta}^\star(\boldsymbol{k}') \rangle = (2\pi)^3 \delta^3(\boldsymbol{k} - \boldsymbol{k}') \mathcal{P}_m(k), \tag{1.12}$$

where the tildes denote Fourier-transformed quantities and $\delta(x)$ is the Dirac δ-function.

With the amount of DM being nearly five times the amount of baryonic matter, it is clear that the dynamics of DM are crucial in translating between the initial

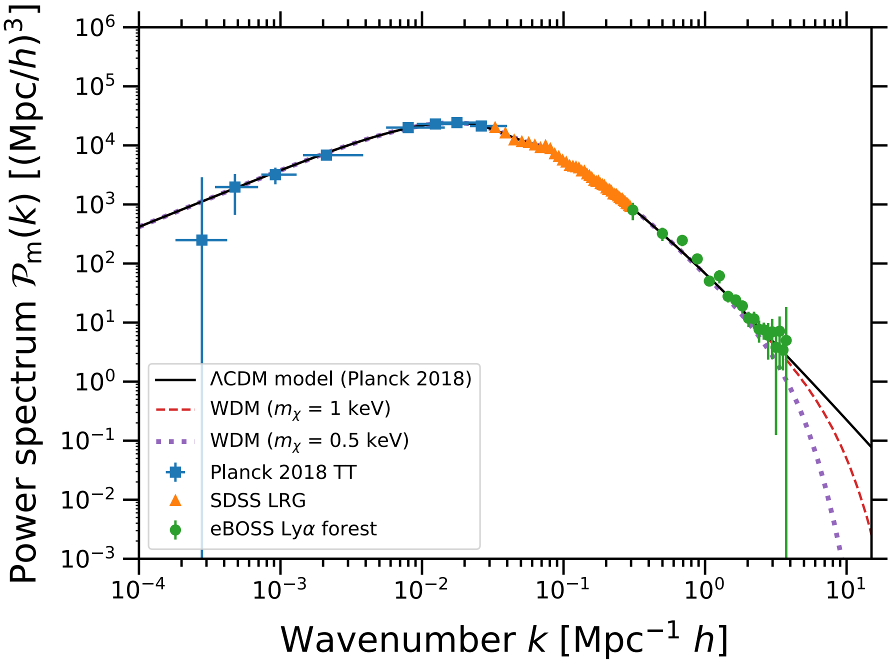

Figure 1-4: Plot of the matter power spectrum $\mathcal{P}_m$ as a function of wavenumber k from different cosmic probes (data from Ref. [60] and references therein). The black line shows the best-fit ΛCDM model from Planck [15]. The dashed (dotted) line indicates a warm-DM thermal relic with $m_\chi = 1\,\mathrm{keV}$ (0.5 keV), using the approximate transfer functions from Ref. [61].

power spectrum[13] $\mathcal{P}_0(k)$ and that observed at later times. In the language of non-relativistic linear perturbation theory (applicable if $\Delta \ll 1$), the evolution of the matter power spectrum is described by [19]

$$\mathcal{P}(k,t) \propto D_+^2(t)\mathcal{P}_0(k), \tag{1.13}$$

where the growth factor D_+ evolves with the cosmic scale factor a as

$$D_+(a) \propto \frac{H(a)}{H_0}\int_0^a \frac{\mathrm{d}x}{[\Omega_m x^{-1} + \Omega_\Lambda x^{-2} - (1-\Omega_m-\Omega_\Lambda)]^{3/2}}. \tag{1.14}$$

---

[13] This is generally assumed to be a power law $\mathcal{P}_0(k) \propto k^{n_s}$, with $n_s = 1$ known as a scale-invariant or Harrison-Zel'dovich spectrum. Recent results from Planck indicate that $n_s = 0.9652(42)$ (68% CI, Ref. [15]).

(Note that D_+ after DM becomes nonrelativistic depends only on the underlying cosmological parameters, not on the DM model itself.)

It is instructive to consider two limiting cases: cold (negligible velocity) and hot (relativistic) DM. In CDM, structure forms from the "bottom up"—numerical simulations indicate that small halos form first and subsequently grow by accretion and merging with other halos (see, e.g., Refs. [62–65] for reviews). On the other hand, DM particles with non-negligible velocities—warm (WDM) to hot (HDM)—are able to freely stream out of primordial overdensities, "washing out" small structures and decreasing the matter power spectrum at small length scales (large k). In particular, simulations with HDM (e.g., eV-scale neutrinos) indicate that structures form from the "top down," since only the most massive overdensities are able to gravitationally retain hot DM. These halos would then gradually fragment into smaller halos over billions of years. Figure 1-4 shows the effect of different "temperatures" of DM on the matter power spectrum. For DM produced in thermal equilibrium, it is generally the case that lower DM masses result in warmer DM, though as we will see in Chapter 2, it is quite possible to obtain cool-to-cold DM with masses as low as the keV scale by other means.

1.2.5 Big Bang nucleosynthesis

The Big Bang nucleosynthesis (BBN) epoch began at $t \sim 1\,\text{minute}$ after the Big Bang, and is the process by which the light elements (isotopes of hydrogen, helium, and lithium) were created from the sea of protons and neutrons (for reviews, see, e.g., Refs. [16, 66, 67]). The onset of BBN depended on two conditions. First, the neutron-proton equilibrium mediated by weak interactions (e.g., $p + e^- \leftrightarrow n + \nu_e$ and $n + e^+ \leftrightarrow p + \bar{\nu}_e$) was broken as the Universe expanded and cooled, allowing protons and neutrons to exist as independent entities. This expansion also reduced the number density and energy of the thermal photons, allowing deuterium (D or ^{2}H, the first step in the BBN reaction network) to form via $p + n \to \text{D} + \gamma$ without being immediately photo-dissociated. As such, one of the central parameters governing BBN is the baryon-to-photon ratio $\eta_b = n_b/n_\gamma$, where n is the number density.

Using terrestrial accelerator measurements of the BBN reaction rates, it is possible to generate predictions for the BBN isotope yields as a function of η_b for one's favorite choice of cosmological parameters using codes such as AlterBBN [68, 69]. (Many cosmological parameters, including the densities Ω_i and the number of relativistic degrees of freedom, affect the expansion rate H of the Universe, and thus its temperature and density evolution.) Deuterium is a particularly attractive probe of BBN, since its low nuclear binding energy ($\sim$2.2 MeV) means that it is easily destroyed in stellar fusion processes; thus, any cosmological measurements of the deuterium-to-hydrogen ratio D/H provide a bound on the BBN value. The primordial deuterium abundance can be determined by spectroscopy of damped Lyman-α systems, dense ($\gtrsim 10^{20}\,\mathrm{cm}^{-2}$) columns of low-metallicity[14] intergalactic gas along the line of sight to quasars. A recent evaluation gives $\mathrm{D/H} = 2.547(25) \times 10^{-5}$ [16], a level of precision at which uncertainties in the nuclear reaction rates affecting deuterium may dominate the BBN error budget [70]. Spectroscopy of extragalactic HII regions also provides a constraint on the cosmic ^{4}He abundance $Y_p \equiv {}^4\mathrm{He/H}$, with a recent evaluation giving $Y_p = 0.245(3)$ [16]. The one isotope whose measured abundance[15] is not consistent with BBN is ^{7}Li, with measurements in extremely low-metallicity stars indicating an abundance $^7\mathrm{Li/H} = (1.6 \pm 0.3) \times 10^{-10}$ [16]. This is nearly a factor of three lower than the BBN concordance model, with a tension approaching 5σ. There are many explanations proposed for this *lithium problem*, ranging from stellar physics (perhaps stars destroy their lithium more quickly than previously believed) to various BSM models (see, e.g., Ref. [72] for a recent review). Excluding the potentially-problematic ^{7}Li, the deuterium and ^{4}He abundances are consistent with a baryon-to-photon ratio $\eta_b^{\mathrm{BBN}} = (6.143 \pm 0.190) \times 10^{-10}$ [16].

---

[14] In astrophysics, any elements with atomic number $Z > 2$ are classified as "metals." Owing to the minuscule production of lithium in BBN, any material with significant metallicity has been processed by stellar fusion.

[15] Note that ^{3}He is generally not used as a probe, since its nuclear behavior in stars is less well-understood (see, e.g., discussions in Refs. [66, 71]).

To connect the BBN constraints with DM, we return to the CMB. First, the measured CMB temperature $T_0 \approx 2.725\,\mathrm{K}$ allows us to convert the baryon-to-photon ratio η_b^{BBN} to the baryon density, giving a value $\Omega_b^{\mathrm{BBN}} h^2 = 0.02244(69)$ [16]. When compared to the Planck telescope's CMB constraints $\Omega_b^{\mathrm{CMB}} h^2 = 0.02233(15)$ assuming a BBN prior on the helium mass fraction Y_p, the level of agreement is excellent [15]. Furthermore, the high precision of the Planck measurements allows a CMB determination of $Y_p = 0.241(25)$ independent of—but consistent with—BBN constraints [15]. The striking agreement between cosmological parameters when the Universe was only a few minutes old, compared to when it was $\sim 3 \times 10^5$ years old, represents one of the strongest tests of the ΛCDM model.

1.3 Particle-physics searches for dark matter

All particle-physics searches for DM must make one crucial assumption: that DM somehow couples to the familiar SM particles by non-gravitational forces. Though there are an uncountable number of different DM models, with a similarly uncountable number of possible DM-SM interactions, DM searches generally fit into one of three classes: collider production, direct detection, and indirect detection.

1.3.1 Collider production

First, DM particles may be produced at colliders via a variety of mechanisms. Owing to their feeble interactions, any DM particles produced in colliders will stream out of the detectors without depositing energy; thus, their presence must be inferred indirectly, e.g., via conservation of energy and momentum. (For recent reviews, see, e.g., Refs. [73–76].) Since the space of collider-accessible DM models is vast, I will briefly describe two largely model-independent search strategies here: Higgs-coupled DM and initial-state radiation (ISR) tagging.

The precise measurement of the Higgs boson's properties is a central aim of current and next-generation energy-frontier experiments. In particular, if DM couples to the Higgs (the so-called “Higgs portal,” e.g., Refs. [77–81]), it may be accessible to colliders in a variety of ways. As one example, consider the so-called *Higgs-strahlung* process $e^+e^- \to ZH \to f_Z\bar{f}_Z\chi\bar{\chi}$, where the Z and H decay into fermion-antifermion pairs (SM particles $f_Z\bar{f}_Z$ and DM particles $\chi\bar{\chi}$, respectively). For a SM Higgs mass $m_H \approx 125\,\text{GeV}$, the Higgs-strahlung cross section peaks at a center-of-mass energy $\sqrt{s} \approx 250\,\text{GeV}$, so we set this as the fiducial energy. The search strategy depends somewhat on how m_χ compares to m_H and $\sqrt{s}$. First, $2(m_\chi + m_{f_Z}) < \sqrt{s}$ for the reaction to be kinematically allowed. Assuming this condition is satisfied, we compare[16] m_χ and m_H. If $m_\chi < \frac{1}{2}m_H$, an on-shell[17] Higgs boson can decay into $\chi\bar{\chi}$ pairs. This is desirable from an experimental point of view, since the reaction probability is enhanced if the H and Z are on-shell. On the other hand, if $m_\chi > \frac{1}{2}m_H$, an on-shell Higgs boson cannot decay into $\chi\bar{\chi}$; thus, the reaction proceeds through an off-shell (virtual) Higgs boson, which greatly suppresses the rate. In the extreme case that $m_\chi > \frac{1}{2}(m_Z + m_H)$, the Z boson can also be forced off-shell, at an even greater cost to the reaction probability. In any case, if the Z boson decay is well-reconstructed (e.g., via $Z \to \mu^+\mu^-$), the production of H can be inferred by energy and momentum conservation, even if the H decay is not observed. Proposed e^+e^- colliders operating up to $\sqrt{s} \approx 350\,\text{GeV}$ (e.g., FCC-ee [82], CEPC [83], CLIC [84], and ILC [85]) may have sensitivity to the branching ratio $\mathcal{B}(H \to \text{invisible})$ at the few-percent level, and since $\mathcal{B}(H \to \text{invisible}) \lesssim 10^{-3}$ in the SM [86–88], any deviation may be indicative of new physics. Of course, this method only applies if the mass of the invisible particle $m_\chi < \frac{1}{2}m_H$, since it involves the search for invisible decays of on-shell Higgs bosons. For larger m_χ produced by off-shell Higgs decays, proton-proton colliders are desirable, with proposed 100-TeV machines (e.g., FCC-hh [89] and and SppC [90]) potentially being sensitive to Higgs-DM interactions for m_χ up to 1 TeV depending on the spin

---

[16] In this discussion, we neglect m_{f_Z}. This is a reasonable assumption since the heaviest SM fermion into which the Z boson can decay is the b quark ($m \sim 4\,\text{GeV}$).

[17] That is, a particle whose energy and three-momentum satisfy $E^2 = \boldsymbol{p}^2 + m^2$. States violating this condition are called “off-shell” or “virtual,” and cannot be detected (they are meaningful only in the context of perturbation theory to map between initial and final states).

of χ and the assumed $H\chi\chi$ coupling strength (e.g., Refs. [91, 92]).

To set the stage for ISR tagging, consider the simplified DM production process at a hadron collider $q\bar{q} \to X^* \to \chi\bar{\chi}$, where X is some new (off-shell) boson with mass m_X that connects to both the SM and DM with couplings g_q and g_χ, respectively (i.e., the cross section $\sigma \propto g_q^2 g_\chi^2$). Since the χ are undetected, this process[18] results in missing transverse momentum[19] $p_{\mathrm{T}}^{\mathrm{miss}}$. The radiation of a single gluon, photon, or $W^\pm/Z$ boson from one of the initial-state quarks serves as a high-momentum "tag" for the event, with single-gluon (→ mono-jet) events dominating at hadron colliders. Most of the SM backgrounds arise from processes such as $q\bar{q} \to gV$, where V is an undetected photon or W/Z boson (either landing outside the detector acceptance or producing neutrinos). Recent results from the ATLAS [93–95] and CMS [96–99] experiments have strongly constrained mono-boson plus $p_{\mathrm{T}}^{\mathrm{miss}}$ production, which have been interpreted in the context of many different DM models.

There are several fundamental challenges for the collider production and detection of DM. First, the DM particles may be too massive or too weakly-coupled to the SM to be efficiently produced at colliders. Furthermore, even if collider experiments someday detect a long-lived, non-interacting BSM particle, such a particle does not necessarily constitute all (or even any) of the DM in the Universe. Ultimately, any DM candidate must also be observed elsewhere in the cosmos by other means, two of which are discussed in the next sections.

1.3.2 Direct detection

Direct-detection experiments seek to detect DM particles in the Galactic halo recoiling against nuclei and/or electrons in detectors (see, e.g., Refs. [100–102] for reviews). As the Sun orbits the Galactic Center, its orbital speed $|\boldsymbol{v}_\odot| \approx 220\,\mathrm{km\,s^{-1}}$

---

[18] By symmetry, the process $q\bar{q} \to X^* \to q\bar{q}$ with cross section $\sigma \propto g_q^4$ is also allowed, leading to a resonance in the di-jet mass spectrum at mass m_X. If the coupling g_q is very small, however, the latter process may be too rare to observe.

[19] Since the colliding particles have negligible momenta transverse to the beam direction, any imbalance in p_{T} in the detector is indicative of undetected particles.

[103, 104] carries it through the Galactic DM halo, inducing an apparent flux of DM particles. If the DM density in the solar neighborhood is $\rho_0 \approx 0.3\,\mathrm{GeV\,cm^{-3}}$ [53–55], the flux of DM particles through the Solar System is approximately

$$7\times 10^{6}\,\mathrm{cm^{-2}\,s^{-1}}\left(\frac{\rho_0}{0.3\,\mathrm{GeV\,cm^{-3}}}\right)\left(\frac{|\boldsymbol{v}_\odot|}{220\,\mathrm{km\,s^{-1}}}\right)\left(\frac{1\,\mathrm{GeV}}{m_\chi}\right). \tag{1.15}$$

Consider a DM particle of mass m_χ and velocity $v \ll c$ striking a nucleus of mass m_A, causing the nucleus to recoil with kinetic energy E_r. If the total sensitive mass of the detector is M, the rate of DM-nucleus scattering events per unit E_r is (see, e.g., Refs. [100, 102])

$$\frac{\mathrm{d}^2N}{\mathrm{d}E_r\,\mathrm{d}t} = \frac{\rho_0 M}{m_A m_\chi}\int_{v_{\mathrm{min}}}^{v_{\mathrm{esc}}} v f(v)\frac{\mathrm{d}\sigma}{\mathrm{d}E_r}\,\mathrm{d}v. \tag{1.16}$$

Here, $v_{\mathrm{min}} = \sqrt{E_{\mathrm{thr}} m_A/(2\mu_A^2)}$ is the minimum DM speed required to induce a detectable recoil energy above the detector energy threshold E_{thr}, and $\mu_A = m_A m_\chi/(m_A + m_\chi)$ is the reduced mass of the DM-nucleus system. The velocity distribution function $f(v)$ is usually chosen to be an isothermal and isotropic Maxwell-Boltzmann distribution, with possible anisotropies due to the so-called "Gaia sausage" of the Milky Way halo [105]. Finally, $v_{\mathrm{esc}} \approx 500\text{–}550\,\mathrm{km\,s^{-1}}$ is the escape velocity of the Milky Way at the solar orbit, above which DM particles are gravitationally unbound [106–108].

All of the particle physics of the interaction is encoded in the differential scattering cross section $\mathrm{d}\sigma/\mathrm{d}E_r$. The form of $\mathrm{d}\sigma/\mathrm{d}E_r$ is of course strongly dependent on the DM model in question, but it is often possible to make some simplifying assumptions. After "integrating out" the (presumably heavy) mediator particle, the differential cross section is (see, e.g., Refs. [100, 102])

$$\frac{\mathrm{d}\sigma}{\mathrm{d}E_r} = \frac{m_A}{2v^2\mu_A^2}\left[\sigma_{\mathrm{SI}}F_{\mathrm{SI}}^2(E_r) + \sigma_{\mathrm{SD}}F_{\mathrm{SD}}^2(E_r)\right], \tag{1.17}$$

where F_{SI} (F_{SD}) is the spin-independent (spin-dependent) scattering form factor and depends on the structure of the nucleus as well as the momentum transfer. To compare experiments using different target materials on an equal footing, it is common for direct-detection experiments to report their constraints in terms of the DM-nucleon cross section σ_N (see, e.g., Refs. [100, 102]) using the relation

$$\sigma_{\rm SI} = \sigma_N \frac{\mu_A^2}{\mu_N^2} \frac{[f_p Z + f_n(A-Z)]^2}{f_n^2}, \tag{1.18}$$

where σ_N is the DM-nucleon cross section, μ_N is the DM-nucleon reduced mass, f_p (f_n) is the relative DM-proton (-neutron) coupling strength, Z is the atomic number, and A is the mass number. In the limit that $f_p = f_n$, the SI cross section $\sigma_{\rm SI} \propto A^2$. On the one hand, this motivates using heavier target nuclei (e.g., Xe), but on the other, the increased nuclear mass reduces the recoil energy (thereby requiring lower-threshold and lower-background detectors). The SD cross section $\sigma_{\rm SD}$ does not benefit from this coherence effect, as it involves the coupling to the spins of unpaired nucleons. The rest of this section will focus on SI searches unless otherwise noted.

There are several classes of direct-detection experiments based on the ways in which the recoil energy E_r is ultimately detected, mainly in the forms of light (scintillation), charge, and/or heat (phonons). This section can only provide a very brief sketch of direct-detection techniques. For more details, see the aforementioned review articles, as well as the individual collaborations' papers. In all cases, extreme care must be taken to reduce radioactive backgrounds[20] as much as possible so as not to mask any DM scattering signals. This requirement is somewhat in tension with the need to make detectors as large as possible to maximize the expected DM scattering rate. The lowest detectable DM mass is determined by the interplay between $E_{\rm thr}$, m_A, and m_χ, since larger m_A and smaller m_χ each result in lower available recoil energies. Energy thresholds of $\mathcal{O}(\text{few keV}_{\rm ee})$ have been achieved in inorganic scintillators such as sodium iodide (e.g., Refs. [109, 110]) and liquid

---

[20] For the nuclear-recoil experiments outlined previously, neutrons (mainly from the uranium and thorium decay chains) are a particularly pernicious background, since they can induce DM-like nuclear recoils.

xenon (e.g., Refs. [111–113]), with values[21] as low as tens of $\mathrm{eV_{ee}}$ reported in cryogenic detectors (e.g., Refs. [114–117]).

A collection of current and projected near-future direct-detection constraints is shown in Fig. 1-5. Several features deserve special mention. First, most of the constraints rapidly weaken at low DM masses m_χ, owing to the low recoil energies produced upon scattering with detectors (thus requiring detectors with ever-lower energy thresholds). For $m_\chi \lesssim 1\,\mathrm{GeV}$, it is sometimes more advantageous to search for DM-electron scattering, owing to the lower mass (and hence greater recoil energy) of electrons compared to nuclei [118–122]. At larger DM masses, the cross section constraints weaken because a smaller DM number density is needed to comprise the same mass density, thus leading to a linear decrease in the expected scattering rate. Lastly, the shaded region at the bottom of the figure represents the *neutrino fog* [123]. For sufficiently low DM cross sections, the background from coherent elastic neutrino-nucleus scattering becomes significant. At lower DM masses ($\lesssim 10\,\mathrm{GeV}$), solar neutrinos are the dominant background, with atmospheric neutrinos dominating at higher m_χ. This region of the parameter space was previously called the *neutrino floor*, since it was assumed that no useful DM constraints could be derived below it. Recent work has shown that progress can still be made, albeit slowly [123]. (In quantitative terms, the neutrino floor/fog represents the transition from signal-dominated to background-dominated DM searches.)

If DM is not discovered prior to encountering the neutrino fog, there are several possibilities to overcome this issue. First, the flux of DM particles is not constant in time; rather, it varies on an annual basis due to the Earth's orbit around the Sun, reaching a maximum (minimum) in June (December) (see, e.g., Refs. [124–126] for reviews). Though this annual modulation is only on the ∼few percent level, it provides a sensitive tool for distinguishing between DM and other terrestrial backgrounds, including neutrinos [127]. Second, some groups are pursuing *directional detection* techniques, in which the arrival direction of the incoming DM particle is

---

[21] To compare different materials (as well as nuclear versus electronic recoils) on an equal basis, it is common to convert nuclear-recoil energies E_{nr} to electronic-recoil equivalent values E_{ee} via $E_{\mathrm{ee}} = \mathcal{Q}(E_{\mathrm{nr}})E_{\mathrm{nr}}$, where $\mathcal{Q}$ is the nuclear quenching factor.

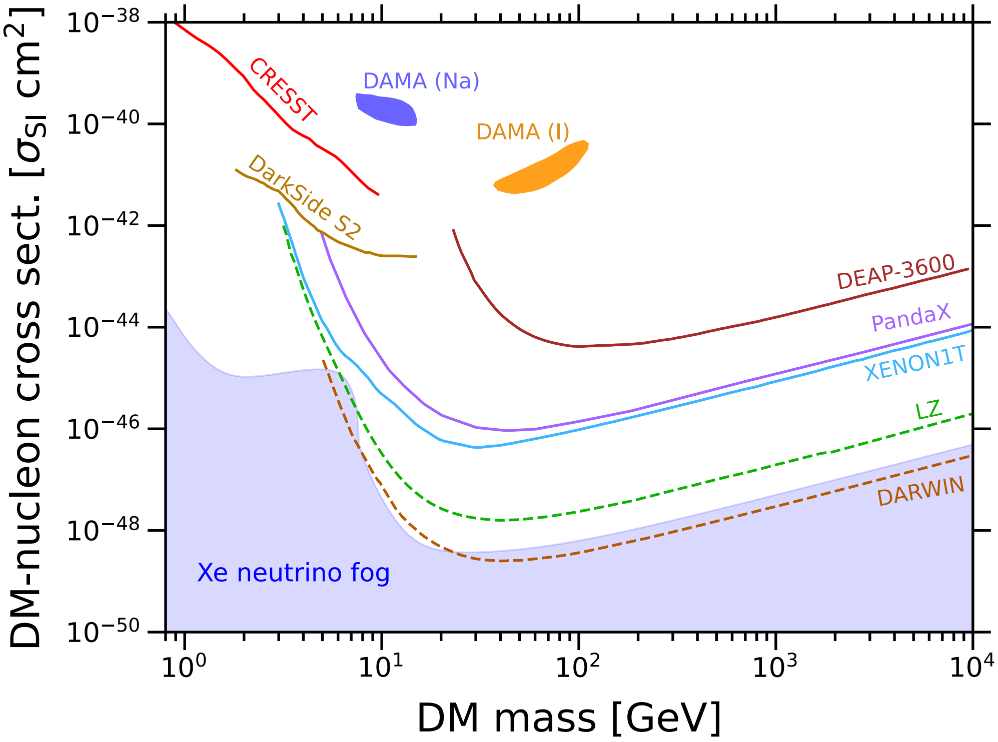

Figure 1-5: Collection of selected current (solid lines) and projected (dashed lines) direct-detection constraints on the spin-independent DM-nucleon cross section σ_{SI}. The closed DAMA contours indicate the 1σ allowed regions for DM-nucleus scattering on sodium (Na) or iodine (I) nuclei. The blue shaded region at the bottom indicates the xenon neutrino fog, where direct-detection searches transition from signal- to background-dominated. Data from Ref. [123] and references therein.

reconstructed from the track left by the struck particle in the detector (see, e.g., Refs. [128, 129] for reviews). The recently-formed CYGNUS collaboration [130] proposes to build a network of He+SF_6 time projection chambers with total volume $>10^3\,\mathrm{m}^3$, with a projected DM sensitivity several orders of magnitude below the xenon neutrino floor. Furthermore, several groups are exploring the use of solid-state detectors, including diamond-crystal defects (see, e.g., Refs. [131–133]) and nuclear emulsions (see, e.g., Refs. [134, 135]) as directional DM detectors.

Before concluding this section, it should be noted that the DAMA/LIBRA group has consistently claimed the observation of an annual modulation in the $<6\,\mathrm{keV_{ee}}$ data from their their sodium iodide scintillators (see, e.g., Ref. [136] for a recent discussion). The statistical significance of this claim now exceeds 10σ, and exhibits the correct annual modulation and single-hit signal morphology expected from DM. Interpreted in the context of spin-independent DM scattering on sodium (iodine) nuclei, these results produce the closed contours shown in Fig. 1-5 near $m_\chi \sim 10\,\mathrm{GeV}$ ($\sim$50 GeV). This signal has not been observed in any other DM experiments, with constraints from many groups disfavoring the DAMA/LIBRA allowed regions by several orders of magnitude. Most of these constraints were derived from liquid noble gas and semiconductor detectors, leading some to suggest "isospin-violating" DM models that preferentially coupled to sodium iodide [137]. Recently, experiments based on the same sodium iodide approach as DAMA/LIBRA (e.g., COSINE-100 and ANAIS) have come online, placing strong tension on the DAMA/LIBRA modulation claim [110, 138–141]. Additional data from these experiments, as well as the forthcoming SABRE [142] and COSINUS [143] experiments, should help to clarify matters.

1.3.3 Indirect detection

The techniques of *indirect detection* rely on only one assumption—that DM particles out in the Universe decay and/or annihilate into SM particles, which may be detectable at Earth (see, e.g., Refs. [144, 145] for reviews). Whether the interaction involves bosons, quark-antiquark pairs, or other unstable SM particles, there are only a few stable final-state particles which would be detectable at Earth: photons, electrons/positrons, (anti)nuclei, and (anti)neutrinos. This section will introduce indirect detection in the context of photons (the formalism for neutrinos is almost identical), with the discussion of charged-particle propagation and detection deferred to Chapter 4.

Consider a small clump of DM with mass M at a distance r from Earth. If DM is composed of a single species of particle with mass m_χ, the number of DM particles in the clump is M/m_χ. If the DM has decay rate Γ_i into some channel—for concreteness, the single-photon channel $\chi \to \gamma + x$, where x is any undetected particle(s)—the flux of photons observed at Earth is

$$F_\gamma \equiv \frac{\mathrm{d}^2 N_\gamma}{\mathrm{d}A\,\mathrm{d}t} = \frac{\Gamma_{\chi\to\gamma x}}{4\pi m_\chi r^2} M. \tag{1.19}$$

If we are interested in the spectral flux $\mathrm{d}F_\gamma/\mathrm{d}E_\gamma$, we need only add a factor $\mathrm{d}N_\gamma/\mathrm{d}E_\gamma$ representing the final-state photon spectrum (normalized to the total number of photons emitted):

$$\frac{\mathrm{d}F_\gamma}{\mathrm{d}E_\gamma} = \frac{\Gamma_{\chi\to\gamma x}}{4\pi m_\chi r^2}\frac{\mathrm{d}N_\gamma}{\mathrm{d}E_\gamma} M. \tag{1.20}$$

In general, we cannot observe single isolated clumps of DM; rather, we observe large regions of galactic halos. To account for this, we replace the mass M with $\int \rho_\chi\, r^2\,\mathrm{d}r\,\mathrm{d}\Omega$, where r is now interpreted as the distance along the line of sight (LOS) through the DM halo and Ω is the solid angle. If ξ is the detection efficiency as a function of solid angle, we thus obtain the specific intensity

$$\mathcal{I}_\gamma \equiv \frac{\mathrm{d}^2 F_\gamma}{\mathrm{d}E_\gamma\,\mathrm{d}\Omega} = \underbrace{\frac{\Gamma_{\chi\to\gamma x}}{4\pi m_\chi}\frac{\mathrm{d}N_\gamma}{\mathrm{d}E_\gamma}}_{\text{particle physics}} \times \underbrace{\frac{1}{\Delta\Omega}\int \xi\,\mathrm{d}\Omega \int_{\mathrm{LOS}} \rho_\chi\,\mathrm{d}r}_{\text{astrophysics}} \qquad \text{(decay)}, \tag{1.21}$$

where $\Delta\Omega \equiv \int_{\mathrm{FOV}} \xi\,\mathrm{d}\Omega$ is the effective solid angle of the instrument. The specific intensity neatly factorizes into two groups of terms. The "particle physics" terms relate to the specific DM model[22] under consideration, whereas the "astrophysics" term depends only on the DM column density within the instrument FOV. In the case of decaying DM, the latter is usually called the $\mathcal{D}$-factor. The mathematical formalism for DM annihilation proceeds in a very similar way, with a few minor changes. First, the relevant particle-physics quantity is the velocity-averaged annihilation cross section $\langle \sigma v \rangle$. Second the annihilation rate scales as ρ_χ^2. Finally, the annihilation rate is suppressed by a further factor of 2 if the DM is not its own antiparticle. Putting everything together, we find

$$\mathcal{I}_\gamma = \frac{\langle \sigma v \rangle}{4\pi a m_\chi^2} \frac{\mathrm{d}N_\gamma}{\mathrm{d}E_\gamma} \times \frac{1}{\Delta\Omega} \int_{\mathrm{FOV}} \xi \, \mathrm{d}\Omega \int_{\mathrm{LOS}} \rho_\chi^2 \, \mathrm{d}r, \qquad \text{(annihilation)} \tag{1.22}$$

where $a = 2$ ($a = 4$) corresponds to DM which is (not) its own antiparticle. In the annihilation case, the astrophysics factor is usually referred to as the $\mathcal{J}$-factor.

In all cases, indirect DM experiments must find a compromise between increasing the expected DM signal and decreasing the expected background. In the case of photon and neutrino detectors, where the particles point directly back to their sources (within the angular resolution of the instrument), this requirement means considering regions where the DM column density is greatest (e.g., the Galactic center, nearby galaxies, and galaxy clusters). In both the decay and annihilation cases, the density profile $\rho_\chi(r)$ represents one of the largest sources of systematic uncertainty when converting between observed flux F_γ and DM properties (e.g., Γ, $\langle \sigma v \rangle$, etc). Additionally, since baryonic matter traces DM to a good approximation, the regions of densest DM correspond to the regions of densest baryonic matter, thus leading to higher backgrounds. The strategies for dealing with these backgrounds depend on the instrument and DM model under consideration. Specific approaches in the context of keV-scale sterile neutrinos and cosmic antinuclei will be discussed in Chapters 3 and 5, respectively.

There are two complications[23] to the indirect-detection approach described for

---

[22] Strictly, the particle physics term also includes a sum over all channels producing the particle of interest, in this case photons, with each channel having its own Γ and $\mathrm{d}N/\mathrm{d}E$.

[23] Discussion of the *detection* of charged particles, e.g., $e^\pm$ and (anti)nuclei, is deferred to Chapter 5.

photons. First, DM may not immediately decay or annihilate into photons, instead proceeding through quark-antiquark, charged-lepton, and gauge-boson production. The hadronization of quarks produces jets of particles, many of which subsequently lead into final states involving photons (e.g., from the decays of π^0 mesons, initial/final-state radiation, internal bremsstrahlung, etc). Fortunately, these processes can (in principle) be calculated using Monte Carlo event generators, allowing the spectra $\mathrm{d}N_\gamma/\mathrm{d}E_\gamma$ to be tabulated for different DM masses (see, e.g., Ref. [146]). Furthermore, the photon flux can be enhanced by a variety of secondary processes (see, e.g., Ref. [147] and references therein). First, in regions of intense magnetic fields, $e^\pm$ (from DM or otherwise) emit synchrotron radiation, leading to a significant flux of radio waves. Second, energetic $e^\pm$ can upscatter infrared photons from the ambient radiation field (from starlight, the CMB, etc) to higher energies (e.g., x-rays) via the inverse-Compton process. Finally, bremsstrahlung emission from $e^\pm$ scattering on the interstellar medium may produce γ-rays in the GeV range for electroweak-scale DM particles.

As one might expect, the search space for indirect DM detection is vast (too vast to summarize in this dissertation). For photons, the detectable parameter space covers more than 20 orders of magnitude in energy, from radio waves to $>$100-TeV γ-rays. The wide array of current and proposed neutrino detectors opens a window in energy from the MeV scale to $>10^{15}$ GeV [148, 149]. Cosmic-ray detectors are sensitive to $e^\pm$ and (anti)nuclei with energies ranging from several hundred MeV to more than 10^{12} GeV [150]. Though most indirect-detection constraints are dependent on the DM model through $\mathrm{d}N/\mathrm{d}E$, the wide range of indirect-detection experiments allows for complementarity in strategy. First, if a DM candidate is detected at colliders or direct-detection experiments, it is very likely accessible to indirect-detection experiments as well. Second, if a DM candidate is observed in one indirect-detection channel (e.g., γ-rays) and is sufficiently massive for several final states to be kinematically accessible, it is quite likely that it will also be detectable in another channel (e.g., lower-energy photons, neutrinos, etc).

1.4 Gravitational constraints

At present, the only knowledge we have of DM concerns its gravitational effects on galaxies, galaxy clusters, and large-scale structures. It is also entirely possible that DM is too feebly-coupled to the SM interactions for it to be detected by the techniques outlined in the previous section. Thus, it pays to say a few words on gravitational DM constraints. This section cannot hope to summarize all of the interesting work[24] being done in this area, so this section is intended to give a sense of some recent directions in the field. I refer the interested reader to the proceedings of the recent Snowmass process (e.g., Refs. [151–156] and the references therein).

1.4.1 Primordial black holes

When they are not feeding on surrounding matter, black holes fit all of the requirements for DM. The extreme energy densities in the early Universe shortly after the Big Bang may lead to the production of many exotic species, including primordial black holes (PBHs). These objects may be formed by a variety of mechanisms (see, e.g., Refs. [157, 158] for recent reviews), with perhaps the simplest being direct gravitational collapse from primordial overdensities. In this scenario, the PBH mass M_{PBH} depends on the formation time $t_\bullet$ following the standard expansion rate in the radiation-dominated epoch:

$$M_{\mathrm{PBH}} \sim \frac{c^3 t_\bullet}{G_{\mathrm{N}}} \sim 10^{15}\,\mathrm{g}\left(\frac{t_\bullet}{10^{-23}\,\mathrm{s}}\right). \tag{1.23}$$

If PBHs are produced by this mechanism, their mass distribution is expected to be "monochromatic"—i.e., peaked sharply near a single value of M_{PBH}. A lower bound of $M_{\mathrm{PBH}} \gtrsim 10^{15}\,\mathrm{g}$ arises from evaporation due to Hawking radiation—any PBHs with masses less than this would not survive to the present day. In the mass

---

[24] For the sake of brevity, this section does not describe constraints arising from galactic rotation curves or galaxy clusters, since those have already been discussed earlier in this chapter. Of course, they will remain crucial tools for the study of DM well into the future.

range $10^{22}\,\mathrm{g} \lesssim M_{\mathrm{PBH}} \lesssim 10^{35}\,\mathrm{g}$, PBHs are constrained by gravitational microlensing, i.e., the brief increase in optical flux occurring when a compact object passes in front of a background star. For PBHs above $10^{33}\,\mathrm{g} \sim 1\,M_{\odot}$, constraints from gravitational-wave telescopes and the disruption of stellar binary systems are most relevant. With the exception of the mass range $\sim 10^{17}$–10^{22} g, which has been difficult to probe, the fractional contribution of PBHs to the observed DM density is generally constrained to be below $\sim 10^{-2}$–10^{-4} (see, e.g., Ref. [159] for a recent collection of constraints).

1.4.2 Gravitational waves

In 2016, the LIGO and VIRGO collaborations announced the first detection of gravitational waves (GWs) from a binary black-hole merger, GW150914 [160], followed shortly by the first multimessenger (gravitational-wave plus electromagnetic) observation of a binary neutron-star merger, GW170817 [161, 162]. With these and subsequent observations, astrophysicists gained a powerful new tool to study the cosmos.

Different DM models may modify the inspirals of compact binary systems in unique (and potentially detectable) ways. As one example, the extreme densities and electromagnetic fields inside neutron stars may result in a significant coupling to light DM fields (e.g., axions[25]), generating a potentially significant force between binary neutron stars and modifying the GW signal during the inspiral phase [165–167]. A potential consequence of ultralight bosonic DM near spinning black holes is the phenomenon of superradiance, i.e., the transfer of large amounts of energy and angular momentum to an ambient field (see, e.g., Ref. [168] for an extensive review). Since superradiance "spins down" black holes over long timescales, the detection of fast-spinning black holes serves as a constraint on the viability of superradiance, and thus on the properties of possible new boson fields (see, e.g., Refs. [169–171]).

---

[25] There have also been proposals to search for axions and dark photons via their effect on the lasers and mirrors inside GW interferometers [163, 164].

As a final example of GW probes of DM, consider pulsar timing arrays (PTAs). Since millisecond pulsars are extremely accurate[26] "clocks," it is possible to use correlations in the arrival times of their radio signals at Earth to search for distortions in the light-travel time across the Galaxy (see, e.g., Refs. [173–175] for reviews). Originally proposed to search for the nHz-to-μHz GW signals from binary supermassive black holes, PTAs are also sensitive to compact substructures (e.g., DM subhalos) expected in a variety of DM models. Current facilities (see, e.g., Refs. [176, 177]) have already delivered novel constraints on ultralight DM ($m \lesssim 10^{-20}$ eV), and it is anticipated that next-generation facilities may be able to deliver constraints on substructures with masses well below $1\,M_\odot$—approaching the scales expected for subhalos in ΛCDM (see, e.g., Refs.[178–180]).

1.4.3 Gravitational lensing

Gravitational lensing—the distortion of light from a background object (the source) by a massive foreground object (the lens)—provides a powerful method for studying the distribution of mass along the line of sight. Gravitational microlensing was previously discussed in the context of PBHs; here, I describe lensing on more cosmic scales.

The prototypical gravitational lensing system is a *strong lens*: a background object (often a quasar) is distorted into an Einstein ring by a foreground galaxy along the line of sight. The pattern of emission depends on the geometry of the source-lens-observer system, as well as the distribution of mass within the lens. Thus, strong gravitational lenses provide a powerful tool for probing the mass distribution of galaxies. In particular, brightness fluctuations in the lensed image can be used to map the distribution of dark substructures in the lensing galaxy. Current facilities (e.g., the Hubble Space Telescope, Keck telescopes, and ground-based radio interferometers) are able to resolve substructures with masses as low as $\sim 10^8\,M_\odot$ (see, e.g., Refs. [181–187]). Pushing below this figure will require im-

---

[26] Many millisecond pulsars have spin-down rates $\dot{P} \lesssim 10^{-19}\,\mathrm{s\,s^{-1}}$, rivaling the accuracy of atomic clocks (see, e.g., Ref. [172]).

provements in the systematics of lens and source modeling (see, e.g., Refs. [188–191]), as well as next-generation observational facilities. Fortunately, upcoming ground-based surveys such as the LSST [192] and space telescopes such as Euclid [193] and Roman [194] should discover tens of thousands of strong-lensing systems for follow-up at other facilities.

At the largest scales, gravitational lensing of CMB photons by massive structures can also provide meaningful constraints on DM parameters (see, e.g., Refs. [195, 196]). This avenue of research is still quite young, with the first detection of CMB anisotropies associated with gravitational structures only arriving several years ago (see, e.g., Refs. [197–199]). Upcoming high-resolution experiments such as CMB-S4 [200] and the Simons Observatory [201], as well as the proposed CMB-HD [202], are well-positioned to take advantage of the lensing maps generated by upcoming optical surveys (e.g., LSST, Euclid, and Roman). Furthermore, these CMB instruments will deliver constraints on other aspects of DM, including its thermal history and potential interactions with electrons and protons (see, e.g., Refs. [203, 204]).

1.5 A few words on modified gravity

There are two possible solutions to the issue of "missing mass" in cosmology: the existence of additional (dark) matter, and/or modifications to the laws of gravity. Though this dissertation focuses on the former, it is instructive to say a few words on the latter.

Perhaps the most popular theory of modified gravity is "modified Newtonian dynamics" (MOND). As proposed by Milgrom [205], MOND amounts to a modification of Newton's second law, which would now read

$$\boldsymbol{F} = I\left(\frac{a}{a_0}\right) m\boldsymbol{a}. \tag{1.24}$$

Here, $\boldsymbol{F}$ is the force applied to a mass m causing the observed acceleration $\boldsymbol{a}$, I is a function which will be described shortly, and a_0 is a fundamental acceleration scale below which the simple $\boldsymbol{F} = m\boldsymbol{a}$ begins to break down. This is equivalent to relating the measured gravitational acceleration $\boldsymbol{g}$ to the Newtonian expectation $\boldsymbol{g}_{\mathrm{N}}$ via[27]

$$\boldsymbol{g}_{\mathrm{N}} = I\left(\frac{g}{a_0}\right)\boldsymbol{g}. \tag{1.25}$$

Milgrom proposed that at small accelerations (i.e., $a \ll a_0$, the so-called "deep MOND" limit), $I \to g/a_0$, whereas for larger accelerations, $I \to 1$, preserving the Newtonian behavior. (At this stage, the function I is purely phenomenological—its exact behavior between the $a \ll a_0$ and $a \gg a_0$ limits depends on the underlying theory.) Assuming that the outskirts of galaxies satisfied the deep MOND limit, Milgrom found that the rotation curves tended toward an asymptotic velocity

$$v = (a_0 G M)^{1/4}. \tag{1.26}$$

With a single value $a_0 \sim 10^{-10}\,\mathrm{m\,s^{-2}}$, MOND can account for a wide variety of galactic dynamics, not just rotation curves (an extensive review can be found in, e.g., Ref. [37]). A particularly striking example is the baryonic[28] Tully-Fisher relation (BTFR, proposed in Ref. [207])

$$\log M_b = \alpha \log v_f - \log \beta. \tag{1.27}$$

Here, $M_b = M_\star + M_{\mathrm{gas}}$ is the total baryonic mass of a galaxy, $M_\star$ is the mass in stars, M_{gas} is the mass in gas, and v_f is the asymptotic rotation speed at large distances from the galaxy. With a single parameter a_0, MOND predicts both the slope ($\alpha = 4$) and intercept ($\beta = a_0 G_{\mathrm{N}}$) of the BTFR across more than four orders of magnitude in M_b. In contrast, reproducing the observed slope and scatter of the BTFR in the context of ΛCDM has historically been difficult. As an additional curiosity, the value of the Hubble parameter $H_0 \approx 70\,\mathrm{km\,s^{-1}\,Mpc^{-1}} \sim 2 \times 10^{-18}\,\mathrm{s^{-1}}$ is quite

---

[27] It should be noted that MOND is a manifestly non-relativistic theory. More about this later.

[28] The original Tully-Fisher paper [206] related $\log L$ to $\log v_f$, where L is the optical luminosity (a reasonable proxy for M_b for massive star-rich galaxies).

close to the inferred value of a_0/c, as well as to the cosmological constant $c\sqrt{\Lambda}$ (see, e.g., Ref. [208]). Whether these are coincidences or indications of a deeper unifying theory remains to be seen.

Despite its successes at galactic scales, MOND struggles at the scale of galaxy clusters. Crucially, MOND cannot account for all of the gravitating matter in galaxy clusters, with the total gravitational mass often several times greater than the observed baryonic mass, particularly in the central regions of clusters [37]. There are several possibilities: (i) missing baryons, (ii) non-baryonic DM, and/or (iii) modifications to MOND (e.g., additional fields or relativistic effects). Finally, it must be noted that MOND is clearly a non-relativistic theory, leading to violations of the strong equivalence principle and a variety of nonlinear effects such as the external field effect (see, e.g., Refs.[209, 210]). A variety of relativistic embeddings of MOND-like theories have been proposed (see, e.g., Ref. [37]), with some recent efforts succeeding in reproducing the CMB and matter power spectra (see, e.g., Ref. [211]).

Perhaps the main question of the MOND-versus-DM debate is, "If the universe is made of cold dark matter, why does MOND get any predictions right?" [212]. MOND makes a remarkably consistent set of predictions for galactic dynamics using only a single parameter, the fundamental acceleration scale a_0, plus the observed baryonic matter distribution. In the absence of a direct detection of DM, or a measured deviation in Newton's second law at low accelerations, choosing between MOND and ΛCDM is somewhat of a matter of personal taste and focus. Cosmologists interested in the large-scale Universe will naturally gravitate to ΛCDM, assuming that the difficulties of modeling baryonic effects in galaxies and clusters will eventually be resolved; conversely, astronomers focusing on the dynamics of galaxies and clusters might be more interested in the wide range of galactic phenomena modeled by MOND, while hoping to find a relativistic parent theory consistent with cosmological data (see the discussion in Ref. [37]).

For the purposes of this dissertation, I will assume the existence of DM. The wealth of precision galactic and cosmological data expected in the coming years may very well settle the issue on the side of MOND, DM, or—perhaps—both. In any case, I hope the instrumentation and data-analysis techniques described in the following chapters may be useful to the broader astroparticle physics community, whether DM ultimately exists or not.

2

Sterile Neutrinos

Without the darkness, how would we recognize the light?

—Tuvok, *Star Trek: Voyager*

In this chapter, I will discuss some of the motivations for sterile neutrinos, a class of hypothetical particles that may account for DM (among many other outstanding issues in particle physics). In Sec. 2.1, I briefly summarize the current three-neutrino picture of the SM. In Sec. 2.2, I describe the experimental and theoretical consequences of the observed nonzero SM neutrino masses. In Sec. 2.3, I describe the neutrino minimal Standard Model (νMSM), a theory which extends the SM by introducing three massive sterile neutrinos and which may simultaneously explain DM, the SM neutrino masses, and the matter-antimatter asymmetry of the Universe. In particular, the keV-scale νMSM sterile neutrino will be the benchmark model for the x-ray analyses described in Chapter 3.

2.1 Neutrinos in the Standard Model

Before considering extensions of the SM, it will be helpful to review the SM neutrino sector as it currently exists. The SM contains three neutrinos whose flavor eigenstates ν_ℓ form weak isospin doublets $\mathsf{L}_\ell = (\nu_\ell, \ell^-)^\mathsf{T}$ with the charged leptons ℓ^-, where $\ell \in \{e, \mu, \tau\}$. The SM Lagrangian for neutrino interactions is the sum of charged-current (CC) and neutral-current (NC) terms [213–215]:

$$\begin{aligned}\mathcal{L}_\nu^{\text{SM}} &= \mathcal{L}_{\text{CC}}^{\text{SM}} + \mathcal{L}_{\text{NC}}^{\text{SM}} \\ &= \left[-\frac{g_w}{\sqrt{2}} \sum_\ell \overline{\nu_{\ell\text{L}}} \gamma^\mu W_\mu^+ \ell_\text{L}^- + \text{h.c.} \right] - \frac{g_w}{2\cos\theta_\text{W}} \sum_\ell \overline{\nu_{\ell\text{L}}} \gamma^\mu Z_\mu \nu_{\ell\text{L}}. \end{aligned} \tag{2.1}$$

Here, g_w is the dimensionless $SU(2)$ coupling, $\theta_\text{W} \simeq \cos^{-1}(m_W/m_Z)$ is the Weinberg angle, γ^μ are the γ-matrices, $\bar{\nu} \equiv \nu^\dagger \gamma^0$, and the subscript "L" denotes the left-handed component of the fields. At this stage, it will be helpful to briefly review a few aspects of neutrino physics in the SM.

2.1.1 Neutrinos are left-handed

In quantum field theories, there are two concepts related to the "handedness" of a particle (see the discussion in, e.g., Refs. [214–216]), here assumed to be a spin-$\frac{1}{2}$ fermion for concreteness. If the fermion has three-momentum $\boldsymbol{p}$ and spin $\boldsymbol{S}$, it is possible to define the *helicity* $\lambda \equiv \boldsymbol{S} \cdot \boldsymbol{p}/|\boldsymbol{p}|$. Fermions with $\lambda < 0$ are said to be left-helical, and those with $\lambda > 0$ are said to be right-helical. The helicity is a constant of motion for free fermions regardless of their mass, but is not Lorentz-invariant for fermions with $m > 0$. The other handedness concept is *chirality*, which relates to the transformation properties of the fermion fields under the Lorentz group. As such, chirality is a Lorentz-invariant quantity, though it is not conserved for particles with $m > 0$. In the $m = 0$ limit (or equivalently in the $E \gg m$ limit), helicity and chirality coincide.

This might have been only an interesting theoretical exercise if it were not for the weak interaction. In 1956, Lee and Yang observed that no one had checked whether the weak interaction violated parity—i.e., symmetry under spatial inversions $\boldsymbol{r} \to -\boldsymbol{r}$ [217]. The first group to report results was led by C. S. Wu, who studied the beta decay of spin-polarized ^{60}Co in a magnetic field $\boldsymbol{B}$. They observed a significant asymmetry in the directions of the outgoing electrons, with a clear preference for electrons to be emitted parallel to the direction of nuclear spin $\boldsymbol{J}$ [218]. Since parity inversion only changed the electron momenta (the pseudovectors $\boldsymbol{B}$ and $\boldsymbol{J}$ were unchanged under parity), this experiment was clear evidence of parity violation in the weak interaction. In time, other groups would find evidence of parity violation in the decays of muons, K^0 mesons, and other particles, and that the parity violation was maximal in the left-handed direction (see, e.g., Ref. [5] for a historical review). This led to the consensus $V - A$ (vector minus axial vector) model of the weak interaction, where each vertex had a factor of the form $\gamma^\mu(1-\gamma^5)$ to account for the left-handed preference of weak decays. Finally, in 1958, Goldhaber *et al* reported the results of an experiment they conducted to determine the helicity of the neutrino [219], based on the electron-capture decay of the metastable isotope $^{152\mathrm{m}}$Eu. By measuring the helicity of the de-excitation γ-ray from the daughter nucleus ^{152}Sm*, the group determined that the helicity of ν_e was negative—i.e., neutrinos are left-helical (and equivalently, antineutrinos are right-helical). It should be noted that in all of these experiments, the outgoing neutrinos were ultrarelativistic, and thus the distinction between helicity and chirality was blurred. This issue will be addressed in the section on neutrino mass models.

2.1.2 There are three light neutrinos, one per charged lepton

How many kinds of neutrinos are there? Over the course of the twentieth century, physicists identified the three charged leptons: the electron e^-, the muon μ^-, and the tau τ^- (and their antiparticles). Following the discovery of what we now know as the electron antineutrino $\bar{\nu}_e$ by Cowan and Reines via the reaction $p + \bar{\nu}_e \to n + e^+$ [220], physicists began to wonder whether there were ν_μ and ν_τ neutrinos as well. The non-observation of processes such as $\mu^\pm \to e^\pm + \gamma$ suggested that there might be conserved lepton numbers L_e, L_μ, and L_τ, and that each charged lepton had an associated neutrino. The first direct evidence of muon-type neutrinos was provided by a group at Brookhaven in 1962 [221], who demonstrated that the neutrinos produced by charged-pion decays (e.g., $\pi^- \to \mu^- + \bar{\nu}_\mu$ + c.c.[1]) could themselves produce muons in detectors (e.g., $p + \bar{\nu}_\mu \to n + \mu^+$ + c.c.). The direct detection of the tau neutrino by the DONuT collaboration [222] came nearly three decades after the discovery of the tau lepton, but by then, physicists had good reason to believe not only that the ν_τ must exist, but that there were only three light neutrinos.

The strongest constraints on the number of neutrinos N_ν come not from direct measurements of charged leptons, but from the measured decay width of the Z boson. Recall that the Z couples to all fundamental fermions in the SM, and decays to any fermion with mass $<\frac{1}{2}m_Z \approx 45.5\,\text{GeV}$ (i.e., to all fermions except the $\sim$175-GeV top quark). In the 1980s–1990s, a concerted effort was launched to study Z-boson production in e^+e^- collisions at the Stanford Linear Collider (SLC) in the United States and at the Large Electron-Positron Collider (LEP) at CERN (see, e.g., Ref. [213] for an extensive review). The number of neutrinos N_ν can be determined from

$$N_\nu = \frac{\Gamma_{\text{inv}}}{\Gamma_\ell}\left(\frac{\Gamma_\ell}{\Gamma_\nu}\right)_{\text{SM}}, \tag{2.2}$$

where $\Gamma_{\text{inv}} = \Gamma_Z - \Gamma_{\text{had}} - \sum_\ell \Gamma_\ell$ is the invisible decay width of the Z boson (i.e., the remainder after accounting for the measured hadronic and leptonic widths Γ_{had} and $\sum_\ell \Gamma_\ell$) and $(\Gamma_\ell/\Gamma_\nu)_{\text{SM}} \approx 2$ is the SM prediction for the ratio of charged-lepton and neutrino decay widths. A recent re-evaluation of LEP data gives $N_\nu = 2.9963(74)$ [223], consistent with there being three light SM neutrinos. If there are more neutrinos, they are heavier than $\sim$45 GeV and/or have suppressed Z-boson couplings. This is also consistent with the Planck CMB limits on the effective number of relativistic fermions $N_{\text{eff}} = 2.92^{+0.36}_{-0.37}$ (95% CI, Ref. [15]).

---

[1] Charge-conjugate, i.e., $\pi^+ \to \mu^+ + \nu_\mu$.

2.2 Neutrino masses

It is often said that the observation of neutrino masses constitutes some of the best evidence for physics beyond the SM. In this section, I describe the theoretical and experimental consequences of massive neutrinos.

2.2.1 Neutrino mass models

If the SM is assumed to contain only the known gauge/Higgs bosons, quarks, charged leptons, and left-handed neutrinos, and to be gauge-invariant, renormalizable, and baryon- and lepton-number conserving at the Lagrangian[2] level, there is an unambiguous prediction for the neutrino masses: they must be zero. To see why, we will attempt to construct neutrino mass terms and see where things go awry.

Dirac mass

We begin by considering the Dirac Lagrangian for a massive fermion field ψ,

$$\mathcal{L}_{\text{Dirac}} = \bar{\psi}(i\gamma^\mu \partial_\mu - m)\psi. \tag{2.3}$$

We can decompose the four-component Dirac spinor into $\psi = \psi_{\text{L}} + \psi_{\text{R}}$, where $\psi_{\text{L}} = \frac{1}{2}(1-\gamma^5)\psi$ and $\psi_{\text{R}} = \frac{1}{2}(1+\gamma^5)\psi$ are the eigenstates of the chiral projection operators. Thus, the Dirac Lagrangian reads (see, e.g., Ref. [216])

$$\begin{aligned}\mathcal{L}_{\text{Dirac}} &= (\overline{\psi_{\text{L}} + \psi_{\text{R}}})(i\gamma^\mu \partial_\mu - m)(\psi_{\text{L}} + \psi_R) \\ &= \overline{\psi_{\text{L}}} i\gamma^\mu \partial_\mu \psi_{\text{L}} + \overline{\psi_{\text{R}}} i\gamma^\mu \partial_\mu \psi_{\text{R}} - m_{\text{D}}(\overline{\psi_{\text{L}}}\psi_{\text{R}} + \overline{\psi_{\text{R}}}\psi_{\text{L}}).\end{aligned} \tag{2.4}$$

This is the Lagrangian for two free Dirac fields ψ_{L} and ψ_{R} with opposite chirality, coupled by the Dirac mass term m_{D}. Note that in the SM, the bare Dirac Lagrangian

---

[2] In the SM, baryon number B and lepton number L are conserved at the classical (Lagrangian) level but not at the quantum (loop) level, where they are anomalous; however, their difference $B - L$ is accidentally conserved (see, e.g., Refs. [224, 225]).

in Eq. 2.3 violates $SU(2)_\mathrm{L} \times U(1)_\mathrm{Y}$ gauge invariance (see, e.g., Refs. [214, 226]). This is compensated by the Higgs mechanism: after electroweak symmetry breaking, the Dirac mass term becomes

$$\mathcal{L}_{\text{Dirac mass}} \to -\frac{Fv}{\sqrt{2}}(\overline{\psi_\mathrm{L}}\psi_\mathrm{R} + \overline{\psi_\mathrm{R}}\psi_\mathrm{L}) - \frac{F}{\sqrt{2}}H(\overline{\psi_\mathrm{L}}\psi_\mathrm{R} + \overline{\psi_\mathrm{R}}\psi_\mathrm{L}). \tag{2.5}$$

Here $v \approx 246\,\text{GeV}$ is the Higgs vacuum expectation value (VEV) [16] and F is the Yukawa coupling of ψ to the Higgs field H. Thus, we can relate the Dirac mass to the Higgs-fermion Yukawa coupling via

$$m_\mathrm{D} = \frac{Fv}{\sqrt{2}}. \tag{2.6}$$

Unfortunately, the Lagrangian in Eq. 2.5 cannot describe neutrino masses using the field content of the Standard Model, as only left-handed neutrinos are known to exist. On a more aesthetic level, the fact that the neutrino Yukawa couplings would need to be extremely small ($F \sim 10^{-12}$), more than six orders of magnitude smaller[3] than the electron's, appears rather unusual (see, e.g., Refs. [227, 228] for discussions).

Majorana mass

As it turns out, there *is* a way to construct neutrino mass terms using only the field content of the Standard Model. Consider the field (see, e.g., Refs. [216, 229])

$$\psi^c \equiv i\gamma^2\gamma^0\psi^\star, \tag{2.7}$$

which is the CP conjugate of ψ. Ignoring the possibility of CP violation for the moment, we immediately realize that $(\psi_\mathrm{L})^c = (\psi^c)_\mathrm{R}$ is a *right-handed* field. Thus,

---

[3] A pattern arises in the quark sector, where the masses range from ∼few MeV ($F \sim 10^{-5}$) for the up and down quarks to ∼175 GeV ($F \sim 1$) for the top quark.

we can write the four-component *Majorana spinor* ψ as

$$\psi = \psi_{\mathrm{L}} + (\psi_{\mathrm{L}})^c = \psi_{\mathrm{L}} + (\psi^c)_{\mathrm{R}}\,, \tag{2.8}$$

using only a single field (in this case, ψ_{L}). In the context of neutrinos, since only the left-handed neutrino fields are known to exist, Eq. 2.8 seems to be exactly what we need. Majorana mass terms arise frequently in effective field theories, e.g., from the dimension-five Weinberg operator (see, e.g., [230, 231])

$$\mathcal{L}_5 \supset \frac{C_5^{\ell\ell'}}{\Lambda}[\overline{L^c_{\ell\mathrm{L}}}\widetilde{H}^\star][\widetilde{H}^\dagger L_{\ell'\mathrm{L}}], \tag{2.9}$$

where $\widetilde{H} \equiv i\sigma_2 H^\star$ is the Higgs conjugate, σ_2 is the second Pauli matrix, C_5 is the Wilson coefficient, $L_{\ell\mathrm{L}}$ is the left-handed lepton doublet of flavor ℓ, and Λ is some (large) energy scale at which new physics appears. This operator generates a Majorana mass

$$M_{\ell\ell'} = \frac{C_5^{\ell\ell'}v^2}{\Lambda}, \tag{2.10}$$

where v is the Higgs VEV. In fact, the operator in Eq. 2.9 can only be generated at tree-level in three ways: (i) by the introduction of three massive singlet (right-handed) fermions, (ii) by the introduction of a new massive scalar triplet, or (iii) by the introduction of a massive fermion triplet. These are called the *type-I*, *type-II*, and *type-III seesaw mechanisms*, respectively (see, e.g., Ref. [232] for a review). In each of these cases, the SM neutrino masses m_ν scale inversely with the mass M of the new (heavy) degrees of freedom — pushing M upward pushes m_ν downward.

Despite its theoretical advantages (no need for separate right-handed fields!), the Majorana mass hypothesis runs into problems with conserved charges. Note that Eq. 2.7 requires $\psi = \psi^c$ — in other words, ψ must be its own antiparticle. One major consequence is that unlike the Dirac case, in which the mass term $m_{\mathrm{D}}\bar{\psi}\psi$ is unaffected by $U(1)$ transformations of the form $\psi \to \psi e^{i\alpha}$, the Majorana mass term $M_{\mathrm{LL}}\overline{\psi_{\mathrm{L}}}(\psi_{\mathrm{L}})^c$ picks up a factor $e^{-2i\alpha}$. In other words, the Majorana mass term does not conserve $U(1)$ quantum numbers (e.g., electric charge, hypercharge, lepton number, etc). This explains why all of the other fundamental fermions cannot have Majorana mass terms—the quarks and charged leptons are, well, electrically-charged! If neutrinos have a nonzero Majorana mass, it would immediately imply violation of lepton number — not just the violation of *individual* lepton flavor numbers, but violation of *total* lepton number L.

A smoking gun for neutrinos having Majorana mass would be the detection[4] of neutrinoless double beta decay ($0\nu 2\beta$). In this process, two neutrons in an atomic nucleus would simultaneously undergo beta decay, producing two protons, two electrons, and no neutrinos at all, thus changing $B - L$ by two units (B is baryon number and L is lepton number). For calorimetric experiments in which both electrons are fully absorbed, the signature of $0\nu 2\beta$ decay would be a peak in the $\mathrm{d}N/\mathrm{d}E_{\beta\beta}$ energy spectrum at $Q_{\beta\beta}$, corresponding to the mass difference between the parent and daughter nuclei. (This is in contrast to the continuous spectrum from "normal" double-beta decay, in which two neutrinos are also released.) A substantial experimental effort has been mounted to search for $0\nu 2\beta$ decays in a variety of isotopes (for a recent review, see Ref. [234]), with lower limits on the half-life exceeding 10^{20} yr in many cases. It is hoped that the coming generation of experiments will be sensitive to $0\nu 2\beta$ decays if the SM neutrinos are in the inverted mass ordering, provided that calculations of the nuclear matrix elements are sufficiently accurate [235]. Since the rate of $0\nu 2\beta$ decay may be unobservably small in the normal-ordering scenario, one might ask whether neutrinos being purely Dirac would have unique and observable consequences. One possibility is neutrinoless *quadruple*[5] beta decay ($0\nu 4\beta$), which can be thought of as two simultaneous $0\nu 2\beta$ processes. Recent work suggests that $0\nu 4\beta$ decays may occur even if neutrinos are purely Dirac [236]. As a possible corollary to the Schechter-Valle theorem, observing $0\nu 4\beta$ decay but not $0\nu 2\beta$ decay may indicate that at least one neutrino is purely Dirac [237].

---

[4] This is a consequence of the Schechter-Valle "black box" theorem [233].

[5] There may also be $0\nu N\beta$ decays for integer $N > 2$, but these will likely be unobservably suppressed due to kinematic and phase-space reasons. The same reasoning applies to $4\nu 4\beta$ decay.

2.2.2 Neutrino masses and oscillations

With the preceding discussion about how neutrinos must be massless in the SM, it comes as somewhat of a surprise that experimental evidence *requires* neutrinos to have mass. This section will briefly review some of the important evidence for neutrino masses—with the full benefit of hindsight (the situation at the time was nowhere near as clear-cut). For historical reviews, see, e.g., Refs. [238, 239].

A brief history

The first piece of evidence for neutrino masses came from the Sun. The fusion reactions in the core of the Sun produce copious amounts of ν_e's, with energies ranging from $\sim$few hundred keV to more than 10 MeV. Since neutrinos are weakly-interacting, they stream freely from the Sun's core, providing a crucial window into the conditions within the Sun. Beginning with the Homestake experiment in the 1960s [240], many solar-neutrino detectors employing different detection techniques consistently recorded $\sim$50% fewer electron-type neutrinos than expected. This discrepancy came to be known as the *solar neutrino problem* (see, e.g., Ref. [241] for a review).

During the latter half of the twentieth century, physicists were also searching for neutrinos produced in cosmic-ray air showers (see, e.g., Ref. [242, 243] for reviews). When cosmic rays (mainly protons) strike nuclei in the upper atmosphere, they produce a hadronic shower consisting of charged pions and other particles. The charged pions decay mainly via $\pi^+ \to \mu^+ + \nu_\mu$ + c.c., with some of the muons also decaying on their way to the surface (mainly via $\mu^+ \to e^+ + \nu_e + \bar{\nu}_\mu$+c.c.). Thus, the flux of these *atmospheric neutrinos* was expected to contain approximately twice as many muon-type neutrinos as electron-type neutrinos. By the 1990s, the IMB [244, 245], Kamiokande [246], Soudan-2 [247], and Super-Kamiokande [248, 249] collaborations were able to measure the rate of muon- and electron-type neutrinos as a function of zenith angle θ. They found that the rate of electron-like events $\mathrm{d}N/\,\mathrm{d}(\cos\theta)$ agreed with expectations for all θ, but there was a substan-

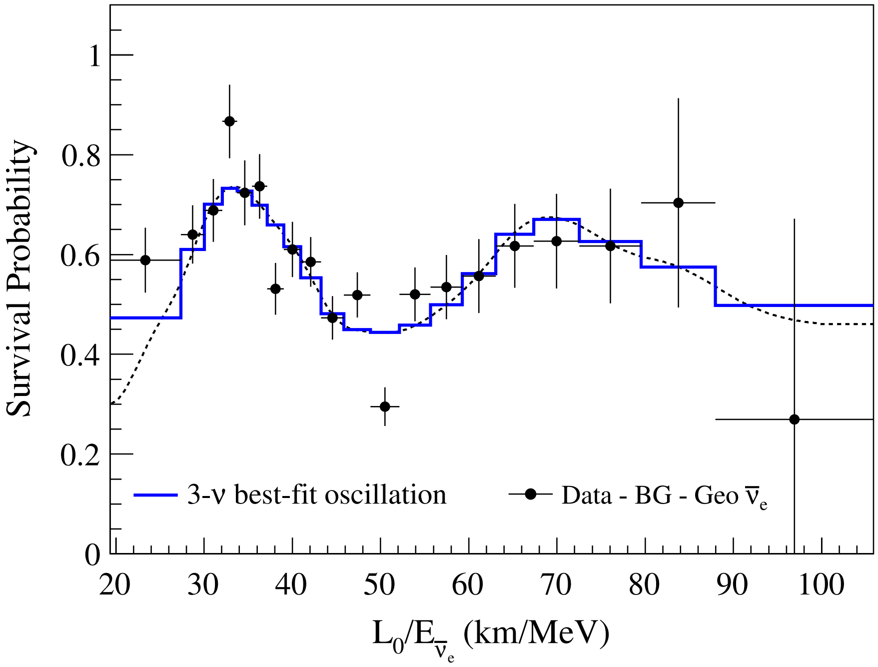

Figure 2-1: Medium-baseline reactor $\bar{\nu}_e$ oscillations observed by the KamLAND experiment after subtracting detector backgrounds ("BG") and geoneutrinos ("Geo $\bar{\nu}_e$"). On the x-axis, $L_0 = 180\,\text{km}$ is the flux-weighted average distance between the detector and reactors. The y-axis shows the survival probability $P(\bar{\nu}_e \to \bar{\nu}_e)$. The solid blue line and black dashed line indicate the binned and smoothed best-fit three-neutrino oscillation model, respectively. Figure reproduced from Ref. [250] with permission. © 2013 by the American Physical Society.

tial deficit of upward-going (i.e., $\cos\theta \lesssim 0$) muon-type events. Essentially, muon neutrinos produced above the detector had the expected rate, but something was depleting the muon neutrinos produced on the other side of the Earth as they propagated through the planet.

To resolve these issues, it was necessary to construct detectors sensitive to all three neutrino flavor eigenstates, and to study neutrinos produced in other contexts. The first such detector was the Sudbury Neutrino Observatory (SNO, Ref. [251]), consisting of ∼1000 tons of heavy water (D_2O) surrounded by photomultiplier tubes. The heavy water allowed SNO to be sensitive to the deuteron breakup reaction $\nu_\ell + \mathrm{D} \to \nu'_\ell + p + n$, with a threshold $E_\nu \gtrsim 2.2\,\mathrm{MeV}$ regardless of neutrino flavor! Since the detector was also sensitive to the processes $\nu_e + n \to e^- + p$ and $\nu_e + e^- \to e^- + \nu_e$, this allowed SNO to perform a simultaneous measurement of the ν_e and $\nu_e + \nu_\mu + \nu_\tau$ solar neutrino fluxes. The ν_e flux was consistent with the previous experiments in showing a deficit compared to the standard solar model, but the all-flavor neutrino flux inferred from deuteron breakup was completely consistent with the standard solar model. These results conclusively demonstrated that the ν_e produced in the Sun's core were oscillating into ν_μ and ν_τ, the latter two of which were undetectable via the reaction $\nu_\ell + n \to p + \ell^-$ owing to the low energies of solar neutrinos. In the coming years, observations of energy-dependent disappearance of reactor $\bar{\nu}_e$ at experiments such as KamLAND [250, 252] supported the oscillation interpretation of solar-neutrino data (see Fig. 2-1). Finally, the early 2000s saw the development of accelerator neutrino experiments capable of generating intense beams of ν_μ or $\bar{\nu}_\mu$ with energies exceeding several hundred MeV. Experiments such as K2K [253] and MINOS [254] found evidence of muon-neutrino disappearance and electron-neutrino appearance, supporting the oscillation interpretation of atmospheric neutrinos.

The modern picture

All of these observations are neatly accommodated using the quantum-mechanical formalism for neutral-particle oscillations (see, e.g., Refs. [215, 255, 256]). In the SM, the three neutrino flavor eigenstates $|\nu_\ell\rangle$ are linear combinations of the three mass eigenstates $|\nu_i\rangle$. These two bases are related by the Pontecorvo-Maki-Nagawa-Sakata (PMNS) matrix U, such that[6]

$$|\nu_\alpha\rangle = \sum_i U^\star_{\alpha i} \, |\nu_i\rangle \, . \tag{2.11}$$

---

[6] For antineutrino oscillations, replace $\mathsf{U}^\star$ with U.

In the SM, the PMNS matrix is a 3×3 unitary matrix containing four independent parameters: three rotation angles $\theta_{12}, \theta_{23}, \theta_{13} \in [0, 90^\circ]$ and a CP-violating phase $\delta_{CP} \in [0, 360^\circ)$. The PMNS matrix is usually written as[7]

$$\mathsf{U} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}, \qquad (2.12)$$

where $c_{ik} \equiv \cos\theta_{ik}$ and $s_{ik} \equiv \sin\theta_{ik}$. Consider a neutrino originally produced in flavor eigenstate $|\nu_\alpha\rangle$ with energy much greater than the neutrino mass, and detected a distance L away. The probability of detecting the neutrino in the flavor eigenstate $|\nu_\beta\rangle$ is given by (neglecting matter effects along the neutrino path) [257]

$$\begin{aligned} P(\nu_\alpha \to \nu_\beta) = \delta_{\alpha\beta} &- 4\sum_{i,k} \mathfrak{Re}[U_{\alpha i}U^\star_{\beta i}U_{\alpha k}U^\star_{\beta k}] \sin^2\left(\frac{\Delta m^2_{ik}L}{4E}\right) \\ &+ 2\sum_{i,k} \mathfrak{Im}[U_{\alpha i}U^\star_{\beta i}U_{\alpha k}U^\star_{\beta k}] \sin\left(\frac{\Delta m^2_{ik}L}{2E}\right) \end{aligned} \qquad (2.13)$$

where $\delta_{\alpha\beta}$ is the Kronecker symbol and $\Delta m^2_{ik} \equiv m_i^2 - m_k^2$ is the mass-squared difference between the two eigenstates. (In the case of antineutrino oscillations, the $\mathfrak{Im}$ term picks up a factor of -1.) The best-fit values of the PMNS neutrino-oscillation parameters are shown in Table 2.1.

(In dense matter, the situation is much more complicated, owing to coherent scattering against electrons and nucleons, generating effective in-media mixing angles θ_m [257]. In media with varying spatially-varying density (e.g., the Sun), this is known as the Mikheyev-Smirnov-Wolfenstein (MSW) effect [258, 259].)

---

[7] If neutrinos have Majorana mass terms, and particularly if there are heavier states (as in the type-I seesaw, for example), more care must be taken, particularly with the Majorana phases (see, e.g., Ref. [16]).

Parameter	Best-fit NO (IO) [$\pm 1\sigma$]	Mainly constrained by ...
Δm^2_{21} [10^{-5} eV^2]	$7.42^{+0.21}_{-0.20}$ ($7.42^{+0.21}_{-0.20}$)	Reactor LBL
$\lvert\Delta m^2_{3k}\rvert$ [10^{-3} eV^2]	$2.517^{+0.026}_{-0.028}$ ($-2.498^{+0.028}_{-0.028}$)	Reactor MBL, accel. LBL
θ_{12} [$^\circ$]	$33.44^{+0.77}_{-0.74}$ ($33.45^{+0.78}_{-0.75}$)	Solar
θ_{23} [$^\circ$]	$49.2^{+0.9}_{-1.2}$ ($49.3^{+0.9}_{-1.1}$)	Accel. LBL
θ_{13} [$^\circ$]	$8.57^{+0.12}_{-0.12}$ ($8.60^{+0.12}_{-0.12}$)	Reactor MBL
δ_{CP} [$^\circ$]	197^{+27}_{-24} (282^{+26}_{-30})	Accel. LBL

Table 2.1: PMNS matrix parameters and global-fit results in the normal (NO) and inverted (IO) mass ordering from NuFIT, including Super Kamiokande atmospheric data [260]. The third column indicates the experimental setup which is most constraining (see Table 14.6 of Ref. [16] for details). MBL is medium baseline ($\gtrsim 1$ km) and LBL is long baseline ($\gtrsim 100$ km). "Accel." means "accelerator."

Remaining questions

Despite the excellent progress in precisely measuring the values of the various PMNS parameters in a variety of experimental conditions, there are several outstanding issues in the three-neutrino[8] picture.

The first question concerns the absolute scale of the neutrino masses. Since oscillations are only sensitive to Δm^2, other methods are needed to determine the actual masses $\{m_1, m_2, m_3\}$. Most recent experiments use the beta decay of tritium (T or ^{3}H), owing to its low decay energy ($Q_\beta \approx 18.6$ keV). Since the effective mass of the neutrino emitted in beta decay $m^2_\beta = \sum_i |U_{ei}|^2 m^2_i$ induces a very small correction to the maximum energy of the outgoing electron, neutrino-mass experiments require extremely sensitive detectors. The first approach has been to use a combination of magnetic spectrometers and electrostatic filters to construct the electron spectrum very near to Q_β. Recent results from the KATRIN experiment give an

---

[8] The possibility of additional neutrino states—of either helicity—is discussed in the next section.

upper limit $m_\beta^2 < 0.8\,\text{eV}^2$ (90% CL, Ref. [261]). A newer technique involves measuring the electron energy via cyclotron resonance emission spectroscopy, with experiments such as Project 8 anticipating $m_\beta{\sim}40$-meV sensitivity in the near future [262, 263].

The second question concerns the ordering of the neutrino mass states. From the best-fit value of $\Delta m_{21}^2 \approx 7.4 \times 10^{-5}\,\text{eV}^2$, it is known that m_1 and m_2 are extremely close and that $m_2 > m_1$. On the other hand, oscillation measurements indicate $|\Delta m_{3i}^2| \approx 2.5 \times 10^{-3}\,\text{eV}^2$—i.e., that m_3 is quite separated from m_1 and m_2, but the sign is not known. In other words, it is not clear from current data whether $m_2 > m_1 \gg m_3$ (the so-called inverted ordering, IO) or $m_3 \gg m_2 > m_1$ (the so-called normal ordering, NO). The mass ordering has dramatic consequences for the detectability of $0\nu\beta\beta$ decay, with the inverted ordering strongly enhancing the expected decay rates. It is hoped that future long-baseline oscillation experiments (e.g., JUNO [264] and DUNE [265]) will be able to resolve the neutrino mass ordering within the coming decade(s).

The final question concerns CP violation. The PMNS matrix admits a CP-violating phase δ_{CP}, which (so long as $\sin\delta_{CP} \neq 0$) would manifest as differences in the oscillation probabilities of $\nu_\alpha \to \nu_\beta$ and $\bar{\nu}_\alpha \to \bar{\nu}_\beta$. At present, δ_{CP} is consistent with 180° at the 1σ level in the normal ordering, and the 3σ level in the inverted ordering. The magnitude of CP violation in the lepton sector is quantified by the Jarlskog invariant

$$\mathcal{J}_{CP} \equiv \mathcal{J}_{CP}^{\text{max}} \sin\delta_{CP} = c_{12}s_{12}c_{23}s_{23}c_{13}^2 s_{13} \sin\delta_{CP}, \tag{2.14}$$

where current oscillation data indicate $\mathcal{J}_{CP}^{\text{max}} = 0.0332 \pm 0.008$ (68% CI, Ref. [260]). This is nearly three orders of magnitude greater than $\mathcal{J}_{CP}^{\text{quarks}}$, indicating that CP violations in the neutrino sector may contribute to the observed matter-antimatter asymmetry of the Universe (discussed in the next section).

2.3 The Neutrino Minimal Standard Model

As discussed previously, experimental evidence is consistent with neutrinos having Dirac and/or Majorana mass terms. The minimal Lagrangian accomplishing this is (e.g., Refs. [266–270])

$$\mathcal{L} = \mathcal{L}_{\rm SM} + i\overline{N_I}\gamma^\mu\partial_\mu N_I - F_{\alpha I}\overline{L_{\alpha \rm L}}\widetilde{H}N_I - \frac{1}{2}M_{IK}\overline{N_I^c}N_K + \text{h.c.}, \tag{2.15}$$

known as the *type-I seesaw*. To the three lepton doublets of the SM, we add $\mathcal{N}$ right-handed singlet neutrinos N with Majorana masses M and Yukawa couplings F. (The conjugate Higgs field $\widetilde{H} = i\sigma_2 H^\star$.) We can write the mass term more compactly as (e.g., Refs. [271, 272])

$$\frac{1}{2}\begin{pmatrix}\overline{\nu_{\rm L}} & \overline{\nu_{\rm R}^c}\end{pmatrix}\mathfrak{M}\begin{pmatrix}\nu_{\rm L}^c \\ \nu_{\rm R}\end{pmatrix} + \text{h.c.} = \frac{1}{2}\begin{pmatrix}\overline{\nu_{\rm L}} & \overline{\nu_{\rm R}^c}\end{pmatrix}\begin{pmatrix}0 & \mathsf{m}_{\rm D} \\ \mathsf{m}_{\rm D}^{\mathsf{T}} & \mathsf{M}\end{pmatrix}\begin{pmatrix}\nu_{\rm L}^c \\ \nu_{\rm R}\end{pmatrix} + \text{h.c.}, \tag{2.16}$$

where $\mathfrak{M}$ is the $(3+\mathcal{N}) \times (3+\mathcal{N})$ seesaw mass matrix, $\mathsf{m}_{\rm D} = v\mathsf{F}$ is the Dirac mass matrix, and M is the Majorana mass matrix. For the rest of this discussion, we will assume that the eigenvalues of M are much greater than those of $\mathsf{m}_{\rm D}$, and also that $\mathcal{N} = 3$. Defining $\theta \equiv -\mathsf{m}_{\rm D}\mathsf{M}^{-1}$, we can write the light (3×3) and heavy $(\mathcal{N} \times \mathcal{N})$ neutrino mass matrices as (e.g., Refs. [271, 272])

$$\begin{cases} \mathsf{m}_{\rm light} = -\theta\mathsf{M}\theta^{\mathsf{T}} \\ \mathsf{M}_{\rm heavy} = \mathsf{M} + \dfrac{1}{2}(\theta^\dagger\theta\mathsf{M} + \mathsf{M}^{\mathsf{T}}\theta^{\mathsf{T}}\theta^\star) \end{cases} \tag{2.17}$$

To uncover the physical neutrino states and their masses, we need only diagonalize the mass matrix $\mathfrak{M}$. This is accomplished with the matrix (e.g., Refs. [271, 272])

$$\mathfrak{U} = \begin{pmatrix} 1 - \frac{1}{2}\theta\theta^\dagger & \theta \\ -\theta^\dagger & 1 - \frac{1}{2}\theta^\dagger\theta \end{pmatrix}\begin{pmatrix}\mathsf{U}_\nu & \\ & \mathsf{U}_N^\star\end{pmatrix} + \mathcal{O}(\theta^3), \tag{2.18}$$

where U_ν is the PMNS matrix for the SM neutrinos and U_N is its analogue for the new singlets. When all is said and done, the flavor eigenstates are (to second-order in θ)

$$\begin{cases} \nu_{\text{light}} \simeq \mathsf{U}_\nu^\dagger \left(1 - \dfrac{1}{2}\theta\theta^\dagger\right) \nu_{\text{L}} - \mathsf{U}_\nu^\dagger \theta N^{\text{c}} \\ \text{N}_{\text{heavy}} \simeq \mathsf{U}_N^\dagger \left(1 - \dfrac{1}{2}\theta^\mathsf{T}\theta^\star\right) N + \mathsf{U}_{\text{N}}^\dagger \theta^\mathsf{T} \nu_{\text{L}}^c. \end{cases} \tag{2.19}$$

Crucially, this means that the light left-handed neutrinos ν_{L} (which we identify as the SM neutrinos) have a small admixture of the right-handed sterile states, and vice-versa. Thus, oscillation-induced mixing opens up a portal between the active and sterile neutrinos, with wide-ranging consequences for particle physics and cosmology.

One particularly interesting realization of the type-I seesaw mechanism is the neutrino minimal Standard Model (νMSM, e.g., Refs. [273, 274]). In this model, the only additions to the SM are three right-handed neutrinos N_1, N_2, and N_3. The masses of N_2 and N_3 are generally chosen to be very large (hundreds of MeV to hundreds of GeV) such that the seesaw mechanism can account for the smallness of the SM neutrino masses. With a mass near the keV scale, the lightest right-handed neutrino N_1 (hereafter symbolized χ) constitutes an ideal DM candidate. Furthermore, the νMSM provides a mechanism for generating the observed matter-antimatter asymmetry in the Universe.

2.3.1 Matter-antimatter asymmetry

As described in Chapter 1, BBN and CMB data indicate the cosmic baryon-to-photon ratio $\eta_b \approx 6 \times 10^{-11}$ (i.e., roughly one part in 10^{11} more matter than antimatter). Generating a net baryon asymmetry (baryogenesis) requires three conditions [275]: (i) C and CP violation, (ii) baryon-number (B) violation, and (iii) non-equilibrium processes.

It turns out that the SM already contains the ingredients for baryogenesis in the electroweak sector. Although baryon number B and lepton number L (as well as $B - L$) are conserved at all orders in perturbation theory, they are violated by nonperturbative effects (see, e.g., Refs. [224, 225, 230] for early work along these lines). The vacuum structure of the SM electroweak sector is quite complex, containing an infinite number of degenerate vacua separated by potential barriers. At low temperatures[9], transitions between adjacent vacua require tunnelling through the barrier (the so-called *instanton* configuration [276]), with the rate scaling as $\exp(-1/g_w^2) \sim \exp(-4\pi \times 29) \sim 10^{-160}$ (see, e.g., Ref. [277]). Since the effective height of the potential barrier is ~10 TeV, high temperatures (comparable to the electroweak scale ~100 GeV) may substantially increase the rate. Configurations with sufficient energy to surmount the potential barrier are known as *sphalerons*, and have become the subject of intense interest (see, e.g., Refs. [278–284]). Unfortunately, this process on its own cannot account for the observed matter-antimatter asymmetry of the Universe. First, the discovery of the SM Higgs boson at $m_H \approx 125\,\text{GeV}$ means that the SM electroweak phase transition is (in the absence of new physics) a continuous transition, rather than first-order[10] (see, e.g., Refs.[285–287]). Second, the magnitude of CP violation in the quark sector appears to be insufficient for reproducing the observed baryon asymmetry (see, e.g., Ref. [288]).

Fortunately, CP-violating neutrino oscillations provide a way out of the aforementioned difficulties [289]. (For quantitative discussions, see, e.g., Refs. [272, 274, 290–294]. This brief overview is based on Ref. [292, 293].) If the heavy neutrinos N_2 and N_3 are nearly degenerate[11] in mass (i.e., $m_{N_2} \approx m_{N_3}$), then resonant CP-violating oscillations may lead to a significant imbalance between neutrinos and antineutrinos at temperatures $T \gtrsim 100\,\text{GeV}$. This lepton asymmetry is converted into the observed baryon asymmetry by the aforementioned sphaleron process. At the end of the electroweak epoch (i.e., when the temperature $T \lesssim 100\,\text{GeV}$),

---

[9] "Low" will be defined shortly.

[10] In first-order EW phase transitions, C and CP asymmetries would be processed into baryon number excesses via sphaleron transitions at the boundaries between the broken (nonzero Higgs VEV) and symmetric (zero Higgs VEV) EW vacua. The abrupt transition at the boundary between the two vacua would prevent the baryon asymmetry from being destroyed by the reverse reactions.

[11] Other works (e.g., Ref. [295]) suggest that this mass degeneracy is not necessary, and that a similar effect can be accomplished via proper modeling of neutrino transport in the primordial plasma.

the sphalerons freeze out. Subsequently, the heavy neutrinos N_2 and N_3 drop out of equilibrium (maintained by reactions such as $\ell^+\ell^- \leftrightarrow \nu N$) at $T \sim$ few GeV and begin to decay. The lepton asymmetry may be further increased by the CP-violating decays of N_2 and N_3 [290], with direct implications for the production of sterile-neutrino DM, discussed next.

2.3.2 Sterile-neutrino DM production

Since sterile-neutrino DM can only interact with the SM particles via oscillation-induced mixing with the lepton neutrinos, it is necessary to consider the effects of the primordial plasma on the active-sterile mixing $\sin^2(2\theta)$. Concretely, the effective in-medium mixing $\sin^2(2\theta_m)$ is given in the two-neutrino picture[12] by (see, e.g., Refs. [272, 297] and references therein)

$$\sin^2(2\theta_m) = \frac{\Delta^2(p)\sin^2(2\theta)}{\Delta^2(p)\sin^2(2\theta) + [\Delta(p)\cos(2\theta) - V_L - V_T]^2}, \tag{2.20}$$

where $\Delta(p) \equiv \Delta m^2/(2p) \approx m_\chi^2/(2p)$. The potentials V_T and V_L are given by

$$\begin{cases} V_T \simeq -\dfrac{8\sqrt{2}}{3} G_\mathrm{F} E_\nu \left[\dfrac{\rho_\nu}{m_Z^2} + \dfrac{\rho_\ell}{M_W^2}\right] \sim 2G_\mathrm{F} T^4 \rho \left[\dfrac{1}{m_Z^2} + \dfrac{r_\ell}{M_W^2}\right] \sim G_\mathrm{eff}^2 T^4 \rho \\ V_L \simeq 0.69 G_\mathrm{F} T^3 L_\nu \end{cases} \tag{2.21}$$

where r_ℓ is a correction factor accounting for the masses of the charged leptons, n_i is the number density of species i, ρ is the energy density, and $L_\nu \equiv (n_\nu - n_{\bar{\nu}})/n_\gamma$ is the lepton asymmetry[13] of the plasma. Significantly, the sterile-neutrino production rate does not increase arbitrarily with temperature due to the thermal potential V_T in the denominator of Eq. 2.20, meaning there is a finite window in the early Universe in which sterile-neutrino DM can be efficiently produced.

---

[12] For full details of the three-neutrino picture (and the underlying finite-temperature QFT) see, e.g., Refs. [272, 296] and references therein.

[13] Note that some authors define the lepton asymmetry in terms of the entropy density s rather than the photon number density n_γ. These are related by $s \approx 1.8 g_{\star s} n_\gamma$, where $g_{\star s}$ is a temperature-dependent factor accounting for the effective relativistic degrees of freedom (see, e.g., Ref. [298]).

If the lepton asymmetry L_ν is not negligible, the term in square brackets in the denominator of Eq. 2.20 may vanish, allowing $\sin^2(2\theta_m) \to 1$ for a time. First proposed by Shi and Fuller [299], this leads to resonant conversion between active and sterile neutrinos, and hence amplified production of the latter. In terms of the rescaled neutrino momentum $x \equiv p/T$, the resonance condition is (e.g., Ref. [297])

$$x_{\text{res}} \approx 0.074 \left(\frac{m_\chi}{1\,\text{keV}}\right) \left(\frac{10^{-3}}{L_\nu}\right) \left(\frac{170\,\text{MeV}}{T}\right)^4, \tag{2.22}$$

where L_ν is normalized by n_γ. As the primordial Universe expands and cools, and active-sterile oscillations deplete the initial lepton asymmetry, both the temperature and lepton asymmetry decrease; thus, the resonance condition shifts to higher neutrino momenta. Crucially, this means that the Shi-Fuller mechanism is most efficient at low momenta, thereby producing a "cool" or even "cold" sterile-neutrino velocity distribution. It should be noted, however, that the resonant production of keV-scale sterile neutrinos peaks at temperatures $T{\sim}100$–$200\,\text{MeV}$, in the same range as the quark-hadron phase transition. Thus, the evolution of L_ν also depends on the neutrino opacity and other transport properties of the hadronic plasma, which is the subject of ongoing research (see, e.g., Ref. [296]).

On the other hand, if the lepton asymmetry L_ν is very small, the denominator of Eq. 2.20 is always much greater than the numerator, meaning that $\sin^2(2\theta_m) \ll 1$. This corresponds to nonresonant production, also known as the Dodelson-Widrow mechanism [300]. In this case, the sterile neutrinos are produced at higher temperatures compared to the resonant-production case, near $T \approx 133\,\text{MeV}(m_\chi/1\,\text{keV})^{1/3}$ [297].

2.3.3 Terrestrial constraints

In addition to cosmological observations, a variety of terrestrial experiments can constrain extensions to the neutrino sector, including the νMSM. This section reviews several such possibilities.

Neutrino experiments

The νMSM scenario outlined previously requires the lightest SM neutrino to have mass $m_1 \lesssim 10^{-6}\,\mathrm{eV}$. Using the Δm^2 values inferred from oscillation experiments, this implies that $m_2 \approx 0.01\,\mathrm{eV}$ and $m_3 \approx 0.05\,\mathrm{eV}$ in the NO, or $m_2 \approx m_3 \approx 0.05\,\mathrm{eV}$ in the IO [272, 291]. Additionally, due to the Majorana mass terms, the νMSM also naturally accommodates $0\nu 2\beta$ decay, with effective masses $m_{\beta\beta}{\sim}1.3$–3.4 meV in the NO and $\sim$13–50 meV in the IO [301]. Next-generation experiments hope to probe the entirety of the IO region, assuming that nuclear matrix element calculations can achieve the requisite level of accuracy [234].

Sterile neutrinos may be produced in nuclear electron-capture or beta decays. Electron-capture nuclei are particularly attractive, as the two-body final state (neutrino and daughter nucleus) particles are mono-energetic. The recoil kinetic energy E_{rec} of the daughter nucleus with mass m_A is given by simple kinematics

$$E_{\mathrm{rec}} = \frac{Q_{\mathrm{EC}}^2 - m_\nu^2}{2(Q_{\mathrm{EC}} + m_A)}, \tag{2.23}$$

where Q_{EC} is the mass difference between the initial- and final-state nuclei. The effect of a massive sterile neutrino would be to produce an additional peak in the E_{rec} spectrum at a slightly lower energy than the SM ($m_\nu \approx 0$) expectation. The number of events in the new peak would be proportional to the coupling $|U_{e4}|^2$. As an example, the BeEST experiment uses ^{7}Be nuclei implanted in a superconducting quantum sensor with few-eV energy resolution [302]. With $Q_{\mathrm{EC}} \approx 860\,\mathrm{keV}$ and $\sim$90% of decays resulting in ground-state ^{7}Li, BeEST has set some of the most competitive limits on sterile neutrinos in the mass range $\sim$100 keV–1 MeV, excluding $|U_{e4}| \gtrsim 10^{-3}$–10^{-4}. Searches for lower-mass sterile neutrinos (in the range $\sim$1–100 keV) are complicated by the difficulty of finding nuclides which (i) purely decay via electron capture, (ii) have low Q_{EC} values, and (iii) have sufficiently long lifetimes to work with. In the keV sterile-neutrino mass range, most proposals have focused on the beta decay of ^{3}H, owing to its relatively small $Q_\beta \approx 17\,\mathrm{keV}$ [303–306].

Colliders

Another constraint on the νMSM scenario comes from colliders. Though the keV-scale DM particle is likely not massive enough to leave a detectable kinematic signature at colliders, the heavier sterile neutrinos[14] might be detectable. Owing to their mixing with the active neutrinos, HNLs can appear at all the same vertices as SM neutrinos (though of course subject to kinematic limits and suppressed by mixing angles).

One avenue for HNL searches at colliders is meson decays. First, consider the charged-current decay of pseudoscalar ($J^P = 0^-$) mesons $P^\pm \in \{\pi^\pm, K^\pm, D^\pm, B^\pm\}$ of the form $P^\pm \to \ell^\pm N$. The decay rate is given by (e.g., Ref. [307])

$$\Gamma(P^\pm \to \ell^\pm N) = \frac{G_\mathrm{F}^2 f_h^2 m_h^2}{8\pi} |V_{ik}|^2 |U_{\ell 4}|^2 \left[y_N^2 + y_\ell^2 - (y_N^2 - y_\ell^2)^2 \right] \lambda^{\frac{1}{2}}(1, y_N^2, y_\ell^2), \quad (2.24)$$

where G_F is the Fermi constant, f_P is the hadronic form factor, m_P is the meson mass, V_{ik} is the relevant CKM matrix element (e.g., V_{ud} for $\pi^\pm$ or V_{us} for $K^\pm$), $y_i \equiv m_i/m_P$, and $\lambda^{\frac{1}{2}}(a,b,c) \equiv \sqrt{a^2 + b^2 + c^2 - 2(ab + bc + ac)}$. In the SM with three light neutrinos ($y_\nu \ll y_\ell$), the electron mode is helicity-suppressed by a factor $\sim 10^4$–10^5 compared to the muon mode, but the presence of an HNL may substantially increase the rate. Since the mixing parameters $U_{\ell 4}$ are of course unknown *a priori*, most experiments usually report their constraints assuming only one of $\{U_{e4}, U_{\mu 4}, U_{\tau 4}\}$ is nonzero. The PIENU and NA62 experiments have provided leading constraints using the decays $\pi^+ \to \ell^+ N$ and $K^+ \to \ell^+ N$, respectively, for masses m_N in the range of several MeV to several hundred MeV [308–311].

Depending on the masses and mixing angles, HNLs may also decay to SM particles within the detector volume. These HNLs may be produced via a number of processes, including the aforementioned meson decays, $W^\pm/Z$ bosons (either on- or off-shell), and many others. The available HNL decay channels are of course set by the mass m_N, with particularly distinctive channels including $N \to \ell_\alpha^+ \ell_\beta^- \nu_\beta$ (where α and β need not be the same flavor). Because the HNL decay rate is sup-

---

[14] Above the MeV scale, these particles are often called “heavy neutral leptons” (HNLs).

pressed by the (potentially very small) mixing parameter $U_{\ell 4}$, an HNL may be expected to travel a distance

$$c\tau_N \sim 1.3\,\mathrm{mm}\left(\sum_\ell |U_{\ell 4}|^2\right)^{-1}\left(\frac{m_N}{1\,\mathrm{GeV}}\right)^{-5}, \tag{2.25}$$

a significant distance even before accounting for the HNL's (likely significant) Lorentz boost [307]. Thus, such decays may be reconstructed as displaced vertices in modern detectors with µm-level track resolution. Experiments such as DELPHI [312], Belle [313, 314], BaBar [315], BESIII [316], MicroBooNE [317], T2K-ND280 [318], CHARM [319–321], ATLAS [322, 323], CMS [324, 325], and LHCb [326] have contributed leading constraints[15] in the HNL mass range from several hundred MeV to nearly 1 TeV. Future fixed-target, collider, and neutrino experiments will continue to constrain the HNL parameter space (see, e.g., Ref. [270, 327] and references therein).

2.3.4 Cosmological constraints

A large part of Chapter 3 will be devoted to searches for decaying sterile-neutrino DM out in the Universe; here, I describe several other cosmological constraints on sterile-neutrino DM in the context of the νMSM.

Early-Universe production

The early-Universe production of sterile-neutrino DM via neutrino mixing must satisfy two principal conditions: it must produce the correct amount[16] of DM, and it must not alter the thermal history of the early Universe or BBN light-element abundances in a detectable way.

---

[15] Note that the references provided are a non-exhaustive list—each of these experiments may have several published results, as well as subsequent re-analyses in terms of different models. See, e.g., Ref. [270] for a recent review.

[16] In the νMSM discussed in this dissertation, the only DM candidate is the sterile neutrino, so this process must produce all of the observed DM density. In other models, of course, DM may contain many different types of particles, which would relax these production constraints.

As discussed in Chapter 1, a wealth of cosmological observations indicate that $\Omega_{\mathrm{DM}}h^2 \approx 0.12$. Since the sterile-neutrino production rate is proportional to the in-medium mixing angle $\sin^2(2\theta_m)$, it stands to reason that increasing the mixing angle increases the amount of DM produced. In the context of nonresonant (Dodelson-Widrow) production, Ref. [328] obtained an approximate relation between the sterile-neutrino mass m_χ and the mixing angle in vacuum $\sin^2(2\theta)$:

$$m_\chi \approx 1.8\,\mathrm{keV}\left[\frac{\sin^2(2\theta)}{10^{-8}}\right]^{-0.62}\left(\frac{\Omega_{\mathrm{DM}}h^2}{0.12}\right)^{0.5}\mathrm{erfc}\left[-1.15\left(\frac{T_{\mathrm{QCD}}}{170\,\mathrm{MeV}}\right)^{2.15}\right]. \quad (2.26)$$

Here, $T_{\mathrm{QCD}} \approx 170\,\mathrm{MeV}$ is the temperature of the quark-hadron phase transition (see Chapter 4). For a given value of $\Omega_{\mathrm{DM}}h^2$, this result provides an upper bound in the mass-mixing angle plane (i.e., values of $\sin^2(2\theta)$ above this bound produce too much DM). In the case of resonant (Shi-Fuller) production via an initial lepton asymmetry, the situation is substantially more complicated, and approximations like Eq. 2.26 are generally not available. Thus, most groups use numerical solvers (e.g., STERILE-DM [296]) to determine the lepton asymmetry required to produce the observed $\Omega_{\mathrm{DM}}h^2 \approx 0.12$ at each point in the mass-mixing angle parameter space. Similarly, if the maximum allowed value of the lepton asymmetry is known, it is possible to set a lower bound on $\sin^2(2\theta)$, below which too little DM is produced. In the νMSM, where the lepton asymmetry is generated by CP-violating oscillations of heavy sterile neutrinos, calculations suggest[17] $L_6 \equiv 10^6(n_\nu - n_{\bar{\nu}})/s \lesssim 700$ (e.g., Refs. [290, 329]). The results of these computations are shown in Fig. 2-2. In any case, the currently-allowed values of L_ν are several orders of magnitude greater than the baryon-to-photon ratio η_b.

A complementary bound on the lepton asymmetry (and thus the mixing angle) arises from BBN. In brief, the introduction of an asymmetry between neutrinos and antineutrinos shifts neutron-proton equilibrium maintained by reactions such as $p + e^- \leftrightarrow n + \nu_e$ and $n + e^+ \leftrightarrow p + \bar{\nu}_e$, thereby modifying the helium-4 abundance

---

[17] In the rest of this dissertation, L_6 will always be $10^6(n_\nu - n_{\bar{\nu}})/s$—i.e., L_6 will never be normalized to n_γ.

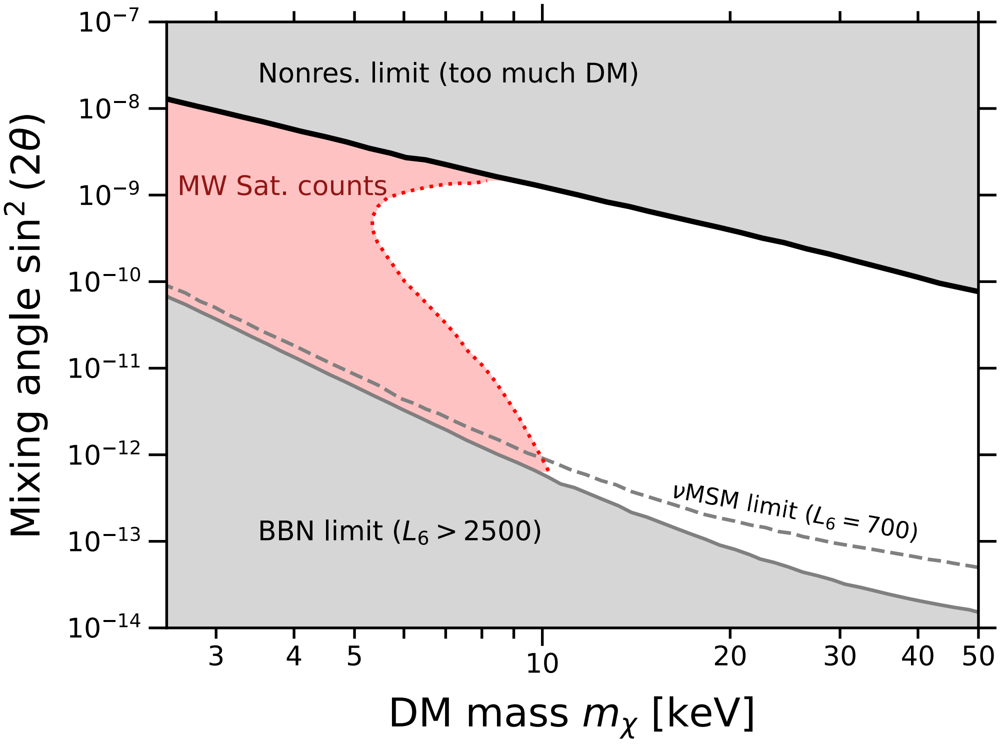

Figure 2-2: Constraints on sterile-neutrino DM in the νMSM from several cosmological sources. The top gray shaded region indicates mixing angles which produce too much DM compared to observations in the nonresonant scenario [300]. The red shaded region is disfavored by the observed number of Milky Way satellite galaxies (here taken to be 47, Ref. [333]). The gray dashed line $L_6 = 700$ is the largest lepton asymmetry achievable in the νMSM, and the bottom gray shaded region corresponds to the BBN limit $L_6 = 2500$ [333].

^{4}He/H $\equiv Y_p$. Thus, measurements of Y_p can be compared to the predictions of standard BBN (i.e., $L_\nu \sim \eta_b \ll 1$) and used to generate a bound on the lepton asymmetry. Previous works find $L_6 \lesssim 2500$ from BBN (see, e.g., Ref. [330]), which is slightly less constraining than the νMSM limit of $L_6 \lesssim 700$. It should be noted that recent measurements of Y_p in extremely metal-poor galaxies deviate from the expectation of zero lepton asymmetry by nearly 3σ, indicating that $L_6 \gtrsim 10^3$ [331, 332]. It is hoped that future CMB experiments (e.g., the Simons Observatory [201] and CMB-S4 [200]) will be able to constrain Y_p at the few-percent level. In the rest of this dissertation, I will adopt the bound $L_6 \lesssim 2500$ to be conservative.

Structure formation

Unlike the canonical cold DM theory, sterile-neutrino DM may have non-negligible momentum at the epoch of cosmic structure formation. Qualitatively, as the sterile-neutrino momentum increases, it more readily escapes from small DM halos. In terms of the matter power spectrum $\mathcal{P}_m(k)$, this implies a suppression of power at large wavenumbers k (i.e., smaller length scales). Comparisons between CDM and alternative DM models are typically made using the transfer function

$$T(k) \equiv \sqrt{\frac{\mathcal{P}_m(k)}{\mathcal{P}_{\mathrm{CDM}}(k)}}. \tag{2.27}$$

For a given DM model (specifying the momentum distribution) and set of initial perturbations (usually taken from the CMB), the matter power spectra $\mathcal{P}_m$ at late times are generally calculated numerically using Boltzmann codes such as CLASS [334] or CAMB [335], or with N-body simulations.

A particularly stringent test for alternative DM models that suppress small-scale structure is the number of Milky Way satellite galaxies. At least 50 confirmed and candidate satellite galaxies are known to orbit the Milky Way, with more discovered each year [336]. Conservatively, any DM model which underestimates the observed number of MW satellites can be ruled out. (In fact, an early motivation for sterile-neutrino DM was the "missing satellites problem"—i.e., the MW appeared to have fewer satellites than predicted by ΛCDM—though the tension may have alleviated recently [337–339].) One complication of this method is the connection between DM halos and the eventual satellite galaxies that form within them—after all, DM halos containing very few stars are extremely hard to see. Furthermore, depending on the treatment of halo-to-galaxy matching, various groups have disfavored sterile-neutrino DM masses below ∼10–30 keV (see, e.g., Refs. [333, 336, 340]). In this dissertation, I will adopt the results of Ref. [333] to be conservative.

An additional nearly-model-independent constraint on the properties of fermion

DM arises from the Pauli exclusion principle. If the DM particle mass is sufficiently small, the number density in the cores of galactic halos may approach the degeneracy limit for a Fermi gas. This *Tremaine-Gunn bound* [341] provides a lower limit on the mass of fermionic DM particles. The central regions of ultra-faint dwarf spheroidal galaxies are particularly attractive probes, owing to their extremely large DM-to-baryon ratios. Recent analyses disfavor $m_\chi \lesssim 2\,\mathrm{keV}$ (see, e.g., Refs. [342, 343] and references therein).

I conclude this section by mentioning two other structure-formation techniques used to constrain sterile-neutrino DM and other alternatives to ΛCDM. The first is the Lyman-α (Lyα) forest: along the line of sight to distant quasars are many galactic halos containing hydrogen, producing absorption lines at a wavelength of $1216\,\mathrm{\AA}$ in the rest frame of the absorber, leading to a "forest" of absorption lines based on the redshift of the quasar and absorbing clouds [344, 345]. As such, the Lyα forest is a powerful probe of the baryonic matter distribution (which is assumed to trace the DM distribution) at large scales. The other technique is based on strong gravitational lensing of quasars by foreground galaxies. The geometry of the lensed quasar image is sensitive to the underlying mass distribution in the lens, including small halos (down to $\sim 10^8\,M_\odot$ in some cases) which may not have visible galaxies within them (see, e.g., Refs. [181, 346–349]). Recent Lyα and lensing constraints on resonantly-produced sterile-neutrino DM are comparable to those obtained from other structure-formation probes (see, e.g., Ref. [350]), though with different systematic uncertainties (in particular, lensing does not depend on the complicated baryonic feedback processes within DM halos).

3

X-Ray Searches for Sterile Neutrinos

There is a crack, a crack in everything /
That's how the light gets in.
—Leonard Cohen, "Anthem"

In this chapter, I describe the process by which x-ray telescope observations are converted into constraints on the decay rate of keV-scale sterile-neutrino DM. In Sec. 3.1, I quantify the expected sterile-neutrino x-ray intensity from a given region of the sky. In Sec. 3.2, I outline the basic statistical procedure by which x-ray spectra are searched for x-ray lines from sterile-neutrino DM. In Sec. 3.3 I (very briefly) describe some of the x-ray instruments which have recently contributed to the DM search. In Sec. 3.4 I detail the design and characteristics of the NuSTAR observatory which make it a valuable tool for DM hunting. In Sec. 3.5, I describe NuSTAR constraints on sterile-neutrino DM, with particular emphasis on my contributions to analyses of data from the M31 galaxy, the Galactic bulge, and the Galactic halo.

3.1 Expected signal

As described in the previous chapter, massive sterile neutrinos decay as a consequence of their oscillation-induced mixing with the SM neutrinos. If the sterile-neutrino mass $m_\chi < 2m_e$, there are two dominant decay modes, shown in Fig. 3-1: the three-neutrino mode $\chi \to \nu_\alpha \nu_\beta \bar{\nu}_\beta$ and the radiative mode $\chi \to \gamma \nu_\alpha$, where α, β may be any lepton flavor. Since it emits a photon, the latter decay mode is much more amenable to indirect-detection searches. When summed over all flavors of the (undetected) neutrino, the radiative decay rate is (e.g., Refs. [351, 352])

$$\Gamma_{\chi \to \gamma \nu} = \frac{9 \alpha_e G_\mathrm{F}^2}{1024 \pi^4} \sin^2(2\theta) m_\chi^5 \tag{3.1}$$

$$\simeq 1.38 \times 10^{-32}\,\mathrm{s}^{-1} \left[\frac{\sin^2(2\theta)}{10^{-10}}\right] \left(\frac{m_\chi}{1\,\mathrm{keV}}\right)^5, \tag{3.2}$$

where $\alpha_e \approx 1/137$ is the electromagnetic fine-structure constant, G_F is the Fermi constant, and θ is the flavor-inclusive active-sterile mixing angle. Since the radiative decay mode contains a two-body final state, the outgoing photon is monoenergetic, with $E_\gamma \simeq \frac{1}{2} m_\chi$ (assuming $m_\chi \gg m_\nu$). For sterile-neutrino masses m_χ at the keV scale, an x-ray photon is emitted, whose spectrum in the sterile-neutrino rest frame is given by

$$\frac{\mathrm{d}N_\gamma}{\mathrm{d}E_\gamma} \simeq \delta(E_\gamma - \tfrac{1}{2} m_\chi). \tag{3.3}$$

The spectrum of Eq. 3.3 may be modified by several effects. First, the velocity of DM particles (of order ~few hundred km s^{-1}) in galactic halos leads to Doppler broadening, leading to a linewidth $\delta E_\gamma / E_\gamma \sim v/c \sim 10^{-3}$. Nearly all operating x-ray telescopes have detector energy resolutions which exceed the Doppler width by at least an order of magnitude, so the δ-function approximation is more than adequate. Second, for galaxies at great distances, the detected photon energy $E_\gamma = \frac{1}{2} m_\chi (1+z)^{-1}$, where z is the redshift. Recalling the results of Eq. 1.21, the specific

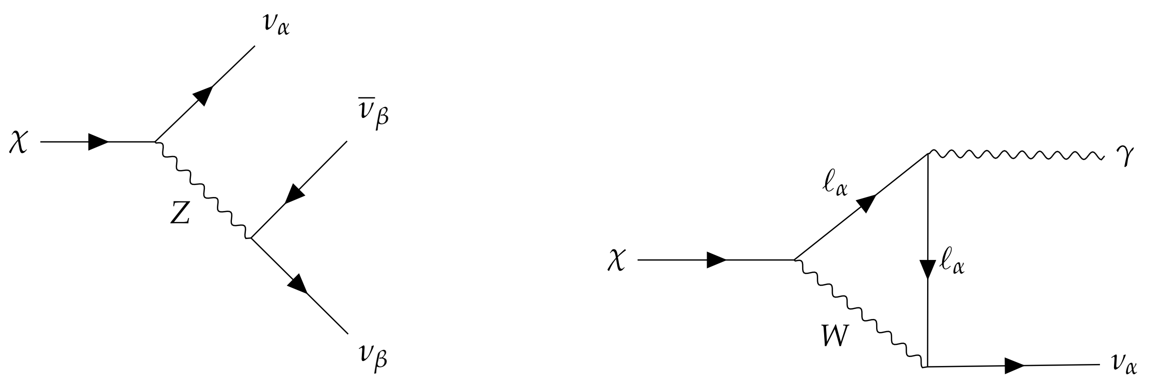

Figure 3-1: **(Left)** Representative Feynman diagram for the $\chi \to 3\nu$ decay, where the neutrino flavors α, β need not be the same. **(Right)** Representative Feynman diagram for the $\chi \to \gamma\nu$ decay mode. Both diagrams were typeset using `tikz-feynman` [353].

x-ray intensity $\mathcal{I}_\gamma \equiv \mathrm{d}^2 F_\gamma / \mathrm{d}E_\gamma \mathrm{d}\Omega$ is given by

$$\mathcal{I}_\gamma = \frac{\Gamma_{\chi\to\gamma\nu}}{4\pi m_\chi}\frac{\mathrm{d}N_\gamma}{\mathrm{d}E_\gamma} \times \underbrace{\frac{1}{\Delta\Omega}\int_{\mathrm{FOV}} \xi \, \mathrm{d}\Omega \int_{\mathrm{LOS}} \rho_\chi \, \mathrm{d}r}_{\langle \mathrm{d}\mathcal{D}/\mathrm{d}\Omega \rangle}, \tag{3.4}$$

where the effective solid angle $\Delta\Omega \equiv \int_{\mathrm{FOV}} \xi \, \mathrm{d}\Omega$ and ξ is the position-detection efficiency across the detector FOV.

(It should be noted here that nearly all of the x-ray analyses discussed in this chapter assume that all of DM is composed of sterile neutrinos, since there is no information to suggest otherwise. If only a fraction f of DM particles are sterile neutrinos, the x-ray constraints derived in this chapter weaken by that same factor.)

3.2 X-ray data format and analysis

In an abstract sense, the purpose of an x-ray telescope is to convert the specific photon intensity $\mathcal{I}_\gamma \equiv \mathrm{d}^2 F_\gamma / \mathrm{d}E_\gamma \, \mathrm{d}\Omega$ of some source[1] on the sky into discrete elec-

---

[1] Strictly, $\mathcal{I}_\gamma$ is only relevant for extended sources (i.e., those which are significantly larger than the angular resolution of the instrument, such that "surface brightness" is a meaningful quantity). For

trical signals. The data output from most x-ray telescopes consists of a list of x-ray events, with each event having at minimum a timestamp t, detector position (x, y), and pulse height I (a proxy for energy). Subsequent data products depend on the parameter of interest: lightcurves depend only on time, spectra depend only on energy, and images depend only on position. The vast majority of x-ray analyses described in this chapter are based on spectra, so it is that formalism I will briefly describe.

Imagine we wish to extract an x-ray spectrum from some extended region of the detector (i.e., some extended region on the sky). The fundamental equation of x-ray spectroscopy is [354]

$$\frac{\mathrm{d}N(I)}{\mathrm{d}t} = \int \mathrm{RMF}(E_\gamma, I)\ \mathrm{ARF}(E_\gamma)\, \mathcal{I}_\gamma \, \mathrm{d}E_\gamma, \tag{3.5}$$

which gives the rate of x-ray events in each pulse-height[2] bin. The RMF is the *redistribution matrix file* which is, as its name suggests, a matrix whose elements encode the probability of a photon with energy E_γ being recorded as a pulse height I in the detector (including the photoabsorption probability and detector energy resolution). On the other hand, the ARF is the *auxiliary response file*, whose primary role is to encode the energy-dependent effective area[3] $A(E_\gamma)$ of the detector and/or optics. (Depending on the instrument, the ARF may also include the throughput of various filters or other optical elements.) For extended sources, the ARF also carries a factor $\Delta\Omega$ representing the solid angle of the region from which the spectrum is extracted. When considered together, the RMF and ARF define the *response* of the detector to incoming x-rays, and are generally distributed in a calibration database (CALDB) maintained by the instrument team. (For more details, including the cases of spatially and/or temporally varying responses, see, e.g., Ref. [354].)

As alluded to previously, the conversion of an x-ray with energy E_γ to a count

---

point sources, $\mathrm{d}F_\gamma/\mathrm{d}E_\gamma$ is more appropriate.

[2] To return to the more physically familiar $\mathrm{d}^2N/\mathrm{d}E_\gamma\mathrm{d}t$, simply divide each bin $N(I)$ in Eq. 3.5 by its associated width in energy.

[3] For extended sources, the ARF also incorporates the effect of vignetting, i.e., the reduction in effective area with increasing distance from the optical axis.

in some detector channel I is an inherently lossy process. As one example, the finite detector energy resolution means it is impossible to uniquely determine the energy of each incident photon: does a count in channel $I = 85$ correspond to a 5-keV photon, a 5.3-keV photon, or the escape peak from a 50-keV photon? Since it is not generally possible to invert Eq. 3.5 to find the incident source spectrum, most x-ray spectral analyses involve the process of *forward folding*. In this scheme, the data analyst defines a model $M(E_\gamma; \boldsymbol{\eta})$ for the incident source spectrum, convolves it with the appropriate[4] ARF and/or RMF, and allows the model parameters $\boldsymbol{\eta}$ to find their best-fit values. To compare the observed counts per bin $\boldsymbol{u} = (u_1, u_2 \cdots u_N)$ to the predicted model counts per bin $\boldsymbol{m}(\boldsymbol{\eta}) = (m_1, m_2, \cdots m_N)$, it is common to use the Poisson likelihood

$$\mathcal{L}(\boldsymbol{\eta}; \boldsymbol{u}) = \prod_i \frac{m_i^{u_i} e^{-m_i}}{u_i!}, \tag{3.6}$$

where the product runs over each bin i in the spectrum. In general, $\mathcal{L}$ may depend on many free parameters $\boldsymbol{\eta}$ whose values we are not principally interested in (e.g., the normalization and slope of background components) but which are nonetheless necessary to define the model. For the sterile-neutrino case, in which we seek to discriminate between models with DM (mixing angle $\theta \neq 0$) and without DM ($\theta = 0$), we define the profile likelihood ratio[5]

$$\Lambda(\theta) \equiv \frac{\mathcal{L}(\theta, \hat{\hat{\boldsymbol{\eta}}})}{\mathcal{L}(\hat{\theta}, \hat{\boldsymbol{\eta}})}, \tag{3.7}$$

which quantifies how the goodness of fit of the model changes with θ. (Single hats denote the unconstrained best-fit values of parameters, whereas $\hat{\hat{\boldsymbol{\eta}}}$ denotes the best-fit values subject to the constraint of fixed θ.) The profile likelihood ratio is an example of a test statistic (TS), a single number which can be used to quantify the goodness of fit. The probability of obtaining a TS at least as great as the observed

---

[4] For example, when modeling spectral components unrelated to the sky (e.g., internal detector backgrounds), only the RMF may be required.

[5] The Neyman-Pearson lemma states that the likelihood ratio is the uniformly most powerful test for distinguishing between two hypotheses at a given confidence level [355].

value Λ_{obs} (assuming the hypothesis $\mathcal{H}_i$ is true) is given by the p-value,

$$p = \int_{\Lambda_{\text{obs}}}^{\infty} f(\Lambda|\mathcal{H}_i)\,\mathrm{d}\Lambda \tag{3.8}$$

where f is the probability density function of Λ for a given hypothesis. For numerical and practical reasons, it is generally more common to work with the negative log-likelihood ratio $\lambda(\theta) \equiv -2\ln\Lambda(\theta)$. Furthermore, the results of Wilks [356] and Wald [357] state that the probability distribution for $\lambda(\theta)$ reduces to the familiar χ^2 distribution for a single degree of freedom, provided certain conditions are met.

Searches for sterile-neutrino DM decay with x-ray telescopes generally take the form of a "line scan." In this procedure, a δ-function line at fixed energy $E_i = \frac{1}{2}m_i$ is added to the model, with normalization set by the incident photon intensity of Eq. 3.4. This line is then convolved with the appropriate instrument response files (ARF and/or RMF). The nuisance parameters $\boldsymbol{\eta}$ are allowed to float[6] to their optimal values while keeping $\mathcal{I}_i$ (and hence the mixing angle θ_i) fixed. since This step is repeated for the next value of $\mathcal{I}(\theta)$ until the full likelihood curve $\lambda(\theta; m_\chi)$ is constructed, after which the line energy is moved to the next value to be tested and the process begins again. With the collection of profile likelihood curves $\lambda(\theta; m_\chi)$, we have everything we need to determine the statistical significance of any detected x-ray lines, or barring that, set upper limits on the mixing angle θ. To accomplish this, we define test statistics t and q for line-detection and upper limits, respectively [358]:

$$t(m_\chi) = \begin{cases} -2\ln(0), & \hat{\theta} \geq 0, \\ 0, & \hat{\theta} < 0 \end{cases} \tag{3.9}$$

$$q(\theta; m_\chi) = \begin{cases} -2\ln(\theta), & \hat{\theta} \leq \theta, \\ 0, & \hat{\theta} > \theta. \end{cases} \tag{3.10}$$

---

[6] This includes the normalizations of background emission lines. In the vicinity of such lines, the allowed flux of DM photons greatly increases (and thus the constraints on the decay rate Γ greatly worsen), since the DM model is allowed to assume the full strength of these background lines.

In terms of the more familiar χ^2 fit statistic applicable in the large-statistics limit, the test statistics are

$$t(m_\chi) \simeq \chi^2(\hat{\theta}) - \chi^2(0) \tag{3.11}$$

$$q(\theta; m_\chi) \simeq \chi^2(\theta) - \chi^2(\hat{\theta}). \tag{3.12}$$

The local statistical significance of any line detection in Gaussian standard deviations is simply $\sqrt{t(m_\chi)}$, and the one-sided 95% upper limit θ_{95} occurs for the largest value[7] of θ where $q = 2.71$. (Note that Wilks' theorem assumes that the true value of θ lies within the parameter space, which becomes problematic if there is a downward fluctuation in the data and we are constraining $\theta \geq 0$. To account for this, the probability distributions of t and q are mixtures of a χ^2 distribution and a spike at $\delta(0)$, a so-called $\frac{1}{2}\chi^2$ distribution [358–360]. The NuSTAR limits described in this chapter assume $\theta \geq 0$ and thus set 95% one-sided upper limits where $q = 2.71$.)

Before turning our attention to specific x-ray instruments, it will be helpful to discuss how the expected DM sensitivity scales with various properties of the telescope. As a figure of merit, consider the signal-to-noise ratio $S/\sqrt{B}$ where S and B are the number of (DM) signal versus background photons, respectively. For a DM particle decaying with rate Γ, the number of DM line photons collected is obtained by integrating Eq. 3.4 over the effective area A, solid angle $\Delta\Omega$, energy E_γ, and exposure time T,

$$S \propto \Gamma A \Delta\Omega T, \tag{3.13}$$

assuming that the intrinsic width of the DM line is much less than the instrument energy resolution ΔE. (As discussed previously, this assumption is nearly always true for keV-scale sterile-neutrino DM, so long as $\Delta E/E \gtrsim 10^{-3}$.) On the other hand, the astrophysical background $\mathcal{I}_{\rm bkg}$ is a continuum, so the number of back-

---

[7] If there is a preference for line emission at some energy, $\lambda(\theta)$ will be parabolic near the minimum, and there may be two values of θ where $q = 2.71$. Since we are setting upper limits, we use the larger value of θ.

ground events B is proportional to ΔE:

$$B \propto \mathcal{I}_{\mathrm{bkg}} A \Delta\Omega T \Delta E. \tag{3.14}$$

Putting both together, the statistical significance of a weak line[8] atop a featureless astrophysical background is (e.g., Ref. [361])

$$\frac{S}{\sqrt{B}} \propto \Gamma \sqrt{\frac{A\Delta\Omega T}{\mathcal{I}_{\mathrm{bkg}}\Delta E}} \quad \text{(astrophysical background).} \tag{3.15}$$

On the other hand, if the background mainly arises from within the detector rather than from astrophysical sources, B is no longer proportional to $A\Delta\Omega$, and instead depends on $R_{\mathrm{bkg}} \equiv \mathrm{d}^2 B/\mathrm{d}E_\gamma \mathrm{d}t$:

$$\frac{S}{\sqrt{B}} \propto \Gamma A\Delta\Omega \sqrt{\frac{T}{R_{\mathrm{bkg}}\Delta E}} \quad \text{(detector background).} \tag{3.16}$$

Lastly, the combination $A\Delta\Omega$ occurs so frequently in these sorts of studies that it is given its own name, the *grasp* (though since most analyses involve integration over some extended region of the sky, we are generally interested in the FOV-averaged grasp $\langle A\Delta\Omega\rangle_{\mathrm{FOV}}$, which takes into account vignetting of the effective area).

3.3 X-ray instruments

The field of x-ray astronomy is vast, and cannot be summarized in one dissertation (for extensive reviews, see, e.g., Refs. [362–364]). In this section, I will briefly summarize the contributions of recent x-ray astronomy missions to the search for keV-scale sterile neutrinos, divided broadly into non-focusing and focusing instruments. Headings with an asterisk (*) denote missions which are currently being built, as well as instruments that are currently (as of early 2023) operating but have

---

[8] Here, "weak" means that the equivalent width of the spectral line is less than ΔE. Note that reducing ΔE may push the DM line from being weak to strong, after which point further improvements in ΔE do not improve the sensitivity.

Instrument	Energy range	$A\Delta\Omega$ [cm^2 deg^2]	E_γ/FWHM	Focusing?
INTEGRAL SPI	18 keV–8 MeV	$\sim 10^4$ @ 100 keV	100–500	✗
Fermi GBM	10 keV–30 MeV	$\sim 10^5$ @ 100 keV$^{(a)}$	$\sim$10	✗
NuSTAR 0-bounce	3–160 keV	$\sim$100	10–200	✗
HaloSat	0.4–7 keV	$\sim$30 @ 1 keV	10–100	✗
Micro-X DM*	0.5–10 keV	$\sim 10^3$	150–3000	✗
XMM-Newton EPIC	0.15–15 keV	$\sim$200 @ 2 keV	20–50	✓
Chandra ACIS-I	0.1–10 keV	$\sim$20 @ 5 keV	10–20	✓
SRG eROSITA*	0.2–10 keV	$\sim 10^3$ @ 1 keV	$\sim$100	✓
SRG ART-XC*	4–30 keV	$\sim$40 @ 8 keV	10–20	✓
XRISM Resolve*	0.4–12 keV	$\sim$0.5 @ 5 keV$^{(a)}$	20–2000	✓
XRISM Xtend*	0.4–12 keV	$\sim$50 @ 5 keV$^{(a)}$	$\sim$few–50	✓
NuSTAR 2-bounce	3–79 keV	$\sim$20 cm^2 @ 10 keV	10–200	✓
Suzaku XIS	0.12–12 keV	$\sim$10 cm^2 @ 5 keV$^{(a)}$	20–50	✓

Table 3.1: Properties of x-ray instruments which have (or will soon, indicated by the asterisk *) contributed to the search for sterile-neutrino DM. The grasp $A\Delta\Omega$ is generally the FOV-averaged value, with $^{(a)}$ representing on-axis values if the FOV-averaged value was not given in the literature. For more details, see Sec. 3.3.

not yet reported sterile-neutrino constraints. It should be noted that with very few exceptions, these instruments were not designed to search for DM, illustrating the flexibility of indirect-detection techniques.

3.3.1 Non-focusing instruments

Unlike lower-wavelength (e.g., optical) photons, x-rays are not easily focused by mirrors or lenses. Thus, many x-ray instruments have been non-focusing devices, analogous to the pinhole camera or camera obscura. These devices are generally simpler to build than their focusing counterparts (discussed next) and are frequently used to study diffuse x-ray emission on large angular scales.

INTEGRAL

Launched in 2002, the International Gamma-Ray Astrophysics Laboratory (INTEGRAL) is Europe's flagship hard x-ray and soft γ-ray observatory. At its core, IN-

TEGRAL is a coded-aperture telescope: incoming photons must first pass through one of the openings drilled into a tungsten mask before striking the detectors. The shadow pattern obtained from viewing the same source at different angles allows for localization of sources on the sky. INTEGRAL's main spectroscopic instrument is (appropriately) the Spectrometer on INTEGRAL (SPI), consisting of an array of 19 high-purity germanium detectors [365]. These detectors provide sensitivity to photon energies $\sim$18 keV–8 MeV, with an energy resolution $\sim$2 keV FWHM for 1-MeV photons. Though SPI is principally used for spectroscopy, its coded mask also allows for imaging, with a FOV $\sim$15–30° in diameter[9] and angular resolution $\sim$2.5° FWHM. With a FOV-averaged grasp $\gtrsim 10^4\,\mathrm{cm}^2\,\mathrm{deg}^2$ across much of its energy range and total exposure time $\gtrsim 1$ Ms across much of the sky, SPI is extremely sensitive to diffuse x-ray and γ-ray line emission. Along with Fermi-GBM (discussed next), INTEGRAL-SPI is responsible for closing the νMSM parameter space for sterile-neutrino masses $m_\chi \gtrsim 40\,\mathrm{keV}$ down to the BBN limit [366, 367]. As of early 2023, the other two instruments—the Imager on-board the INTEGRAL Satellite (IBIS, Ref. [368]) and the Joint European X-ray Monitors (JEM-X, Ref. [369])—have not reported sterile-neutrino DM constraints.

Fermi

Launched in 2009, the Fermi Gamma-Ray Space Telescope (hereafter "Fermi") is NASA's flagship high-energy γ-ray observatory. It consists of two main instruments, the Large Area Telescope (LAT) and the Gamma-ray Burst Monitor (GBM). The LAT is a pair-conversion telescope [370], consisting of stacks of high-Z foils to convert incoming γ-rays into e^+e^- pairs, which are subsequently tracked by silicon strip detectors and absorbed in cesium iodide calorimeters. The LAT energy range of $\sim$20 MeV–300 GeV energy range is far too high for νMSM sterile-neutrino DM searches, but has contributed to many searches for GeV-to-TeV-scale DM (see, e.g., Refs. [371–374] for a very non-exhaustive list of recent work). Fortunately, the GBM

---

[9] Depending on whether one takes the fully-coded FOV or the FOV out to zero response, respecively.

is more well-suited to keV-scale DM searches, consisting of twelve sodium iodide (NaI) scintillators and two bismuth germanium oxide (BGO) scintillators. The NaI scintillators are sensitive to photon energies of $\sim$few keV–few MeV, whereas the BGO scintillators cover the range $\sim$150 keV–30 MeV. (Both detectors have energy resolutions $\lesssim 10\%$ FWHM.) With an on-axis effective area peaking near $\sim 100\,\mathrm{cm}^2$ for both NaI and BGO [375], an instantaneous FOV $\Delta\Omega \sim 2 \times 10^4\,\mathrm{deg}^2$, and a substantial cumulative exposure time, the sterile-neutrino sensitivity of GBM is comparable to SPI. Prior to NuSTAR, the all-sky sterile-neutrino constraints from Fermi-GBM were some of the strongest in the mass range $\sim$25–40 keV [376].

HaloSat

HaloSat was a CubeSat[10] launched in 2018 to (somewhat ironically) search for the "missing" baryons in the Milky Way halo by means of their soft x-ray emission [377]. The science payload consisted of three silicon drift detectors (SDDs) with energy range 0.4–7 keV and energy resolution $\sim$85–130 eV FWHM. The instrument's 0.3-cm^2 effective area above 1 keV was compensated by an effective FOV $\Delta\Omega \approx 115\,\mathrm{deg}^2$, giving HaloSat a grasp comparable to much larger x-ray telescopes without the systematic uncertainties associated with focusing optics or detector pixelization. Despite its short orbital lifetime (2018–2021), HaloSat provided constraints on $m_\chi = 7\,\mathrm{keV}$ sterile-neutrino DM that were comparable to instruments with orders of magnitude more exposure time [378].

Micro-X*

Micro-X differs from the other x-ray instruments in this section in that it is not a satellite—it is a sounding rocket [379–381]. As such, the instrument only has a few hundred seconds at the apogee of each flight to make its measurement; thus, sensitivity is crucial. At the core of Micro-X is a $\sim$1-cm^2 array of $\sim$120 cryogenic transition edge sensor (TES) bolometers with energy resolution better than 10 eV

---

[10] A small satellite optimized for simplicity and cost.

FWHM. For the upcoming launch dedicated to searching for DM, the Micro-X focusing optics will be removed to increase the detector FOV from 0.2° to nearly 30°, greatly increasing the expected DM signal. With only several hundred seconds of exposure time near the MW Galactic Center, it is anticipated that Micro-X will achieve constraints in the sterile-neutrino mass range $m_\chi \lesssim 10\,\mathrm{keV}$ that meet or exceed those of XMM-Newton and NuSTAR [379]. In 2022, Micro-X completed its first successful (focused) observation of the Cassiopeia A supernova remnant, and the follow-up DM-dedicated flight will hopefully occur by 2023–2024.

3.3.2 Focusing telescopes

In contrast to the previous instruments, some telescopes use x-ray focusing optics to increase the collecting area and improve angular resolution, though generally at the cost of requiring a smaller field of view. (For more details on focusing optics in the context of NuSTAR, see Sec. 3.4.)

XMM-Newton

Launched in 1999, XMM-Newton is the European Space Agency's flagship x-ray observatory. The spacecraft is built around three Wolter-I x-ray telescopes, responsible for focusing x-rays onto the Reflection Grating Spectrometers (RGS[11], Ref. [382]) and the European Photon Imaging Cameras (EPIC). The EPIC detectors are silicon pixel devices, consisting of two arrays of front-illuminated metal-oxide semiconductor (EPIC-MOS, Ref. [383]) and one array of back-illuminated pn-junction (EPIC-pn, Ref. [384]) detectors. All three EPIC modules share the same $\sim$30-arcmin-diameter FOV (i.e., $\Delta\Omega \approx 0.8\,\mathrm{deg}^2$), and have similar spectral resolutions ($\sim$70 eV FWHM at 1 keV to $\sim$150 eV FWHM at 6.4 keV) and bandpasses (0.15–12 keV). When considered together, the three EPIC detectors have a combined FOV-averaged grasp of $\sim 100\,\mathrm{cm}^2\,\mathrm{deg}^2$ in the energy range 1–5 keV, making

---

[11] The RGS are dispersive spectrometers which provide high-resolution ($\Delta E/E \lesssim 10^{-2}$ FWHM) spectroscopy in the energy range 0.33–2.5 keV. They are not generally used for DM searches since their grasp is quite low.)

XMM-Newton a leading player in the search for keV-scale DM for many years (see, e.g., Refs. [385–390]). Unfortunately, the decline in mirror reflectivity and detector efficiency for x-ray energies above $\sim$10 keV limits XMM-Newton's ability to search for higher-mass DM.

Chandra

In many ways, the Chandra X-ray Observatory (CXO) is NASA's counterpart to XMM-Newton. Also launched in 1999, CXO consists of a single Wolter-I x-ray telescope feeding several focal plane instruments. The primary instrument[12] is the silicon CCD-based Advanced CCD Imaging Spectrometer (ACIS, Ref. [394]), which contains arrays for imaging (ACIS-I) and grating spectroscopy (ACIS-S). The ACIS-I array provides high-resolution imaging ($\lesssim$1 arcsec) across a 17×17 arcmin^2 FOV ($\Delta\Omega \approx 0.08\,\text{deg}^2$), with an energy resolution $\sim$100-200 eV FWHM for photon energies 1–5 keV. The FOV-averaged grasp of ACIS-I is $\sim$10–20 cm^2 deg^2 in the x-ray energy range $\sim$1–5 keV—somewhat lower than XMM-Newton, mainly owing to CXO's smaller FOV. Despite this, CXO has still contributed extensively to the search for sterile-neutrino DM, owing to its large catalogue of observations and excellent angular resolution for excluding contamination from x-ray point sources (see, e.g., Refs. [395–397]).

Suzaku

The Suzaku observatory was launched in 2005 as a joint venture between the Japan Aerospace Exploration Agency (JAXA) and NASA, and operated until 2015. Suzaku consisted of three instruments, the X-Ray Spectrometer (XRS, Ref. [398]), the X-ray Imaging Spectrometer (XIS, Ref. [399]), and the Hard X-ray Detector (HXD, Ref. [400]). The XRS consisted of a cryogenic microcalorimeter array with $\sim$6-eV FWHM energy resolution in the energy range 0.3–12 keV; unfortunately,

---

[12] The other CXO focal-plane instrument is the High Resolution Camera (HRC, Ref. [391]), a microchannel plate detector with superior angular resolution (though worse energy resolution) compared to ACIS. Both ACIS and HRC can also be used in conjunction with low- and medium-energy grating spectrometers similar to XMM-Newton's RGS [392, 393].

the liquid helium coolant system failed shortly after launch, rendering the XRS inoperable. The other focusing x-ray instrument aboard Suzaku was the XIS, consisting of four arrays of silicon CCD arrays observing the same 19×19 arcmin^2 ($\Delta\Omega \sim 0.1\,\mathrm{deg}^2$) FOV and sensitive to 0.2–12-keV photons. Despite its somewhat small peak effective area ($\sim$200 cm^2), XIS provided some of the few sterile-neutrino constraints in the mass range $\sim$10–20 keV prior to NuSTAR (see, e.g., Refs. [401–405]). Finally, the HXD was a non-imaging collimated detector consisting of silicon *p-i-n* diodes and gadolinium silicate scintillators surrounded by anticoincidence detectors. The energy bandpass of HXD was 10–600 keV, with a peak effective area $\sim$200 cm^2 near 100 keV and a solid angle $\Delta\Omega \sim 0.3\,\mathrm{deg}^2$ below 100 keV to $\sim$20 deg^2 above 100 keV.

Swift

The Neil Gehrels Swift (hereafter Swift) observatory was launched into low-Earth orbit in 2004 to study γ-ray bursts [406]. Swift's science payload contains two x-ray instruments, the X-Ray Telescope (XRT, Ref. [407]) and the Burst Alert Telescope (BAT, Ref. [408]). The XRT is a grazing-incidence telescope with energy range 0.2–10 keV and peak on-axis effective area $\sim$100 cm^2 at 1.5 keV, coupled to a silicon CCD detector with energy resolution $\sim$0.2 keV. Despite its small solid angle ($\Delta\Omega \sim 0.05\,\mathrm{deg}^2$), the extensive observation catalog of XRT ($>$ 70 Ms) and small low-energy background rate allowed the authors of Ref. [397] to obtain sterile-neutrino constraints that were comparable to CXO's. As of early 2023, there have been no DM searches using the non-focusing BAT.

Spektr-RG*

The Spektr-RG mission was launched in 2019 as a joint venture between the (Russian) Space Research Institute (IKI) and the (German) Max Planck Institute for Extraterrestrial Physics (MPE). The satellite contains two primary science instruments, the soft x-ray telescope eROSITA [409] managed by MPE and the hard x-ray telescope ART-XC [410] managed by IKI. Both telescopes are designed to survey the entire sky at least eight times.

The eROSITA instrument consists of seven telescope modules (TMs 1–7), each consisting of an x-ray mirror assembly and a detector. The x-ray optics are grazing-incidence Wolter-I telescopes consisting of gold-plated nickel, each observing the same $\sim 1^\circ$-diameter focused FOV. The detectors are silicon CCD arrays with energy resolution ranging from $\sim$0.06 keV FWHM at 0.3 keV to $\sim$0.17 keV FWHM at 10 keV. The total grasp of the system ranges from $\sim 10^3\,\mathrm{cm}^2\,\mathrm{deg}^2$ at 1 keV to $\sim 10^2\,\mathrm{cm}^2\,\mathrm{deg}^2$ at 5 keV, a substantial increase over previous x-ray telescopes. Therefore, it is anticipated that eROSITA will be integral[13] in searching for keV-scale DM, including sterile neutrinos in the Galactic halo (see, e.g., Refs. [411, 412] for sensitivity projections). Following the Russian Federation's invasion of Ukraine in early 2022, MPE placed eROSITA in safe mode. At that time, eROSITA had completed four of its planned eight all-sky passes. It is unknown when (or whether) eROSITA will be brought back online, or what condition the instrument will be in if it is.

The ART-XC instrument will enable the first all-sky imaging survey in hard x-rays, with energy bandpass 4–30 keV slightly overlapping with eROSITA. Like eROSITA, ART-XC consists of seven TMs, consisting of x-ray optics and detectors. The ART-XC mirrors, however, consist of iridium-plated nickel/cobalt to provide increased reflectivity at higher energies. The seven ART-XC TMs observe the same 36-arcmin-diameter ($\Delta\Omega \sim 0.3\,\mathrm{deg}^2$) FOV, achieving a grasp of nearly $\sim 40\,\mathrm{cm}^2\,\mathrm{deg}^2$ at 8.1 keV [410]. This is comparable to NuSTAR's 0-bounce sensitivity while also providing arcminute-level angular resolution. The detectors are solid-state CdTe double-sided strip devices, which provide the necessary photoabsorption efficiency for hard x-rays. The projected sterile-neutrino constraints of ART-XC are discussed in Ref. [411], and an analysis of the first two years of ART-XC data is described in Ref. [413]. At the time of writing (early 2023), ART-XC remains online.

---

[13] Pardon the pun.

Hitomi and XRISM*

The Hitomi observatory was a JAXA-led project launched to low-Earth orbit in 2016, containing three x-ray instruments: the Soft X-ray Spectrometer (SXS, Ref. [414]), the Soft X-ray Imager (SXI, Ref. [415]), and the Hard X-ray Imager (HXI, Ref. [416]), each fed by focusing telescopes with peak on-axis effective areas of $\sim$several hundred cm^2. The SXI was a silicon CCD pixel detector with energy resolution $\sim$0.16 keV, energy range 0.4–12 keV, and solid angle $38 \times 38\,\mathrm{arcmin}^2$ ($\Delta\Omega \sim 0.4\,\mathrm{deg}^2$). The HXI consisted of four layers of silicon strip detectors and one layer of CdTe strip detectors, providing a solid angle $9' \times 9'$ and peak effective area $\sim 10^3\,\mathrm{cm}^2$ for 5–80-keV photons. The SXS was perhaps the most technically ambitious instrument on the observatory, consisting of a 36-pixel array of HgTe absorbers and silicon thermistors cooled to 50 mK. Despite its small solid angle ($\Delta\Omega \sim 0.003\,\mathrm{deg}^2$), the SXS offered an unparalleled energy resolution of $\lesssim 7\,\mathrm{eV}$ in the x-ray energy range 0.4–10 keV. It was hoped that the exquisite SXS energy resolution would allow any sterile-neutrino feature to be distinguished from neighboring atomic spectral lines. Sadly, the destruction of the Hitomi observatory only one month after launch meant that only a single target (the Perseus galaxy cluster) could be observed for sterile-neutrino DM during commissioning [417, 418].

Fortunately, this story does not end with the destruction of Hitomi, and the X-Ray Imaging and Spectroscopy Mission (XRISM, previously called the X-ray Astronomy Recovery Mission [419, 420]) was quickly approved. Owing to the hard x-ray performance of NuSTAR, the HXI from Hitomi will not carry over to XRISM. Rather, the science payload of XRISM will consist of Resolve, a near-clone of SXS [421]; and Xtend, an analogue of SXI with improved energy resolution [422]. At the time of writing (early 2023), it is hoped that XRISM will launch this year.

3.3.3 The 3.5-keV excess

In 2014–2015, two independent groups reported evidence of a line-like excess of x-ray photons near 3.5 keV in energy, in the stacked spectra of galaxy clusters [386] and in the Milky Way, M31, and Perseus cluster [387, 423]. In both cases, the excess was consistent with a decaying $\sim$7-keV sterile neutrino with mixing angle $\sin^2(2\theta) \simeq (0.2\text{–}2) \times 10^{-10}$ depending on the choice of DM density profile, though other models were also considered. This excess was subsequently observed by CXO [396] and NuSTAR [424, 425], though in the latter case, it was consistent in strength and position with a known instrument background feature [426]. This anomaly provided a natural target for x-ray telescopes, and eventually almost all operating instruments would weigh in.

The following years saw a flurry of competing detection and exclusion claims, far too many to list here (see Refs. [427, 428] for recent reviews). In addition to the intrinsic faintness of the excess, there were several additional aspects that complicated matters. First, uncertainties in the DM density profile can substantially modify the inferred mixing-angle constraints. Even within the Milky Way, different density profile choices can change the DM constraints by at least tens of percent, to say nothing of the substantially greater uncertainties when more distant galaxies and clusters (for which kinematic data may be sparse) are considered. Second, the faintness of the line means that small variations in the treatment of the instrument background can have large effects on the projected sensitivity. In particular, Ref. [388] generated substantial discussion owing to their use of novel spline-based background-modeling techniques in the context of XMM-Newton [429–431], though follow-up work obtained similar constraints using a more traditional parametric background model [390]. In any case, constraints from a variety of telescopes searching for the line in a variety of astrophysical targets are necessary to fully understand the nature of this feature.

3.4 The NuSTAR observatory

Launched in 2012, the Nuclear Spectroscopic Telescope Array (NuSTAR) observatory is the first high-energy focusing telescope in orbit [432]. Before discussing NuSTAR's contributions to DM searches (and my contributions to those efforts [1–

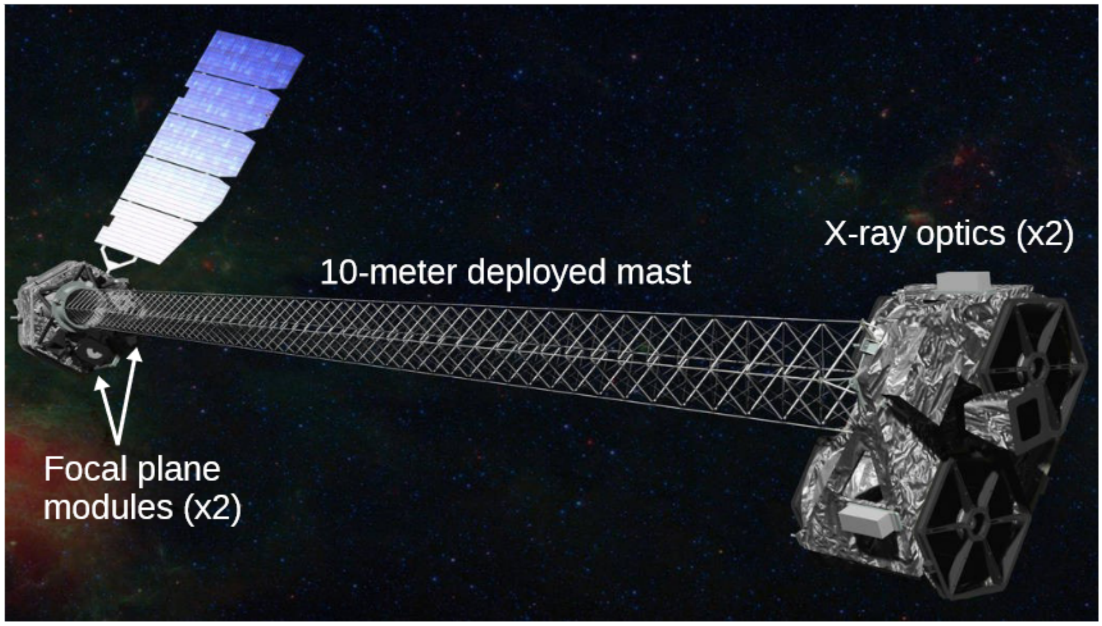

Figure 3-2: Artist's impression of the NuSTAR observatory in orbit, with major components labeled. © by NASA/JPL-Caltech.

3]), it will be helpful to describe the design of the observatory itself. A diagram of the observatory is shown in Fig. 3-2.

3.4.1 X-ray optics modules

The NuSTAR optics modules (OMs) contain the mirrors which focus x-rays onto the focal-plane detectors. The mirrors are conical approximations to the Wolter-I design, consisting of hyperbolic and parabolic segments. Incoming photons with energies $E_\gamma = 3\text{–}79\,\text{keV}$ reflect twice within the OMs, once off the hyperbolic segments, and once off the parabolic segments, earning these focused x-rays the name "2-bounce" (2b) photons.

Focusing x-ray telescopes rely on the properties of photons at *grazing incidence* (since x-ray photons hitting mirrors at *normal incidence* tend to go straight through). Defining ϕ to be the angle tangent to the surface, Snell's law relates the angles of incidence and refraction (ϕ_i and ϕ_r, respectively) [354]:

$$\frac{\cos\phi_\mathrm{r}}{\cos\phi_\mathrm{i}} = \frac{1}{1-\delta}. \tag{3.17}$$

For angles $\phi \ll 1$, there is a critical angle $\phi_\mathrm{crit} \simeq \sqrt{2\delta}$ above which there is no refracted wave and total reflection occurs. For a material with electron number density n_e, mass density ρ, atomic number Z, and mass number A, the optical constant δ is a decreasing function of photon energy E_γ [354]:

$$\delta = \frac{r_e n_e}{2\pi}\left(\frac{hc}{E_\gamma}\right)^2 \approx 0.0025\left(\frac{Z}{A}\right)\left(\frac{\rho}{1\,\mathrm{g\,cm^{-3}}}\right)\left(\frac{1\,\mathrm{keV}}{E_\gamma}\right)^2. \tag{3.18}$$

Since $\phi_\mathrm{c} \propto \sqrt{Z\rho}/E_\gamma$, large values of Z and/or ρ are needed to ensure reasonable reflection angles, particularly at high energies. For example, the CXO and XMM-Newton telescopes use mirrors coated with iridium and gold, respectively. Since the mirrors are only reflective for grazing-incidence x-rays, a single mirror with physical area A_0 presents a very small effective area $A_\mathrm{eff} \sim A_0\phi$. To increase the collecting area, x-ray telescopes generally use many nested concentric mirrors.

NuSTAR's mirrors required technological innovations well beyond those of previous telescopes. For one, NuSTAR's maximum energy of 79 keV is nearly an order of magnitude above that of CXO and XMM. To counteract the steep decrease in x-ray reflectivity with increasing energy while keeping the optics compact and lightweight, the 133 concentric NuSTAR mirrors are composed of novel depth-graded multilayers of varying high- and low-density materials—i.e., platinum/silicon carbide (Pt/SiC) and tungsten/silicon (W/Si) deposited onto glass bases [433]. Each mirror segment contains $\sim$200 multilayer pairs, and the constructive interference they generate greatly enhances the x-ray reflectivity at high energies for angles $\phi > \phi_\mathrm{c}$. The combined on-axis effective area of the two NuSTAR OMs is $\sim 10^3\,\mathrm{cm}^2$ at 10 keV, giving a peak FOV-averaged grasp of $\sim 20\,\mathrm{cm}^2\,\mathrm{deg}^2$ for 2-bounce photons.

3.4.2 Focal plane modules

At the opposite end of the 10-meter carbon-fiber mast from the OMs are the focal plane modules (FPMs). The FPMs consist of three principal components: a beryllium window, a solid-state detector array, and an anticoincidence detector.

Before entering the detectors, x-ray photons must first pass through ∼100-μm-thick beryllium (Be) window, one per FPM. These windows serve two purposes. First, they block lower-energy optical and ultraviolet photons from striking the detectors while allowing $E_\gamma \gtrsim 3\,\text{keV}$ photons to pass through. The transmission efficiency $\mathcal{E}_{\text{Be}}$ of the windows reaches ∼70% by 3 keV and >95% by 8 keV. (The opacity of the windows to optical photons means that NuSTAR is capable of observing the Sun, a capability other observatories (e.g., CXO and XMM-Newton) lack.) Second, the windows shield the detectors against small micrometeoroid impacts.

At the heart of the NuSTAR FPMs are the solid-state cadmium zinc telluride (CdZnTe or CZT) hybrid arrays sensitive to photons with energies ∼3–160 keV with energy resolution ∼0.4–0.9 keV FWHM [434, 435]. Each FPM array contains four detector chips, with each chip (labeled DET0–DET3) consisting of a single $(20 \times 20 \times 2)\,\text{mm}^3$ CZT crystal. A monolithic platinum contact defines the cathode, whereas the platinum anode of each chip is segmented into a 32 × 32 grid of pixels (i.e., each FPM array contains 1024 pixels). Each pixel is equipped with its own signal-processing electronics (preamplifier, shaper, discriminator, etc), and a custom ASIC handles the analog-to-digital conversion[14] and digital readout. Alhough the ASIC is primarily set up to detect positive (electron-generated[15]) pulses, it is also sensitive to negative (hole-generated) pulses, allowing the software to determine the depth of interaction (internal detector backgrounds generally occurring deeper than astrophysical x-rays). The use of CZT presents several advantages for NuSTAR compared to previous silicon- or germanium-based detectors. First, CZT offers a substantially greater quantum efficiency for detecting high-energy pho-

---

[14] NuSTAR carries 12-bit ADCs, giving 4096 0.04-keV-wide bins in the energy range 1.6–165.4 keV.

[15] Since electrons have greater mobility than holes in CZT.

tons compared to silicon and germanium, owing to its high atomic number[16] ($\sim$50 for CZT versus 14 for silicon and 32 for germanium). Second, CZT is operable near room temperature with acceptable energy resolution and leakage current owing to its large band gap ($\gtrsim 1.5\,\mathrm{eV}$), whereas silicon and germanium x-ray spectrometers generally must be cooled to approximately 170 K and 77 K, respectively (to say nothing of cryogenic calorimeters at the $\lesssim 1\,\mathrm{K}$ level). Finally, CZT is much easier to segment into pixels than germanium, allowing for high-resolution imaging.

Finally, the CZT detector arrays are surrounded by anticoincidence shields. These shields consist of cesium iodide (CsI) crystals which fluoresce when struck by ionizing radiation (particularly protons or Compton-scattered γ-rays). Events in the CZT which coincide in time with energy depositions in the CsI are rejected, thereby decreasing the total instrument background. The event rate registered by the anticoincidence shields is crucial for defining the good time intervals (GTIs), ensuring that the data are not contaminated by solar flares, geomagnetic storms, or observatory passages through the South Atlantic Anomaly (SAA, a region of intense charged-particle flux in low-Earth orbit). Somewhat ironically, the CsI itself generates x-ray fluorescence and activation lines as a consequence of the space radiation environment, producing a series of x-ray lines near 30 keV.

(Each FPM is also equipped with ^{155}Eu calibration sources with activity $\sim$10 μCi at launch. The sources are located on retractable arms that can extend into the optical path, uniformly illuminating the detector arrays with x-ray lines at energies between 5–90 keV [436, 437]. To avoid the possibility of the source becoming "stuck" in front of the detectors, the arm has not been deployed since 2015. Energy-scale calibration of the detectors instead relies on emission lines from the Kepler supernova remnant and detector fluorescence lines [438].)

---

[16] The photoelectric absorption cross section scales with atomic number approximately as Z^4.

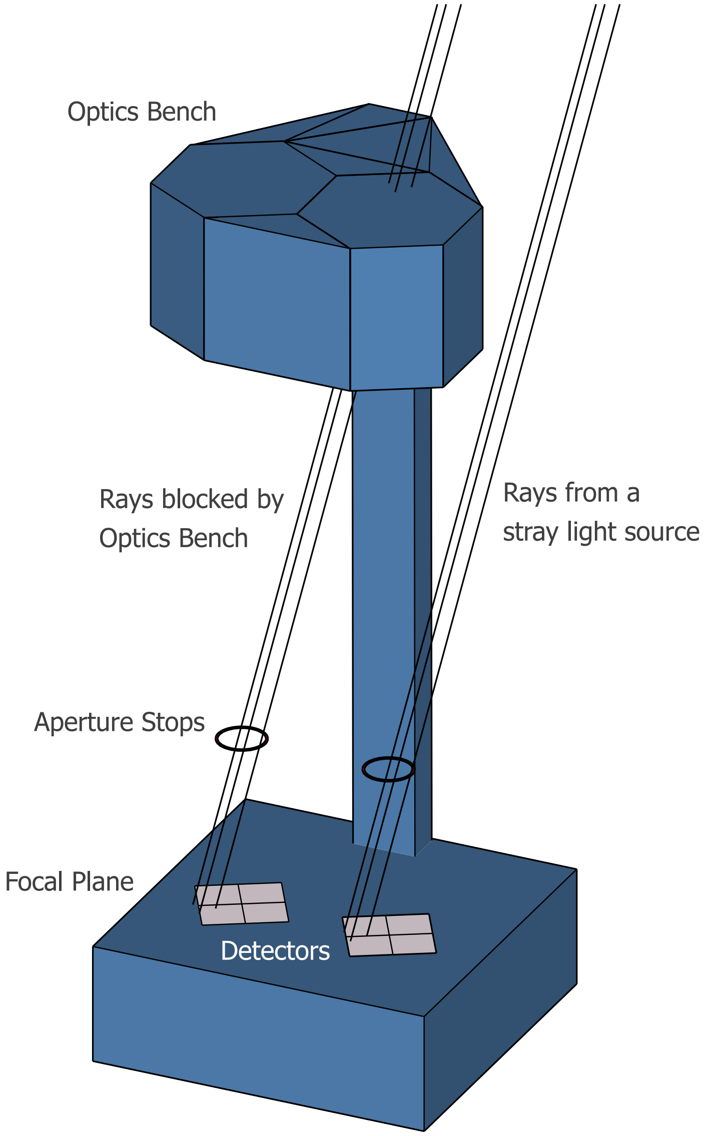

Figure 3-3: A schematic (not to scale) of the NuSTAR instrument design, illustrating the 0-bounce concept. The circular aperture stops above each FPM define regions of $\sim$3.5° radius on the sky (some of which is blocked by the optics bench) from which 0-bounce photons can strike the detectors. Figure reproduced from Ref. [425] with permission. © 2017 by the American Physical Society.

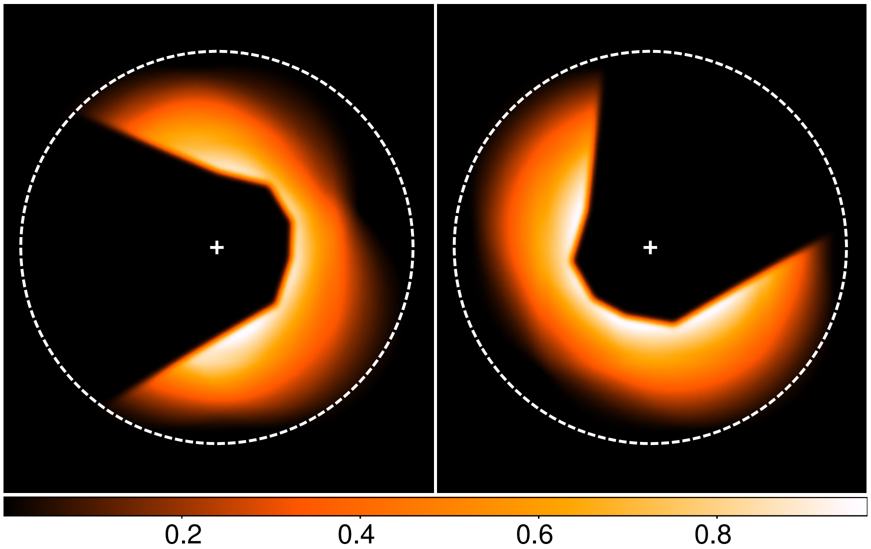

Figure 3-4: Regions of the sky from which 0-bounce photons may reach NuSTAR's FPMA (left) and FPMB (right). The white cross indicates the location of the optical axis, and the dashed circle indicates the 3.5° radius of the aperture stops. The colorbar is proportional to the efficiency $\xi_{0\mathrm{b}}$ with which photons from each patch of solid angle can be detected. Figure reproduced from Ref. [425] with permission. © 2017 by the American Physical Society.

3.4.3 0-bounce photons

Unlike previous telescopes like CXO and XMM-Newton, the optical path between the NuSTAR optics and detectors is largely unshielded and open to the sky (Fig. 3-3). Thus, off-axis x-rays may reach the detectors without interacting with the focusing optics, hence their alternative name of "0-bounce" (0b) photons. The regions of sky from which 0-bounce photons may strike each FPM array are shown in Fig. 3-4. The outer radius of $\sim$3.5° is set by the geometry of the aperture stops in front of the detectors, with the optics bench blocking approximately half of the circle to form a crescent or "Pac-Man" shape. Since each NuSTAR detector pixel subtends a slightly different solid angle on the sky, we define an effective solid angle

$$\Delta\Omega_{0\mathrm{b}} = \int_{\mathrm{FOV}} \xi_{0\mathrm{b}} \, \mathrm{d}\Omega, \tag{3.19}$$

where ξ_{0b} is the position-dependent efficiency of detecting a 0-bounce photon in a given element of solid angle. Each FPM array subtends an effective solid angle $\Delta\Omega_{0b} \approx 4.5\,\mathrm{deg}^2$, nearly two orders of magnitude larger than the $(13\,\mathrm{arcmin})^2 \approx 0.05\,\mathrm{deg}^2$ solid angle for 2-bounce photons. If a bright x-ray source lands within the 0-bounce FOV of a given NUSTAR observation, large regions of the detector plane may be covered by "stray light" from that source. Even if there are no such bright sources, every NuSTAR observation contains an irreducible contribution from the diffuse cosmic x-ray background (CXB) coming through both the 0-bounce and 2-bounce apertures. As discussed previously, the number of photons received from diffuse sources scales with the telescope's grasp $A\Delta\Omega$ averaged over the FOV. The 0-bounce grasp for each FPM array is $12.5\,\mathrm{cm}^2 \times 4.5\,\mathrm{deg}^2 \approx 56\,\mathrm{cm}^2\,\mathrm{deg}^2$, compared to the maximum 2-bounce grasp $\sim$10 $\mathrm{cm}^2\,\mathrm{deg}^2$. For analyses of x-ray point sources, this stray light is often a hindrance; however, as we shall see, these unfocused x-rays provide NuSTAR unparalleled sensitivity to sterile-neutrino DM in certain regions of the parameter space. Furthermore, since the 0-bounce sensitivity is not constrained by the reflectivity of the optics, these analyses can use the full 3–160-keV energy range of the detectors, rather than the 3–79-keV range of the optics.

3.5 NuSTAR sterile-neutrino searches

To date, there have been six published NuSTAR analyses searching for sterile-neutrino DM. I will discuss three previous NuSTAR analyses (Bullet cluster, blank-sky, and Galactic center) before discussing my contributions to analyses of the M31 galaxy, the Milky Way bulge, and the Milky Way halo. The first three constraints are plotted in Fig. 3-5.

3.5.1 Previous NuSTAR analyses

The first NuSTAR search for DM used observations of the Bullet cluster [440]. This analysis differs from the rest in that it used only 2-bounce photons, since the angu-

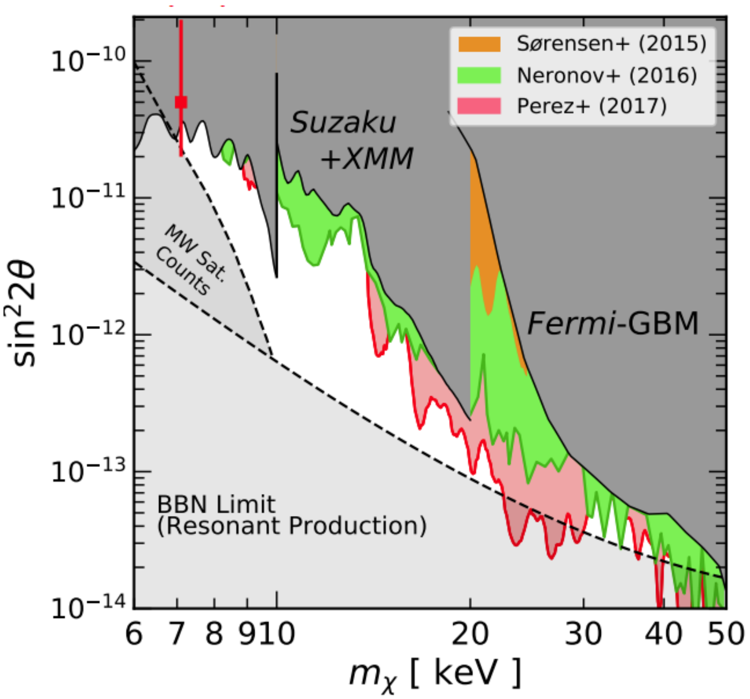

Figure 3-5: Constraints on the νMSM parameter space as of 2017, illustrating the unique capabilities of NuSTAR. The dark gray shaded regions at the top indicate constraints from previous x-ray telescopes (see, e.g., Refs. [366, 376, 395, 404, 439]). The Milky Way satellite counts and BBN constraints are from Ref. [333], assuming $N_{\text{sat}} = 47$ and $L_6 = (n_\nu - n_{\bar{\nu}})/s < 2500$, respectively. For additional information on the latter two constraints, see Chapter 2. NuSTAR constraints are as follows: Sørensen+(2015) [440], Neronov+(2016) [441], and Perez+(2017) [425]. The red 3.5-keV point ($m_\chi \approx 7\,\text{keV}$) is from Ref. [387]. Figure generated from code supplied by K. C. Y. Ng in the style of Ref. [425].

lar extent of the cluster is comparable to NuSTAR's focused FOV. The Bullet cluster is an attractive candidate for x-ray DM searches, owing to its large DM content and the spatial separation of the x-ray emitting gas from the densest regions of DM. Unfortunately, the use of Bullet cluster data comes with several disadvantages. First, the Bullet cluster is bright in thermal x-ray emission, creating a background component that must be modeled or subtracted; in either case, this astrophysical background weakens the DM constraints. Second, the redshift $z \sim 0.3$ of the Bullet cluster means that the $E_\gamma = 3\,\text{keV}$ lower limit of the NuSTAR detectors corresponds to a sterile-neutrino mass limit $m_\chi \sim 7.8\,\text{keV}$. Despite these challenges, the NuSTAR Bullet cluster constraints are still able to exclude some of the "window" in the νMSM parameter space between $m_\chi \sim 10$–$20\,\text{keV}$.

In 2016, two groups realized that the 0-bounce photons landing on the NuSTAR detectors could substantially increase its sensitivity to decaying DM. The first was based on stacked blank-sky observations of the COSMOS and ECDFS extragalactic survey fields [442], whereas the second used survey data of the MW Galactic Center region [425]. In many respects, these analyses are complementary, and derive similar constraints on the νMSM parameter space. The blank-sky observations had longer exposure time and lower astrophysical background but lower DM density, whereas the GC observations had higher DM density but shorter exposure time and higher astrophysical backgrounds. In particular, the GC analysis was limited by line and continuum emission from the Galactic ridge (GRXE), believed to result from accreting white dwarfs (see, e.g., Refs. [443–446]), with the 6.4-keV line from ionized iron constituting a major irreducible background. Despite these challenges, both analyses were instrumental in beginning to close the sensitivity gap between $m_\chi{\sim}10$–$30\,\text{keV}$ shown in Fig. 3-5.

3.5.2 M31 galaxy

At a distance ${\sim}780\,\text{kpc}$, the M31 galaxy is the nearest large galaxy to the Milky Way. Previous works used observations from XMM-Newton [385] and CXO [395] searched for sterile neutrinos with masses $m_\chi \lesssim 20\,\mathrm{keV}$. In our NuSTAR analysis [1], we present the first hard x-ray constraints from M31, searching for sterile-neutrino decays in the mass range 10–220 keV (energy range 5–110 keV).

Our NuSTAR analysis is based on eight epochs of M31 observations between 2012–2016, with each FPM recording $\sim 0.6\,\mathrm{Ms}$ exposure following SAA and solar-flare filtering. This dataset is a small subset of NuSTAR's M31 data, owing to several selection cuts. To avoid the intense x-ray backgrounds near the central regions of M31, we select observations including at most two bright (0.35–8-keV flux above 2×10^{-13} erg cm^{-2} s^{-1}) resolved x-ray point sources. This leaves two regions of M31, one of which was observed for two epochs and the other for six. We remove all detector pixels within a $60''$ radius around CXO 004429.57+412135.1 and CXO 004527.34+413253.5, and a $100''$ radius around CXO 004545.57+413941.5. We then extract x-ray spectra and the associated response files (RMFs and 2-bounce ARFs) from the full detector planes for each observation using NUSTARDAS V1.5.1. This produces 16 spectra (eight epochs, each of which had data from both FPMA and FPMB). We use the public NUSKYBGD tool [426] to calculate the corresponding 0-bounce effective areas $A_{0\mathrm{b}}$ and solid angles $\Delta\Omega_{0\mathrm{b}}$ for both FPMs, accounting for dead pixels and masked point sources. As shown in Fig. 3-6, the 0-bounce FOVs largely avoid x-ray emission from the visible disk of M31.

We separately fit each x-ray spectrum to a background model based on Ref. [426] in the energy range 5–110 keV (owing to systematic difficulties with the background model in the 3–5 keV energy range described in Ref. [425]). This model consists of four main components: a broken powerlaw continuum representing Compton scattering within the telescope (with fixed spectral indices), a forest of Lorentzian fluorescence and activation lines (with fixed position and width), a thermal-plasma APEC model[17] representing the Solar background (with fixed temperature), and a powerlaw spectrum representing the diffuse cosmic x-ray back-

---

[17] We also consider replacing the APEC model with a powerlaw whose spectral index is fixed for each spectrum using Earth-occulted data (i.e., when the telescope is behind the Earth and no astrophysical x-rays are expected). We find no significant difference in the fit quality of our $E_\gamma > 5$ keV dataset between these two models.

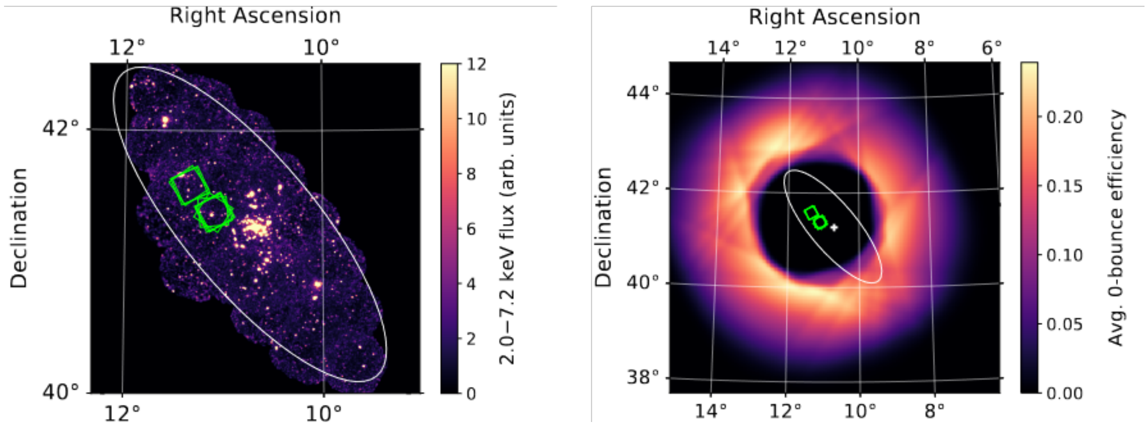

Figure 3-6: **(Left)** Mosaic of XMM-Newton observations of M31 [447, 448], showing the 2-bounce FOVs of the NuSTAR observations (green squares) and the outline of the M31 disk (gray ellipse). **(Right)** Zoomed-out version of the previous, where the colorbar now shows the average 0-bounce coverage from both FPMs and all eight observations. Reprinted with permission from Ref. [1]. © 2019 by the American Physical Society.

ground (CXB, with fixed normalization and spectral index). To account for diffuse x-ray emission within the M31 disk, we also include a powerlaw model with spectral index and normalization both free to fit. (Since the 0-bounce FOVs avoid the M31 disk, we do not include a 0-bounce M31 emission component.) We define custom 2-bounce response files for each observation by multiplying the NUSTARDAS-extracted ARFs[18] by the corresponding solid angle ($\sim$0.047 deg^2 minus the areas removed around the point sources). Similarly, we derive 0-bounce response files for each observation by multiplying the beryllium window transmission function by $A_{0\mathrm{b}}$ and $\Delta\Omega_{0\mathrm{b}}$. We assign these custom respose files to the 0-bounce CXB, 2-bounce CXB, and 2-bounce M31 components, as appropriate. We also include the attenuation of the CXB, M31, and DM emission in the interstellar medium by means of the TBABS model in XSPEC, though the equivalent optical depth $\tau \lesssim 10^{-2}$ is not significant in the energy range of our analysis [449]. (The CXB, M31, and DM emission are also attenuated by the platinum contact coating and layer of inactive CZT on each detector, which we incorporate by means of the NUABS model.) An

---

[18] These files already include the beryllium-window transmission efficiency.

example spectral fit is shown in Fig. A-1.

Our DM search strategy is based on that described in Sec. 3.2, wherein we grid the mass range 10–220 keV into 100 evenly-spaced bins per decade in $\log_{10} m_\chi$ and search each for evidence of a DM decay line at each value of $E = \frac{1}{2} m_\chi$ using the profile likelihood. The DM signal in these NuSTAR observations contains contributions from both M31 and the MW, in both the 0-bounce and 2-bounce apertures. Since M31 is far from the MW Galactic center, the choice of MW DM profile does not contribute significantly to the error budget, and we adopt the shallow NFW profile discussed in Ref. [425]. We model the M31 DM distribution with the NFW profile discussed in Ref. [450], with scale radius $r_s = 16.5\,\mathrm{kpc}$, virial radius $R_{200} = 207\,\mathrm{kpc}$, and scale density $\rho_\mathrm{s} = 0.418\,\mathrm{GeV\,cm^{-3}}$. After taking the geometry of the 0-bounce and 2-bounce FOVs into account and averaging over all observations, we find that the 0-bounce MW component contributes $\sim$50% of the expected signal, the 0-bounce M31 component $\sim$30%, the 2-bounce MW component $\sim$5%, and the 2-bounce M31 component $\sim$15%. To account for the additional sensitivity afforded by the 2-bounce aperture, we scale our 0-bounce DM constraints by the *boost factor*

$$f_\mathrm{boost}(E_\gamma) = 1 + \frac{A_\mathrm{2b}(E_\gamma)}{A_\mathrm{0b}} \frac{\Delta\Omega_\mathrm{2b}}{\Delta\Omega_\mathrm{0b}} \frac{D_\mathrm{2b}}{D_\mathrm{0b}} \tag{3.20}$$

where D is the FOV-averaged DM column density.

Unlike previous NuSTAR DM searches (e.g., Refs. [424, 425] which stack x-ray spectra from different regions of the sky into a single spectrum, our M31 analysis derives constraints from each individual spectrum, which are then statistically combined. Specifically, we separately fit each of the 16 x-ray spectra and search for evidence of DM, collecting likelihood curves $\chi_i^2(\theta; m_\chi)$ where i labels the spectra. We then construct the quantity

$$X^2(\theta; m_\chi) = \sum_i \chi_i^2(\theta; m_\chi) \tag{3.21}$$

by summing over all of the individual observations. The corresponding test statistics for detection and upper limits are obtained simply by replacing χ^2 with X^2 in

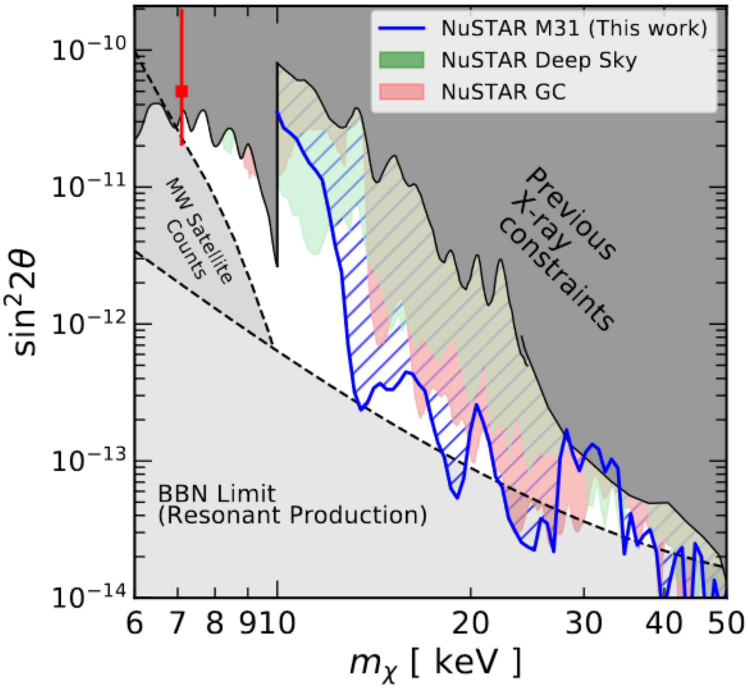

Figure 3-7: Constraints on sterile-neutrino DM in the νMSM from NuSTAR observations of M31 (blue hatched region, Ref. [1]), deep sky fields (green shaded region, Ref. [442]), and the Galactic Center (red shaded region, ref. [425]). The dark gray shaded regions indicate previous x-ray constraints from CXO [395], NuSTAR Bullet cluster [440], INTEGRAL [366], and Fermi-GBM [376] observations. The BBN and MW satellite limits are the same as Fig. 2-2. For further details, see Ref. [1], from which this figure has been reproduced with permission. © 2019 by the American Physical Society.

Eqs. 3.11–3.12, where the best-fit signal strength $\hat{\theta}$ is now interpreted as the joint best-fit value accounting for all of the spectra. We compare these results to those derived from a more "traditional" analysis in which we stack the x-ray spectra from the different fields into a single spectrum, and find that the two approaches give very similar constraints. In both cases, we find no evidence for DM emission lines.

Compared to previous analysis techniques, our statistical combination method provides substantial gains in flexibility. First, it allows for the possibility that different spectra may have very different x-ray background models. For example, it would be extremely difficult to stack spectra from the Milky Way Galactic center with spectra from M31 and expect the resulting spectrum to be described by a single background model. Our statistical combination method handles this case easily. Another approach would be to fit all of the individual spectra simultaneously, introducing a DM model with the same signal strength in all of them and recording the combined $\chi^2 = \sum_i (y_i - m_i)/\sigma_i^2$, where i now indexes each individual data point across all of the spectra. This would allow each spectrum to have its own background model and would provide equivalent results to the statistical-combination approach (since a product of likelihoods is equivalent to the sum of log-likelihoods, and hence the sum of χ^2), but at the cost of computational resources required to simultaneously fit all of the spectra. Furthermore, the simultaneous-fitting approach is not persistent—if the analyst wished to add a new observation to the analysis, it would be necessary to refit everything. On the other hand, our statistical-combination approach produces one file of $\chi_i^2(\theta)$ for each m_χ, so incorporating constraints from new observations simply requires calculating $\chi^2(\theta; m_\chi)$ for that observation (using the same grid of m_χ values) and adding it to Eq. 3.21.

As shown in Fig. 3-7, our NuSTAR M31 sterile-neutrino DM constraints represent a significant improvement over previous analyses. In particular, the lack of bright x-ray line emission from the Galactic ridge significantly strengthens our DM constraints in the mass range ∼10–20 keV. Unfortunately, the x-ray continuum emission from M31 still weakens our constraints near $m_\chi \sim 10\,\mathrm{keV}$.

3.5.3 Milky Way Galactic bulge

For our next analysis, we return to the central region of the Milky Way, where the expected signal from DM decay is greatest. A full accounting of this analysis can be found in Ref. [2], and proceeds very similarly to the M31 case described previously.

To reduce the x-ray backgrounds from bright point sources as well as diffuse

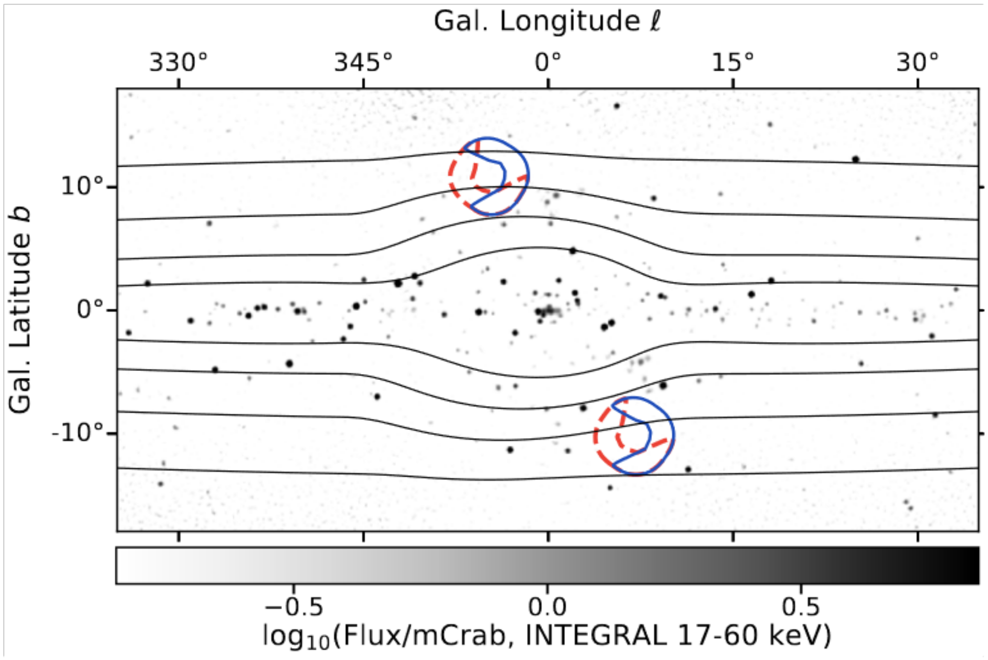

Figure 3-8: The MW Galactic bulge region in hard x-rays. The dashed red and solid blue "Pac-Man" contours indicate the 0-bounce FOVs of NuSTAR's FPMA and FPMB, respectively. The thin black contours indicate the expected GRXE flux (proportional to infrared surface brightness from stars), which decreases away from the Galactic plane [451, 452]. The colorbar indicates the 17–60 keV flux measured by INTEGRAL, showing many bright x-ray point sources along the Galactic plane [453]. Reproduced from Ref. [2] with permission. © 2020 by the American Physical Society.

emission from the Galactic ridge, we requested two dedicated NuSTAR observations in the Galactic bulge region shown in Fig. 3-8, centered at Galactic coordinates $(355.2^\circ, +10.9^\circ)$ and $(7.0^\circ, -10.2^\circ)$. These fields were chosen to remain close to the Galactic center (i.e., where the DM is densest) while their high latitudes place them outside most of the Galactic ridge emission. In one of the observations, the RS CVn-type stellar binary HD 152178 lands within the 2-bounce FOV, so we exclude all detector pixels within a $75''$ radius around that source's Gaia coordinates. We then extract spectra from the full detector planes using NUSTARDAS and define custom 0-bounce and 2-bounce response files as described previously, accounting for the effective area and solid angle of each aperture.

Owing to the low level of astrophyical x-ray emission in these fields, it is important to define the background model with care. We parametrize the CXB with the spectral model of Ref. [454], freezing the spectral index but allowing the normalization to float by ±10% around its nominal value measured by INTEGRAL [455] (though we require the 0-bounce and 2-bounce CXB components to have the same incident flux per solid angle). Since we find that the APEC solar model provides a poor fit to the 5–10-keV data, we follow the prescription of Refs. [1, 446] and replace the APEC model with a simple powerlaw with a fixed spectral index Γ and free normalization. (We determine Γ for each observation by fitting the Earth-occulted spectra to a model containing only the internal detector components—i.e., no CXB and GRXE—allowing Γ to vary. Subsequently, we freeze the value of Γ to its best-fit value in each of the four individual occulted-mode spectra.) We model the GRXE with an APEC model with fixed $kT = 8\,\text{keV}$ [446]. Since the GRXE appears in both the 0-bounce and 2-bounce FOVs of each observation, we require both apertures to have the same metallicity and incident flux per solid angle. In all spectra, the event rate from the GRXE is approximately one order of magnitude below that of the CXB and internal detector background. Since the NuSTAR sensitivity in the energy range 20–95 keV is limited by the forest of instrumental lines (and because INTEGRAL, Fermi-GBM, and NuSTAR have already strongly constrained those energy ranges), we only use data in the energy ranges 5–20 keV and 95–110 keV, the latter being necessary to constrain the normalization of the internal continuum. To account for the smaller exposure time, each spectrum is binned to 100 bins per decade in $\log_{10} E_\gamma$, providing ∼10% statistical uncertainty per bin. The fit quality is somewhat worse than the M31 analysis, ranging from $\chi^2/60 \approx 1.0$ to 1.39 depending on the spectrum, suggesting that the low astrophysical x-ray emission is revealing systematic issues with the NuSTAR background model.

To search for evidence of DM decay lines, we follow the approach developed in our M31 analysis: introduce a DM emission line at fixed energy $E_i = \frac{1}{2} m_i$, convolve the line with the 0-bounce instrument response, and extract likelihood curves $\chi^2(\theta; m_i)$ from each spectrum separately. We consider several DM density

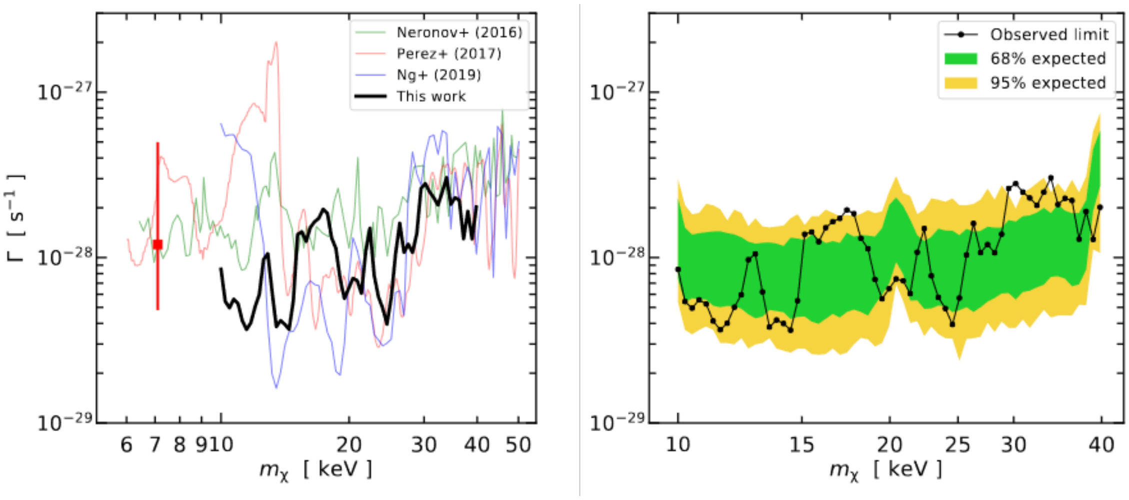

Figure 3-9: **(Left)** Constraints on the single-photon decay rate $\Gamma_{\chi\to\gamma x}$ versus DM mass m_χ for several NuSTAR analyses ("Neronov+" from Ref. [442], "Perez+" from Ref. [425], "Ng+" from Ref. [1]). The thick black line labeled "this work" corresponds to the Galactic bulge analysis of Ref. [2]. **(Right)** Same quantity as the left, comparing the observed limit (black) to the expected 68% (green) and 95% (yellow) containment bands derived from simulations. Reprinted with permission from Ref. [2]. © 2020 by the American Physical Society.

profiles: a "traditional" NFW profile (Eq. 1.7), a "cored" NFW profile with constant density within 1.5 kpc, and a shallow NFW profile 1.10 with $\gamma = 0.7$. In all cases, we adopt a scale radius $r_s = 20\,\text{kpc}$ and local DM density $\rho_0 = 0.4\,\text{GeV}\,\text{cm}^{-3}$ at the solar position. We also consider a Burkert profile (Eq. 1.8) with scale radius $r_s = 8\,\text{kpc}$ and local DM density $\rho_0 = 0.5\,\text{GeV}\,\text{cm}^{-3}$. We ultimately adopt the shallow NFW profile as a reasonable "median" choice (the NFW and coreNFW profiles give DM densities ~20% higher, and the Burkert profile is ~10% lower).

To assess whether our derived DM limits are consistent with the expected statistical power of our data, we generate 100 mock realizations of each spectrum using the `fakeit` tool in XSPEC. These mock spectra are based on the best-fit no-DM spectral models with Poisson noise corresponding to the exposure time. We then pass these mock spectra through the spectral-fitting and DM-limit-setting machinery described previously, resulting in an ensemble of 100 expected upper limits on the mixing angle θ at each sterile-neutrino mass. From this, we construct the 68%

and 95% containment bands around our expected limit, and find that our data-derived limit is consistent with expectations. We also consider the effect of adding a flat 7.5% systematic uncertainty to each bin, sufficient to bring $\chi^2/60 \lesssim 1$ for all spectra. Even though this is an extremely conservative approach, we find that our eventual DM limits only weaken by $\sim$50% in the mass range 10–12 keV, but are still $\sim$50% more constraining than the previous limits in this mass range [442].

3.5.4 Milky Way halo

For our final NuSTAR sterile-neutrino search [3], we present two analyses based on disjoint datasets, using novel treatments of the NuSTAR instrument background to achieve improved sensitivity compared to previous work.

Spatial-gradient analysis

The NuSTAR analyses described up to this point have been traditional parametric analyses: extract all of the x-ray events from some region of the detector (sky) and fit the resulting spectrum to a model consisting of the various instrumental and astrophysical components. Such analyses use the 0-bounce information in a very "coarse" way—the distinctive intensity pattern shown in Fig. 3-4 is condensed into a single number, the effective solid angle $\Delta\Omega_{\mathrm{0b}}$. Could there be a way of more effectively utilizing the spatial distribution of x-rays on the detector plane?

The NuSTAR analysis of Ref. [456] presents a method for separating 0-bounce photons from internal detector backgrounds on the basis of their intensity pattern on the detectors. In brief, x-rays from diffuse sources (e.g., the CXB or decaying DM) would have "Pac-Man"-shaped intensity patterns on the detectors, whereas the internal detector backgrounds (e.g., Compton scattering or activation/fluorescence lines) have an approximately flat spatial distribution. The total number of counts N_{pix} in the ith pixel recorded during exposure time T in the vicinity of energy E_γ is given by [3, 456]

$$N_{\mathrm{pix},i}(E_\gamma) = \left[M_{\mathrm{int}} \left(\frac{\mathrm{d}N}{\mathrm{d}t}\right)_{\mathrm{int}} + R_{\mathrm{pix}} A \Omega \mathcal{E}_{\mathrm{tot}} \left(\frac{\mathrm{d}F}{\mathrm{d}\Omega}\right)_{0\mathrm{b}}\right]_i T. \tag{3.22}$$

The internal detector background rate $(\mathrm{d}N/\mathrm{d}t)_{\mathrm{int}}$ is expected to be approximately uniform within each of the eight detector chips, with variations between detectors accounted for in the response M_{int} (taken from the NuSTAR CALDB). The 0-bounce flux per solid angle $(\mathrm{d}F/\mathrm{d}\Omega)_{0\mathrm{b}}$ (here assumed to be constant across the FOV) is modulated by the physical pixel area $A_i = (0.6\,\mathrm{mm})^2$, the solid angle Ω_i calculated by NUSKYBGD, the energy-dependent transmission efficiency $\mathcal{E}_{\mathrm{tot}}$ (incorporating the beryllium shield and detector surface layers), and the pixel response R_{pix} packaged in the NuSTAR CALDB. (We do not include a 2-bounce spatial template in Eq. 3.22 since the expected rate of 2-bounce CXB photons is at or below the level of the internal detector background.)

Since all parameters except $(\mathrm{d}N/\mathrm{d}t)_{\mathrm{int}}$ and $(\mathrm{d}F/\mathrm{d}\Omega)$ are known, we construct the Poisson likelihood

$$\mathcal{L} = \prod_i \left[\frac{N_{\mathrm{pix}}^n \exp\left[-N_{\mathrm{pix}}\right]}{n!}\right]_i \tag{3.23}$$

in each energy bin, where n_i is the observed number of counts in the ith pixel (here i runs over all pixels in all of the observations). By minimizing $-2\ln\mathcal{L}$ in each of the 100 energy bins between 3–20 keV with respect to $(\mathrm{d}F/\mathrm{d}\Omega)_{0\mathrm{b}}$, we obtain 0-bounce spectra and response files for FPMA and FPMB. By applying this technique to data from the NuSTAR extragalactic survey fields (in which the contribution from astrophysical x-ray sources is negligible), we obtain spectra for FPMA and FPMB with exposure time $\sim$7 Ms per FPM. These spectra are the same as those presented in Ref. [456], though ours have a much narrower binning.

Owing to the purity of the 0-bounce spectra, we only include two components in the background model. The first is the CXB parametrization of Ref. [454],

$$\frac{\mathrm{d}^2F}{\mathrm{d}E_\gamma\mathrm{d}\Omega} = 0.0024\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}\,\mathrm{deg}^{-2}\,\mathrm{keV}^{-1}\left(\frac{E_\gamma}{1\,\mathrm{keV}}\right)^{-\Gamma_{\mathrm{CXB}}}\exp\left[-\frac{E_\gamma}{E_{\mathrm{fold}}}\right] \tag{3.24}$$

where we fix the values of $\Gamma_{\mathrm{CXB}} = 1.29$ and $E_{\mathrm{fold}} = 41.13$ to the values shown to provide good fits to these spectra in Ref. [456]. (We allow the incident CXB flux to vary as a check on the stability of the model, and obtain consistent results between the FPMs.) The second component is a powerlaw of the form $E_\gamma^{-\Gamma_{\mathrm{solar}}}$ to account for residual solar emission making it into the 0-bounce aperture. The normalization and spectral index of the solar component is allowed to vary independently between the FPMs, though Γ_{solar} is constrained to be > 2 to ensure it does not interfere with the CXB. Owing to the low ISM column density in the direction of these fields, we do not include an x-ray attenuation component. The best-fit spectral models for the two FPMs are shown in Fig. 3-10.

To search for evidence of DM in these two spectra, we follow the same procedure discussed previously in Refs. [1, 2]. For each of 200 trial masses m_i in the range 6–40 keV, we introduce a Gaussian line with energy $E_i = \frac{1}{2}m_i$ and construct likelihood curves $\chi^2(\theta; m_i)$. With the exception of two bins[19] near $m_\chi \sim 40\,\mathrm{keV}$, we find no evidence of x-ray lines, and thus set one-sided 95%-confidence upper limits on the model-independent decay rate $\Gamma_{\chi\to\gamma x}$ and the active-sterile mixing angle θ. (Owing to the ~90–110° separation of these fields from the GC, the choice of DM density profile only modifies our constraints by ±10%.) To ensure that our limits are consistent with the expected statistical fluctuations, we generate 10^3 mock spectra per FPM using the `fakeit` tool in XSPEC. Each mock spectrum is based on the best-fit no-DM spectral model with Poisson noise appropriate for the 7-Ms exposure time. We pass these mock spectra through the same data-analysis pipeline as the real data, obtaining 68% and 95% containment bands around the median expected upper limit. With the exception of the two highest-mass bins discussed previously, our results are entirely consistent with the expected statistical fluctu-

---

[19] A DM interpretation for this excess is disfavored by previous NuSTAR analyses [1, 2, 425, 442].

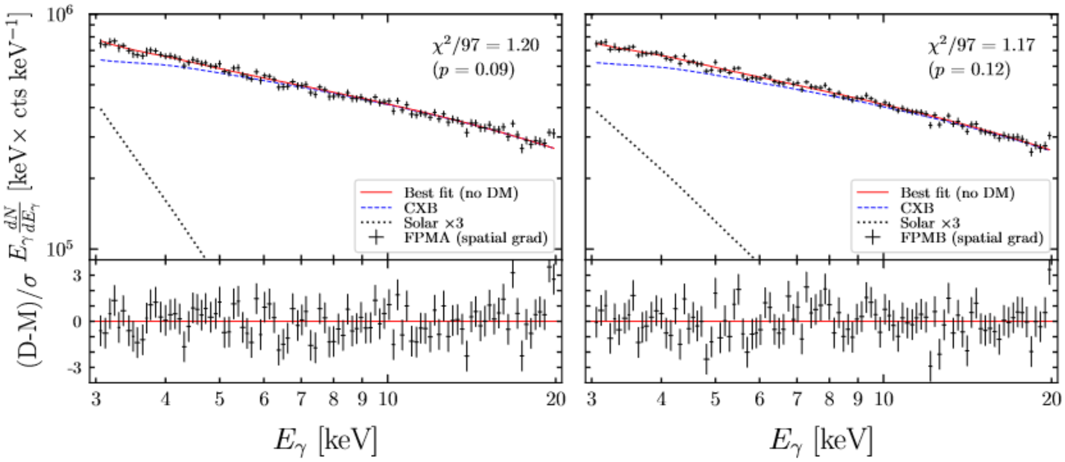

Figure 3-10: Blank-sky NuSTAR spectra from FPMA (left) and FPMB (right) extracted using the spatial-gradient technique described in Sec. 3.5.4. The solid red lines indicate the best-fit no-DM model, including contributions from the 0-bounce CXB (blue dashed) and solar (black dotted) components. The solar component has been multiplied by a factor of three to aid visibility. The bottom panels show the residuals (Data – Model) divided by the bin-wise statistical uncertainty σ. Figure reproduced from Ref. [3] with permission. © 2023 by the American Physical Society.

ations, as shown in Fig. 3-11. The nondetection of the 3.5-keV line in this dataset provides compelling evidence that this feature is indeed part of the NuSTAR internal detector background, rather than an astrophysical emission line.

Parametric analysis

For the second analysis, we exploit NuSTAR's large "back catalogue" of observations scattered throughout the MW halo (more than 2000 observations between 2012–2017). To avoid diffuse and point x-ray sources near the Galactic plane, we exclude all observations with latitude $|b| < 15^\circ$. From this set of $\sim$600 observations, we implement a data-driven process to exclude x-ray point sources and charged particle flares, described fully in Ref. [3]. For each observation, we obtain eight x-ray spectra (one per NuSTAR detector chip) with the associated 0-bounce and 2-bounce response files. To alleviate the computational difficulty of fitting

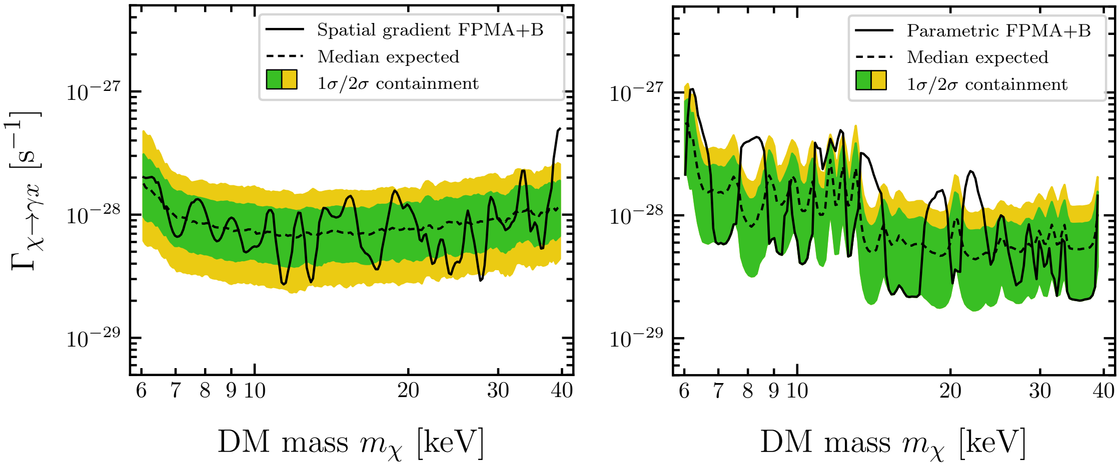

Figure 3-11: **(Left)** Constraints on the two-body DM decay rate to a single photon $\Gamma_{\chi\to\gamma x}$ for the spatial-gradient analysis of Sec. 3.5.4. The black solid line shows the 95%-confidence upper limit, the dashed line shows the median expected limit, and the green (yellow) bands show the 1σ (2σ) containment. **(Right)** Same as previous, for the parametric analysis. The data-derived limit is power-constrained at the -1σ level. For more details, see Ref. [3], from which this figure has been reproduced with permission. © 2023 by the American Physical Society.

all ∼8 × 600 spectra, we create eight stacked x-ray spectra, one per detector chip. We obtain custom 0-bounce and 2-bounce response files in a similar manner to Refs. [1, 2] by constructing the exposure-time-weighted average of the response from each individual observation.

To match the exquisite statistical power of the NuSTAR data (∼20 Ms per spectrum), we update the background model proposed by Ref. [426]. A full accounting of this updated model is found in Ref. [3]; here, we note several improvements from the previous version. First, unlike previous NuSTAR analyses which combine data from each FPM's four detector chips, we construct individual spectral models for each detector to account for their different responses. Second, we model the solar x-rays with a powerlaw plus several Lorentzian lines between 3–7 keV, rather than an APEC. The powerlaw index for each spectrum is fixed to the best-fit value obtained by fitting the corresponding Sun-illuminated Earth-occulted data. We find that this approach provides improved flexibility, since the normalizations

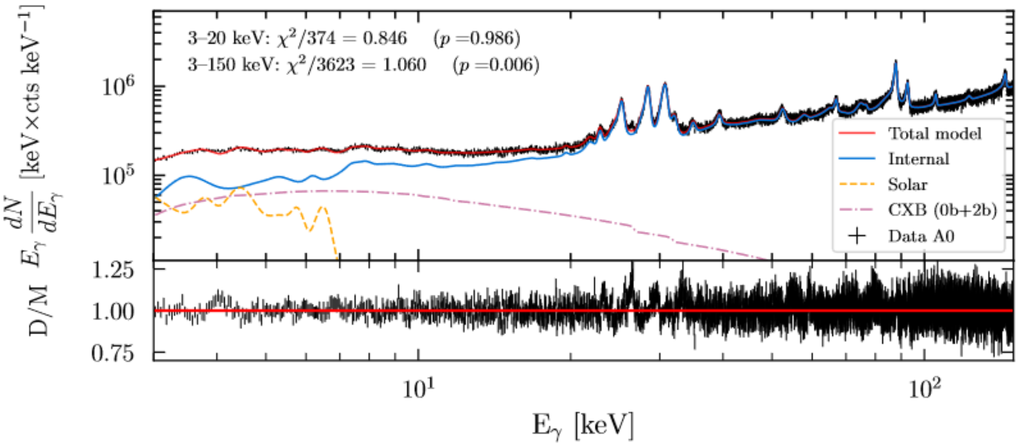

Figure 3-12: Best-fit no-DM model for NuSTAR detector A0 (DETA0) as described in the parametric analysis of Sec. 3.5.4. A flat 2.5% systematic has been added in quadrature to each bin. The bottom panel shows the ratio of data to model (D/M). Corresponding figures for the other seven detectors may be found in Appendix A. Reproduced from Ref. [3] with permission. © 2023 by the American Physical Society.

of the lines are now free to vary independently. Third, we introduce several broad x-ray lines in the energy range 10–20 keV, motivated by the broad "plateau" observed in the spectra of Earth-occulted data that may be a result of x-rays reflected from Earth's atmosphere. The other model components (CXB, internal continuum, etc) are the same as Ref. [426]. A representative spectrum from detector 0 of FPMA (i.e., DETA0) is shown in Fig. 3-12. The fit quality is excellent across most of the 3–160-keV energy range, except in some isolated regions near known x-ray lines. To account for this, we conservatively add a flat 2.5% systematic uncertainty to all bins. Furthermore, we allow the background lines in the 20–95-keV energy range to find their best-fit values in the no-DM hypothesis, but freeze their normalizations when constructing the profile likelihood. This ensures that the low-energy tails of these lines are properly modeled, while reducing the computational complexity of the model at negligible cost to our DM constraints in the 3–20-keV energy range of interest. (The background line normalizations in the 3–20-keV and 95–160-keV energy ranges are free to vary at all stages of our analysis.)

To search these spectra for evidence of DM, we follow the same profile-likelihood procedure as our M31 and Galactic bulge analyses [1, 2], deriving constraints for each of the eight spectra independently and combining the profile-likelihood curves to obtain the final limits. We consider the same DM density profiles as the spatial-gradient analysis discussed above, with different profile choices affecting our final constraints by only $\sim$20%. To assess whether our DM constraints are consistent with statistical fluctuations around the best-fit no-DM model, we use the "Asimov" procedure [358]. Rather than generating many random realizations of each NuSTAR spectrum and scanning each for evidence of DM (which would be computationally expensive, owing to the complexity of the background model), we instead generate a single Monte Carlo spectrum for each NuSTAR detector based on the best-fit background model (including the 2.5% systematic). We pass these eight spectra through the same profile-likelihood pipeline as the real data, obtaining the $\pm N\sigma$ bands around the median expected 95% upper limit where X^2 defined in Eq. 3.21 increases from its minimum by $(1.64 \pm N)^2$ (see, e.g., Ref. [457]). Since the -2σ band is undefined in this procedure, and to ensure our DM constraints are conservative, we power-constrain our DM limits—i.e., we do not allow our DM constraints to become stronger than the (median-1σ) expectation from the Asimov simulations [458]. The joint constraints from all eight spectra are shown in Fig. 3-13.

3.6 Closing thoughts

The search for x-rays from decaying keV-scale sterile-neutrino DM has been ongoing for several decades, involving nearly every major x-ray telescope. Since the early 2000s, the the νMSM has been a compelling target, owing to its ability to account for DM, neutrino masses, and the observed cosmic matter-antimatter asymmetry. With the possible exception of the 3.5-keV anomaly, no putative DM signal has withstood the test of time. Assuming sterile-neutrino DM is produced by mixing with active neutrinos, only a few small openings in the mass-mixing

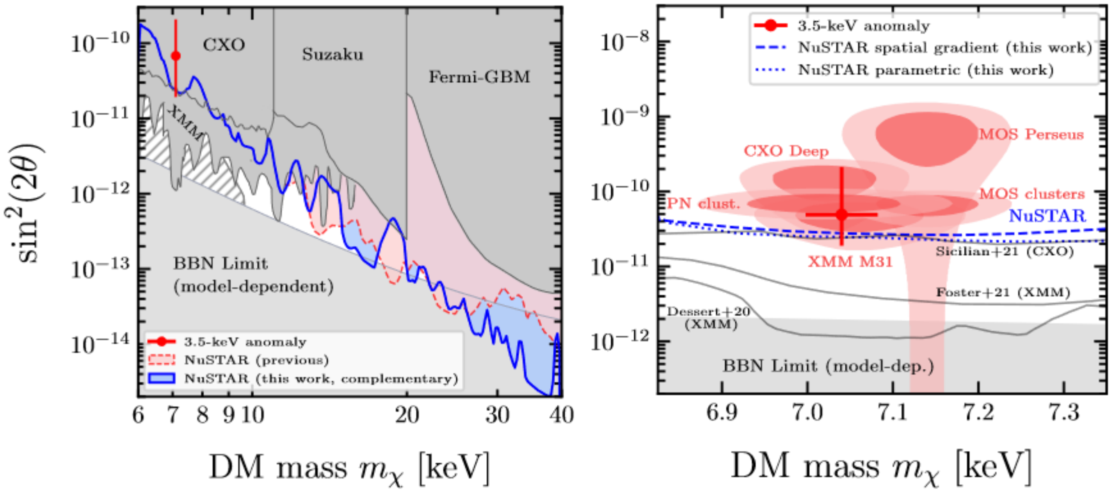

Figure 3-13: **(Left)** Coverage of the νMSM parameter space as of early 2023. The blue NuSTAR constraint comes from the MW halo analysis of Sec. 3.5.4. Previous NuSTAR constraints [1, 3, 425, 442] are shown in the light red shaded region. Dark gray shaded regions indicate other x-ray constraints [376, 390, 397, 401–404]. The red point is the 3.5-keV detection claim of Ref. [387]. **(Right)** Magnified view of the region near $m_\chi = 7\,\text{keV}$. The dark (light) red contours show the 1σ (2σ) detections of the 3.5-keV anomaly as shown in Ref. [297]. For more details, see Ref. [3], from which this figure has been reproduced with permission. © 2023 by the American Physical Society.

angle parameter space remain. Upcoming data from SRG, Micro-X, XRISM, and Athena will be essential for searching for anomalous x-ray lines, and for fully testing the νMSM. Of course, it may be the case that only a fraction of the DM is composed of sterile neutrinos, in which case the parameter space would open up once again. Finally, the νMSM is not the only model capable of producing keV-scale sterile-neutrino DM, with other possibilities including the decay of heavy scalar fields (see, e.g., Ref. [427] for a review of other production mechanisms). In these alternative scenarios, the thermal history of the DM would be quite different, and the BBN and satellite-count constraints may be modified.

Beyond sterile neutrinos, indirect x-ray searches for DM will continue as long as there are x-ray telescopes. Though they are beyond the scope of this dissertation, axions and axion-like particles (ALPs) are another well-motivated class of DM candidates with a coupling to photons. Recent analyses of data from CXO [459–463], XMM-Newton [464], and NuSTAR [465–467] have constrained the $a\gamma\gamma$ coupling in astrophysical magnetic fields, producing constraints complementary to terrestrial experiments [468]. Finally, x-ray telescopes are not limited to keV-scale DM, since the decay or annihilation of heavier (MeV, GeV, or above) particles can produce a significant x-ray flux from inverse-Compton scattering, bremsstrahlung, and final-state radiation (see, e.g., Refs. [469, 470]). In any case, x-rays will remain a powerful tool for indirect DM detection well into the future.

4

Cosmic Rays and Antinuclei

On what can we now place our hopes of solving the many riddles which still exist as to the origin and composition of cosmic rays?

—Victor F. Hess

In astrophysics, *cosmic rays* are energetic charged particles[1] (protons, electrons, heavier nuclei, and their antiparticles) continuously bombarding Earth from space. In Sec. 4.1, I briefly review some of the terminology and phenomenology of these particles. In Sec. 4.2, I motivate the discussion of antinuclei, extremely rare particles that may serve as ideal messengers for DM, with a view towards both terrestrial and cosmic production. In Sec. 4.3, I describe some of the recent cosmic-ray experiments which have contributed to the study of these rare and enigmatic cosmic messengers, motivating the discussion of the GAPS experiment which forms the rest of this dissertation.

---

[1] From here on, I will use "cosmic rays" to refer to protons, nuclei, and their antiparticles, unless otherwise noted (i.e., not including electrons/positrons).

4.1 Introduction to cosmic rays

The field of cosmic-ray physics is vast, and I cannot hope to summarize it in this dissertation. Thus, this section will serve as a high-level overview of the field, motivating some of the crucial concepts and terminology to follow.

4.1.1 Cosmic-ray observables

At Earth, the fundamental observable for cosmic rays is their intensity (flux per energy per solid angle) $\mathcal{I} \equiv \mathrm{d}^2F/\mathrm{d}E\mathrm{d}\Omega$. As shown in Fig. 4-1, the observed cosmic-ray spectra range over many orders of magnitude in energy (for a compilation of data, see, e.g., Ref. [471]). Collectively, these experiments have measured the spectra of cosmic-ray nuclei ranging from protons ($Z = 1$) to nickel ($Z = 28$). The spectra of cosmic-ray nuclei generally follow a powerlaw of the form $\mathcal{I} \propto E^{-2.7}$ in the energy range of several GeV to several hundred TeV, naturally accommodated in the supernova-remnant scenario (discussed next).

At this stage, it is necessary to discuss the units of cosmic-ray energy. The most obvious choice is to simply use kinetic energy, or—in the case of nuclei—kinetic energy per nucleon. Another common choice is the rigidity $\mathcal{R} \equiv pc/Ze$, where p is the magnitude of the three-momentum, c is the speed of light, and Ze is the charge expressed as a multiple of e. As will be described later in this chapter, the rigidity is a natural choice for detectors based on magnetic spectrometers, since the bending radius r in a uniform magnetic field $\boldsymbol{B}$ perpendicular to the momentum is simply $r = \mathcal{R}/B$. Additionally, the acceleration of cosmic rays at their production site is more naturally expressed in terms of $\mathcal{R}$ (see, e.g., Ref. [472]). If the electric charge Ze can be measured by other means (e.g., by the energy deposited in a known thickness of detector), a measurement of $\mathcal{R}$ is equivalent to a measurement of momentum. Furthermore, if the velocity $\beta = v/c$ can be determined, it is possible to determine the mass (and hence identity) of the particle using the relation $p = \gamma\beta cm$, where $\gamma = (1-\beta^2)^{-1/2}$ is the usual Lorentz factor.

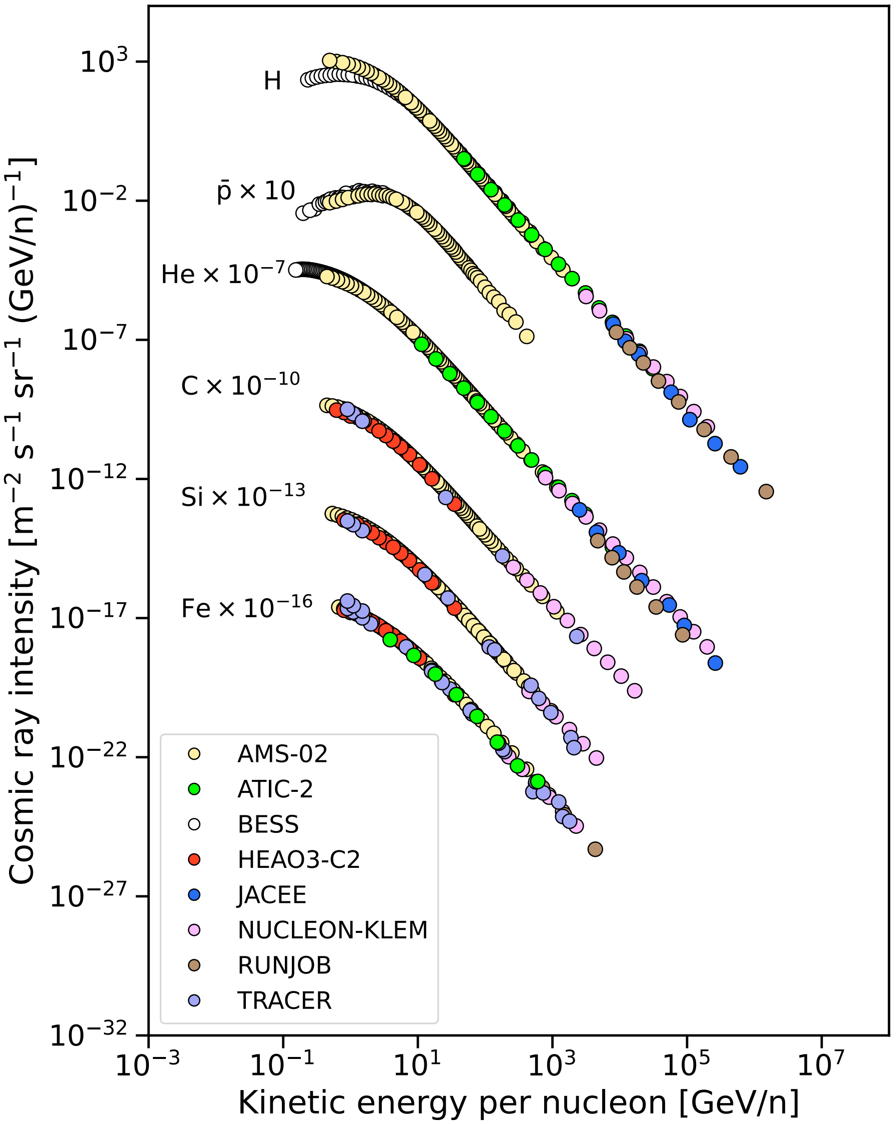

Figure 4-1: Spectra of a variety of cosmic-ray nuclei observed near Earth. Only a small number of experiments are plotted for simplicity, and errorbars are suppressed so as not to clutter the plot. Data taken from Ref. [471]. References are as follows: AMS-02 [473–476], ATIC-2 [477], BESS [478, 479], HEAO3-C2 [480], JACEE [481], NUCLEON-KLEM [482, 483], RUNJOB [484], and TRACER [485–487]. Inspired by Fig. 30.1 of Ref. [16].

4.1.2 Sources

In the standard picture, hadronic[2] cosmic rays with energies up to the PeV (10^{15} eV) scale are produced in diffusive shock waves surrounding Galactic supernova remnants (see, e.g., Refs. [488–490] for reviews). In this scenario, first proposed by Fermi [491], cosmic rays gain energy by reflecting from magnetic fields entrained in the blast wave of the supernova as it strikes the surrounding interstellar medium. Diffusive-shock acceleration naturally produces a cosmic-ray spectrum proportional to E^{-2}, with the somewhat softer observed $E^{-2.7}$ spectrum possibly resulting from propagation in the Galaxy [489]. Assuming $\sim$2–3 Galactic supernovae per century and $\mathcal{O}(10\%)$ efficiency in transferring the supernova energy to cosmic rays [492, 493], this mechanism can account for the observed energy density of cosmic rays in the Galaxy. Since the arrival directions of cosmic rays are scrambled by interstellar magnetic fields, a host of telescopes across the electromagnetic spectrum have extensively studied particle acceleration in supernova remnants via synchrotron, inverse-Compton, and π^0-decay emission, though further work is needed to distinguish between leptonic and hadronic acceleration (see, e.g., Refs. [494, 495]). The next generation of radio, x-ray, and γ-ray telescopes will be extremely helpful in this regard.

The change in spectral index at energies of a few PeV (the so-called “knee”) and again at a few EeV (10^{18} eV) suggests a different acceleration mechanism is at play at these ultra-high energies. Furthermore, cosmic rays with energies greater than a few PeV have magnetic gyroradii in excess of several kiloparsecs—in other words, these ultrahigh-energy cosmic rays cannot be confined to the Galaxy. These facts suggest that cosmic rays with energies $\gtrsim$ few PeV are likely extragalactic in origin.

---

[2] Consisting of protons and heavier nuclei, as opposed to leptonic cosmic rays (electrons and positrons). The latter are outside the scope of this dissertation.

4.1.3 Propagation

Cosmic rays subsequently diffuse into the Galaxy, propagating along turbulent magnetic fields in the interstellar medium. The general propagation equation for species i with momentum $\boldsymbol{p}$ is (e.g., Refs. [488, 496])

$$
\begin{aligned}
\frac{\partial \psi_i}{\partial t} &= q_i(\boldsymbol{r}, \boldsymbol{p}, t) && \text{(source)} \\
&+ \boldsymbol{\nabla} \cdot (D_{xx} \boldsymbol{\nabla} \psi_i - \boldsymbol{V} \psi) && \text{(spatial convection)} \\
&+ \frac{\partial}{\partial p} p^2 D_{pp} \frac{\partial}{\partial p} \frac{\psi_i}{p^2} && \text{(diffusive reacceleration)} \\
&- \frac{\partial}{\partial p} \left[\dot{p} \psi_i - \tfrac{1}{3} p (\boldsymbol{\nabla} \cdot \boldsymbol{V}) \psi_i) \right] && \text{(momentum convection)} \\
&- (\Gamma_i^{\mathrm{col}} + \Gamma_i^{\mathrm{dec}}) \psi_i && \text{(depletion by collision or decay)}
\end{aligned}
\tag{4.1}
$$

Here, $\psi_i(\boldsymbol{r}, \boldsymbol{p}, t) = 4\pi p^2 f_i(\boldsymbol{r}, \boldsymbol{p}, t)$ is the momentum-space distribution, f_i is the phase-space distribution function, q_i includes all of the sources of species i (injection, spallation, decay, etc), D_{xx} is the spatial diffusion coefficient, $\boldsymbol{V}$ is the bulk convective velocity, D_{pp} is the momentum diffusion coefficient, $\dot{p} = \mathrm{d}p/\mathrm{d}t$ is the rate of change of momentum, and Γ_i^{col} and Γ_i^{dec} are the rates at which i are depleted due to collisions and decay, respectively. Owing to its complexity, Eq. 4.1 is generally solved numerically for some assumed distribution of sources in a model of the Galaxy using codes such as GALPROP [497] or DRAGON [498]. The fluxes and flux ratios of the various cosmic-ray species are then calculated at Earth, and the parameters of the cosmic-ray propagation model are tuned until the breadth of experimental results can be consistently reproduced.

Cosmic rays can be divided along many axes, but a particularly useful distinction is between *primary* and *secondary* cosmic rays. Primary cosmic rays are those which have remained the same species from their production/acceleration until their detection at Earth. In contrast, secondary cosmic rays are the product of collisions of primary cosmic rays with the interstellar medium along their trajectory. As one example, it has been known since the 1950s that the observed abundances of lithium, beryllium, and boron cannot be the result of stellar fusion reactions, as proton bombardment in stellar environments rapidly converts these elements to helium (see, e.g., Refs. [71, 499]). Rather, these elements result from the fragmentation and spallation of heavier cosmic-ray nuclei (e.g., carbon and oxygen) in hard-scattering interactions. In particular, the cosmic-ray abundances of beryllium and boron are sensitive to the confinement time of cosmic rays within the Galaxy, since ^{10}Be undergoes beta-decay to ^{10}B with a half-life of approximately 1.4×10^6 yr (see, e.g., Refs. [488, 490, 500] and references therein). Another major class of secondary cosmic rays is the antiparticles, some of which will be discussed later in this chapter.

Upon entering the solar system, cosmic rays encounter the Sun's magnetic field and experience *solar modulation* (see, e.g., Ref. [501] for a review). Since the Sun continuously emits a stream of charged particles (the solar wind), the magnetic field within the solar system varies with the 11- and 22-year solar activity cycles. Owing to the complexity of modeling the spatial and temporal variations of the solar magnetic field throughout the solar system, it is common to invoke the so-called *force-field approximation* [502], wherein all of the complexity of the problem is rolled into a single potential ϕ which is allowed to change with solar activity. In particular, the spectrum $\mathcal{I}_{\mathrm{E}}$ measured near Earth[3] is related to the local interstellar spectrum $\mathcal{I}_{\mathrm{LIS}}$ by [502, 503]

$$\mathcal{I}_{\mathrm{E}}(T_{\mathrm{E}}) \approx \frac{T_{\mathrm{E}} + 2m}{T_{\mathrm{LIS}} + 2m} \frac{T_{\mathrm{E}}}{T_{\mathrm{LIS}}} \mathcal{I}_{\mathrm{LIS}}(T_{\mathrm{LIS}}), \tag{4.2}$$

where T is the kinetic energy, m is the particle mass, and $T_{\mathrm{LIS}} = T_{\mathrm{E}} + Ze\phi$. The effect of the solar modulation is to reduce the flux of low-energy ($\lesssim 1\,\mathrm{GeV/n}$) cosmic rays entering the solar system. Since determining ϕ requires knowledge of the local interstellar spectrum, recent measurements by the Voyager spacecraft outside the heliosphere (as well as measurements by other observatories over several solar cycles) have been invaluable inputs (see, e.g., Refs. [504, 505]).

---

[3] Though neglecting the Earth's magnetic field.

In the case of the Earth, the magnetic field is generally more temporally stable compared to the Sun's. The effect of the magnetic field results in a position-dependent *geomagnetic cutoff*, since particles heading toward different positions on Earth encounter different magnitudes and directions of magnetic fields. The transport of charged particles in the Earth's magnetic field is a complicated problem well outside the scope of this dissertation (see, e.g., Ref. [506] for reviews), so it is typical to describe the geomagnetic shielding effect in terms of the vertical cutoff rigidity (i.e., the rigidity below which downgoing cosmic rays are not detected). The vertical cutoff rigidity is lowest near the Earth's magnetic poles, since incoming charged particles are traveling nearly parallel to the magnetic field lines (and thus experience minimal Lorentz forces). The issue of geomagnetic cutoff will be particularly important for the low-energy cosmic antinuclei searches discussed later in this chapter.

4.2 Antinuclei

In particle physics, an antinucleus is a bound state of antiprotons ($\bar{p}$, consisting of $\bar{u}\bar{u}\bar{d}$ valence quarks) and antineutrons ($\bar{n}$, consisting of $\bar{u}\bar{d}\bar{d}$). Like all antiparticles, antinuclei have the same masses, magnetic moments, and other properties as their "normal" matter counterparts, with only their charges reversed. As such, antinuclei are exotic and rare particles—on Earth, they are found only in the subatomic debris of high-energy particle collisions. Their extreme rarity in cosmic rays is itself an indication that the Universe is predominantly composed of matter, rather than containing large amounts of relic antimatter. Before discussing the possible DM signatures in cosmic antinuclei, it will be helpful to briefly review how antinuclei are produced in high-energy nuclear collisions on Earth.

4.2.1 Experimental history of antinuclei

The first antinucleus to be discovered was the antiproton, produced at the Bevatron[4] accelerator in 1955 [507] The minimal reaction for antiproton production is

$$p + p \to p + p + p + \bar{p}. \tag{4.3}$$

In the fixed-target configuration (i.e., where one of the initial-state protons is at rest), Eq. 4.3 has a kinetic-energy threshold[5] of $6m_p \approx 5.6\,\mathrm{GeV}$. The antiprotons were produced with forward momentum $\gtrsim 1\,\mathrm{GeV}/c$, so a series of magnetic spectrometers was used to select particles by the desired charge (negative) and momentum (1.19 GeV/c). Since antiprotons with this momentum would be traveling much slower than the dominant π^- background ($\beta \sim 0.75$ versus $\sim$0.99, respectively) a series of plastic scintillators and glass Čerenkov counters were used to measure the velocity of outgoing particles. The discovery of the antineutron came the following year, also at the Bevatron, by the charge-exchange reaction [508]

$$p + \bar{p} \to n + \bar{n}. \tag{4.4}$$

Unlike antiprotons, however, antineutrons are electrically neutral, so they cannot be selected using magnetic spectrometers. To account for this, the Bevatron group used the other major property of antiparticles—their annihilation. The annihilation of an $\bar{n}$ with a nucleus would produce a shower of $\pi^{\pm}$ and nuclear fragments, with no incoming ionization track.

The following decades would see further studies of antiproton and antineutron production, as well as attempts to produce ever-heavier antinuclei. The next logical candidate was the antideuteron ($\overline{\mathrm{D}}$ or $^{2}\overline{\mathrm{H}}$), the bound state of one $\bar{p}$ and one $\bar{n}$. In 1965, two groups reported the first evidence of antideuterons produced in the collisions of protons with beryllium targets, using 30-GeV protons from the

---

[4] The BeV (for "billion eV") was an older term for the GeV.

[5] In the original Bevatron experiment, the energy threshold was somewhat lowered by the Fermi momentum of nucleons in the copper target.

Brookhaven Alternating Gradient Synchrotron [509] and 19-GeV protons from the CERN Proton Synchrotron [510]. The following years would see the discovery of antitritons ($\overline{\mathrm{T}}$ or $^{3}\overline{\mathrm{H}}$, Ref. [511]), antihelium-3 ($^{3}\overline{\mathrm{He}}$, Ref. [512]), and antihelium-4 ($^{4}\overline{\mathrm{He}}$, Ref. [513]) nuclei. Around the same time, particle physicists put antiprotons to work in their accelerators, eventually culminating in high-energy proton-antiproton colliders such as the CERN Super Proton-Antiproton Synchroton and the Fermilab Tevatron.

4.2.2 Production models

With the wealth of (anti)nuclear production data in proton-proton (pp), electron-positron ($e^{+}e^{-}$), proton-nucleus (pA), and nucleus-nucleus (AA) collisions, it has become possible to precisely interrogate the formation mechanisms of these exotic particles. There are two general models for the formation of light antinuclei in high-energy collisions: the coalescence model and the statistical hadronization model.

Coalescence model

The central assumption of the coalescence model is that (anti)nucleons produced sufficiently close in momentum space will coalesce (see, e.g., Refs. [514–516]). In this model, the production rate of an antideuteron, say, with momentum $\boldsymbol{k}_{\overline{\mathrm{D}}}$ at a center-of-mass energy $\sqrt{s}$ is given by (e.g., Ref. [517])

$$\frac{\mathrm{d}N_{\overline{\mathrm{D}}}(\sqrt{s},\boldsymbol{k}_{\overline{\mathrm{D}}})}{\mathrm{d}\boldsymbol{k}_{\overline{\mathrm{D}}}^{3}} = \int \mathrm{d}^{3}\boldsymbol{k}_{\bar{p}}\,\mathrm{d}^{3}\boldsymbol{k}_{\bar{n}}\frac{\mathrm{d}^{2}N_{\bar{p}\bar{n}}(\sqrt{s},\boldsymbol{k}_{i})}{\mathrm{d}\boldsymbol{k}_{\bar{p}}^{3}\,\mathrm{d}\boldsymbol{k}_{\bar{n}}^{3}}C(\sqrt{s},\boldsymbol{k}_{\bar{p}},\boldsymbol{k}_{\bar{n}})\delta^{3}(\boldsymbol{k}_{\overline{\mathrm{D}}}-\boldsymbol{k}_{\bar{p}}-\boldsymbol{k}_{\bar{n}}). \tag{4.5}$$

The function $C(\sqrt{s},\boldsymbol{k}_{\bar{p}},\boldsymbol{k}_{\bar{n}})$ is known as the coalescence function, and encodes the probability of $\bar{p}$ and $\bar{n}$ with relative momentum $\Delta\boldsymbol{k} = \boldsymbol{k}_{\bar{p}} - \boldsymbol{k}_{\bar{n}}$ coalescing into $\overline{\mathrm{D}}$. To simplify the problem, many analyses assume that C does not depend on $\sqrt{s}$, and that $C(\Delta\boldsymbol{k}) = \Theta(\Delta\boldsymbol{k}^{2} - p_{0}^{2})$—i.e., any nucleons with relative momentum greater than the coalescence momentum p_{0} will coalesce with 100% probability (Θ

is the unit step function). With the further approximation that the $\bar{p}$ and $\bar{n}$ are produced uncorrelated, the yield of any (anti)nucleus with mass A—composed of Z (anti)protons and N (anti)neutrons—can be expressed as (e.g., Ref. [518])

$$E_A \left(\frac{\mathrm{d}N_A}{\mathrm{d}\boldsymbol{k}_A^3} \right) \approx B_A \left(E_{\bar{p}} \frac{\mathrm{d}N_{\bar{p}}}{\mathrm{d}\boldsymbol{k}_{\bar{p}}^3} \right)^Z \left(E_{\bar{n}} \frac{\mathrm{d}N_{\bar{n}}}{\mathrm{d}\boldsymbol{k}_{\bar{n}}^3} \right)^N . \tag{4.6}$$

The coalescence factor $B_A \equiv A[4\pi p_0^3/(3m_p)]^{A-1}$ depends on the coalescence momentum p_0, the only free parameter of this simplified model. The inferred values of p_0 for antideuteron production in a variety of experiments depend on $\sqrt{s}$, the type of colliding particles, and on the choice of Monte Carlo event generators[6] needed to model the hadronic physics. Perhaps unsurprisingly for such a simple model, this naïve coalescence model is likely parametrizing several other effects, including the finite volume of the collision region and correlations in the final-state particles (see, e.g., Refs. [519, 520] for additional discussion). Recent work including finite-size and momentum-correlation corrections [518, 521] and multiphase transport modeling [522] has helped to make the models more realistic, but more data is clearly needed (see Sec. 4.2.3).

Statistical hadronization model

The statistical hadronization model (SHM) was developed in response to the wealth of high-energy nucleus-nucleus collision data that became available after the 1980s (for extensive theoretical and experimental reviews, see, e.g., Refs. [523–525]). In these collisions, with nucleon-nucleon center-of-mass collision energies $\sqrt{s_{\mathrm{NN}}}$ ranging from several GeV to several TeV, the high energy density ($>1\,\mathrm{GeV\,fm^{-3}}$) causes the quarks and gluons (collectively referred to as *partons*) to deconfine from the hadrons. These partons rapidly enter a state of approximate thermal and chem-

---

[6] These are also required to simulate the yields of stable SM particles (e.g., antinuclei) from the initial DM decay or annihilation (e.g., $\chi\chi \to b\bar{b} \to \text{hadrons} + \cdots$). Since these generators are tuned using different datasets and employ different algorithms (especially in their treatment of nonperturbative hadronic processes), the choice of generator represents a significant source of systematic uncertainty.

ical equilibrium, as the internal pressure of the plasma causes it to expand and cool. Below the chemical freezeout temperature T_{cf}, the system is once again described in terms of a weakly interacting gas of hadrons. The number density n_i of a hadronic species with mass m_i, degrees of freedom g_i, and chemical potential μ_i is simply the integral over $p_i = |\boldsymbol{p}_i|$ (e.g., Ref. [526])

$$n_i = \frac{g_i}{2\pi^2} \int_0^\infty \frac{p_i^2 \, \mathrm{d}p_i}{e^{(E_i - \mu_i)/T} \pm 1}, \tag{4.7}$$

where $E_i = \sqrt{p_i^2 + m_i^2}$ and the $+1$ (-1) refers to fermions (bosons). As shown in Fig. 4-2, the SHM provides an excellent fit to more than 20 particle species produced in high-energy nuclear collisions with only three free parameters (chemical freezeout temperature T_{cf}, baryon chemical potential μ_b, and volume V). Owing to the extremely high collision energies and small baryonic chemical potentials achievable at heavy-ion colliders such as RHIC and the LHC, matter and antimatter are produced in approximately equal amounts, making these heavy-ion experiments ideal for the study of composite antinuclei (i.e., antideuterons and heavier).

It should be noted that although the SHM has been a boon to the heavy-ion community, its utility for predicting cosmic antinuclei fluxes (either from DM or from standard astrophysical sources) is likely limited (see, e.g., the discussion in Ref. [527]). First, as will be discussed next, most antinuclei produced by cosmic-ray collisions (i.e., secondary antinuclei) are produced at low energies near threshold, where the interacting system is not likely to reach thermal and chemical equilibrium. Second, unlike the coalescence model, the SHM does not predict the momentum distribution of the final-state antinuclei. Thus, most cosmic antinucleus work uses the coalescence model, and this dissertation will do the same.

4.2.3 Secondary antinucleus production in cosmic rays

Secondary production of antinuclei in cosmic rays proceeds through the same general steps as terrestrial accelerators: a high-energy particle strikes a target nucleus

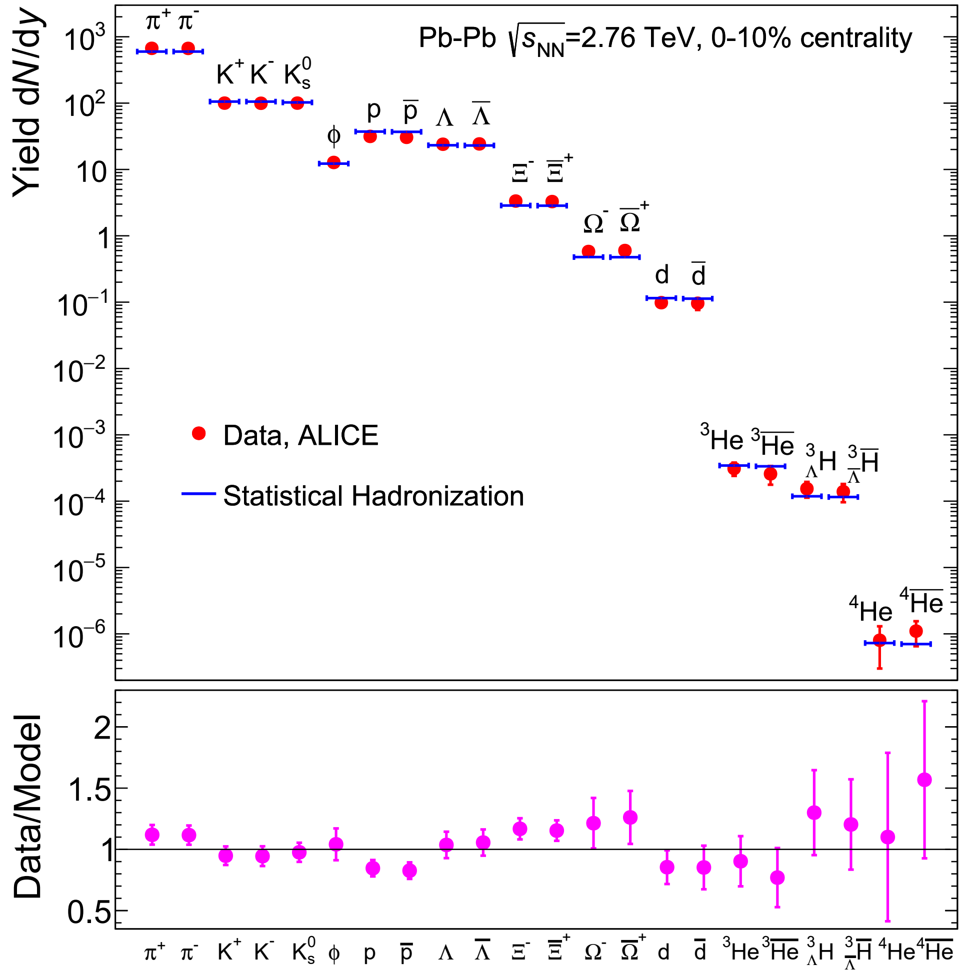

Figure 4-2: Measured yields of hadrons per unit rapidity y in central lead-lead collisions at ALICE. The data (red) correspond to mid-rapidity particles at a center-of-mass collision energy $\sqrt{s_{\mathrm{NN}}} = 2.76\,\mathrm{TeV}$ per nucleon pair. The blue lines indicate the statistical hadronization model predictions for $T_{\mathrm{cf}} = 156.5 \pm 1.5\,\mathrm{MeV}$, $\mu_{\mathrm{b}} = 0.7 \pm 3.8\,\mathrm{MeV}$, and $V = 5280 \pm 410\,\mathrm{fm}^3$ [528]. Figure reproduced from Ref. [529] with permission. © 2018 by Springer Nature.

at rest. Consider a cosmic-ray nucleus $i \in \{p, \mathrm{He}, \bar{p}, \cdots\}$ colliding with a stationary nucleus $k \in \{p, \mathrm{He}\}$. The production rate $q_{\overline{\mathrm{D}}}^{\mathrm{sec}}$ of secondary antideuterons, for example, with energy E is given by (e.g., Ref. [530])

$$q_{\overline{\mathrm{D}}}^{\mathrm{sec}}(\boldsymbol{r}, E) = 4\pi \sum_{i,k} n_k(\boldsymbol{r}) \int_{E_{ik}^{\mathrm{min}}}^{\infty} \frac{\mathrm{d}\sigma(ik \to \overline{\mathrm{D}}X; E)}{\mathrm{d}E_i} \Phi_i(\boldsymbol{r}, E_i)\, \mathrm{d}E_i. \tag{4.8}$$

Here, n is the number density in the ISM, $\mathrm{d}\sigma(ik \to \overline{\mathrm{D}}X; E)/\mathrm{d}E_i$ is the differential cross section producing antideuterons with energy E, and Φ_i is the flux of incident cosmic-ray particles. The threshold E_{ik}^{min} is set by kinematics—for example, the minimal reaction $p + p \to p + p + p + \bar{p} + n + \bar{n}$ has a threshold kinetic energy of $17m_p$, assuming one of the incident protons is at rest. Since the cosmic-ray flux Φ_i is a sharply decreasing function of energy, the production rate of secondary antideuterons is further suppressed.

One of the dominant sources of systematic uncertainty for secondary antinucleus production is the production cross section $\mathrm{d}\sigma/\mathrm{d}E_i$. Since the flux of cosmic rays decreases rapidly with increasing energy, lower-energy cosmic rays have a significant impact on the secondary production of antinuclei. Unfortunately, there is comparatively little data for $pp \to \bar{p}X$ production for incident proton kinetic energies $10\,\mathrm{GeV} \lesssim E_p \lesssim 100\,\mathrm{GeV}$ (center-of-mass collision energies $5\,\mathrm{GeV} \lesssim \sqrt{s} \lesssim 25\,\mathrm{GeV}$), and the data for low-energy $\overline{\mathrm{D}}$ production is even more sparse (see, e.g., Ref. [517]). Fortunately, there are several experimental efforts to remedy this issue. The fixed-target NA61/SHINE detector [531, 532] is a fixed-target magnetic spectrometer designed to study hot QCD matter in a variety of colliding nuclei with beam energies ranging from tens of GeV/n to a few hundred GeV/n. Recent results on antiproton production in these mid-energy nuclear collisions has been a welcome input to cosmic-ray physics (see, e.g., Refs. [533, 534]). Another CERN experiment contributing to this effort is LHCb—since the collision region can be filled with a small amount of gas, it is possible to study antiproton production in fixed-target pA collisions, including the first measurements of antiproton production in pHe collisions [535]. Finally, the upcoming PANDA experiment at the

Facility for Antiproton and Ion Research (FAIR) will study the collisions of 1.5–15-GeV/c antiprotons with nuclei [536], which will improve modeling of antiproton propagation in the Galaxy as well as tertiary production (e.g., $\bar{p}\mathrm{A} \to \overline{\mathrm{D}}X$).

4.2.4 Antideuterons as a DM probe

Consider a hypothetical DM particle with mass $m_\chi \gtrsim 10\,\mathrm{GeV}$ which decays or annihilates with another DM particle via some BSM process. Depending on the mass of the particle, there may be many final states available: $\nu\bar{\nu}$, $\ell^+\ell^-$, gg, $q\bar{q}$, W^+W^-, ZZ, HH, and so on (see, e.g., Ref. [146]). With the exception of e^+e^- and $\nu\bar{\nu}$, all of these final states are unstable, and all of them can lead to high-energy quarks and/or gluons which form the basis of this discussion. As has been conclusively demonstrated in terrestrial experiments, quarks and gluons cannot exist on their own—rather, they must be confined into hadrons (see, e.g., Ref. [5]). The process by which an initial quark or gluon is converted into a jet of color-neutral hadrons is known as *hadronization*, and proceeds by gluon radiation and $q\bar{q}$ pair production until the available energy is too low to continue generating new particles. At this point, the system consists of a jet of (mostly unstable) hadrons, which subsequently decay into (anti)protons, (anti)neutrons, and electrons/positrons on timescales of $\lesssim 1\,\mu\mathrm{s}$. Some of these antiprotons and antineutrons will coalesce into antideuterons, as described in the previous section, and subsequently make their way to Earth. (At this stage, it should also be emphasized that unlike photons, which can be localized to a particular area of sky, the arrival directions of cosmic rays are scrambled by magnetic fields. Though this means a loss of directional information, it also means that large volumes of the Galactic halo can contribute to the observed signal.)

The reader may wonder why so much effort has been spent considering the production and coalescence of antideuterons, when such DM models would clearly produce antiprotons much more abundantly. The answer comes down to expected backgrounds. As described previously, antiprotons are constantly being produced

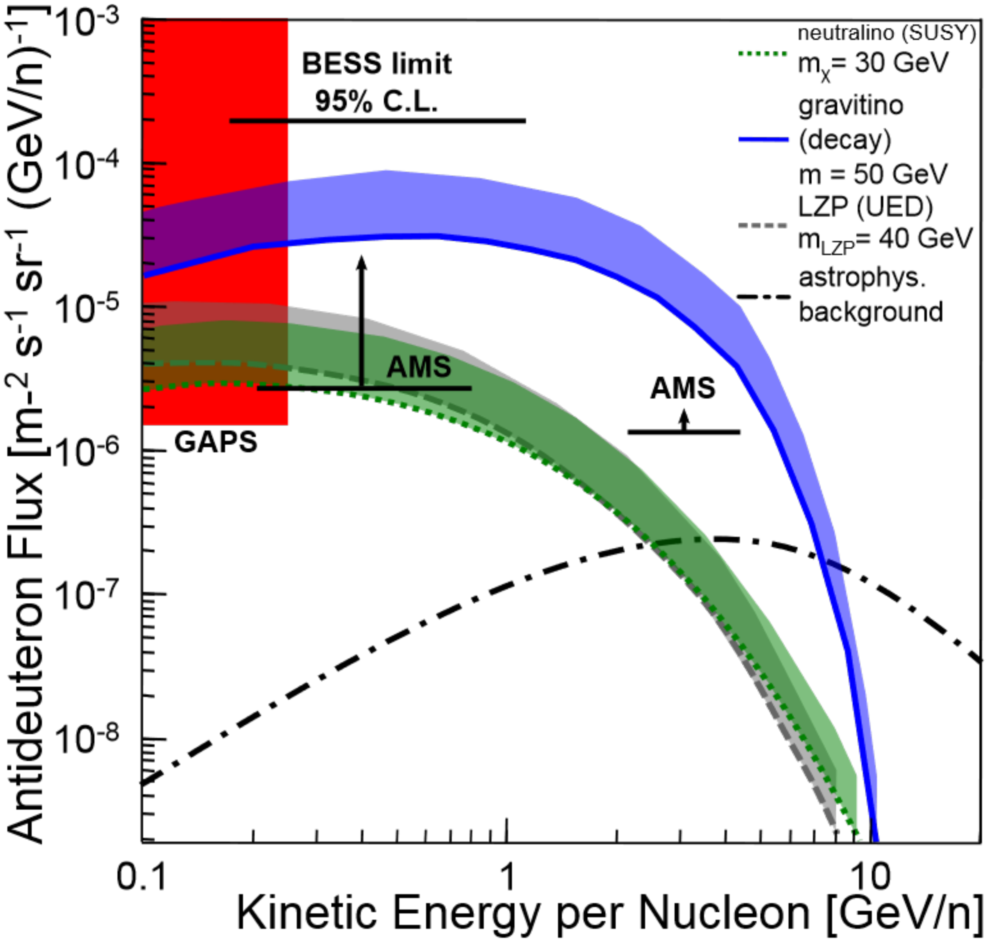

Figure 4-3: Expected antideuteron flux from a variety of DM models (blue, gray, and green shaded bands indicating Galactic propagation uncertainties) compared to the expectation from secondary astrophysical production (black dashed-dotted line). The red shaded region indicates the projected GAPS upper limit (3σ) after three 35-day Antarctic flights. The black lines show the BESS (95%, Ref. [537]) and projected AMS-02 (3σ, Refs. [519, 538]) upper limits. For more details, see Ref. [539], from which this figure has been reproduced with permission.

via the collisions of cosmic rays with the interstellar medium; thus, determining the DM contribution to the antiproton spectrum would require knowing the astrophysical processes to extremely (and perhaps prohibitively) high confidence (see, e.g., Ref. [540] for an early discussion). The expected rate of secondary antideuteron production, however, is greatly suppressed by the heightened kinetic energy threshold compared to $\bar{p}$ production ($17m_p$ versus $7m_p$, respectively, for pp collisions), meaning that relatively few[7] primary cosmic rays have sufficient energy to produce antideuterons. As shown in Fig. 4-3, low-energy ($\lesssim 1\,\mathrm{GeV/n}$) antideuterons represent a nearly background-free channel for DM searches in a variety of models, including supersymmetry (SUSY), weakly-interacting massive particles (WIMPs), primordial black holes, and models with extra dimensions (see, e.g., Refs. [519, 540–545]) . Detection of even a handful of low-energy antideuterons would be a "smoking gun" for new physics, either in the context of DM or modifications to the standard picture of cosmic-ray propagation.

4.3 Cosmic antinuclei measurements

Following the first detection of cosmic antiprotons in 1979 [546], a variety of experiments have searched for antiprotons and heavier antinuclei. Since the atmosphere very efficiently shields the Earth's surface from incoming cosmic rays, these experiments must be balloon- or space-based (altitudes of $\gtrsim$30 km and $\gtrsim$100 km, respectively). In this section I will discuss the two experiments with leading sensitivity to cosmic antideuterons, BESS and AMS-02.

4.3.1 BESS

The Balloon-borne Experiment with a Superconducting Spectrometer (BESS) was—as its name suggests—a balloon experiment with ten successful flights between

---

[7] Note that the reaction $p + \bar{p} \to p + \bar{p} + n + \bar{n}$ has a threshold kinetic energy comparable to that of $p + p \to p + p + p + \bar{p}$—thus, antiproton collisions may be a significant contributor to antideuteron production despite the small $\bar{p}/p$ flux ratio.

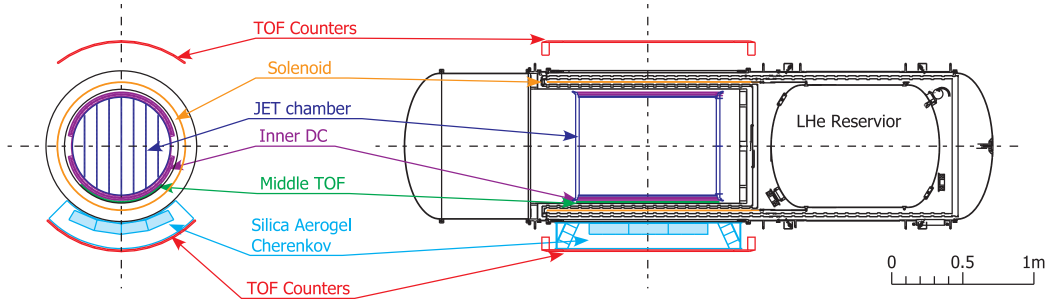

Figure 4-4: Schematic of the BESS-Polar II instrument in its final configuration. The detector subsystems are described in Sec. 4.3.1. Reproduced from Ref. [548] with permission. © 2012 by the American Physical Society.

1993–2008, optimized for the detection of low-energy cosmic antinuclei. BESS flew in several configurations over its life, so this section will describe the final (and most sensitive) configuration, BESS-Polar II [547].

As shown in Fig. 4-4, the central component of BESS-Polar II was the solenoidal spectrometer with acceptance $\sim 0.2\,\mathrm{m}^2\,\mathrm{sr}$. A 0.8-T solenoidal field provided the bending power for incoming charged particles, whose trajectories were tracked in gas-filled jet-type (JET) and inner drift chambers (DC). The ~140-µm (per point) spatial resolution provided a rigidity resolution $\delta\mathcal{R}/R \sim 0.4\%$ at 1 GV. The time-of-flight (TOF) counters consisted of plastic scintillator paddles placed above and below the solenoid, providing $\mathrm{d}E/\mathrm{d}x$ and timing information and serving as the initial event trigger. The ~100-ps timing resolution of the TOF translated to a ~2.5% measurement of the velocity $\beta = v/c$. Using these quantities, the mass of the incoming particle could be determined from simple kinematics:

$$m^2 = (Ze)^2\mathcal{R}^2(\beta^2 - 1). \tag{4.9}$$

To veto the orders-of-magnitude more abundant protons and muons, BESS also included a silica aerogel Čerenkov counter with refractive index $n = 1.03$ below the solenoid (since antiprotons would only emit Čerenkov light at $\mathcal{R} \gtrsim 3.8\,\mathrm{GV}$). Finally, the instrument included an additional TOF scintillator paddle within the magnet bore to ensure that the lowest-energy antiprotons (those most prone to stopping or multiple scattering) would still be detected. Since it flew from the South Pole during the solar minimum, BESS-Polar II was able to measure antiprotons with kinetic energies as low as $\sim$200 MeV. During its one-month flight in 2007–2008, BESS produced the first precision measurements of cosmic antiprotons below $\sim$1 GeV in kinetic energy, as well as leading limits on antihelium in the rigidity range 1.6–14 GV. Analysis of the antideuteron sensitivity is still ongoing, but it is anticipated that BESS-II Polar will substantially improve upon the flux limit of $1.9 \times 10^{-4}\,\mathrm{m^{-2}\,s^{-1}\,sr^{-1}\,(GeV/n)^{-1}}$ (95% CL) set by previous BESS flights in the energy range 0.17–1.5 GeV/n [537].

4.3.2 AMS-02

The second[8] Alpha Magnetic Spectrometer (AMS-02) experiment is a cosmic-ray detector that was mounted to the International Space Station (ISS) in 2011, and has been operating[9] ever since. AMS-02 is a general-purpose cosmic-ray detector, capable of recording particles with rigidity from $\lesssim$1 GV to several TV and charges up to $|Z| = 28$. Owing to its 51.6° orbital inclination with respect to the equator, the ISS orbit ranges between $\pm 51.6^\circ$ latitude, where the vertical cutoff rigidity often exceeds several GV.

As shown in Fig. 4-5, the AMS-02 detector consists of a variety of subsystems to measure different properties of cosmic rays. Like BESS, AMS-02 is a magnetic spectrometer, where a 0.15-T dipole field is responsible for bending the trajectories of incoming charged particles. The rigidity is measured by nine planes of double-sided silicon strip detectors (seven planes within the magnet bore, one plane above, and one plane below), providing position resolution better than 10 μm in the bending direction. Four planes of plastic-scintillator TOF panels (two above the magnet and two below) measure the velocity β with resolution $\sim$4% for $|Z| = 1$

---

[8] The first AMS experiment (AMS-01) flew in the payload bay of the Space Shuttle Discovery for ten days in 1998. Its magnet would later form the core of AMS-02.

[9] Assuming no unforeseen instrument failures, the end-of-life for AMS-02 will likely be the same as the ISS, currently scheduled to de-orbit in $\sim$2030.

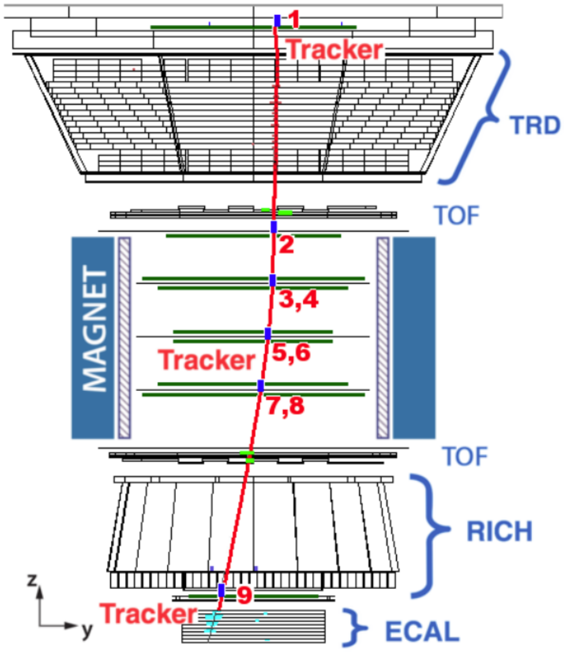

Figure 4-5: Schematic of the AMS-02 experiment on the International Space Station, showing a charged particle (red) entering from the top of the instrument. The detector subsystems are described in Sec. 4.3.2. Reproduced from Ref. [549] with permission. © 2015 by Elsevier B.V.

particles, a figure which improves with increasing $|Z|$ [550]. Since the mass resolution scales as $(\delta m/m)^2 = (\delta\mathcal{R}/\mathcal{R})^2 + \gamma^4(\delta\beta/\beta)^2$, AMS-02 also includes ring-imaging Čerenkov detectors to provide $\delta\beta/\beta \sim 10^{-3}$ for high-velocity ($\beta \approx 1$) particles [551]. At the base of the detector is an electromagnetic calorimeter (ECAL) consisting of 17 radiation lengths of lead sheets and scintillating fibers. The ECAL measures the energies of incoming electrons and positrons to better than ~3% above 100 GeV, and allows the detector to distinguish between energetic electrons/positrons and hadrons [552].

The large sample of cosmic-ray events recorded by AMS-02 since 2011 has been a boon to cosmic-ray studies. The experiment's contributions are too numerous to list here (see, e.g., Ref. [475] for a recent review of results), so I will focus primarily on those results related to cosmic antinuclei. In 2016, the collaboration published precision measurements of the proton and antiproton fluxes in the rigidity range 1–450 GV (kinetic energy range 0.5–445 GeV), extending the measurements by BESS and PAMELA. One interesting result was the discovery of an unexpected excess of antiprotons near 10 GV rigidity, which several groups assessed to be consistent with DM annihilation for masses $m_\chi \sim 50$–$100\,\mathrm{GeV}$ and cross sections $\langle \sigma v \rangle \sim 10^{-26}\,\mathrm{cm}^3\,\mathrm{s}^{-1}$—both compatible with new electroweak-scale physics (e.g., WIMPs) and with the Galactic center γ-ray excess (see, e.g., Ref. [553] and references therein). The status of the antiproton excess remains unresolved, owing to possible systematic effects and bin-to-bin correlations which are known to be significant at these low rigidities [554].

Perhaps more surprising, however, was the announcement in 2015 that AMS-02 had observed several candidate events with charge $Z = -2$ and masses near ~3–4 GeV—i.e., *antihelium* ($\overline{\mathrm{He}}$, Refs. [555, 556]). Owing to the extremely low event rate (the inferred $\overline{\mathrm{He}}/\mathrm{He}$ ratio is less than 10^{-8}), an understanding of possible systematic effects is crucial. If these events turn out to be genuine antihelium nuclei, the possible implications[10] for cosmic-ray and DM physics are enormous. Owing to the wide range of predicted antideuteron and antihelium fluxes from DM or standard astrophysical processes, a complementary experiment with different systematics is very much needed. One such experiment, the General Antiparticle Spectrometer (GAPS), comprises the rest of this dissertation.

---

[10] Some groups even proposed that these events might result from residual clouds of primordial antimatter, or even entire *antistars* (see, e.g., Ref. [557]).

5

The GAPS Experiment

Particles of raw inspiration sleet through the universe all the time.
Every once in a while one of them hits a receptive mind [...]
—Terry Pratchett

The General Antiparticle Spectrometer (GAPS) balloon experiment will be the first detector optimized for low-energy cosmic antinuclei. GAPS is an international collaboration, with principal funding support from the NASA Astrophysics Research and Analysis (APRA) program, the Institute for Space and Aeronautical Science (ISAS) at the Japan Aerospace Exploration Agency (JAXA), the Agenzia Spaziale Italiana[1] (ASI), and the Istituto Nazionale di Fisica Nucleare[2] (INFN). In Sec. 5.1, I briefly review the science goals of the GAPS experiment *vis-à-vis* cosmic antinuclei. In Sec. 5.2, I describe the novel exotic atom detection technique at the core of the GAPS concept. In Sec. 5.3, I outline the various GAPS payload subsystems. I conclude in Sec. 5.4 with a brief description of recent and ongoing GAPS integration and testing work to prepare for launch.

---

[1] Italian Space Agency

[2] (Italian) National Institute of Nuclear Physics

5.1 GAPS science goals

The GAPS experiment has two principal science goals: construct the cosmic antiproton spectrum in an unprecedentedly low energy range, and set leading limits on low-energy antideuterons and heavier antinuclei.

In each $\sim$30-day flight, GAPS will detect several hundred low-energy antiproton events. These antiproton events will demonstrate the sensitivity of the exotic-atom technique, while also serving as controls for various systematics common to the antideuteron and antihelium searches (see Ref. [558] for further discussion). In particular, the low-energy antiproton spectrum at the top of the atmosphere (TOA) is not the same as that recorded at the top of the instrument (TOI), owing to the small amount of residual atmosphere above the $\sim$37-km float altitude. First, this residual atmosphere slows incoming charged particles, with greater losses coming at larger angles θ from the zenith (since particles coming from larger angles must traverse more atmosphere). For antiprotons in the range $0.75 < \cos\theta < 1$, a particle with $\beta \approx 0.4$ ($\beta \approx 0.68$) at TOA corresponds[3] to $\beta \approx 0.25$ ($\beta \approx 0.65$) at TOI. Second, low-energy charged particles can be completely absorbed by the atmosphere: only $\sim$40% of cosmic antiprotons with $\beta = 0.4$ are expected survive from TOA to TOI, compared to $\sim$70% of antiprotons with $\beta = 0.7$. Third, the collisions of cosmic rays with nuclei in the upper atmosphere can produce low-energy antiprotons via the same reactions that produce antiprotons in the interstellar medium. These atmospheric antiprotons constitute an irreducible background to the total antiproton spectrum, exceeding the flux of cosmic antiprotons for $\cos\theta \lesssim 0.5$ ($\theta \gtrsim 60^\circ$). Thus, it is anticipated that antiproton events with $\cos\theta < 0.5$ will be used to study atmospheric effects common to all GAPS antinuclei searches, reserving $\cos\theta > 0.5$ for the cosmic antinucleus analysis. These measurements will also serve to validate the treatment of the geomagnetic and heliospheric magnetic fields.

GAPS will also deliver leading limits on low-energy cosmic antideuterons and antihelium. With a target sensitivity of $2 \times 10^{-6}\mathrm{m}^{-2}\,\mathrm{s}^{-1}\,\mathrm{sr}^{-1}\,(\mathrm{GeV/n})^{-1}$ (99% CL)

---

[3] For context, the GAPS antiproton energy range is $\sim$0.07–0.21 GeV/n at TOA, corresponding to $\sim$0.03–0.19 GeV/n at TOI.

for antideuterons and $(0.06\text{–}6) \times 10^{-6}\mathrm{m}^{-2}\,\mathrm{s}^{-1}\,\mathrm{sr}^{-1}\,(\mathrm{GeV/n})^{-1}$ (95% CL) for $^{3}\overline{\mathrm{He}}$ after ∼100 days of total flight time[519, 558, 559], GAPS will complement previous searches by AMS-02 and BESS (discussed in Chapter 4). At these levels, GAPS is expected to be sensitive to a variety of DM scenarios shown in Fig. 4-3, e.g., tens-of-GeV neutralinos and gravitinos from supersymmetric models or the lightest Z_3 particle (LZP) from warped extra-dimensional models (see, e.g., Refs. [519, 542–544, 560] and references therein). In these cases, the predicted antideuteron or antihelium flux is several orders of magnitude greater than the expected background from cosmic-ray collisions in the interstellar medium, meaning that the identification of only a few antideuterons or antihelium-3 nuclei would be groundbreaking evidence of DM and/or new cosmic-ray physics.

5.2 The GAPS exotic-atom technique

The GAPS experiment will employ a novel detection technique based on the formation and destruction of *exotic atoms*, wherein one of the electrons is replaced with a heavier particle. The stopping of antinuclei to form exotic atoms within the detector, as well as the subsequent de-excitation and annihilation of the exotic atoms, produces a variety of distinctive signatures which enable GAPS to achieve the required sensitivity and background rejection for low-energy antinucleus studies. The process occurs in three principal stages—stopping, cascade, and annihilation—discussed next.

5.2.1 Stopping

Consider a low-energy antinucleus entering the GAPS detector from above, as shown in Fig. 5-1 for an antiproton. The incoming particle first passes through two layers of plastic-scintillator time-of-flight (TOF) panels, in which its velocity is precisely measured. For a heavy particle with charge Z and velocity $\beta = v/c$, the rate of energy loss per unit length $\mathrm{d}E/\mathrm{d}x$ is given by the Bethe-Bloch formula (see,

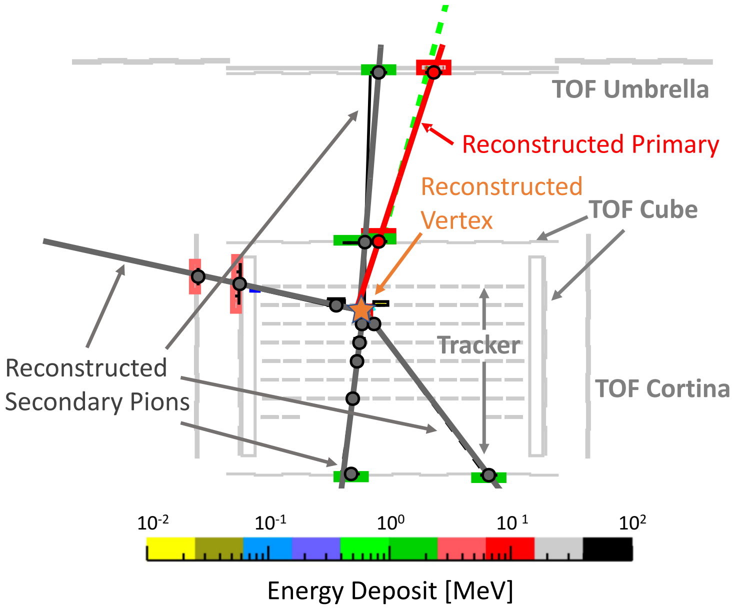

Figure 5-1: A simulated low-energy antiproton event in the GAPS experiment. The antiproton (green dotted) enters from the upper right, passing through two layers of TOF panels and two tracker layers before stopping and annihilating into a shower of secondary pions. These secondary pions are tracked through the tracker and TOF system. For more details, see Secs. 5.2 and 5.3. Reproduced from Ref. [558] with permission. © 2023 by Elsevier B.V.

e.g., Chapter 34 of Ref. [16])

$$\frac{1}{\rho}\frac{\mathrm{d}E}{\mathrm{d}x} \approx -0.3\,\mathrm{MeV}\,\mathrm{cm}^{-2}\,\mathrm{g}^{-1}\left(\frac{Z_A}{A}\right)\left(\frac{Z^2}{\beta^2}\right)\left[\ln\left(\frac{2\beta^2\gamma^2 m_e c^2}{I}\right)-\beta^2\right], \tag{5.1}$$

where Z_A is the atomic number of the target material, A is the mass number, ρ is the mass density, $\gamma = (1-\beta^2)^{-1/2}$ is the usual Lorentz factor and I is the electronic excitation energy of the target material. For particles with the same charge (e.g., antiprotons and antideuterons), the value of $\mathrm{d}E/\mathrm{d}x$ (and hence the distance traveled in the detector before stopping) depends only on β. Since particles with the same kinetic energy per nucleon but different masses have the same β, the particle with the greater mass (e.g., an antideuteron) will travel further in the tracker, since it has greater total kinetic energy to lose. Thus, low-energy antiprotons and antideuterons can be separated by their stopping distance in the detector, provided their β can be determined. Since stopping power is proportional to Z^2, higher-charged particles (e.g., antihelium or heavier cosmic-ray nuclei) are easily distinguished by their large energy depositions and short ranges [560].

5.2.2 Cascade

Eventually, the antinucleus (hereafter called "cascader" for reasons which will soon become clear) comes to rest in the tracker and replaces an atomic electron. For concreteness, assume that the nucleus has atomic number Z_A and mass m_A. The principal quantum number n into which the cascader is originally captured scales as $\mu^{1/2}$, where $\mu \equiv m_A m/(m_A + m)$ is the reduced mass of the cascader and nucleus. This exotic atom is inherently unstable, since the cascader does not experience Pauli blocking from atomic electrons; thus, it rapidly[4] cascades to lower-energy orbitals as shown in Fig. 5-2, first by emitting Auger electrons and subsequently x-rays. Ignoring the shielding effects of the electrons and the finite size of

---

[4] In this discussion, the cascader is assumed to be a stable particle (e.g., an antiproton or antideuteron) or at least long-lived with respect to the timescales of atomic transitions.

the nucleus, the energies of the exotic-atom orbitals are (see, e.g., Ref. [561])

$$E_{n,j} \approx \frac{(\alpha Z_A Z)^2}{2n^2}\mu c^2 - \frac{(\alpha Z_A Z)^4}{2n^4}\left[\frac{n}{j+\frac{1}{2}} - \frac{3}{4}\right]\mu c^2 + \cdots \tag{5.2}$$

Here, $\alpha \approx 1/137$ is the electromagnetic fine-structure constant, $\mu \equiv m_A m/(m_A + m)$ is the reduced mass defined previously, n is the principal quantum number, and j is the total angular momentum (spin plus orbital) of the cascader. For a heavy particle with charge -1 in a state (n, j), the energies of the exotic orbitals are greater than the corresponding electron orbitals by a factor proportional to m/m_e. Thus, the photons emitted as the exotic atom lowers its energy are in the $\mathcal{O}(1\text{–}100\,\text{keV})$ energy range—i.e., x-rays—even for relatively low n. Using a cascade model to calculate the occupancies and transition probabilities, it is possible to predict the x-ray energies and yields of these exotic atoms (see, e.g., Refs. [562–566] and references therein). These transitions have been extensively studied in elements across the entire periodic table for exotic atoms containing μ^-, π^-, and K^- (see, e.g., Refs. [561, 567–574] for reviews). Recent work studying the x-rays emitted from exotic atoms of hydrogen, helium, nitrogen, carbon, fluorine, aluminum, sulfur, chlorine, and bromine have demonstrated the applicability of the technique for antiprotons (see, e.g., Refs. [566, 574]). Since antideuterons are approximately twice as massive as antiprotons, the x-rays released from antideuteronic exotic atoms will be more energetic than the corresponding antiprotonic x-rays, as shown in Fig. 5-3. Furthermore, since x-rays can only be emitted in the decays of exotic atoms formed from negatively-charged particles, requiring x-ray detection serves to veto the orders-of-magnitude more-abundant positively-charged cosmic-ray protons and nuclei. Distinguishing x-rays from antiproton and antideuteron captures requires detectors with energy resolutions better than 4 keV FWHM, the implications of which will be discussed in Chapter 6.

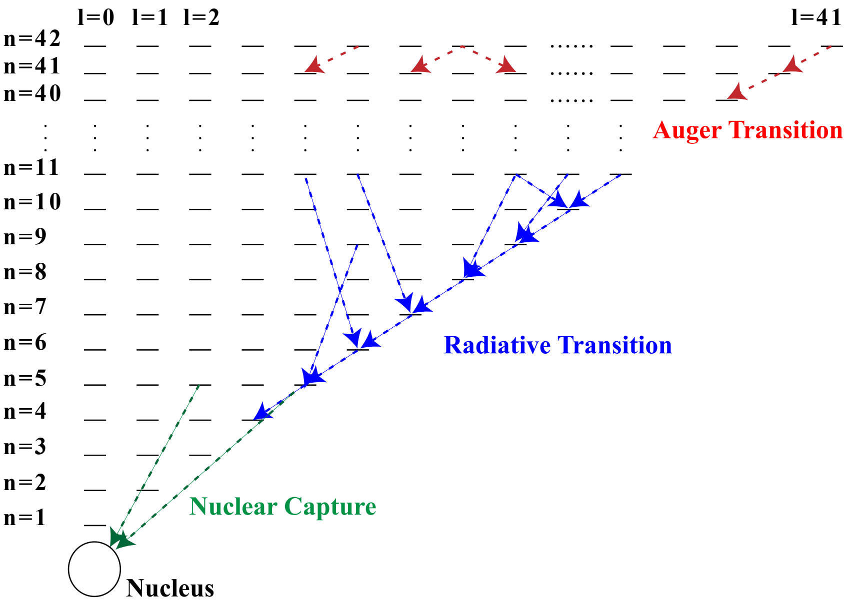

Figure 5-2: Schematic of an exotic-atom cascade beginning in a state with high principal quantum number n and orbital quantum number l. The atom subsequently de-excites by Auger (electron-emitting, red) and radiative (x-ray-emitting, blue) transitions, eventually being captured by the nucleus in an $n = 5$ state. Reproduced from Ref. [566] with permission. © 2013 by Elsevier B.V.

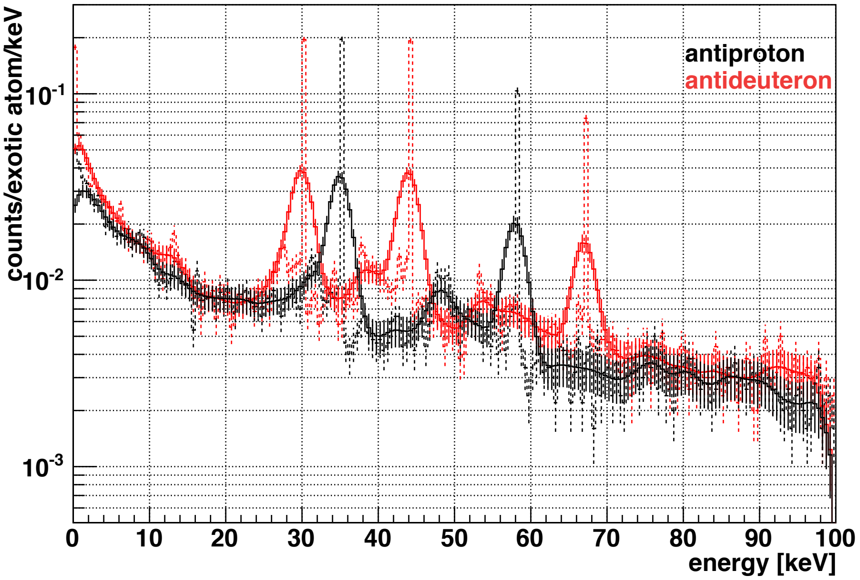

Figure 5-3: X-ray spectra from a GEANT4 simulation of antiproton (black) and antideuteron (red) capture in silicon. The dashed lines show the initial energy distributions, and the solid lines show the effect of 3-keV FWHM detector energy resolution. Reproduced from Ref. [519] with permission. © 2016 by Elsevier B.V.

5.2.3 Annihilation

Replacing an atomic electron with a heavy particle has another consequence—the Bohr radii of the atomic orbitals scale approximately as m^{-1}. The low-n orbitals thus bring the particle extremely close to the nucleus, until it eventually interacts with the nucleons. Antiprotons and heavier antinuclei will annihilate with the nucleus, producing an "annihilation star" of mesons, protons, and other nuclear fragments as shown in Fig. 5-1. The Crystal Barrel experiment at CERN, for example, studied low-energy proton-antiproton annihilation in liquid hydrogen to many final states involving mesons [575]. Many of these final states include extremely unstable hadronic resonances (e.g., ρ, ω, η, and ϕ mesons), which subsequently decay to $\pi^{\pm}$ and $K^{\pm}$ mesons[5] that live long enough to travel through the detector without decaying. The aforementioned Crystal Barrel data indicate that the average charged-pion multiplicity in $p\bar{p}$ annihilation, for example, is $\sim$3. Since antideuteron annihilations are too rare to study at colliders, models such as the intranuclear cascade model (see, e.g., Refs. [576–578]) are used to predict the interactions of antideuterons with nuclei. Roughly speaking, each additional antinucleon contributes $\sim$3 more charged pions to the annihilation, so the number of charged pions provides a measurement of the mass of the incoming antinucleus. Additionally, the fragmentation of the nucleus emits protons with kinetic energy spectrum [579]

$$\frac{\mathrm{d}N_p}{\mathrm{d}E_p} = 2\langle N_p\rangle \left(\frac{E_p}{\pi T^3}\right)^{1/2} \exp\left[-\frac{E_p}{T}\right], \tag{5.3}$$

where $\langle N_p\rangle$ is the average proton multiplicity (0.86 ± 0.05 for silicon) and $T \sim 40\,\mathrm{MeV}$ is the "nuclear temperature." Compared to $\pi^{\pm}$, these protons have much lower energies, and can thus be identified by their much larger energy deposits as they traverse the Si(Li) detectors (see, e.g., Ref. [519]). The pions and protons are a distinctive and unique signature of antinucleus annihilation at rest, providing both a measurement of the antinucleus mass and background rejection against the am-

---

[5] These annihilations also produce π^0 mesons, which decay almost instantly to two photons. Since these photons are extremely energetic ($\gtrsim 70\,\mathrm{MeV}$) and thus difficult to detect in a balloon payload, they are not included in the rest of this discussion.

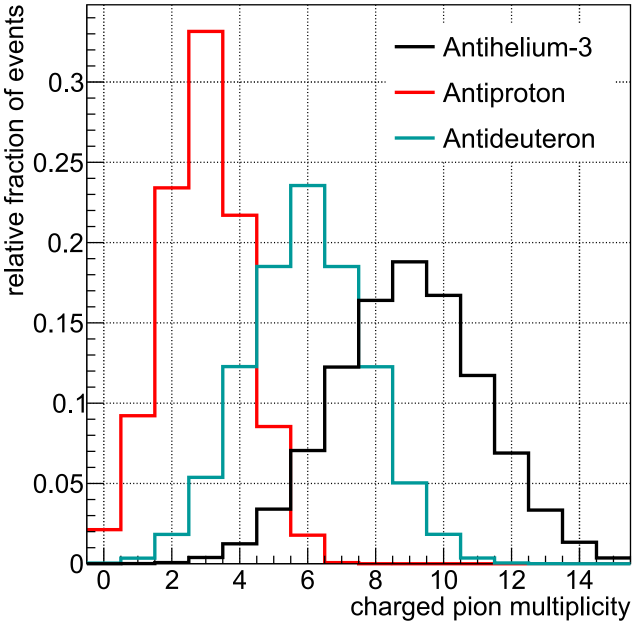

Figure 5-4: Multiplicity distributions of charged pions following (from left to right) antiproton, antideuteron, and antihelium-3 annihilation on silicon, simulated using GEANT4. Reproduced from Ref. [560] with permission. © 2021 by Elsevier B.V.

bient cosmic rays (assuming the outgoing charged particles can be traced back to a common vertex). Simulated charged-pion multiplicity distributions for antiproton, antideuteron, and antihelium-3 annihilation in silicon are shown in Fig. 5-4.

5.3 GAPS subsystems

With a total mass of nearly 2500 kg, GAPS will be one of the most complex Antarctic balloon missions ever to fly. Thus, it seems fitting to briefly outline each of the GAPS subsystems, illustrated in Fig. 5-5.

5.3.1 Silicon tracker

The silicon tracker (hereafter simply "tracker") is the heart of the GAPS detector, both conceptually and literally. Mainly developed by MIT, Columbia University

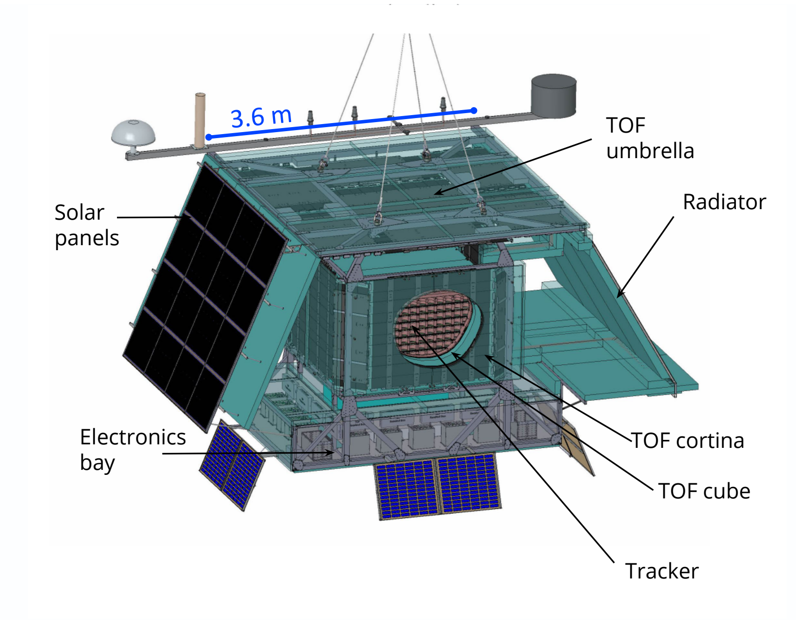

Figure 5-5: Model of the GAPS balloon experiment in its flight configuration, showing the positions of the various subsystems. For more details, see Sec. 5.3. Reproduced from Ref. [558] with permission. © 2023 by Elsevier B.V.

(CU), and the University of Hawai'i at Mānoa (UHM), the tracker fills several roles simultaneously. First, it serves as a large-area ($\sim$10 m^2) stopping target for incoming antinuclei. Second, it must be able to measure the positions and energy deposits from the incoming antinucleus and from the charged particles in the annihilation shower, as well as tracking these particles to a common vertex. Third, the tracker must be able to distinguish x-rays from antiproton and antideuteron exotic atoms, requiring energy resolutions better than 4 keV FWHM. Finally, the tracker must be operable at relatively warm temperatures near –40 °C to eliminate the need for heavy cryogenic systems.

To meet these challenging requirements, the GAPS tracker will consist of more than 1100 lithium-drifted silicon [Si(Li)] detectors manufactured by Shimadzu Corporation. Each circular Si(Li) detector has a diameter of 10 cm and a thickness of 2.5 mm, and is segmented into eight approximately-rectangular strips. A full account of the GAPS Si(Li) detector design and performance will be given in Chapter

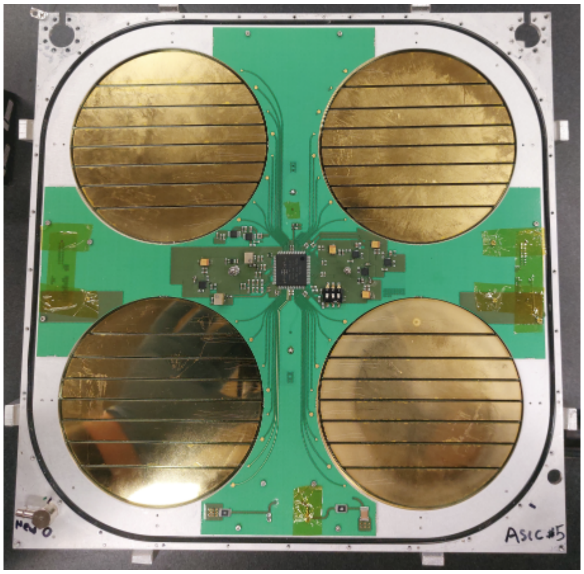

Figure 5-6: Photograph of the GAPS Si(Li) detector module with top cover removed, showing the four detectors (gold), the front-end board (green) and the SLIDER32 ASIC (black chip in the center). For flight, the traces between the wirebond pads (adjacent to each detector) and the ASIC will be covered with copper shields to reduce electromagnetic interference. Photo courtesy of Mengjiao Xiao.

6. Here, it suffices to note that the lithium-drifting process allows the economical production of large-area silicon detectors capable of maintaining good energy resolution at the relatively warm temperatures expected in flight.

The basic functional unit of the tracker is the Detector Module (hereafter simply "module"), consisting of four Si(Li) detectors mounted to a $9.5'' \times 9.5'' \times 0.125''$ aluminum frame as shown in Fig. 5-6. Each module carries a front-end board (FEB) to which each of the 32 Si(Li) strips are wirebonded and which facilitates power and communication between the modules and the rest of the instrument. Each module is enclosed by two aluminized polypropylene windows to provide a barrier against electromagnetic interference and atmospheric contaminants (e.g., water vapor). Each module is also connected via low-outgassing fluorosilicone tubing to a common manifold through which dry N_2 gas can be pumped during integration and testing on the ground, as an additional layer of protection against atmospheric contamination.

At the center of the FEB is the custom SLIDER32 ASIC chip developed by INFN, which facilitates communication between the Si(Li) detectors and the flight computer (see, e.g., Refs. [580–582]). Each of the 32 analog inputs of the SLIDER32 carries a charge-sensitive preamplifier, shaper, and sample-and-hold coupled to an 11-bit ADC. To maintain acceptable energy resolution over the full dynamic range of the tracker (tens of keV for exotic-atom x-rays to tens of MeV for annihilation products), the SLIDER32 includes both low- and high-gain modes and dynamic signal compression [583]. To avoid placing too much stress on the cooling and power systems, each chip consumes <10 mW per channel. The noise performance of the SLIDER32 will be discussed further in Chapter 6.

Within the tracker, six modules are chained together to form each row, and six rows are placed side-by-side to form each layer. The modules rest on polyethylene foam, providing both mechanical support and thermal insulation while minimizing the mass of passive material presented to incoming cosmic rays. The full tracker consists of ten layers (i.e., spaces for 1440 detectors), but only the upper seven layers will consist of functional Si(Li) detectors for the first flight. (The bottom three layers will contain a mix of non-functional silicon detectors and aluminum disks, both serving as stopping targets for highly-penetrating antinuclei.) The modules on alternating layers are rotated by 90 degrees to provide $\sim$1-cm track resolution for throughgoing particles.

5.3.2 Time-of-flight

The time-of-flight (TOF) system is the other particle detector on GAPS (see, e.g., Refs. [584, 585] for recent discussions). Principally contributed by UCLA, the TOF fulfills two primary roles: measuring the velocity $\beta = v/c$ of incoming and outgoing charged particles, and serving as a trigger for the rest of the detector to begin recording data. As shown in Fig. 5-5, the TOF consists of three functional units, each composed of a single layer of plastic scintillator paddles. Immediately outside the silicon tracker volume is the “cube,” surrounding the tracker on all six

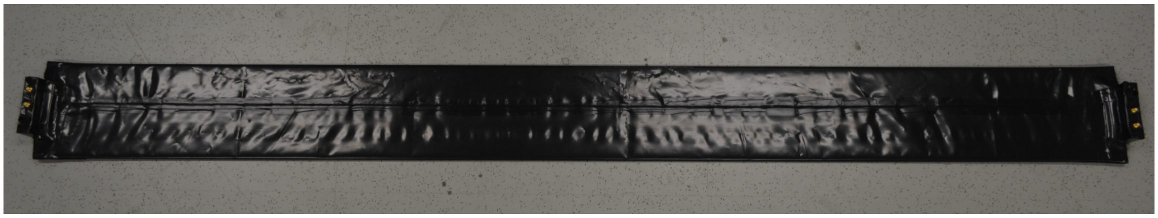

Figure 5-7: A 180-cm-long TOF paddle for the GAPS experiment, wrapped in black vinyl. The SiPMs and preamplifier boards are enclosed in the rectangular structures at each end. Photo courtesy of Takeru Hayashi.

sides. At a distance of ∼30 cm from the cube is the "cortina," surrounding the tracker on five sides (i.e., all but the bottom). Finally, at a distance of ∼90 cm above the tracker is the single-panel "umbrella." Combined, these panels ensure that particles from the upper half of the solid angle must pass through at least two TOF paddles, thereby allowing a measurement of their velocity.

The TOF system is composed of more than 180 plastic scintillator paddles manufactured by Eljen Technology. The paddles are composed of polyvinyl toluene, an organic polymer with high scintillation light yield and fast (few ns) pulse response. Each paddle is wrapped with a layer of aluminum foil to facilitate total internal reflection of the scintillation photons, as well as a layer of black vinyl to eliminate ambient light leaks. The ends of each paddle are coupled to six $6 \times 6\,\mathrm{mm}^2$ S13360-6050VE silicon photomultiplier (SiPM) detector arrays manufactured by Hamamatsu Photonics. These SiPMs have a peak photon detection efficiency of ∼40% at 450 nm, well-matched to the EJ-200 scintillator. The SiPMs are themselves mounted to a preamplifier board which sums the signals from the six SiPMs and outputs shaped high-gain (HG) and low-gain (LG) pulses. A finished paddle is shown in Fig. 5-7.

The TOF readout scheme is optimized to reduce the rate of cosmic-ray triggers from several hundred kHz to a more manageable ∼few hundred Hz, and consists of three layers: readout-and-trigger boxes (RATs), the master trigger board, and the TOF computer. Each RAT carries one power board, two readout boards, and one local trigger board (LTB), and manages the power and signals from 16 preamplifier boards (i.e., eight paddles). The LG signals are sent to the LTB, which makes the initial decision whether that group of paddles has registered an "interesting" event in $\lesssim$ 200 ns. The LTBs are connected to the master trigger board by gigabit Ethernet links, and it is the master trigger board which synthesizes the data from all of the LTBs and sends the trigger signal (as well as the clock synchronization signal) to the full TOF and Si(Li) tracker. The readout boards take the HG signals as inputs and use the DRS4 ASIC to sample the waveforms at $>$2 GS per second [586]. (Each readout board also contains a 14-bit ADC and electronics for onboard timing and amplitude calibration.) Communication between the readout boards, master trigger board, and the flight computer is handled by a dedicated TOF computer, again using dedicated gigabit Ethernet links. Tests of the TOF system using atmospheric muons demonstrate a timing resolution near 300 ps between the two ends of the paddles, providing a spatial resolution of several cm along the length of the paddle. Overall, it is anticipated that the full TOF will provide velocity resolutions better than $\sim$0.5–1% in the range $0.25 \lesssim \beta \lesssim 0.7$ (see, e.g., Refs. [584, 585]).

5.3.3 Thermal

The mass and electrical power constraints of a balloon mission place substantial limits on the GAPS thermal control system. As discussed previously, the Si(Li) detectors comprising the GAPS tracker must be maintained at temperatures near –40 °C to be operable and to provide good energy resolution. Additionally, the cooling system must not generate electromagnetic interference (EMI) which might degrade the detectors' performance, nor may it present a large material thickness which might degrade the reconstruction of low-energy charged particles.

To meet these requirements, GAPS will use a novel oscillating heat pipe (OHP) cooling system [587] developed by collaborators at JAXA. The OHP consists of 32 loops of stainless-steel capillary tube (inner diameter 1 mm, outer diameter 3 mm) running between the tracker and a large radiator mounted to the Sun-opposing side of the experiment. The tubes contain a liquid refrigerant held close to its boiling point, with each OHP loop thermally coupled to one detector module in each tracker layer by means of anodized[6] aluminum collars mounted to the corner of each module. Heat generated by the detector modules is conducted to the OHP tubes, cooling the detectors and causing a small amount of the refrigerant to vaporize. Since the vapor has a lower density than the liquid refrigerant, the former rises up through the capillary tubes, eventually reaching the radiator. The warm vapor deposits its heat into the radiator, returning to the liquid phase and beginning the cycle again. The OHP system is completely passive, requiring no mechanical pumps (though small heaters are mounted at the base of each loop and on the reservoir to control the temperature of the whole system).

5.3.4 Electronics systems

Electronics bay

Below the Tracker and TOF Cube sits the electronics bay, which houses the equipment needed to control the various instrument subsystems. At the center of the electronics bay is the flight computer, responsible for issuing commands to (and receiving data from) the various subsystems. The flight computer merges the data packets from the Si(Li) Tracker (via the backend electronics) and the TOF computer into *events* (e.g., an incoming antinucleus and its outgoing x-rays and charged particles would be associated to a single event) and performs the initial reconstruction to determine velocity, arrival direction, etc. The flight computer is also responsible for determining how much data to save for each event. Since the projected TOF trigger rate is $\sim$300 Hz, most events will be stored in a "reduced" form, containing only the Si(Li) and TOF energy deposits and the TOF timing information. For the small subset of "interesting" events (e.g., antinucleus candidates), the lower rate will allow some initial reconstruction to be performed and a more extensive set of

---

[6] To ensure that the metal OHP loops (which are essentially giant antennae) do not conduct electromagnetic noise into the Si(Li) modules.

information to be saved. As of early 2023, the GAPS data pipeline is still being finalized, but it is anticipated that these “interesting” events will also contain the full TOF waveforms and some subset of reconstructed physics quantities (velocity, arrival direction, number of secondaries, etc). Owing to telemetry limitations (discussed next), most of the data from GAPS will be saved to solid-state drives (SSDs) recovered after each flight. The electronics bay also contains the low- and high-voltage power supplies for the Si(Li) Tracker and the TOF, and their associated backend electronics. All of these systems will be powered by solar panels (capable of providing up to 1.6 kW peak power) feeding a set of five lithium iron phosphate batteries.

CSBF equipment

The Columbia Scientific Balloon Facility (CSBF) manages flight operations for the NASA long-duration balloon (LDB) program. For GAPS, there are several key pieces of CSBF equipment. The first is the communication system, which facilitates uplink (commands) and downlink (data) between the instrument and the ground stations. For most of the Antarctic balloon flight, GAPS will be out-of-range from any radio stations on the ground; thus, all communication with the payload must go through the Tracking and Data Relay Satellite System (TDRSS) or Iridium networks. The bandwidth of these systems is $\sim$few tens of Kbps, limiting the amount of data that can be downlinked during flight and thus necessitating on-board storage (the aforementioned SSDs). CSBF also provides a rotator from which the entire payload is hung to provide azimuthal pointing (e.g., for solar-power or thermal control) and a GPS system to provide near-real-time position and altitude information. Finally, CSBF provides the equipment to detach the payload from the balloon at the end of the flight[7] and the parachutes to ensure a (hopefully) soft landing on the Antarctic ice.

---

[7] Or if the balloon begins to drift off-course, e.g., toward the ocean.

5.4 GAPS status

From the first proposal[8] in 2002 [588] and antiproton beam tests in 2004–2005 [589], GAPS has been a nearly two-decade effort. In this section, I give a (very brief) timeline of some relevant milestones.

The first test of the GAPS architecture [TOF plus Si(Li) detectors] occurred in June 2012 with the flight of the prototype GAPS (pGAPS) balloon experiment. The pGAPS payload contained six Si(Li) detectors in two columns, as well as three planes of plastic-scintillator TOF panels and a prototype OHP cooling system. During its approximately seven-hour flight near Japan (three hours at $\sim$30 km altitude), pGAPS recorded more than 10^6 cosmic-ray events and demonstrated stable performance of all relevant subsystems [590]. Following the flight of pGAPS, further development work in the USA, Japan, and Italy demonstrated the viability of the GAPS concept. The project was approved for funding by NASA in 2016, and by international partner agencies shortly thereafter. During the next five years, GAPS collaborators developed prototype (and eventually flight-ready) Si(Li) detectors, TOF counters, electronics, and other hardware/software.

Integration of the GAPS experiment began at the MIT Bates Research and Engineering Center in late 2021 with the construction of the GAPS Functional Prototype (GFP). As shown in Fig. 5-8, the GFP was a sub-scale version of the full GAPS payload, consisting of 36 Si(Li) detector modules, two layers of TOF paddles, a prototype OHP cooling system, and the associated front-end and back-end electronics for each of these subsystems. The Si(Li) detector modules were arranged in three layers, with each layer consisting of two rows of six modules per row. The two layers of TOF panels were placed 0.4 m and 1.4 m above the top layer of Si(Li) detectors, providing a sufficient distance for the reconstruction of atmospheric muons up to $\beta = 1$. The GFP served as the first end-to-end test of the GAPS detector systems by tracking atmospheric muons through the TOF panels

---

[8] GAPS was originally proposed to include a gas ionization counter as the exotic-atom target. When the design shifted to silicon detectors, the name was changed to the General Antiparticle Spectrometer, but the acronym was kept (thus making GAPS somewhat of a "backronym").

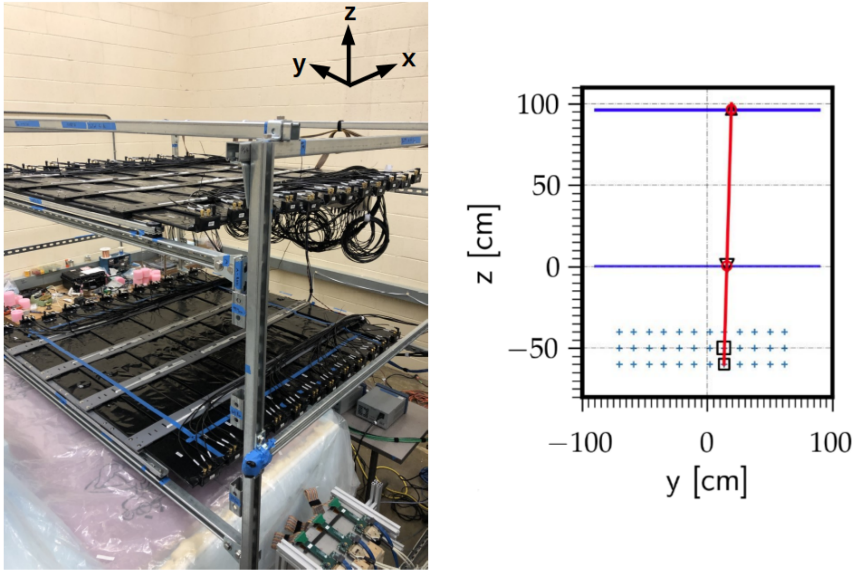

Figure 5-8: **(Left)** Photograph of the GFP at MIT Bates showing the two panels of TOF counters (black) above the Si(Li) detector modules (inside the pink/blue foam at bottom). Photo courtesy of Mengjiao Xiao. **(Right)** Display of an atmospheric muon event in the GFP. The horizontal blue lines near $z = 0$ and $z = 100\,\text{cm}$ indicate the TOF panels, and each blue cross shows one Si(Li) detector. The black squares and triangles show energy deposits, and the red line shows the reconstructed track. Courtesy of Achim Stoessl.

and then the Si(Li) detectors.

In early 2022, the integration and testing of the full GAPS experiment began at MIT Bates with the delivery of the goldola (shown in Fig. 5-9), the assembly of the OHP system, and the integration and testing of the bottom six layers[9] of Si(Li) detectors. In October 2022, the payload was shipped to the UC Berkeley Space Sciences Lab (SSL), where the final layers of Si(Li) detectors were mounted. In November 2022, integration of the TOF panels commenced. In parallel, an extensive campaign of testing the various subsystems in their flight configuration has been underway. In mid-2023, the full GAPS payload will undergo thermal vacuum testing at National Technical Systems in Los Angeles to verify functionality in the atmospheric conditions expected at float altitude. It is anticipated that further testing of the full payload will occur at the CU Nevis Laboratories in New York throughout the remainder of 2023 followed by balloon compatibility testing at the Columbia Scientific Balloon Facility in Palestine, Texas, in mid-2024. The GAPS Collaboration is working toward our first Antarctic balloon flight, currently planned for late 2024.

---

[9] Again, the bottom three layers of the tracker will consist of non-functional silicon detectors and aluminum disks for the first flight. It is hoped that future flights will include all ten layers of functional Si(Li) detectors.

Figure 5-9: The GAPS gondola (large white basket-shaped structure) at MIT Bates prior to integration. The vertical blue foam structure in the back is the cold plate (coupled to methanol chillers) used to cool the radiator during ground testing. Photo courtesy of Brandon Roach.

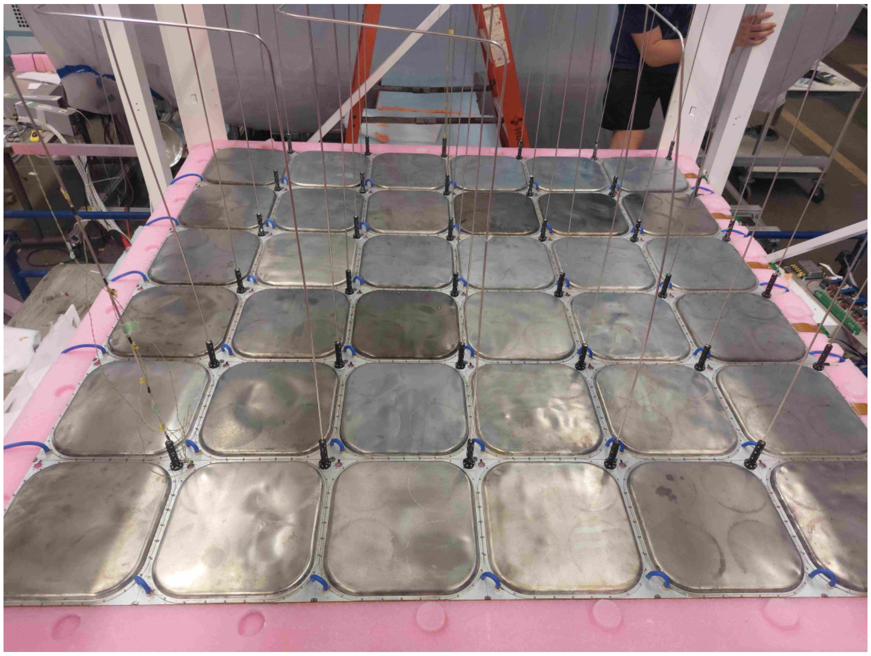

Figure 5-10: A layer of 36 GAPS Si(Li) detector modules in the tracker. The black cylinders at the corner of each module are the aluminum cooling collars through which the OHP tubes run vertically. Photo courtesy of Brandon Roach.

6

GAPS Si(Li) Detector Testing

Whether we are based on carbon or on silicon makes
no fundamental difference; we should each be
treated with appropriate respect.
—Arthur C. Clarke, *2010: Odyssey Two*

In this chapter, I describe in more detail the lithium-drifted silicon [Si(Li)] detectors that are at the core of the GAPS experiment. In Sec. 6.1, I give a brief introduction to the physics and operational principles of semiconductor (specifically silicon) detectors. In Sec. 6.2, I describe the design and fabrication of the GAPS Si(Li) detectors themselves. In Sec. 6.3, I detail the procedures and results from x-ray tests of the GAPS Si(Li) detectors at flight temperatures, with particular emphasis on the factors affecting x-ray energy resolution. I conclude in Sec. 6.5.

6.1 Introduction to silicon detectors

In particle physics, silicon detectors are a type of ionization counter, in which the energy deposited by subatomic particles is converted into an electronic signal by means of electron-ion pairs generated in the detector medium. In this section, I briefly review some elements of semiconductor device physics before moving on to the specifics of the GAPS detector design.

6.1.1 Silicon semiconductors

More than perhaps any other element, silicon has shaped the trajectory of the modern world. It is the second most-abundant element by mass in the Earth's crust after oxygen, with which it combines to form a great variety of silicate minerals. In its pure form, silicon is a bluish-gray semiconductor, having electrical resistivity intermediate between conductors and insulators. The behavior of semiconductors can be understood in terms of their band structure: in the crystal, the uncountably many electron energy levels form continuous bands. In the valence band, electrons are tied to their atoms, but in the conduction band, electrons can propagate through the crystal, thereby conducting current. Most silicon semiconductors used in devices ranging from computer microprocessors to particle detectors are generally *doped* with impurities to alter their electrical characteristics. Silicon dopants are generally divided into two types: *p*-type atoms with three valence electrons (e.g., boron or aluminum) or *n*-type atoms with five valence electrons (e.g., phosphorus or arsenic). Since silicon has four valence electrons, *n*-type atoms can "donate" their valence electron to the crystal. Similarly, *p*-type atoms can be thought of as having a *hole* in their valence shell which can be filled by an electron. (Note that the dopants are still neutral atoms, and thus the overall crystal is neutral as well.)

6.1.2 *p*-*n* junctions

When a *p*-type and *n*-type semiconductor are joined together, a *p*-*n* junction is formed. Electrons from the *n*-type material diffuse into the *p*-type material, leaving the *n* side with a slight positive charge and the *p* side with a slight negative charge. A *depletion region* with no free charge carriers forms between the *p*- and *n*-type materials, and an electric field $\boldsymbol{E}$ pointing from *n* to *p* across the depletion region eventually resists the further flow of charge, and an equilibrium is established. Any charge carriers generated in the depletion region (e.g., by ionizing radiation) are swept out of the region by the ***E***-field, making the *p*-*n* junction a promising radiation detector. The depth of the depletion region is (e.g., Ref. [591])

$$x = x_n + x_p = \sqrt{\frac{2\epsilon V_0}{q} \frac{(n_a + n_d)}{n_a n_d}}, \tag{6.1}$$

where ϵ is the permittivity of the material, $V_0 \approx 1\,\mathrm{V}$ is the bias generated by thermal diffusion near room temperature in silicon, q is the elementary charge, and n_a and n_d are the number densities of acceptor (*p*-type) and donor (*n*-type) atoms, respectively. If the *n*-side of the junction is more heavily doped than the *p*-side (i.e., $n_d \gg n_a$), then the depletion region is almost entirely on the *p*-side, having a width[1] (e.g., Ref. [591])

$$x \approx 0.32\,\mu\mathrm{m} \sqrt{\left(\frac{\rho_p}{1\,\Omega\,\mathrm{cm}}\right)\left(\frac{V_0}{1\,\mathrm{V}}\right)}, \tag{6.2}$$

where ρ_p is the resistivity of the *p*-type semiconductor. If $\rho_p \sim 10^3\,\Omega\,\mathrm{cm}$ (already a reasonably high resistivity), the depletion region is only ~10 μm wide, which leads to diodes with poor detection efficiency.

To remedy this issue, most *p*-*n* junction detectors are operated in reverse-bias mode. In this configuration, the positive terminal of a voltage supply is connected to the *n*-type material, with the other terminal connected to the *p*-type. The electric field induced by the voltage supply serves to "pull" the depletion region wider—simply replace $V_0 \approx 1\,\mathrm{V}$ in Eq. 6.1 with the appropriate bias voltage. Unfortunately, a variety of physical and practical limitations prevent V_0 (and hence the depletion depth) from becoming arbitrarily large. First, increasing the bias voltage runs the risk of dielectric breakdown in the bulk material—i.e., the flow of (large, damaging) currents even in the absence of ionizing radiation. This effect can be somewhat mitigated with high-resistivity ($\rho \gtrsim 10^3\,\Omega\,\mathrm{cm}$) silicon crystals, though generally at

---

[1] For a heavily *p*-doped semiconductor (i.e., $n_a \gg n_d$), replace ρ_p with ρ_n and 0.32 with 0.52 in Eq. 6.1.

a significant monetary cost. Furthermore, the environment of balloon payloads is generally not favorable to the high voltages ($\gtrsim 1\,\mathrm{kV}$) needed to operate such detectors. On the one hand, generating these high voltages on a balloon payload is difficult from an electrical-engineering perspective. On the other hand, the small amount of residual atmosphere at float altitude means that electrical arcs can develop between sensitive components if the voltage is too high (see, e.g., Ref. [592]). Taken together, all of these factors favor GAPS detectors which can be fully depleted with a bias of only $\sim$250 V.

6.1.3 Compensation and lithium drifting

Producing high-resistivity *intrinsic* (i.e., non-doped) semiconductor material is extremely challenging. Trace amounts of impurities (e.g., oxygen or carbon at or below the parts-per-billion level) can dramatically affect the conductivity of the material. Despite these challenges, it is possible to produce high-resistivity semiconductors via the process of *compensation*—i.e., doping an already-doped semiconductor with the opposite kind of atom (e.g., doping a *p*-type material with *n*-type atoms). Junctions produced in such a way are called *p-i-n* junctions, where the *i* refers to the compensated[2] material.

A particularly attractive dopant for semiconductor compensation is lithium. Owing to its $\sim$0.03-eV effective ionization energy in silicon [593], lithium serves as a donor (*n*-type) atom. The idea of diffusing lithium atoms into semiconductors—a process known as *drifting*—has been studied since the 1960s (see, e.g., Refs. [594–597] for early work, and Refs. [598, 599] for reviews). Since lithium has a high mobility in silicon and germanium[3], it can be drifted through *p*-type semiconductors to form large compensated regions, up to several mm thick. These lithium-drifted silicon—Si(Li)—detectors have since found widespread use in many nu-

---

[2] Somewhat confusingly, *i* does not mean "intrinsic"!

[3] Lithium-drifted germanium—Ge(Li)—is outside the scope of this dissertation. It is worth noting, however, that lithium atoms are *so* mobile in germanium that allowing a Ge(Li) detector to warm above liquid-nitrogen temperatures is very likely to ruin the lithium-drifted region. Modern high-purity germanium (HPGe) detectors made of intrinsic Ge do not have this issue.

clear physics, particle physics, and astrophysics experiments (see, e.g., Refs. [600–613] for some examples). Early Si(Li) detectors were small (a few cm^2), but improvements in production methods eventually allowed large-area detectors to be deployed whilst maintaining good energy resolution. It should be noted, however, that many of these large-area Si(Li) detectors had to be cooled—some to liquid-nitrogen temperatures—to reduce leakage current (a dominant contribution to detector noise, and hence energy resolution).

6.2 GAPS Si(Li) detectors

Si(Li) detectors are ideal for the GAPS tracker described in Chapter 5, owing to their large sensitive areas, thick depletion regions, and reasonable bias-voltage requirements. However, the constraints of a balloon experiment require these thick, large-area detectors to provide $<$4-keV FWHM energy resolutions at relatively warm temperatures near $-40\,^\circ\text{C}$, more than 100 K above cryogenic detectors. To produce the more than 1100 Si(Li) detectors needed for the first flight, the GAPS Collaboration contracted with Shimadzu Corporation of Japan, with whom we developed a cost-effective mass-production protocol. Detailed descriptions of the process are given in Refs. [559, 614–616]; this section gives a brief review of important aspects.

6.2.1 Si(Li) detector design

The design of the GAPS Si(Li) detector is shown in Fig. 6-1. Each circular detector has a 10-cm diameter and a 2.5-mm thickness, with eight $\sim 8\,\text{cm}^2$ strips defined by grooves cut into the top surface. Near the outer edge of the detector is a circular *guard ring*, which prevents leakage current from flowing into the active strips, thereby reducing noise (see, e.g., Refs. [599, 618]). The depletion region (i.e., sensitive volume) of the detectors is nearly 2.3 mm thick, whilst requiring only 250 V to fully bias.

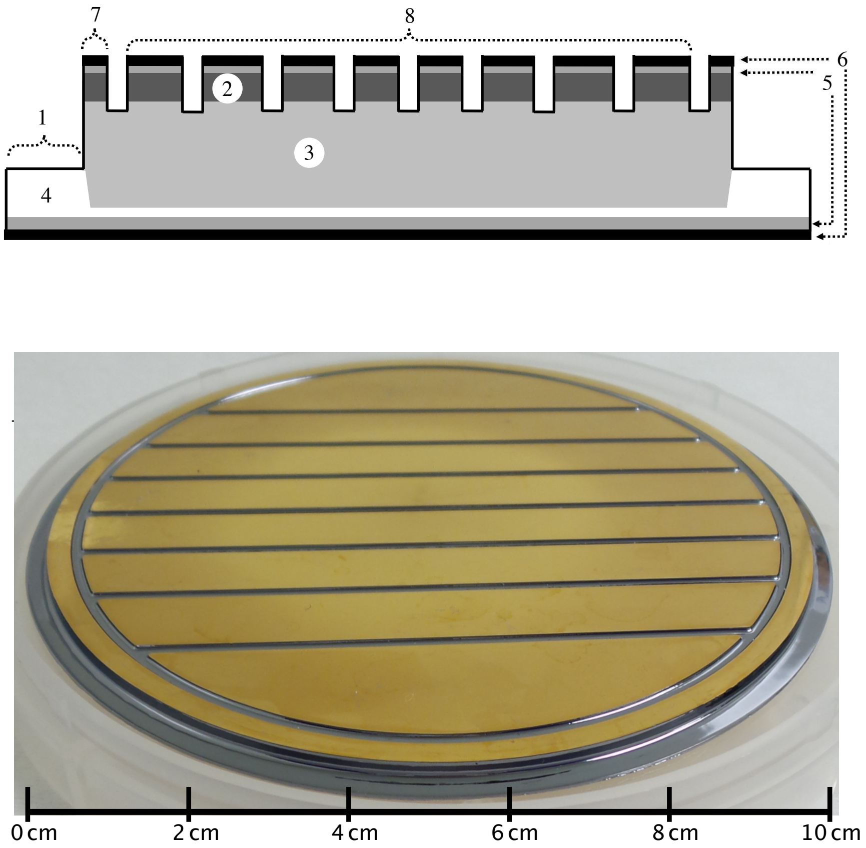

Figure 6-1: **(Top)** Cutaway view of a GAPS Si(Li) detector (not to scale). Different features are as follows: (1) top-hat brim, (2) n^+ layer, (3) lithium-drifted region, (4) undrifted lithium layer, (5) nickel layer, (6) gold layer, (7) guard ring, and (8) active strips. **(Bottom)** Photograph of a GAPS Si(Li) detector prior to passivation. The gold-colored surfaces are the gold electrical contacts, and the grayish surfaces in the grooves and top-hat brim are exposed silicon. This figure was reproduced from Ref. [617] with permission. © 2019 by IOP Publishing.

The GAPS Si(Li) detector design is a compromise between several factors. First, the volume of the depletion region requires special attention to be paid to controlling leakage current, particularly at the relatively warm GAPS operating temperatures. Second, it is advantageous for the area of each strip to be as small as possible, since the strip capacitance (which contributes to noise) is proportional to area. (Fortunately, the strip capacitance decreases as the depletion depth increases.) Furthermore, small strips provide improved position resolution for ionizing radiation events, and thus improved track-reconstruction efficiency. On the other hand, the complexity of detector fabrication and readout increases as the strip size decreases, with the latter being a particularly significant constraint for balloon payloads such as GAPS. The final eight-strip design is a compromise between these factors, with strip capacitance generally between 37–39 pF and strip width $\sim$1 cm. Additionally, the GAPS detectors are operable at temperatures exceeding –40 $^\circ$C either at atmospheric pressure or in vacuum, eliminating the need for large and heavy cooling systems and pressure vessels.

6.2.2 Si(Li) detector fabrication

A full accounting of GAPS Si(Li) detector production is given in Refs. [559, 615], from which this section is summarized. GAPS detector production begins with high-resistivity ($\rho \sim 10^3\,\Omega\,\mathrm{cm}$) p-type silicon supplied by SUMCO Corporation, grown in the $\langle 111 \rangle$ crystal orientation using the float-zone process to keep atomic impurity concentrations (e.g., carbon and oxygen) below $(1\text{–}2) \times 10^{16}\,\mathrm{cm}^{-3}$ (i.e., below $\sim$0.1 ppm). The resulting 10-cm-diameter ingots are diced into 2.5-mm-thick wafers, from which the Si(Li) detectors are made. The wafers are then cleaned with a solution of methanol, xylene, and acetone to remove any organic contaminants, followed by an immersion in an etchant consisting of hydrofluoric acid, nitric acid, and acetic acid to remove oxides.

The n^+ contact and top-hat brim

The top (strip side) n^+ electrode begins by placing the wafer into a thermal vacuum chamber (pressure $< 10^{-4}$ Pa, temperature $\sim$280 °C) and passing electric current through a small sample of lithium. The evaporated lithium is deposited onto the top surface of the wafer. After $\sim$1 minute the lithium evaporation is terminated, but the chamber temperature and pressure are maintained for another $\sim$1 minute to allow further lithium drift. The detector is then returned to room temperature over $\sim$150 minutes to allow the lithium to diffuse into a $\sim$100-µm-thick layer, forming the top n^+ layer. The surface is then cleaned to remove any oxides. To produce a conductive contact atop the n^+ layer, the vacuum chamber is again pumped down to $< 10^{-4}$ Pa and an 18-nm-thick layer of nickel is evaporated onto the top n^+ layer, followed by a 120-nm-thick layer of gold. Gold was chosen for its electrical conductivity and corrosion resistance, with the nickel chosen for its high adhesion to silicon. An ultrasonic impact grinder is used to machine away the outer circumference of the wafer, leaving a 1-mm-thick "top-hat brim" of p-type silicon. Finally, the n^+ and p contacts are coated in Apiezon® wax dissolved in xylene to protect them from the etchant used to clean the top-hat.

Lithium drifting

Achieving a uniform lithium drift within each detector is crucial for low-leakage-current (and thus low-noise) operation. In collaboration with Shimadzu Corporation, the GAPS Collaboration developed a custom lithium-drifting apparatus, consisting of a heater, voltage supply[4], and current monitor. Over the course of $\sim$3 hours, the detector bias voltage is increased from zero to 600 V in 100-V steps. The width x of the drifted region scales as

$$x = \sqrt{2\mu_{\mathrm{Li}} V t}, \tag{6.3}$$

---

[4] During lithium drifting, the positive voltage is applied to the n^+ side, rather than a negative voltage being applied to the opposite side.

where μ_{Li} is the lithium mobility (itself a function of temperature), V is the bias voltage, and t is time. During the drift, the heater keeps the detector at 100 °C, though its output decreases over time as resistive losses in the ever-deepening drifted layer increase. The drifting procedure is terminated after ∼6200 minutes or once the leakage current begins to increase beyond ∼25 mA, whichever occurs first. This keeps a thin (∼100-μm) layer of undrifted silicon between the drifted region and the *p*-type contact, which further reduces leakage current. Detector conditions during the drifting process (temperature, leakage current, time, etc) are recorded for each detector for later analysis.

Groove machining and the *p* contact

Following lithium drifting, an ultrasonic impact grinder is used to machine additional grooves into the n^+ side of the detector. The first is a ∼0.3-mm-deep and 1-mm-wide groove defining the guard ring, as shown in Fig. 6-1. When the detector is reverse-biased, the high resistivity of the guard ring isolates the entire n^+ layer from the edge of the detector (particularly the top-hat, whose exposed surfaces are a significant source of leakage current), significantly reducing noise. Additional grooves are cut through the n^+ contact to form eight approximately equal-area strips, ensuring that each strip is separated from its neighbors by high-resistivity lithium-drifted silicon. The detector surface is cleaned to remove any debris and coated with etch-resistant wax, after which the top-hat and *p*-side are etched. A nickel-gold contact is then evaporated onto the *p*-side using the same technique as for the n^+ contact. This monolithic *p*-side contact is used to supply the detector bias voltage.

Final etching

After covering the n^+ and *p* electrodes with wax (i.e., leaving the grooves and top-hat bare), the detector is cleaned with a methanol-based solvent to remove wax and to set the exposed silicon to be lightly *n*-type. The detector then undergoes a final series of etchings to clean and smooth the exposed grooves, which also causes the leakage current to decrease by 2–3 orders of magnitude at cold temperatures. A total etching time of 15 minutes (one 10-minute etching plus an additional 5-minute etching) was determined to be sufficient to reduce the leakage current while not removing too much material from the grooves.

Initial electrical measurements

Once a Si(Li) detector is produced, it undergoes several electrical measurements at JAXA. Of particular importance are the strip capacitance C_{strip} and leakage current I_{leak}, since they are major contributors to the noise budget (see, e.g., the discussion in Sec. 6.4). Additionally, since $C_{\text{strip}} = \epsilon A/x$, where $\epsilon \sim 1.1 \times 10^{-4}\,\text{pF}\,\mu\text{m}^{-1}$ is the permittivity of silicon, A is the strip area, and x is the depletion depth, the uniformity of C_{strip} at the operating bias of –250 V is a check on the uniformity of x. As described in Ref. [616], more than 95% of the GAPS Si(Li) strips have C_{strip} in the range 38–40 pF at room temperature[5], with the most probable value being 39 pF. Room-temperature leakage currents for each strip and guard ring were also measured for reverse-bias voltages from zero to –500 V, with the most-probable value at –250 V being $\sim$2 μA. Full details of these and other measurements may be found in Ref. [616]. Following this initial testing at JAXA, each Si(Li) detector is enclosed in a plastic case, placed into an antistatic bag with silica desiccant (to suppress humidity), and vacuum-sealed for shipment to CU.

Passivation

At this stage, the GAPS Si(Li) detectors are usable as ionizing radiation detectors, delivering the required x-ray energy resolution, but are extremely vulnerable to environmental effects. For one, microscopic metal flakes from the electrodes can land in the grooves between adjacent strips, or between strips and the guard ring, drastically increasing the leakage current. For another, water vapor and organic

---

[5] C_{strip} is expected to decrease by $\sim$1–2 pF at the GAPS operating temperature near –40 °C.

compounds can deposit onto the bare silicon surfaces of the grooves and top-hat brim. Since the GAPS detectors must remain inside the payload for many months during integration, testing, and flight, it is not feasible to continuously clean the detector surfaces.

The *passivation* process led by CU provides a barrier against degradation of the GAPS detectors, ensuring that the Si(Li) detectors are well-protected from the environment following production. Unlike previous works using thermal deposition[6] of silicon oxides or nitrides, the GAPS detectors are passivated with polyimides, a type of organic polymer. (For full details, see Refs. [559, 619].) Each detector was first cleaned with acetone, methanol, and deionized water, and the bare silicon was etched with a mixture of hydrofluoric acid, nitric acid, and glacial acetic acid. To promote adhesion of the polyimide passivation layer, a 0.1% solution of γ-Aminopropyltriethoxysilane (APS) in deionized water was applied to the detectors and heated at 110 °C for 30 minutes. The silicon surfaces were then coated in the polyimide precursor[7], and the detectors were placed into an oven which was brought to 210 °C at a rate of $5\,^\circ\mathrm{C\,min^{-1}}$. The detectors were removed from the oven after 60 minutes at 210 °C. We note that the room-temperature leakage current of each strip and guard ring was measured before and after passivation, to verify that the performance had not degraded. Room-temperature leakage-current measurements for each strip and guard ring were performed before and after passivation using a Keithley 487 picoammeter/voltage source in 10-V increments from 0 V to –400 V. When measuring the leakage current of a given strip, the other seven strips plus guard ring were grounded together, and the detector was placed inside a dark box to avoid photovoltaic effects. (These measurements also verified that the passivation process did not significantly increase the leakage current.) Finally, accelerated room-temperature environmental testing equivalent to 14 days at 50% relative humidity, or 6 months of exposure to outgassing from organic materials

---

[6] Since the drifted lithium is vulnerable to decompensation at temperatures $\gtrsim 250\,^\circ\mathrm{C}$, these processes are not suitable for GAPS detectors.

[7] We used VTEC PI-1388, a mixture of diamines and dianhydrides in *N*-Methyl-2-pyrrolidone (NMP) solvent. Specifically, the final protocol combines VTEC PI-1388 and additional NMP in a 1:1 ratio, which also enables (much easier) application to the detectors by pipette.

(including the fluorosilicone, G10, and printed-circuit-board materials in the GAPS detector modules) showed no significant degradation in x-ray energy resolution. The long-term storage of the GAPS detectors is discussed later in this chapter.

6.3 GAPS Si(Li) testing campaign

From late 2019 to early 2022, groups at MIT[8] and UHM tested each of the more than 1100 Si(Li) detectors that were produced by Shimadzu. This testing served two purposes: to ensure that the detectors were operable in the full GAPS instrument, and to characterize their x-ray energy resolution.

6.3.1 Si(Li) quality and pre-selection modeling

To categorize the performance of the Si(Li) strips and detectors, we create three broad classes—x-ray, tracking, and non-usable—as shown in Tables 6.1 and 6.2. At the operating temperature of –37 °C, there are two factors governing the performance of the Si(Li). The first is a constraint on the leakage current: functional strips must have $I_{\mathrm{leak}} < 50\,\mathrm{nA}$ (above which the ASIC readout saturates), and functional detectors must have $I_{\mathrm{tot}} < 1\,\mu\mathrm{A}$ (above which the flight power system can no longer fully bias the detectors). A non-functional strip with high leakage current will simply not be read out in flight, but non-functional detectors cannot be biased (and are thus never connected to the flight power supply). The second constraint on detector quality arises from x-ray energy resolution. As described in Chapter 5, an energy resolution of 4 keV FWHM is needed to distinguish between x-rays from antiprotonic and antideuteronic exotic atoms. Thus, we define *x-ray-quality* strips to have anticipated x-ray energy resolutions $<$ 4 keV FWHM using the flight ASIC readout. Strips with energy resolutions $>$ 4 keV FWHM are still usable for charged-particle tracking, so they are called *tracking-quality*. We also cat-

---

[8] I was actively involved in the testing and analysis at MIT, supervised by our postdoctoral researcher Mengjiao Xiao. I am also grateful to Ian Bouche (postbaccalaureate researcher), Anika Katt (undergraduate researcher), and Kelsey Yee (graduate researcher) for their time and dedication to this effort. It was a genuine pleasure working with all of you.

Strip quality	Pre-selection [room-temp.]	Cold [–37 °C]	Yield
X-ray	$I_{\mathrm{leak}} < 2\,\mu\mathrm{A}$	ASIC FWHM < 4 keV and $I_{\mathrm{leak}} < 50\,\mathrm{nA}$	7090
Tracking	$2\,\mu\mathrm{A} < I_{\mathrm{leak}} < 12.5\,\mu\mathrm{A}$	ASIC FWHM < 4 keV and $I_{\mathrm{leak}} > 50\,\mathrm{nA}$	944
Non-usable	$I_{\mathrm{leak}} > 12.5\,\mu\mathrm{A}$	Unstable waveform or $I_{\mathrm{leak}} > 50\,\mathrm{nA}$	854
Total	—	—	8888

Table 6.1: Definitions of Si(Li) strip quality used in the lab testing described in Sec. 6.3. Strip leakage currents I_{leak} are referenced at –250 V. This table will appear in the published version of Ref. [4].

egorize the usable detectors into x-ray and tracking quality grades, based on how many functional and x-ray-quality strips they have.

To group detectors which will likely have similar performance into modules (thereby minimizing "reshuffling" of detectors between modules after initial x-ray testing), we devise a series of pre-selection criteria based on room-temperature strip (Table 6.1) and detector (Table 6.2) leakage currents measured at room temperature following passivation. Note that all detectors were tested at –37 °C, regardless of their pre-selection quality. The effects of different temperatures on performance will be discussed in Sec. 6.4.3.

6.3.2 Detector storage and shipping

Before describing the x-ray testing setup, it is important to briefly discuss protocols for detector storage and shipping. For long-term storage at CU, the Si(Li) detectors are placed in freezers maintained at –20 °C to slow the rate of lithium drift. Before being placed in the freezer, each detector is placed in a vacuum-sealed antistatic bag along with a silica desiccant pack to maintain relative humidity RH $<$ 5%. For short-term storage (e.g., for testing at MIT or UHM), detectors are kept in room-temperature dryboxes maintained at RH $\lesssim$ 10–15% and covered with blackout

Detector quality	Pre-selection [room-temp.]	Cold [–37 °C]	Yield
X-ray	$\geq$7 x-ray strips and $I_{\text{tot}} < 300\,\mu$A	$\geq$7 x-ray strips and $I_{\text{tot}} < 1\,\mu$A	776
Tracking	$I_{\text{tot}} < 300\,\mu$A	ASIC FWHM < 4 keV and $I_{\text{tot}} < 1\,\mu$A	235
Non-usable	$I_{\text{tot}} > 300\,\mu$A	Unstable waveform or $I_{\text{tot}} > 1\,\mu$A	100
Total	—	—	1111

Table 6.2: Same as Table 6.1 but for Si(Li) detector quality. Note that detectors flagged as non-usable during pre-selection were still tested at –37 °C. I_{tot} refers to the total current drawn by the detector at –250 V. This table will appear in the published version of Ref. [4].

cloths to prevent photodegradation.

The protocol for Si(Li) module shipment[9] between the various testing sites is optimized to provide mechanical and environmental protection for the modules. Each module is placed into an antistatic bag with silica desiccant packs and lightly vacuum-sealed, before being sealed in another antistatic bag. The bagged module is then placed into a custom antistatic foam insert for shipping.

We have also studied the stability of performance (i.e., x-ray energy resolution) for several dozen GAPS Si(Li) detectors during the course of 2019–2022. Ref. [559] reported no significant change in x-ray energy resolution for detectors kept in cold storage between September 2019 and September 2020. A more comprehensive stability analysis for all $>$ 1100 detectors is ongoing and will likely be published later, but preliminary results suggest no significant changes in performance for timescales as long as two to three years, either in cold storage or in low-humidity room-temperature dryboxes.

---

[9] We thank FedEx for their support in this effort.

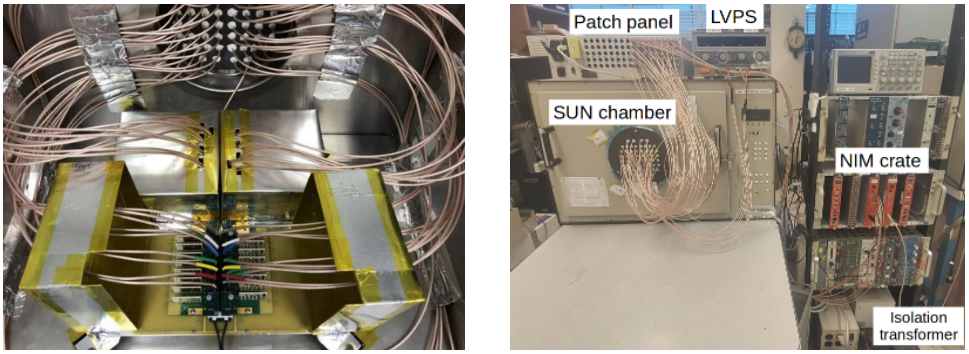

Figure 6-2: **(Left)** Photograph of a GAPS Si(Li) module mounted to the SUN chamber door, showing the cable feedthroughs, Faraday cages, discrete preamplifier holder boards (green), and FR4 top board (yellow). **(Right)** Photograph of the Si(Li) module testing system at MIT. For more details, see Sec. 6.3.3. This figure will appear in the published version of Ref. [4].

6.3.3 Testing setup

The setup used to test GAPS Si(Li) detector modules at cold temperatures is shown in Figs. 6-2 and 6-3, and shares many similarities with the single-detector testing setup described in Ref. [617]. Two module testing systems—one at MIT and one at UHM—were built using the same commercially-available components. (Much of this material can also be found in Ref. [4], of which I was a main author.)

Charge-sensitive discrete preamplifiers

Since the SLIDER32 ASIC was still under development during the Si(Li) testing campaign, we use custom charge-sensitive discrete[10] preamplifiers (hereafter simply *preamplifiers*) to process the signals from the Si(Li) strips. A simplified preamplifier design is shown in Fig. 6-4, and consists of three main components: an inverting operational amplifier (op-amp), a feedback capacitor C_F, and a feedback resistor R_F. Consider a voltage V_{in} (e.g., from the accumulation of charge on the strip capacitance C_{strip}) presented to the input of an ideal op-amp with gain $\mathrm{d}V_{\text{out}}/\mathrm{d}V_{\text{in}} = -A$. Assuming no other parallel capacitance in the readout chain,

---

[10] Here, “discrete” refers to the use of discrete electronic components (e.g., resistors, capacitors, and transistors), as opposed to integrated circuits.

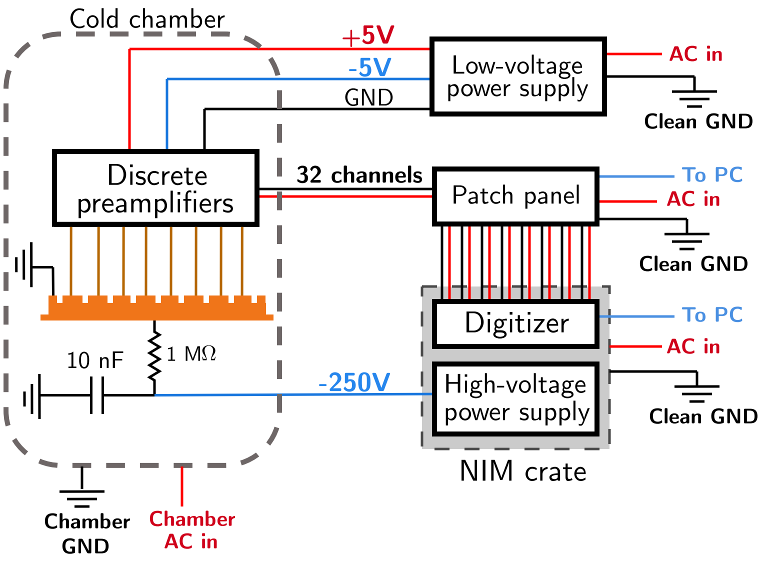

Figure 6-3: Schematic of the electrical connections for the GAPS Si(Li) testing campaign described in Sec. 6.3.3. This figure will appear in the published version of Ref. [4].

the ratio of the charge Q_{F} on C_F to the initial charge Q_{strip} on the Si(Li) strip is [620]

$$\frac{Q_{\mathrm{F}}}{Q_{\mathrm{strip}}} = \left(1 + \frac{C_{\mathrm{strip}}}{C_F(A+1)}\right)^{-1}. \tag{6.4}$$

Thus, assuming $A \gg C_{\mathrm{strip}}/C_F$, essentially all of the charge Q_{strip} will land on C_F. This produces a voltage $V_{\mathrm{F}} = (A+1)V_{\mathrm{in}}$ across C_F, which is discharged through the parallel feedback resistor R_F. This generates a voltage tail pulse at the output, proportional to e^{-t/τ_F}, where the time constant $\tau_F = R_F C_F$.

The discrete preamplifiers used for GAPS detector testing are based on the architecture of Ref. [621]. Each Si(Li) strip is read out by its own preamplifier, incorporating a 100-MΩ feedback resistor and a 0.5-pF feedback capacitor (i.e. $\tau_F = 50\,\mu\mathrm{s}$). A low-noise n-channel junction field effect transistor (JFET) is included at the input of each amplifier. To read out an entire Si(Li) detector, eight preamplifiers can be placed into a single holder board, providing signal input, low-voltage power (±5 V and ground), and output connections. Spring-loaded

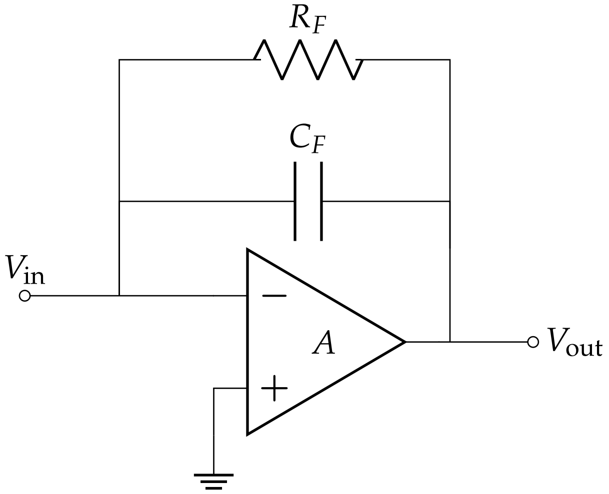

Figure 6-4: Sketch of a simple charge-sensitive preamplifier based on an inverting op-amp with gain A, feedback capacitance C_F, and feedback resistance R_F. For simplicity, the op-amp bias voltages are not shown.

conductive pins are used to couple each Si(Li) strip to its associated readout channel. A custom low-noise gain stage is also attached to the output of the preamplifiers, increasing the output voltage (and hence the dynamic range) by a factor of five.

Si(Li) detector module

The Si(Li) detector modules used in laboratory testing are very similar in design to the flight modules described in Chapter 5, with several modifications to accommodate the discrete preamplifier (as opposed to ASIC) readout. Most significantly, the top aluminized plastic window is replaced with a custom FR4 board shown in Fig 6-2. This board includes a terminal for the –250-V detector bias, with a 1-MΩ resistor mounted on the bias line for overcurrent protection. A 10-nF capacitor between the bias line and ground serves to filter noise on the input side. The underside of the FR4 board includes spring-loaded conductive pins which couple the Si(Li) strips to conductive pads on the top surface of the board. Finally, the board includes mounts for the preamplifier holder boards described previously, whose conductive pins couple to the pads on the FR4 board. To shield the Si(Li) detectors and preamplifiers from electromagnetic interference, we place a custom $\sim 10 \times 10 \times 10\,\mathrm{cm}^3$ Faraday cage (composed of 0.6-mm-thick aluminum sheets) atop each detector as shown in Fig. 6-2, grounding each cage to the aluminum module frame (and thus to the common ground).

Environmental chamber

To simulate the approximately –40 °C temperature expected in flight, we test the Si(Li) detector modules in a SUN Electronics EC13 environmental chamber (hereafter "SUN chamber"). The chamber includes inlets for liquid nitrogen to facilitate cooling. To ensure a low-humidity environment throughout testing, we also continuously flush a small amount of dry nitrogen gas through the chamber. The temperature within the chamber is controlled by the onboard microcontroller, with a built-in fan ensuring a uniform atmosphere within the chamber. The SUN chamber doors include an attached tray, to which the detector module is mounted (though ensuring the module is electrically isolated from the door). As will be discussed in Sec. 6.3.4, the fan, cooling, and heating systems are turned off while performing the x-ray spectral measurements to reduce electromagnetic noise.

Cable feedthroughs

To provide low-voltage and high-voltage inputs, as well as signal outputs from the 32 Si(Li) strips, we modify the SUN chamber doors to include custom cable feedthroughs. These feedthroughs are based on custom-machined black Teflon flanges, with holes drilled through to accommodate the RG316 coaxial cables (32 cables for Si(Li) strip readout, eight cables for preamplifier low-voltage input (one +5 V and one –5 V cable per holder board), one cable for high-voltage detector bias, and several spare cables for redundancy). The coaxial cables were potted to the flange using Torr Seal®, a low-outgassing and vacuum-safe epoxy suitable for low-temperature work.

Cadmium-109 radioactive source

To generate the x-rays[11] used to test the energy resolution of the Si(Li) detectors, we use a ^{109}Cd source with initial activity 1 mCi. This isotope has several distinct advantages. First, ^{109}Cd undergoes electron-capture decay with a half-life of $\sim$463 days to the first excited nuclear state of ^{109}Ag, the de-excitation of which produces an 88-keV photon [622]. As described in Chapter 5, this is in the expected 20–100-keV energy range for exotic-atom x-rays in silicon. Additionally, since ^{109}Cd only decays via electron capture (and ^{109}Ag is stable), there are no other charged particles[12] that might cause unwanted backgrounds in the detectors during testing. Finally, ^{109}Cd also emits lower-energy x-rays between $\sim$22–25 keV, allowing us to test the low-energy detector response (though this is outside the scope of this dissertation).

External electronics

The external electronics consist of several layers. The first is the Tripp-Lite® 1-kW power isolation transformer, which provides low-noise AC power to the rest of the electronics. The signals from all 32 Si(Li) strips are connected to a CXAR-64/MF patch panel, allowing us to select up to eight strips (i.e., one full detector) to send to the CAEN N6725 digitizer. The digitizer is placed in a NIM crate alongside the CAEN 1471A high-voltage power supply (HVPS), the latter of which supplies the –250-V detector bias voltage and can monitor supply currents as high as 300 μA with a resolution of 5 nA. We note that the onboard cooling fan in the NIM crate was disabled and replaced with a separate fan in front of the crate to reduce electromagnetic noise pickup. The digitizer and patch panel are controlled by a desktop computer using custom software developed by collaborators at UHM [623].

---

[11] By "x-rays" I mean "photon with energy less than $\sim$100 keV." In nuclear physics, these 88-keV photons are classified as γ-rays, since they come from nuclear (rather than atomic) processes.

[12] Strictly speaking, ^{109}Cd can also undergo internal conversion, transferring some of its decay energy to orbital electrons. However, the low energies of these electrons (up to a few tens of keV) mean that they are very easily stopped by passive materials above the detectors (aluminum Faraday cages, FR4 board, etc).

6.3.4 Testing procedure

After mounting the Si(Li) detector module to the SUN chamber door and connecting it to the electrical systems, we cool the chamber to a conservative –37 °C over the course of $\sim$45 minutes (though some detectors were also tested at –43 °C to study the temperature dependence of the leakage current and A_f, as described in the next section). Once the chamber has stabilized at the desired temperature, we apply the high-voltage detector bias, ramping from zero to –250 V at a rate of $\sim$3 V s^{-1}. Once the current (as monitored by the CAEN HVPS module) has stabilized, we record x-ray events for each detector for three minutes, with the chamber heating, cooling, and fan disabled to reduce electromagnetic noise pickup. The rest of this section describes how we convert the voltage pulses from the preamplifiers into x-ray energy spectra.

Pulse shaping

As will be described extensively in the next section, the energy resolution of a semiconductor detector depends on how the signals are *shaped*. All amplifiers shape pulses to some extent (e.g., the tail pulses produced from the aforementioned charge-sensitive preamplifiers), but further shaping is often pursued to achieve additional goals. These include maximizing the signal-to-noise ratio, decreasing pileup[13], and so on. As described in Sec. 6.4, the noise characteristics of a semiconductor detector depend on the manner and timescale over which the pulse is shaped. A common measurement is the *peaking time* τ_p, the time required for a pulse to rise from its baseline value to its maximum.

Experimental measurements of the noise characteristics of an x-ray detector generally proceed in the following way. First, a given value of τ_p is chosen[14], and an x-ray energy spectrum is collected. The energy spectrum is then fit to extract the energy resolution (e.g., the FWHM of a Gaussian peak for a mono-energetic

---

[13] Signals with long tails are prone to overlapping at high rates, distoring amplitude measurements.

[14] In the case of the CAEN N6725 digitizer, this is typically accomplished using the included MC^{2}ANALYZER software, though we use an alternative approach discussed next.

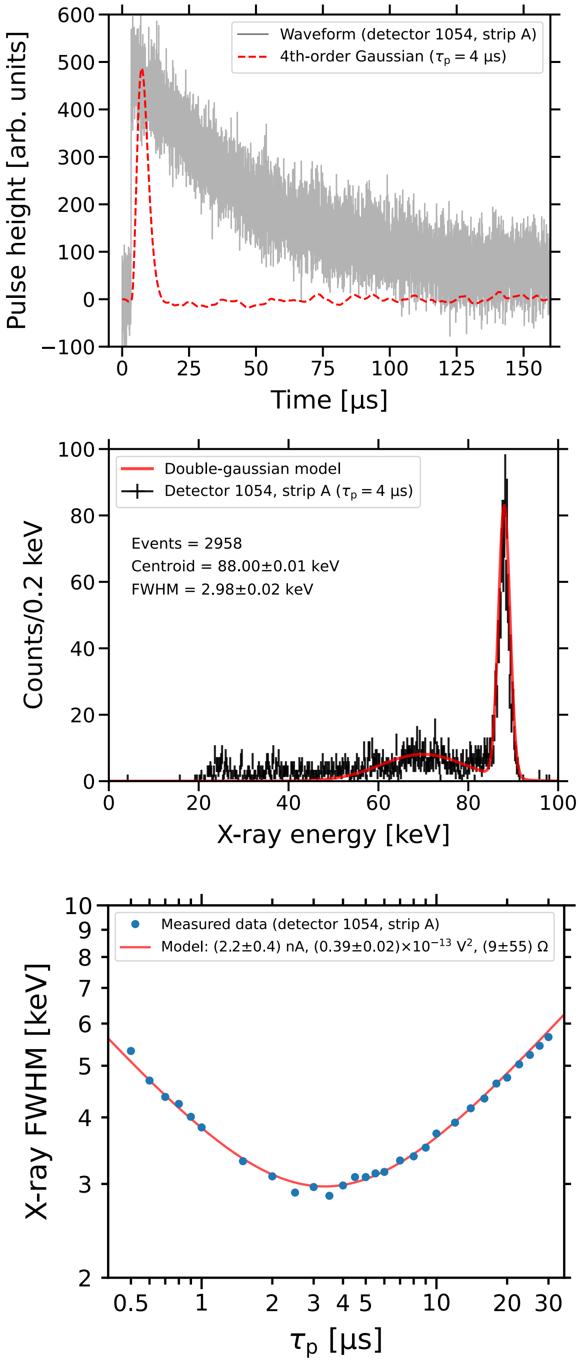

Figure 6-5: The data-analysis workflow for strip A of GAPS Si(Li) detector 1054 tested at –37 °C. **(Top)** A single baseline-subtracted waveform and the corresponding $\tau_p = 4$ µs Gaussian-shaped pulse. **(Middle)** The corresponding ^{109}Cd x-ray energy spectrum, binned to 0.2 keV/bin for presentation. **(Bottom)** The best-fit x-ray noise model from Eq. 6.23 as a function of τ_p. FWHM uncertainties are comparable to the size of the data points. This figure will appear in the published version of Ref. [4].

source) at that τ_p, and the process begins again for the next τ_p. To fully characterize the noise performance of each GAPS Si(Li) strip, we require at least 10 values of τ_p between 0.5–30 µs. Since each setting of τ_p requires $\sim$3 minutes to accumulate sufficient x-ray counts, and the CAEN N6725 digitizer can process eight Si(Li) strips (i.e., one full detector) at once, it would require *at minimum* two hours to test each detector module. After including the time required to re-cool the module before and after each test, the time required nearly doubles.

To increase the throughput of our testing system, we use custom software to extract voltage waveforms directly from the CAEN N6725 digitizer and store them as ROOT files, applying the shaper in software rather than hardware. Each waveform covers a time interval of 160 µs with 4 ns between samples, including 4 µs before the peak to define the baseline voltage. As before, we can record several thousand x-ray events per strip in $\sim$3 minutes. Crucially, however, our software-shaping approach requires only a single set of waveforms for each strip, so acquiring a full dataset for each Si(Li) detector requires only $\sim$3 minutes. This nearly *halves* the testing time for each detector module, requiring only $\sim$2 hours from start (removing the module from storage) to finish (returning the module to storage). We use a fourth-order Gaussian shaper [624] for our default analysis. (Here we note another advantage of our offline waveform approach: we can reprocess the raw data with another shaper at a later date, should the need arise.)

In a sense, this offline-processing approach trades the time required for data-acquisition for storage space and processing time. Fortunately, disk storage is relatively inexpensive, and the full set of GAPS Si(Li) waveforms for one detector (several thousand per strip) requires $\lesssim$ 1.5 GB when stored as a ROOT file. Additionally, the waveform-shaping can be easily multithreaded. For context, the GAPS server at MIT carries 48 GB of RAM and an Intel X5650 CPU (clock speed 2.67 GHz and 48 threads), and requires $\sim$10 minutes to fully process the data from each detector using the entire thread pool.

X-ray spectra

To generate x-ray spectra for each strip, we convolve the waveforms with the fourth-order Gaussian filter implemented in software. An example waveform before and after shaping is shown in Fig. 6-5. The maximum value of the shaped pulse is proportional to the energy deposited in the Si(Li) detector strip, *modulo* the effects of electronic noise discussed in the next section. We process each strip's set of waveforms using 29 different values of τ_p from 0.5–30 μs, generating one spectrum for each peaking time. To convert the somewhat arbitrary units of the shaped pulse height to the more-familiar energy, we first fit the x-ray spectrum to a model consisting of the sum of two Gaussian peaks, representing the x-ray full-absorption peak and the Compton shoulder. We allow the normalization, width, and centroid of both peaks to find their best-fit values, and use the centroid of the full-absorption peak (88 keV for ^{109}Cd) to convert[15] the x-ray spectrum to energy. An example is shown in Fig. 6-5. We also record the other fit parameters, particularly the FWHM of the full-absorption peak, since it is this value that quantifies the x-ray energy resolution at each peaking time.

6.4 Noise sources in Si(Li) detectors

In an ideal world, the only source of variance in energy measurements from semiconductor detectors would be statistical fluctuations in the number of electron-ion pairs generated. Unfortunately, as philosophers and theologians remind us, we do not live in an ideal world, so we must deal with various sources of electronic noise that degrade the signals from our detectors. This section serves as a brief introduction to the modeling of noise sources in semiconductor detectors systems.

---

[15] Previous work using ^{109}Cd and ^{241}Am sources has shown that the GAPS Si(Li) detector response is linear in the energy range of interest for this study (see, e.g., Ref. [617]).

6.4.1 Noise fundamentals

In detector systems, there are two types of noise we are interested in: voltage noise v_n and current noise i_n. Specifically, we are interested in their power spectral distributions[16] $\mathrm{d}v_n^2/\mathrm{d}f$ and $\mathrm{d}i_n^2/\mathrm{d}f$. If the amplifier has a gain $A(f)$, the output noise voltages and currents are (see, e.g., Ref. [620])

$$\begin{cases} v_{n,\mathrm{out}}^2 = \displaystyle\int_0^\infty \frac{\mathrm{d}v_n^2}{\mathrm{d}f}[A(f)]^2\,\mathrm{d}f \\ i_{n,\mathrm{out}}^2 = \displaystyle\int_0^\infty \frac{\mathrm{d}i_n^2}{\mathrm{d}f}[A(f)]^2\,\mathrm{d}f \end{cases} \tag{6.5}$$

This section describes the sources of noise in GAPS Si(Li) detectors.

Thermal noise

One common source of noise[17] in circuit elements is thermal (Johnson-Nyquist) noise [625, 626]. Recall that the power spectrum radiated by a blackbody at temperature T is given by Planck's formula,

$$\frac{\mathrm{d}P}{\mathrm{d}f} = \frac{hf}{\exp[hf/kT]-1}. \tag{6.6}$$

In the limit $hf \ll kT$, this reduces to

$$\frac{\mathrm{d}P}{\mathrm{d}f} \approx kT. \tag{6.7}$$

Imagine an ideal voltage source v_n and resistor R connected in series to a load resistor R_L. If R has a thermal voltage noise v_L and current noise i_n, this corresponds to a power

$$P_L = \frac{v_L^2}{R_L} = \frac{v_n^2 R_L}{(R+R_L)} = i_n^2 R_L \tag{6.8}$$

---

[16] Ref. [620], from which much of this section is summarized, defines $e_n^{(2)} \equiv \mathrm{d}v_n^{(2)}/\mathrm{d}f$ and $i_n^{(2)} \equiv \mathrm{d}i_n^{(2)}/\mathrm{d}f$. The latter especially is extremely confusing (is i_s a current or the power spectral density associated with that current?), so I will explicitly write $\mathrm{d}v_n^{(2)}/\mathrm{d}f$ and $\mathrm{d}i_n^{(2)}/\mathrm{d}f$.

[17] The derivation in this section is reproduced from Ref. [620].

dissipated in the load resistor. Maximum power dissipation occurs when $R = R_L$, giving

$$\frac{\mathrm{d}P}{\mathrm{d}f} = 4kT \quad \text{(thermal)}. \tag{6.9}$$

Thermal noise is a type of *white noise*, since the power spectrum in Eq. 6.9 is independent of frequency (i.e., it contains all frequencies equally, much like white light). The corresponding voltage and current spectral distributions of thermal noise in a resistor R at temperature T are

$$\frac{\mathrm{d}v_n^2}{\mathrm{d}f} = 4kTR \quad \text{and} \quad \frac{\mathrm{d}i_n^2}{\mathrm{d}f} = \frac{4kT}{R}. \tag{6.10}$$

A related form of noise is the so-called *kTC noise*, where C is the total capacitance in parallel with the detector. If there is also a resistance R_p in parallel with (and including) the detector, the noise voltage spectral distribution is given by [620]

$$\frac{\mathrm{d}v_n^2}{\mathrm{d}f} = \frac{4kTR_p}{1 + (2\pi f R_p C)^2}, \tag{6.11}$$

which integrated over all frequencies gives

$$v_n^2 = \frac{kT}{C}. \tag{6.12}$$

Shot noise

In semiconductors, *shot noise* refers to the random processes of charge carriers being generated in the crystal lattice by thermal fluctuations. In the time domain, shot noise arises from δ-function pulses of the form $\delta(t - t_k)$, where t_k is the time at which each charge quantum is generated. In the frequency domain, the shot-noise current has a spectral distribution[18]

$$\frac{\mathrm{d}i_n}{\mathrm{d}f} = 2q, \tag{6.13}$$

---

[18] The derivation in this section is reproduced from Ref. [620].

where q is the magnitude of the electron charge. This means that the rms value of the current is $i_n = q\sqrt{2}$. Imagine a material generating electrons at a rate r, corresponding to a leakage current $I_{\text{leak}} = rq$. Consider a current $i = i_p + i_q$, where i_p and i_q have the same magnitude but are separated by a phase offset ϕ. The mean-square current $\langle i^2 \rangle$ is given by

$$\langle i^2 \rangle = i_p^2 + i_q^2 + 2 i_p i_q \cos\phi. \tag{6.14}$$

Extending the analogy to an arbitrary sum of currents with random phases, the ϕ-dependent terms average to zero, giving

$$\langle i^2 \rangle = 2rq^2. \tag{6.15}$$

Thus, the shot-noise current spectral density is given by

$$\frac{\mathrm{d}i_n^2}{\mathrm{d}f} = 2qI_{\text{leak}}. \tag{6.16}$$

Like thermal noise, shot noise is *white*, since its power spectrum is independent of frequency.

1/f noise

A ubiquitous source of noise in detectors is the so-called $1/f$ noise or *flicker noise*. This phenomenon has been observed in many condensed-matter systems (see, e.g., Refs. [627–632] and references therein for reviews), taking power spectral densities of the form[19] $1/f^\alpha$ ($\alpha \sim 0.5$–2). The power spectral density of $1/f$ noise is

$$\frac{\mathrm{d}P}{\mathrm{d}f} = \frac{S_f}{f}, \tag{6.17}$$

---

[19] $1/f^\alpha$ noise also comes in an array of "colors" depending on the value of α. Noise with $\alpha = 2$ is *red* (since it increases at the low-frequency end), $\alpha = 0$ is *white* as discussed previously, and $\alpha = 1$ is *pink* since it is intermediate between red and white. The rest of this section will assume $\alpha = 1$, since this is the relevant noise source for GAPS Si(Li) detectors.

where S_f describes the amplitude of the fluctuations. In terms of noise voltage,

$$\frac{\mathrm{d}v_n^2}{\mathrm{d}f} = \frac{A_f}{f}. \tag{6.18}$$

To illustrate one possible source of $1/f$ noise[20], consider a material carrying a current I in which charge carriers are trapped after a characteristic lifetime τ. In frequency space, this corresponds to the power spectral density

$$\frac{\mathrm{d}P}{\mathrm{d}f} \propto \frac{\tau}{1 + (2\pi\tau f)^2}, \tag{6.19}$$

which is approximately flat for small f and scales as $1/f^2$ for large f. However, by combining the noise distributions in Eq. 6.19 with values of τ separated by several orders of magnitude, it is possible to obtain $1/f$ distributions over several decades in frequency—see, e.g., Refs. [633–636] for early work along these lines. (A similar effect arises in astrophysics, where the superposition of the synchrotron spectra emitted by a collection of electrons with a powerlaw energy distribution is itself a powerlaw—see, e.g., Refs. [472, 637].)

6.4.2 Noise modeling

In this section, I will present the noise model used for the GAPS Si(Li) detector testing setup described previously. In the context of charge-sensitive amplifiers like those used for GAPS Si(Li) detector testing, a useful figure-of-merit for quantifying the electronic noise in the system is the *equivalent noise charge* (ENC). By definition, a signal with collected charge equal to the ENC would have a signal-to-noise ratio of one.

To construct the noise model (ENC versus peaking time τ_p) for the detector readout system shown in Fig. 6-6, we use the results of Sec. 6.4.1. Current noise sources in parallel to the preamplifier input (i.e., shot noise in the detector and

---

[20] This derivation is reproduced from Ref. [620].

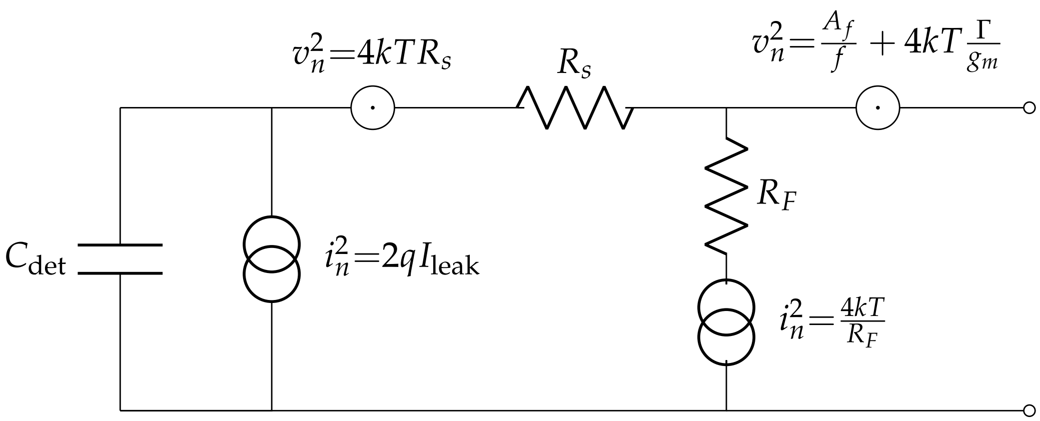

Figure 6-6: Schematic of the readout system for a single GAPS Si(Li) strip, showing the various noise sources. The open circles at the right-hand side of the figure represent the terminals for the discrete preamplifier. Figure adapted from Refs. [619, 638].

thermal noise in the preamplifier feedback resistor) are called *parallel noise*, and have an ENC contribution (see, e.g., Ref. [639])

$$\mathrm{ENC}_{\mathrm{par}}^2 \propto \int_{-\infty}^{\infty} [w(t)]^2 \, \mathrm{d}t \propto F_i \tau_p \tag{6.20}$$

where $w(t)$ is the weighting function of the shaper (see, e.g., Refs. [638, 639]). The quantity F_i is known as the *form factor* or *noise integral* which depends[21] on the choice of shaper, and τ_p is the peaking time. Similar results may be obtained for the *series noise* (voltage sources in series with the preamplifier input) and $1/f$ noise:

$$\mathrm{ENC}_{\mathrm{ser}}^2 \propto \int_{-\infty}^{\infty} [w'(t)]^2 \, \mathrm{d}t \propto \frac{F_v}{\tau_p} \tag{6.21}$$

$$\mathrm{ENC}_{1/f}^2 \propto A_f F_{vf} \tag{6.22}$$

Note that $1/f$ noise also appears in series with the preamplifier input, but carries a different form factor and time dependence owing to its noise power spectrum.

Putting all of these pieces together in quadrature, the full noise model for the

---

[21] The values of the form factors also change depending on how one chooses to parametrize time—peaking time τ_p, shaping time τ_s, etc.

GAPS Si(Li) detectors is [617]

$$\mathrm{ENC}^2 = \left(2qI_{\mathrm{leak}} + \frac{4kT}{R_F}\right) F_i \tau_p + 4kTC_{\mathrm{tot}}^2 \left(R_s + \frac{\Gamma}{g_m}\right) \frac{F_v}{\tau_p} + 2\pi A_f C_{\mathrm{tot}}^2 F_{vf}. \quad (6.23)$$

Here, q is the elementary charge, I_{leak} is the strip leakage current, k is the Boltzmann constant, R_F is the 100-MΩ feedback resistance in the preamplifier, C_{tot} is the total capacitance in parallel with the input (defined shortly), R_s is the series resistance between the Si(Li) strip and the preamplifier, $\Gamma \approx 1$ is a constant related to the JFET channel in the preamplifier, $g_m = 18\,\mathrm{mS}$ is the room-temperature transconductance of the JFET, and A_f is the coefficient of $1/f$ noise. The total parallel capacitance $C_{\mathrm{tot}} = C_{\mathrm{strip}} + C_{\mathrm{int}} + C_{\mathrm{FET}} + C_{\mathrm{stray}}$, where $C_{\mathrm{strip}} \approx$ 37–39 pF is the Si(Li) strip capacitance, C_{int} is the inter-strip capacitance, $C_{\mathrm{FET}} \approx 10\,\mathrm{pF}$ is the capacitance of the input FET, and C_{stray} is any stray capacitance. Previous testing of single Si(Li) detectors found $C_{\mathrm{int}} + C_{\mathrm{stray}} \approx 20\,\mathrm{pF}$, so we freeze $C_{\mathrm{tot}} = 70\,\mathrm{pF}$ when analyzing data from the strips. (Note also from Eq. 6.23 that C_{tot} is degenerate with A_f and R_s, so it is necessary to freeze C_{tot} to constrain the latter two parameters.) The form factors F_i, F_v, and F_{vf} depend on the pulse shaper. For the fourth-order Gaussian shaper used in our fiducial analysis, $F_i = 0.45$, $F_v = 1.02$, and $F_{vf} = 0.52$. To extract the best-fit values of I_{leak}, A_f, and R_s, we fit the x-ray FWHM versus τ_p data for each strip to Eq. 6.23 using the `Minuit` algorithm in ROOT. The conversion between ENC and FWHM energy resolution in silicon is given by

$$\mathrm{FWHM} \simeq 2.355\,\varepsilon_{\mathrm{Si}} \frac{\mathrm{ENC}}{q}, \quad (6.24)$$

where $\varepsilon_{\mathrm{Si}} \approx 3.6\,\mathrm{eV}$ is the average electron-hole pair generation energy in silicon at the GAPS operating temperature (see, e.g., Refs. [640, 641] for recent measurements) and q is the elementary charge.

An example noise model for strip A of detector 1054 is shown in Fig. 6-5. When tested with the discrete preamplifiers, GAPS Si(Li) strips exhibit a characteristic "U" shape, with the best energy resolution typically being achieved near $\tau_p \sim 4\,\mu\mathrm{s}$.

As described in Eq. 6.23, the low-τ_p end is dominated by resistances, either in series with the preamplifier or within the preamplifier itself (note that $\Gamma/g_m \sim 5.5\,\Omega$). For intermediate τ_p, the energy resolution is determined mainly by A_f, and for large τ_p, by the leakage current (for $I_{\rm leak} = 1\,$nA, the $2qI_{\rm leak}$ is a factor $\sim$2.5 greater than the $4kT/R_F$ term). For a strip with negligible $I_{\rm leak}$, A_f, and R_s, the minimum achievable x-ray energy resolution using these discrete preamplifiers is $\sim$2 keV FWHM, as set by R_F, $C_{\rm tot}$, and Γ/g_m. It is worth noting that for low values of $I_{\rm leak}$, A_f, and/or R_s, the noise model in Eq. 6.23 may encounter difficulties in fitting these parameters for several reasons. First, low values of these parameters mean that the x-ray energy resolution is approaching the preamplifier noise floor, and thus the contributions from the different terms in Eq. 6.23 become difficult to distinguish. Second, freezing $C_{\rm tot}$ (necessary to break the degeneracy with A_f and R_s) may cause difficulties if $C_{\rm tot}$ differs from the assumed 70 pF.

As described previously, the flight readout for the GAPS instrument will use a custom ASIC and front-end electronics, rather than the discrete charge-sensitive preamplifiers used in our laboratory testing. The ASIC noise model is (see, e.g., Ref. [580])

$$\mathrm{ENC}^2 = 2q(I_{\rm leak} + I_k)F_i\tau_p + 4kTR_{\rm eq}C_{\rm eq}^2\frac{F_\nu}{\tau_p} + 2\pi A_f C_{\rm eq}^2 F_{\nu f}, \tag{6.25}$$

where $I_k \approx 2.5\,$nA is the effective current in the analog channel, $R_{\rm eq} \approx 40\,\Omega$ is the equivalent input resistance, and $C_{\rm eq} = C_{\rm strip} + C_{\rm amp}$ is the equivalent capacitance including contributions from both the Si(Li) strip ($C_{\rm strip} \approx 37\,$pF) and preamplifier input ($C_{\rm amp} \approx 5\,$pF). The form factors for the ASIC unipolar (RC)–(CR)^2 shaper are $F_i = 0.64$, $F_\nu = 0.85$, $F_{\nu f} = 0.54$. Since the available ASIC peaking times τ_p range from 0.25–1.8 μs, the noise contributions from reasonable values of $I_{\rm leak}$ are generally lower than the A_f contribution, and the (generally small) series resistance R_s does not play a significant role. By using the values of $I_{\rm leak}$ and A_f extracted from x-ray data taken using the discrete preamplifiers, it is possible to use Eq. 6.25 to predict the x-ray energy resolution that a strip may achieve with the flight ASIC readout. It is these ASIC-projected values that we use to define the quality grade of each strip and detector at –37 °C, as shown in Tables 6.1 and 6.2. Further testing is ongoing to validate the ASIC noise model in Eq. 6.25 in the full payload, and will be published at a later date. To be conservative, we set $A_f = \max(0.87 \times 10^{-13}\,\mathrm{V}^2,\ A_{f,\mathrm{fit}})$, where the former value is derived from simulations of the ASIC [580] and the latter is the best-fit value from Eq. 6.23.

6.4.3 Results

As shown in Tables 6.1 and 6.2, the GAPS testing campaign at MIT and UHM examined the performance of 1111 detectors (i.e., 8888 strips). Of these, nearly 70% of detectors (and 80% of strips) were deemed suitable for x-ray spectroscopy, and nearly 90% (of both detectors and strips) were deemed suitable for charged-particle tracking. This substantially exceeds the requirements (and expectations) set prior to production. In this section, I describe several results of this testing, based on analyses I performed with the MIT, UHM, and CU groups [4].

Detector yield rate over time

The production of GAPS detectors by Shimadzu was separated into sixteen batches, with each batch lasting approximately one month. Thus, it is instructive to examine the yield rate of each detector quality grade over time. The results are shown in Fig. 6-7. For most batches, the yield rate of x-ray-quality detectors exceeded 60–70%, though there are several deviations that are worth further consideration. In particular, batches 1 and 2 have noticeably lower yield rates of x-ray-quality detectors than the following batches. We attribute this to a change in the lithium-drifting procedure. For batches 1 and 2, lithium-drifting was only stopped when the current exceeded $\sim$25 mA (or began increasing exponentially), or when the heater output dropped to zero. This led to some detectors being drifted for more than 6500 minutes [616]. From batch 3 onward, an additional limit on the drift time of 6200 minutes was imposed, and the yield rate remained high through the rest

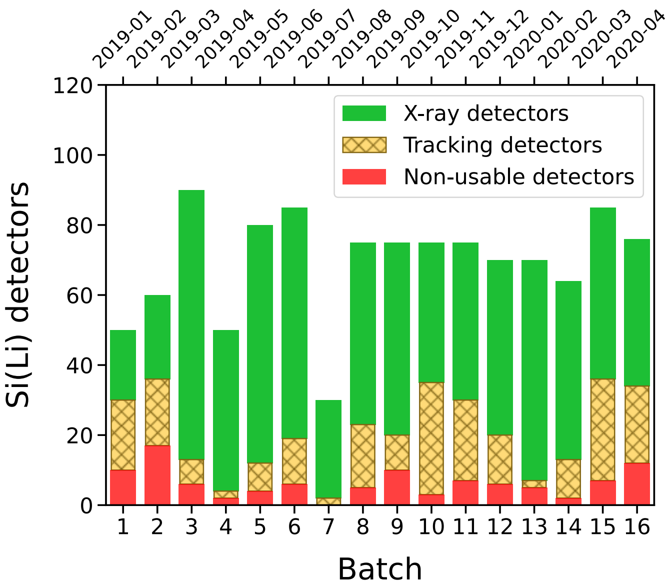

Figure 6-7: Stacked bar chart of the Si(Li) detector quality per production batch, as described in Sec. 6.4.3. This figure will appear in the published version of Ref. [4].

of production.

Detector noise parameters

Since the ASIC noise model in Eq. 6.25 is sensitive primarily to the noise parameters I_{leak} and A_f, it is these values we are most interested in extracting[22] from the discrete-preamplifier testing. The distributions of these parameters for strips on x-ray-quality and tracking-quality detectors are shown in Fig. 6-8. As expected, the I_{leak} and A_f distributions for x-ray-quality detectors peak at lower values than their tracking-quality counterparts, with the latter having longer tails. However, the x-ray-quality distributions also have tails extending beyond $I_{\text{leak}} \approx 10\,\text{nA}$ and $A_f \approx 1.5 \times 10^{-13}\,\text{V}^2$, above which strips tend to be tracking-quality. This further reinforces our observation that otherwise x-ray-quality detectors can still have

---

[22] It was not possible to measure I_{leak} directly for all >1100 GAPS detectors at cold temperatures due to time constraints. However, direct measurements of I_{leak} on ∼100 strips near –40 °C are consistent with the results described here and with Eq. 6.26. Unfortunately, is not possible to directly measure A_f—it can only be extracted from the noise model.

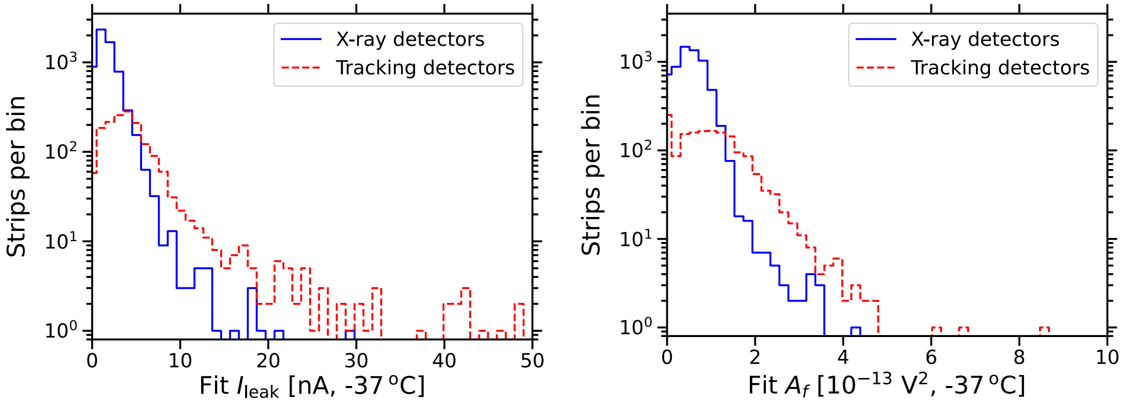

Figure 6-8: **(Left)** Histogram of best-fit leakage currents I_{leak} for the GAPS Si(Li) strips extracted using the noise model in Eq. 6.23. Strips on X-ray-quality detectors are denoted with a blue solid line, and those on tracking-quality detectors with a red dashed line. **(Right)** Same as left, for the best-fit values of A_f. This figure will appear in the published version of Ref. [4].

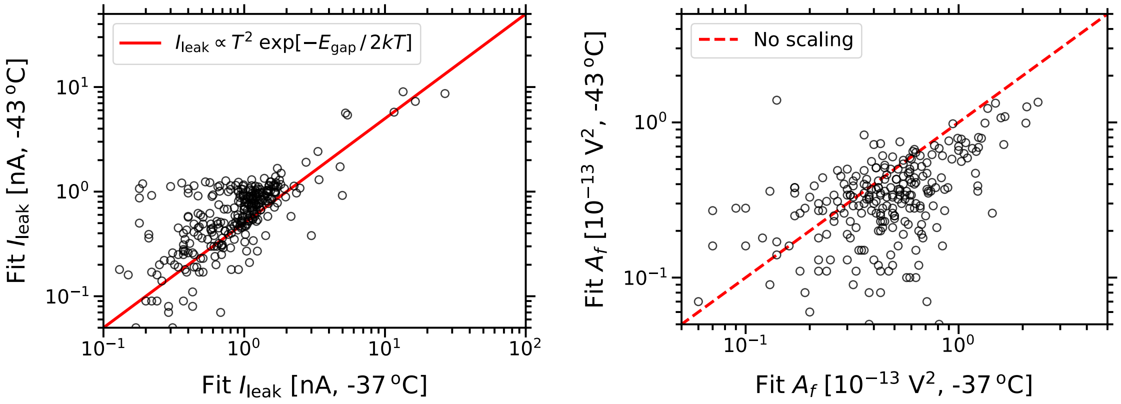

Figure 6-9: **(Left)** Scatter plot showing the best-fit leakage currents for the 37 detectors (296 strips) tested at both –37 °C and –43 °C. For clarity, errorbars on individual points are not shown, but are generally at the $\pm(10\text{–}20)\%$ level for most strips. The predicted temperature scaling from Eq. 6.26 is shown with the red solid line. **(Right)** Same as left, for the best-fit values of A_f. The red dashed line indicates the expectation if A_f does not change with temperature. This figure will appear in the published version of Ref. [4].

tracking-quality (or even non-usable) strips, without negatively affecting the performance of their neighbors.

We also examine the temperature dependence of the x-ray energy resolution. Since the temperature of the Si(Li) detectors in flight will depend on time and position within the instrument, proper interpretation of the results (particularly for any x-ray event candidates) relies on knowledge of the energy resolution at that time. To study this temperature dependence, we tested a sample of 37 detectors at –43 °C in addition to the usual –37°C, and extracted their noise parameters in the same way as before using Eq. 6.23 (though with the appropriate change to T). Plots of I_{leak} and A_f at these two temperatures are shown in Fig. 6-9. We did not expect A_f to change appreciably with temperature (and indeed the best-fit values of A_f are weakly correlated with temperature), but we did expect I_{leak} to decrease. Quantitatively, the leakage current generated by a bulk[23] semiconductor is observed to scale with absolute temperature T as (see, e.g., Ref. [642])

$$I_{\text{leak}} \propto T^2 \exp\left[-\frac{E_{\text{gap}}}{2kT}\right], \tag{6.26}$$

where $E_{\text{gap}} \approx 1.14\,\text{eV}$ is the silicon band gap [643] and k is the Boltzmann constant. From Eq. 6.26, a useful rule-of-thumb is that reducing the temperature by 6 °C near the GAPS operating temperature of –40 °C causes I_{leak} to decrease by a factor $\sim$2. This also means that decreasing the temperature from 25 °C (near room temperature) to –37 °C reduces the leakage current by a factor $\gtrsim$500, further motivating the room-temperature criteria in Tables 6.1 and 6.2.) The fit I_{leak} data at –37 °C and –43 °C generally agree with Eq. 6.26, though the scaling at lower values of I_{leak} is somewhat weaker. As described in Ref. [4], this may be the result of surface currents (which do not scale as Eq. 6.26) dominating over bulk currents, or difficulties encountered in fitting Eq. 6.23 to small values of I_{leak}. To quantify the temperature effect on the GAPS x-ray energy resolution, we use the ASIC noise model in Eq. 6.25 to predict the x-ray energy resolution of these strips at –37°C

---

[23] As opposed to, e.g., surface effects.

and –43°C using the best-fit values of I_{leak} and A_f. We then scale each I_{leak} to the other temperature[24] using Eq. 6.26 while keeping A_f constant, and use Eq. 6.25 to predict the x-ray energy resolution at both temperatures. We find that the FWHMs using the best-fit and extrapolated values of I_{leak} are consistent at the $\sim$0.1-keV level, so a measurement of the module temperature (facilitated by a sensor integrated into each module's front-end board) is sufficient to predict any temperature dependence in detector performance. (This is helped by the fact that short ASIC peaking times do not allow I_{leak} to have much effect on the noise level.)

6.5 Closing thoughts

In this chapter, I described my contributions to the GAPS Si(Li) detector testing effort, which ran from 2019–2022 and involved collaborators from around the world. These results confirm that large-area Si(Li) detectors for x-ray spectroscopy and charged-particle tracking can be economically mass-produced and deployed at relatively warm temperatures. The noise modeling I performed directly influenced the placement of detectors within the GAPS instrument to optimize sensitivity (better-performing detectors being placed near the top, where antinuclei are more likely to stop).

The GAPS program began nearly two decades before my arrival at MIT in 2017, and will hopefully continue for many years into the future. The first Antarctic flight of GAPS is currently scheduled for late 2024, annd the Collaboration is actively working to test the performance of all of the instrument systems. In particular, tests of the tracker energy response to cosmic-ray muons and laboratory x-ray sources will increase our confidence in the performance of the integrated flight instrument.

---

[24] That is, if I_{leak} was extracted at –37°C, we use Eq. 6.26 to predict its value at –43°C, and vice-versa.

7

Coda

The pursuit of knowledge is without end,
for knowledge is never a thing complete.
—Sir Gideon Ofnir, *Elden Ring*

Evidence that most of the matter in the Universe is composed of dark (i.e., invisible and nonbaryonic) degrees of freedom dates back nearly a century. Though cosmological observations have routinely confirmed this discrepancy between visible and gravitating mass, the nature of this dark matter (DM) remains unresolved. The field of indirect detection offers a plethora of techniques to search for DM via its decay and/or annihilation to stable particles (photons, cosmic rays, neutrinos, etc), using data from astrophysical observatories.

In this dissertation, I described my contributions to the field of indirect DM detection over the last six years. Using x-ray data from the NuSTAR space telescope, I developed analysis techniques to fully leverage its sensitivity to decaying keV-scale DM (e.g., sterile neutrinos) in the Galactic halo and the M31 galaxy. Though these analyses did not find any evidence of sterile neutrinos, we were able to set world-leading constraints on the allowed sterile-neutrino mixing angle in the context of the νMSM. This work also spawned a "side project" which recently presented the first NuSTAR constraints on the disintegration of heavier (MeV-to-GeV-scale) DM candidates using x-rays emitted as final-state radiation or inverse-Compton scattering. The second part of my dissertation focused on the laboratory testing and analysis of the more than 1100 Si(Li) detectors that will fly in the GAPS Antarctic balloon instrument, designed to search for low-energy cosmic antiprotons, antideuterons, and other antinuclei. This work demonstrates that large-area Si(Li) detectors can be economically mass-produced, while still offering good x-ray spectral resolutions at non-cryogenic temperatures. These results directly informed the placement of detectors, and will be critical in assessing the sensitivity of the flight instrument.

The coming years will bring a wealth of new cosmological and astroparticle data to scrutinize for signs of DM. The GAPS payload is currently scheduled for launch in late 2024, and many new and ongoing space-based experiments will deliver data in the coming years (e.g., SRG, Micro-X, AMS, XRISM, and ATHENA, to name but a few). In parallel, the search for DM will continue at collider and direct-detection experiments. Perhaps most tantalizing are the upcoming cosmological surveys (e.g., LSST, DES, Euclid, and Roman), which will study the gravitational effects of dark matter in millions of galaxies across cosmic time. This data will allow precision tests of many DM models using only their gravitational signatures on large-scale structure, and to significantly reduce key systematic uncertainties for direct- and indirect-detection experiments.

A

Appendix A: Supplementary Figures

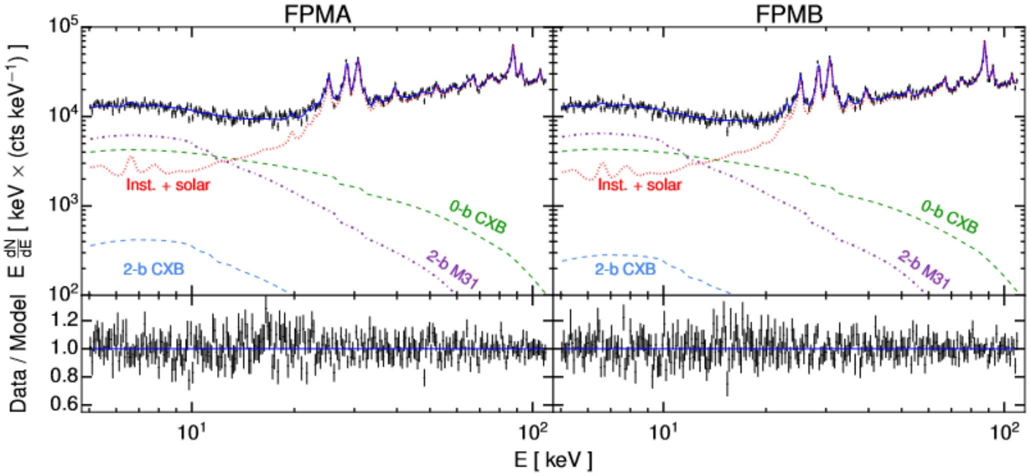

Figure A-1: Fit to the 5–110-keV NuSTAR data from M31 field 50026002003 described in Sec. 3.5.2, showing the 0-bounce CXB (green), 2-bounce (blue) CXB, instrumental+solar background (red), and the 2-bounce M31 component (purple). Reproduced from Ref. [1] with permission. © 2019 by the American Physical Society.

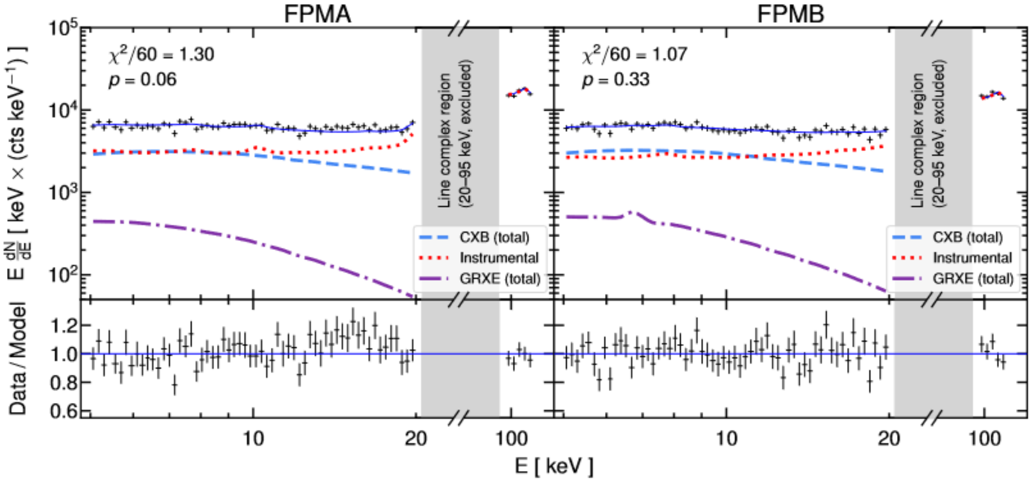

Figure A-2: Fit to the 5–20-keV and 95–110-keV NuSTAR data from Galactic bulge field 40410001002 described in Sec. 3.5.3. The blue dashed curve shows the total (0-bounce + 2-bounce) CXB contribution, the purple dash-dotted line the total (0-bounce + 2-bounce) GRXE, and the red dotted line the instrumental background. Reproduced from Ref. [2] with permission. © 2020 by the American Physical Society.

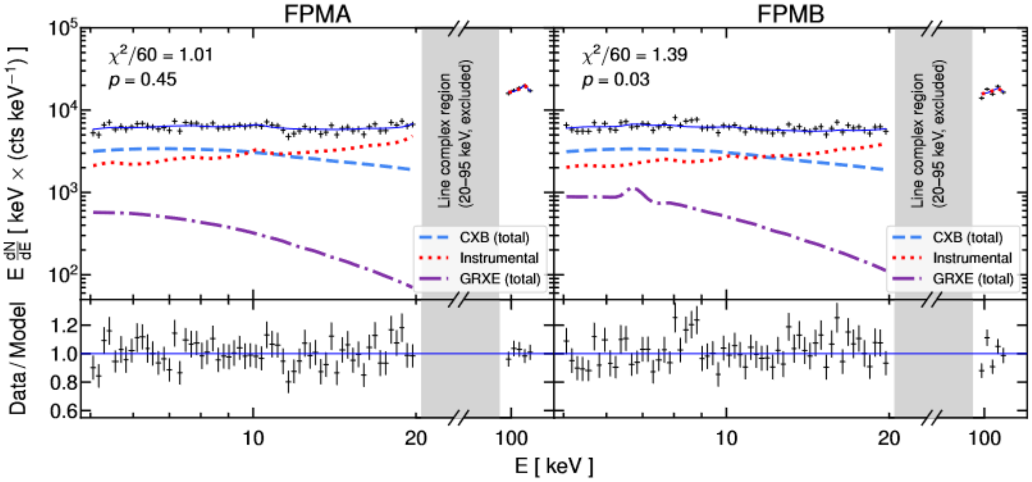

Figure A-3: Same as previous, for Galactic bulge field 40410002002. Reproduced from Ref. [2] with permission. © 2020 by the American Physical Society.

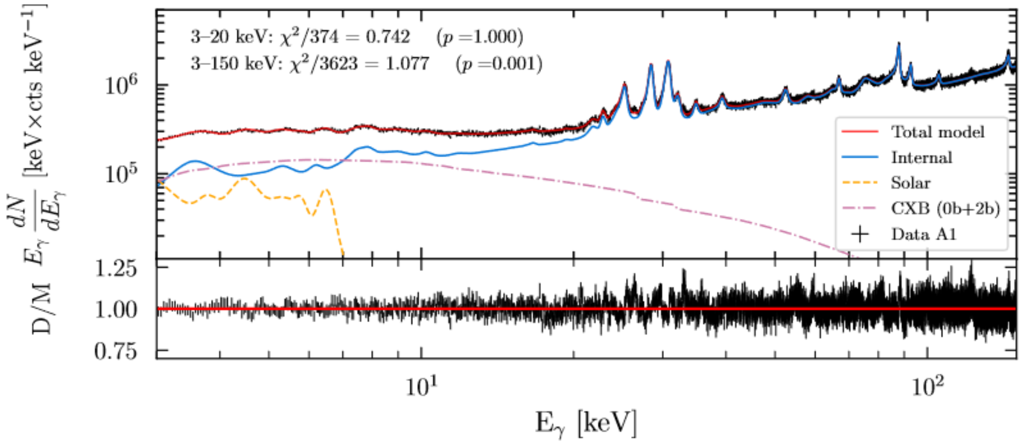

Figure A-4: Same as Fig. 3-12, but for DETA1. Reproduced from Ref. [3] with permission. © 2023 by the American Physical Society.

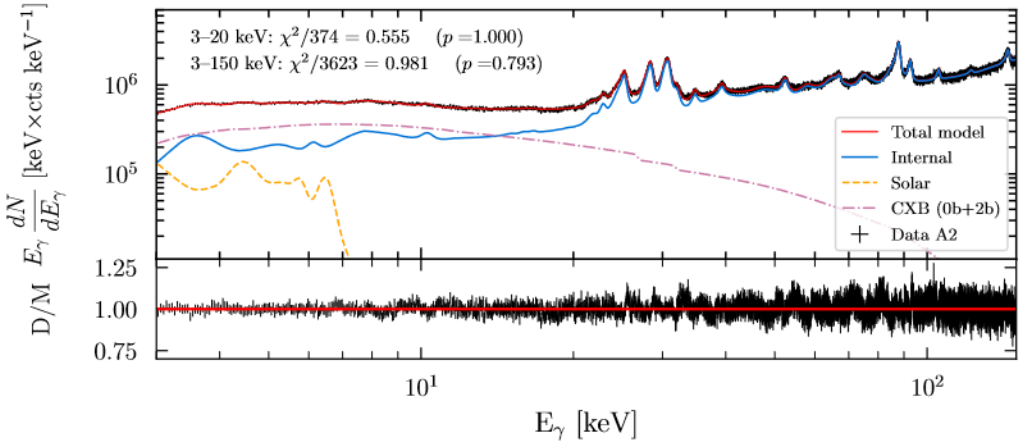

Figure A-5: Same as Fig. 3-12, but for DETA2. Reproduced from Ref. [3] with permission. © 2023 by the American Physical Society.

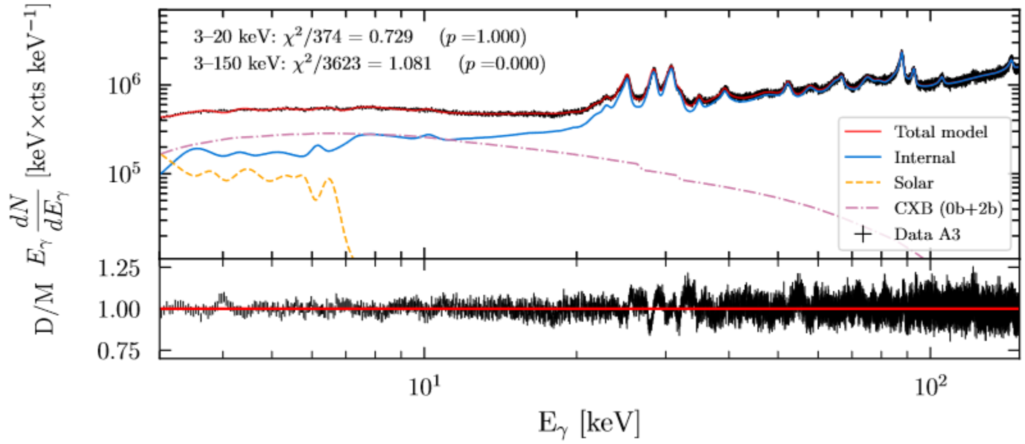

Figure A-6: Same as Fig. 3-12, but for DETA3. Reproduced from Ref. [3] with permission. © 2023 by the American Physical Society.

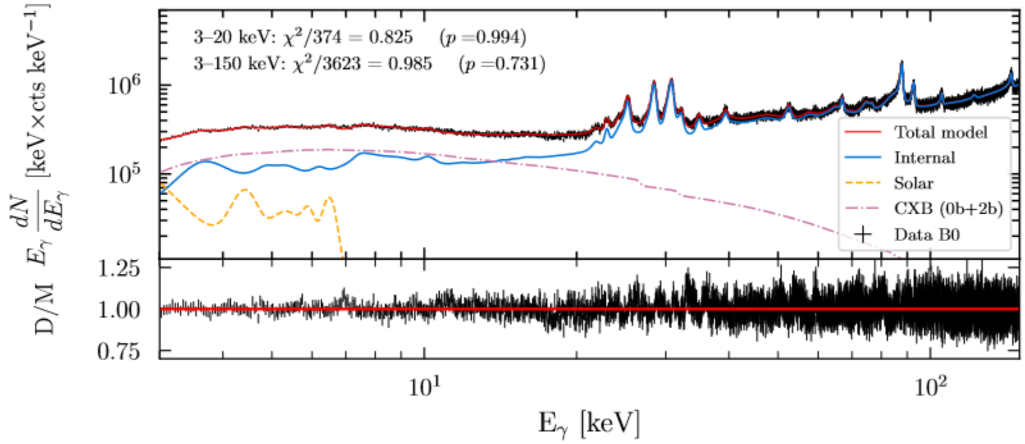

Figure A-7: Same as Fig. 3-12, but for DETB0. Reproduced from Ref. [3] with permission. © 2023 by the American Physical Society.

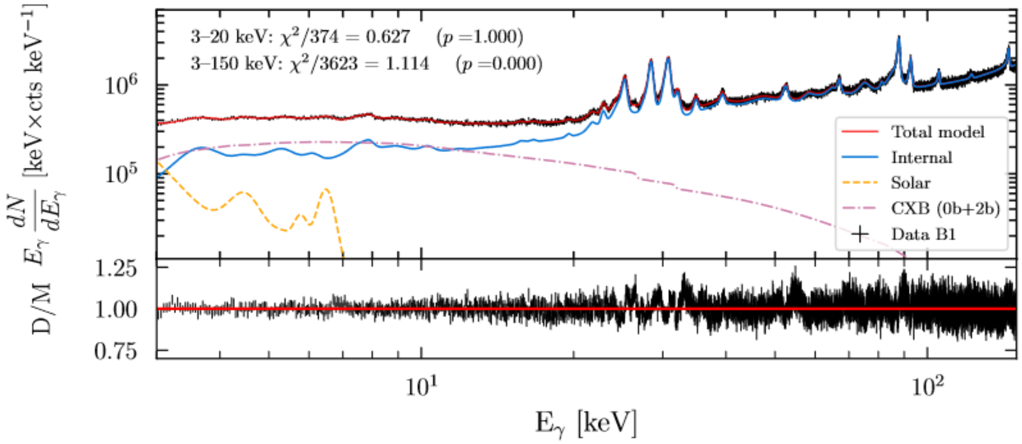

Figure A-8: Same as Fig. 3-12, but for DETB1. Reproduced from Ref. [3] with permission. © 2023 by the American Physical Society.

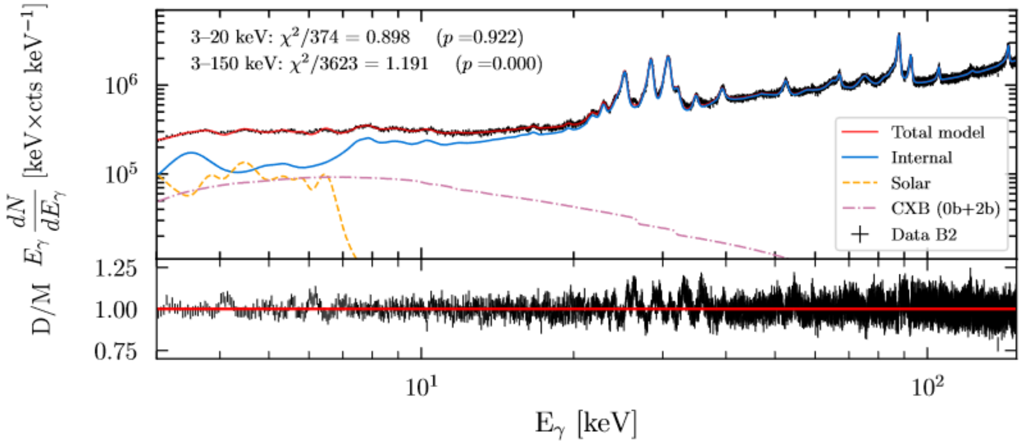

Figure A-9: Same as Fig. 3-12, but for DETB2. Reproduced from Ref. [3] with permission. © 2023 by the American Physical Society.

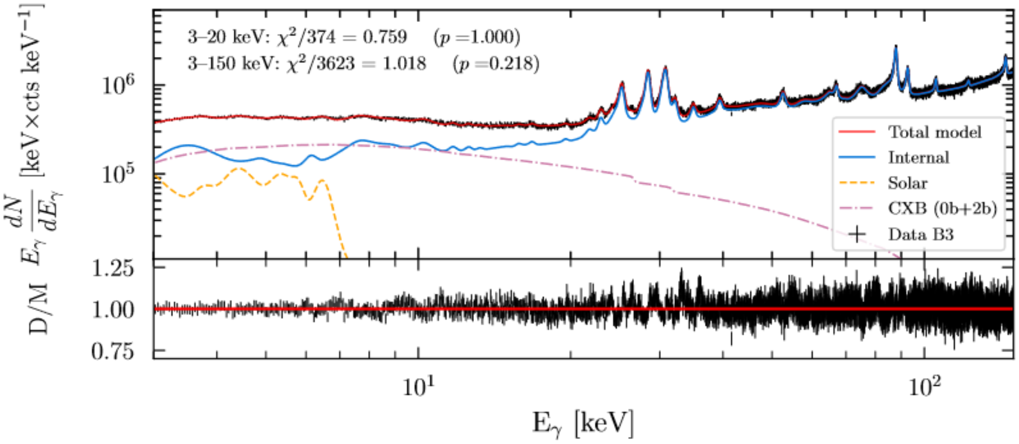

Figure A-10: Same as Fig. 3-12, but for DETB3. Reproduced from Ref. [3] with permission. © 2023 by the American Physical Society.

Bibliography

[1] K. C. Y. Ng, B. M. Roach, K. Perez, J. F. Beacom, S. Horiuchi, R. Krivonos, and D. R. Wik, "New constraints on sterile neutrino dark matter from NuSTAR M31 observations," Phys. Rev. D **99**, 083005 (2019), arXiv:1901.01262 [astro-ph.HE].

[2] B. M. Roach, K. C. Y. Ng, K. Perez, J. F. Beacom, S. Horiuchi, R. Krivonos, and D. R. Wik, "NuSTAR tests of sterile-neutrino dark matter: New Galactic bulge observations and combined impact," Phys. Rev. D **101**, 103011 (2020), arXiv:1908.09037 [astro-ph.HE].

[3] B. M. Roach, S. Rossland, K. C. Y. Ng, K. Perez, J. F. Beacom, B. W. Grefenstette, S. Horiuchi, R. Krivonos, and D. R. Wik, "Long-exposure NuSTAR constraints on decaying dark matter in the Galactic halo," Phys. Rev. D **107**, 023009 (2023), arXiv:2207.04572 [astro-ph.HE].

[4] M. Xiao *et al.*, "Large-scale detector testing for the GAPS Si(Li) Tracker," (2023), in review at IEEE Trans. Nucl. Sci., arXiv:2305.00283 [physics.ins-det].

[5] R. N. Cahn and G. Goldhaber, *The Experimental Foundations of Particle Physics* (Cambridge Univ. Press, Cambridge, 2009).

[6] S. Raby, *Introduction to the Standard Model and Beyond: Quantum Field Theory, Symmetries and Phenomenology* (Cambridge Univ. Press, Cambridge, 2021).

[7] R. P. Woodard, "How far are we from the quantum theory of gravity?" Rep. Prog. Phys. **72**, 126002 (2009), arXiv:0907.4238 [gr-qc].

[8] S. Carlip, D.-W. Chiou, W.-T. Ni, and R. Woodard, "Quantum gravity: A brief history of ideas and some prospects," Int. J. Mod. Phys. D **24**, 1530028 (2015), arXiv:1507.08194 [gr-qc].

[9] A. Addazi *et al.*, "Quantum gravity phenomenology at the dawn of the multi-messenger era—A review," Prog. Part. Nucl. Phys. **125**, 103948 (2022), arXiv:2111.05659 [hep-ph].

[10] A. Ashtekar and E. Bianchi, "A short review of loop quantum gravity," Rep. Prog. Phys. **84**, 042001 (2021), arXiv:2104.04394 [gr-qc].

[11] A. Friedmann, "On the curvature of space," Gen. Relativ. Gravit. **31**, 1991 (1999), [Z. Phys. **10**, 377 (1922)].

[12] G. Lemaître, "A homogeneous universe of constant mass and increasing radius accounting for the radial velocity of extra-galactic nebulæ," Mon. Not. R. Astron. Soc. **91**, 483 (1931), [Ann. Soc. Sci. Bruxelles A **47**, 49 (1927)].

[13] H. P. Robertson, "Kinematics and world-structure," Astrophys. J. **82**, 284 (1935).

[14] A. G. Walker, "On Milne's theory of world-structure," Proc. London Math. Soc. **42**, 90 (1937).

[15] N. Aghanim *et al.* (Planck Collaboration), "Planck 2018 results. VI. Cosmological parameters," Astron. Astrophys. **641**, A6 (2020), [Erratum: A&A **652**, C4 (2021)], arXiv:1807.06209 [astro-ph.CO].

[16] R. L. Workman *et al.* (Particle Data Group), "Review of particle physics," Prog. Theor. Exp. Phys. **2022**, 083C01 (2022).

[17] A. Chudaykin, K. Dolgikh, and M. M. Ivanov, "Constraints on the curvature of the Universe and dynamical dark energy from the full-shape and BAO data," Phys. Rev. D **103**, 023507 (2021), arXiv:2009.10106 [astro-ph.CO].

[18] H. Goldstein, C. Poole, and J. Safko, *Classical Mechanics*, 3rd ed. (Addison Wesley, 2001).

[19] P. Schneider, *Extragalactic Astronomy and Cosmology: An Introduction*, 2nd ed. (Springer Berlin, Heidelberg, 2015).

[20] F. Zwicky, "Republication of: The redshift of extragalactic nebulae," Gen. Relativ. Gravit. **41**, 207 (2009), [Helv. Phys. Acta **6**, 110 (1933)].

[21] G. Bertone and D. Hooper, "History of dark matter," Rev. Mod. Phys. **90**, 045002 (2018), arXiv:1605.04909 [astro-ph.CO].

[22] P. Rosati, S. Borgani, and C. Norman, "The evolution of x-ray clusters of galaxies," Annu. Rev. Astron. Astrophys. **40**, 539 (2002), arXiv:astro-ph/0209035.

[23] M. Arnaud, "X-ray observations of clusters of galaxies," in *International Enrico Fermi School of Physics: Background Microwave Radiation and Intracluster Cosmology* (2005) arXiv:astro-ph/0508159.

[24] J. R. Peterson and A. C. Fabian, "X-ray spectroscopy of cooling clusters," Phys. Rep. **427**, 1 (2006), arXiv:astro-ph/0512549.

[25] H. Boehringer and N. Werner, "X-ray spectroscopy of galaxy clusters," Astron. Astrophys. Rev. **18**, 127 (2010), arXiv:0907.4277 [astro-ph.CO].

[26] N. Ota, "X-ray spectroscopy of clusters of galaxies," Res. Astron. Astrophys. **12**, 973 (2012), arXiv:1211.0679 [astro-ph.CO].

[27] S. Ettori, A. Donnarumma, E. Pointecouteau, T. H. Reiprich, S. Giodini, L. Lovisari, and R. W. Schmidt, "Mass profiles of galaxy clusters from x-ray analysis," Space Sci. Rev. **177**, 119 (2013), arXiv:1303.3530 [astro-ph.CO].

[28] M. Cataneo and D. Rapetti, "Tests of gravity with galaxy clusters," Int. J. Mod. Phys. D **27**, 1848006 (2018), arXiv:1902.10124 [astro-ph.CO].

[29] R. A. Sunyaev and Ya. B. Zeldovich, "Microwave background radiation as a probe of the contemporary structure and history of the universe," Annu. Rev. Astron. Astrophys. **18**, 537 (1980).

[30] S. Bocquet *et al.* (SPT Collaboration), "Mass calibration and cosmological analysis of the SPT-SZ galaxy cluster sample using velocity dispersion σ_v and x-ray Y_X measurements," Astrophys. J. **799**, 214 (2015), arXiv:1407.2942 [astro-ph.CO].

[31] J. P. Dietrich *et al.* (SPT Collaboration), "Sunyaev–Zel'dovich effect and x-ray scaling relations from weak lensing mass calibration of 32 South Pole Telescope selected galaxy clusters," Mon. Not. Roy. Astron. Soc. **483**, 2871 (2019), arXiv:1711.05344 [astro-ph.CO].

[32] K. Umetsu *et al.*, "Weak lensing analysis of x-ray-selected XXL galaxy groups and clusters with Subaru HSC data," Astrophys. J. **890**, 148 (2020), arXiv:1909.10524 [astro-ph.CO].

[33] D. Akino *et al.*, "HSC-XXL: Baryon budget of the 136 XXL groups and clusters," Publ. Astron. Soc. Jpn. **74**, 175 (2022), arXiv:2111.10080 [astro-ph.CO].

[34] I.-N. Chiu, M. Klein, J. Mohr, and S. Bocquet, "Cosmological constraints from galaxy clusters and groups in the eROSITA final equatorial depth survey," Mon. Not. R. Astron. Soc. **522**, 1601 (2023), arXiv:2207.12429 [astro-ph.CO].

[35] D. Paraficz, J. P. Kneib, J. Richard, A. Morandi, M. Limousin, E. Jullo, and J. Martinez, "The Bullet cluster at its best: Weighing stars, gas, and dark matter," Astron. Astrophys. **594**, A121 (2016), arXiv:1209.0384 [astro-ph.CO].

[36] D. Clowe, M. Bradac, A. H. Gonzalez, M. Markevitch, S. W. Randall, C. Jones, and D. Zaritsky, "A direct empirical proof of the existence of dark matter," Astrophys. J. Lett. **648**, L109 (2006), arXiv:astro-ph/0608407.

[37] B. Famaey and S. McGaugh, "Modified Newtonian dynamics (MOND): Observational phenomenology and relativistic extensions," Living Rev. Relativ. **15**, 10 (2012), arXiv:1112.3960 [astro-ph.CO].

[38] M. Milgrom, "Marriage à-la-MOND: Baryonic dark matter in galaxy clusters and the cooling flow puzzle," New Astron. Rev. **51**, 906 (2008), arXiv:0712.4203 [astro-ph].

[39] G. W. Angus, B. Famaey, and A. Diaferio, "Equilibrium configurations of 11 eV sterile neutrinos in MONDian galaxy clusters," Mon. Not. R. Astron. Soc. **402**, 395 (2010), arXiv:0906.3322 [astro-ph.CO].

[40] J. Binney and S. Tremaine, *Galactic Dynamics*, 2nd ed., Princeton Series in Astrophysics (Princeton Univ. Press, Princeton, 2008).

[41] C. V. Barbieri, F. Fraternali, T. Oosterloo, G. Bertin, R. Boomsma, and R. Sancisi, "Extra-planar gas in the spiral galaxy NGC 4559," Astron. Astrophys. **439**, 947 (2005), arXiv:astro-ph/0504534.

[42] F. Lelli, S. S. McGaugh, and J. M. Schombert, "SPARC: Mass models for 175 disk galaxies with Spitzer photometry and accurate rotation curves," Astron. J. **152**, 157 (2016), arXiv:1606.09251 [astro-ph.GA].

[43] P. Li, F. Lelli, S. McGaugh, and J. Schombert, "A comprehensive catalog of dark matter halo models for SPARC galaxies," Astrophys. J. Suppl. Ser. **247**, 31 (2020), arXiv:2001.10538 [astro-ph.GA].

[44] Y. Sofue and V. Rubin, "Rotation curves of spiral galaxies," Annu. Rev. Astron. Astrophys. **39**, 137 (2001).

[45] J. F. Navarro, C. S. Frenk, and S. D. M. White, "A universal density profile from hierarchical clustering," Astrophys. J. **490**, 493 (1997), arXiv:astro-ph/9611107 [astro-ph].

[46] S. Tulin and H.-B. Yu, "Dark matter self-interactions and small scale structure," Phys. Rep. **730**, 1 (2018), arXiv:1705.02358 [hep-ph].

[47] A. Del Popolo and M. Le Delliou, "Review of solutions to the cusp-core problem of the ΛCDM model," Galaxies **9**, 123 (2021), arXiv:2209.14151 [astro-ph.CO].

[48] A. Burkert, "The structure of dark matter halos in dwarf galaxies," Astrophys. J. Lett. **447**, L25 (1995), arXiv:astro-ph/9504041 [astro-ph].

[49] J. Einasto, "On the construction of a composite model for the Galaxy and on the determination of the system of Galactic parameters," Trudy Astrofiz. Inst. Alma-Ata **5**, 87 (1965).

[50] E. Retana-Montenegro, E. Van Hese, G. Gentile, M. Baes, and F. Frutos-Alfaro, "Analytical properties of Einasto dark matter haloes," Astron. Astrophys. **540**, A70 (2012), arXiv:1202.5242 [astro-ph.CO].

[51] M. Pato, F. Iocco, and G. Bertone, "Dynamical constraints on the dark matter distribution in the Milky Way," J. Cosmol. Astropart. Phys. **1512**, 001 (2015), arXiv:1504.06324 [astro-ph.GA].

[52] N. C. Relatores *et al.*, "The dark matter distributions in low-mass disk galaxies. II. The inner density profiles," Astrophys. J. **887**, 94 (2019), arXiv:1911.05836 [astro-ph.GA].

[53] J. I. Read, "The local dark matter density," J. Phys. G: Nucl. Part. Phys. **41**, 063101 (2014), arXiv:1404.1938 [astro-ph.GA].

[54] P. F. de Salas, K. Malhan, K. Freese, K. Hattori, and M. Valluri, "On the estimation of the local dark matter density using the rotation curve of the Milky Way," J. Cosmol. Astropart. Phys. **10**, 037 (2019), arXiv:1906.06133 [astro-ph.GA].

[55] P. F. de Salas and A. Widmark, "Dark matter local density determination: Recent observations and future prospects," Rep. Prog. Phys. **84**, 104901 (2021), arXiv:2012.11477 [astro-ph.GA].

[56] Y. Sofue, "Rotation curve of the Milky Way and the dark matter density," Galaxies **8**, 37 (2020), arXiv:2004.11688 [astro-ph.GA].

[57] D. J. Fixsen, E. S. Cheng, J. M. Gales, J. C. Mather, R. A. Shafer, and E. L. Wright, "The cosmic microwave background spectrum from the full COBE FIRAS data set," Astrophys. J. **473**, 576 (1996), arXiv:astro-ph/9605054.

[58] W. Hu and S. Dodelson, "Cosmic microwave background anisotropies," Annu. Rev. Astron. Astrophys. **40**, 171 (2002), arXiv:astro-ph/0110414.

[59] W. Hu and M. J. White, "A CMB polarization primer," New Astron. **2**, 323 (1997), arXiv:astro-ph/9706147.

[60] S. Chabanier, M. Millea, and N. Palanque-Delabrouille, "Matter power spectrum: From Lyα forest to CMB scales," Mon. Not. R. Astron. Soc. **489**, 2247 (2019), arXiv:1905.08103 [astro-ph.CO].

[61] M. Viel, K. Markovic, M. Baldi, and J. Weller, "The non-linear matter power spectrum in warm dark matter cosmologies," Mon. Not. R. Astron. Soc. **421**, 50 (2012), arXiv:1107.4094 [astro-ph.CO].

[62] E. Bertschinger, "Simulations of structure formation in the Universe," Annu. Rev. Astron. Astrophys. **36**, 599 (1998).

[63] J. Diemand and B. Moore, "The structure and evolution of cold dark matter halos," Adv. Sci. Lett. **4**, 297 (2011), arXiv:0906.4340 [astro-ph.CO].

[64] R. S. Somerville and R. Davé, "Physical models of galaxy formation in a cosmological framework," Annu. Rev. Astron. Astrophys. **53**, 51 (2015), arXiv:1412.2712 [astro-ph.GA].

[65] M. Vogelsberger, F. Marinacci, P. Torrey, and E. Puchwein, "Cosmological simulations of galaxy formation," Nat. Rev. Phys. **2**, 42 (2020), arXiv:1909.07976 [astro-ph.GA].

[66] R. H. Cyburt, B. D. Fields, K. A. Olive, and T.-H. Yeh, "Big Bang nucleosynthesis: Present status," Rev. Mod. Phys. **88**, 015004 (2016), arXiv:1505.01076 [astro-ph.CO].

[67] B. D. Fields, K. A. Olive, T.-H. Yeh, and C. Young, "Big-Bang nucleosynthesis after Planck," J. Cosmol. Astropart. Phys. **03**, 010 (2020), [Erratum: JCAP **11**, E02 (2020)], arXiv:1912.01132 [astro-ph.CO].

[68] A. Arbey, "AlterBBN: A program for calculating the BBN abundances of the elements in alternative cosmologies," Comput. Phys. Commun. **183**, 1822 (2012), arXiv:1106.1363 [astro-ph.CO].

[69] A. Arbey, J. Auffinger, K. P. Hickerson, and E. S. Jenssen, "AlterBBN v2: A public code for calculating Big-Bang nucleosynthesis constraints in alternative cosmologies," Comput. Phys. Commun. **248**, 106982 (2020), arXiv:1806.11095 [astro-ph.CO].

[70] O. Pisanti, G. Mangano, G. Miele, and P. Mazzella, "Primordial deuterium after LUNA: Concordances and error budget," J. Cosmol. Astropart. Phys. **04**, 020 (2021), arXiv:2011.11537 [astro-ph.CO].

[71] C. Iliadis, *Nuclear Physics of Stars*, 2nd ed. (Wiley-VCH, Weinheim, 2015).

[72] B. D. Fields, "The primordial lithium problem," Annu. Rev. Nucl. Part. Sci. **61**, 47 (2011), arXiv:1203.3551 [astro-ph.CO].

[73] O. Buchmueller, C. Doglioni, and L.-T. Wang, "Search for dark matter at colliders," Nat. Phys. **13**, 217 (2017), arXiv:1912.12739 [hep-ex].

[74] F. Kahlhoefer, "Review of LHC dark matter searches," Int. J. Mod. Phys. A **32**, 1730006 (2017), arXiv:1702.02430 [hep-ph].

[75] B. Penning, "The pursuit of dark matter at colliders—An overview," J. Phys. G: Nucl. Part. Phys. **45**, 063001 (2018), arXiv:1712.01391 [hep-ex].

[76] A. Boveia and C. Doglioni, "Dark matter searches at colliders," Annu. Rev. Nucl. Part. Sci. **68**, 429 (2018), arXiv:1810.12238 [hep-ex].

[77] N. Craig, H. K. Lou, M. McCullough, and A. Thalapillil, "The Higgs portal above threshold," J. High Energy Phys. **02**, 127 (2016), arXiv:1412.0258 [hep-ph].

[78] J. Baglio, A. Djouadi, and J. Quevillon, "Prospects for Higgs physics at energies up to 100 TeV," Rep. Prog. Phys. **79**, 116201 (2016), arXiv:1511.07853 [hep-ph].

[79] G. Arcadi, A. Djouadi, and M. Kado, "The Higgs-portal for dark matter: Effective field theories versus concrete realizations," Eur. Phys. J. C **81**, 653 (2021), arXiv:2101.02507 [hep-ph].

[80] S. Argyropoulos, O. Brandt, and U. Haisch, "Collider searches for dark matter through the Higgs lens," Symmetry **13**, 2406 (2021), arXiv:2109.13597 [hep-ph].

[81] M. Cepeda, S. Gori, V. M. Outschoorn, and J. Shelton, "Exotic Higgs decays," Annu. Rev. Nucl. Part. Sci. **72**, 119 (2021), arXiv:2111.12751 [hep-ph].

[82] A. Abada *et al.* (FCC Collaboration), "FCC-ee: The lepton collider," Eur. Phys. J. Spec. Top. **228**, 261 (2019).

[83] M. Dong *et al.* (CEPC Study Group), "CEPC conceptual design report: Volume 2—Physics & detector," (2018), arXiv:1811.10545 [hep-ex].

[84] T. K. Charles *et al.* (CLIC Collaboration and CLICdp Collaboration), "The Compact Linear Collider (CLIC)—2018 summary report," CERN Yellow Rep. **2/2018** (2018), 10.23731/CYRM-2018-002, arXiv:1812.06018 [physics.acc-ph].

[85] "The International Linear Collider technical design report—Volume 2: Physics," (2013), arXiv:1306.6352 [hep-ph].

[86] O. Cerri, M. de Gruttola, M. Pierini, A. Podo, and G. Rolandi, "Study the effect of beam energy spread and detector resolution on the search for Higgs boson decays to invisible particles at a future e^+e^- circular collider," Eur. Phys. J. C **77**, 116 (2017), arXiv:1605.00100 [hep-ex].

[87] Y. Tan *et al.*, "Search for invisible decays of the Higgs boson produced at the CEPC," Chinese Phys. C **44**, 123001 (2020), arXiv:2001.05912 [hep-ex].

[88] K. Mekala, A. F. Zarnecki, B. Grzadkowski, and M. Iglicki (CLICdp Collaboration), "Sensitivity to invisible scalar decays at CLIC," Eur. Phys. J. Plus **136**, 160 (2021), arXiv:2002.06034 [hep-ex].

[89] A. Abada *et al.* (FCC Collaboration), "FCC-hh: The hadron collider," Eur. Phys. J. Spec. Top. **228**, 755 (2019).

[90] J. Tang, Y. Zhang, Q. Xu, J. Gao, X. Lou, and Y. Wang, "Study overview for Super Proton-Proton Collider," in *Snowmass 2021* (2022) arXiv:2203.07987 [hep-ex].

[91] D. Goncalves, T. Plehn, and J. M. Thompson, "Weak boson fusion at 100 TeV," Phys. Rev. D **95**, 095011 (2017), arXiv:1702.05098 [hep-ph].

[92] G. Arcadi, A. Djouadi, and M. Raidal, "Dark matter through the Higgs portal," Phys. Rep. **842**, 1 (2020), arXiv:1903.03616 [hep-ph].

[93] M. Aaboud *et al.* (ATLAS Collaboration), "Search for dark matter in events with a hadronically decaying vector boson and missing transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector," J. High Energy Phys. **10**, 180 (2018), arXiv:1807.11471 [hep-ex].

[94] G. Aad *et al.* (ATLAS Collaboration), "Search for dark matter in association with an energetic photon in pp collisions at $\sqrt{s}$ = 13 TeV with the ATLAS detector," J. High Energy Phys. **02**, 226 (2021), arXiv:2011.05259 [hep-ex].

[95] G. Aad *et al.* (ATLAS Collaboration), "Search for new phenomena in events with an energetic jet and missing transverse momentum in pp collisions at $\sqrt{s} = 13\,\text{TeV}$ with the ATLAS detector," Phys. Rev. D **103**, 112006 (2021), arXiv:2102.10874 [hep-ex].

[96] A. M. Sirunyan *et al.* (CMS Collaboration), "Search for dark matter produced with an energetic jet or a hadronically decaying W or Z boson at $\sqrt{s} = 13\,\text{TeV}$," J. High Energy Phys. **07**, 014 (2017), arXiv:1703.01651 [hep-ex].

[97] A. M. Sirunyan *et al.* (CMS Collaboration), "Search for new physics in final states with a single photon and missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV," J. High Energy Phys. **02**, 074 (2019), arXiv:1810.00196 [hep-ex].

[98] A. M. Sirunyan *et al.* (CMS Collaboration), "Search for supersymmetry in proton-proton collisions at $\sqrt{s} = 13$ TeV in events with high-momentum Z bosons and missing transverse momentum," J. High Energy Phys. **09**, 149 (2020), arXiv:2008.04422 [hep-ex].

[99] A. Tumasyan *et al.* (CMS Collaboration), "Search for new particles in events with energetic jets and large missing transverse momentum in proton-proton collisions at $\sqrt{s}$ = 13 TeV," J. High Energy Phys. **11**, 153 (2021), arXiv:2107.13021 [hep-ex].

[100] R. J. Gaitskell, "Direct detection of dark matter," Annu. Rev. Nucl. Part. Sci. **54**, 315 (2004).

[101] D. Akimov, "Techniques and results for the direct detection of dark matter (review)," Nucl. Instrum. Methods Phys. Res. A **628**, 50 (2011).

[102] M. Schumann, "Direct detection of WIMP dark matter: Concepts and status," J. Phys. G: Nucl. Part. Phys. **46**, 103003 (2019), arXiv:1903.03026 [astro-ph.CO].

[103] R. Schönrich, "Galactic rotation and solar motion from stellar kinematics," Mon. Not. R. Astron. Soc. **427**, 274 (2012), arXiv:1207.3079 [astro-ph.GA].

[104] C. R. Hayes, D. R. Law, and S. R. Majewski, "Constraining the solar Galactic reflex velocity using Gaia observations of the Sagittarius stream," Astrophys. J. Lett. **867**, L20 (2018), arXiv:1809.07654 [astro-ph.GA].

[105] N. W. Evans, C. A. J. O'Hare, and C. McCabe, "Refinement of the standard halo model for dark matter searches in light of the Gaia Sausage," Phys. Rev. D **99**, 023012 (2019), arXiv:1810.11468 [astro-ph.GA].

[106] H. H. Koppelman and A. Helmi, "Determination of the escape velocity of the Milky Way using a halo sample selected based on proper motion," Astron. Astrophys. **649**, A136 (2021), arXiv:2006.16283 [astro-ph.GA].

[107] L. Necib and T. Lin, "Substructure at high speed. II. The local escape velocity and Milky Way mass with Gaia eDR3," Astrophys. J. **926**, 189 (2022), arXiv:2102.02211 [astro-ph.GA].

[108] Z. Prudil *et al.*, "Milky Way archaeology using RR Lyrae and Type II Cepheids. II. High-velocity RR Lyrae stars and Milky Way mass," Astron. Astrophys. **664**, A148 (2022), arXiv:2206.00417 [astro-ph.GA].

[109] R. Bernabei *et al.*, "First model independent results from DAMA/LIBRA-phase2," Nucl. Phys. At. Energy **19**, 307 (2018), arXiv:1805.10486 [hep-ex].

[110] G. Adhikari *et al.* (COSINE-100 Collaboration), "Strong constraints from COSINE-100 on the DAMA dark matter results using the same sodium iodide target," Sci. Adv. **7**, abk2699 (2021), arXiv:2104.03537 [hep-ex].

[111] M. Schumann, L. Baudis, L. Bütikofer, A. Kish, and M. Selvi, "Dark matter sensitivity of multi-ton liquid xenon detectors," J. Cosmol. Astropart. Phys. **10**, 016 (2015), arXiv:1506.08309 [physics.ins-det].

[112] E. Aprile *et al.* (XENON Collaboration), "Search for light dark matter interactions enhanced by the Migdal effect or bremsstrahlung in XENON1T," Phys. Rev. Lett. **123**, 241803 (2019), arXiv:1907.12771 [hep-ex].

[113] Y. Meng *et al.* (PandaX-4T Collaboration), "Dark matter search results from the PandaX-4T commissioning run," Phys. Rev. Lett. **127**, 261802 (2021), arXiv:2107.13438 [hep-ex].

[114] R. Agnese *et al.* (SuperCDMS Collaboration), "Search for low-mass dark matter with CDMSlite using a profile likelihood fit," Phys. Rev. D **99**, 062001 (2019), arXiv:1808.09098 [astro-ph.CO].

[115] A. H. Abdelhameed *et al.* (CRESST Collaboration), "First results from the CRESST-III low-mass dark matter program," Phys. Rev. D **100**, 102002 (2019), arXiv:1904.00498 [astro-ph.CO].

[116] A. Juillard *et al.*, "Low-noise HEMTs for coherent elastic neutrino scattering and low-mass dark matter cryogenic semiconductor detectors," J. Low Temp. Phys. **199**, 798 (2019), arXiv:1909.02879 [physics.ins-det].

[117] E. Armengaud *et al.* (EDELWEISS Collaboration), "Search for sub-GeV dark matter via the Migdal effect with an EDELWEISS germanium detector with NbSi transition-edge sensors," Phys. Rev. D **106**, 062004 (2022), arXiv:2203.03993 [astro-ph.GA].

[118] R. Essig, M. Fernandez-Serra, J. Mardon, A. Soto, T. Volansky, and T.-T. Yu, "Direct detection of sub-GeV dark matter with semiconductor targets," J. High Energy Phys. **05**, 046 (2016), arXiv:1509.01598 [hep-ph].

[119] L. Barak *et al.* (SENSEI Collaboration), "SENSEI: Direct-detection results on sub-GeV dark matter from a new Skipper-CCD," Phys. Rev. Lett. **125**, 171802 (2020), arXiv:2004.11378 [astro-ph.CO].

[120] D. W. Amaral *et al.* (SuperCDMS Collaboration), "Constraints on low-mass, relic dark matter candidates from a surface-operated SuperCDMS single-charge sensitive detector," Phys. Rev. D **102**, 091101 (2020), arXiv:2005.14067 [hep-ex].

[121] E. Aprile *et al.* (XENON Collaboration), "Emission of single and few electrons in XENON1T and limits on light dark matter," Phys. Rev. D **106**, 022001 (2022), arXiv:2112.12116 [hep-ex].

[122] Z. Y. Zhang *et al.* (CDEX Collaboration), "Constraints on sub-GeV dark matter–electron scattering from the CDEX-10 experiment," Phys. Rev. Lett. **129**, 221301 (2022), arXiv:2206.04128 [hep-ex].

[123] C. A. J. O'Hare, "New definition of the neutrino floor for direct dark matter searches," Phys. Rev. Lett. **127**, 251802 (2021), arXiv:2109.03116 [hep-ph].

[124] A. K. Drukier, K. Freese, and D. N. Spergel, "Detecting cold dark-matter candidates," Phys. Rev. D **33**, 3495 (1986).

[125] K. Freese, J. Frieman, and A. Gould, "Signal modulation in cold-dark-matter detection," Phys. Rev. D **37**, 3388 (1988).

[126] K. Freese, M. Lisanti, and C. Savage, "Colloquium: Annual modulation of dark matter," Rev. Mod. Phys. **85**, 1561 (2013), arXiv:1209.3339 [astro-ph.CO].

[127] J. H. Davis, "Dark matter vs. neutrinos: The effect of astrophysical uncertainties and timing information on the neutrino floor," J. Cosmol. Astropart. Phys. **03**, 012 (2015), arXiv:1412.1475 [hep-ph].

[128] F. Mayet *et al.*, "A review of the discovery reach of directional dark matter detection," Phys. Rep. **627**, 1 (2016), arXiv:1602.03781 [astro-ph.CO].

[129] S. E. Vahsen, C. A. J. O'Hare, and D. Loomba, "Directional recoil detection," Annu. Rev. Nucl. Part. Sci. **71**, 189 (2021), arXiv:2102.04596 [physics.ins-det].

[130] S. E. Vahsen *et al.*, "CYGNUS: Feasibility of a nuclear recoil observatory with directional sensitivity to dark matter and neutrinos," (2020), arXiv:2008.12587 [physics.ins-det].

[131] S. Rajendran, N. Zobrist, A. O. Sushkov, R. Walsworth, and M. Lukin, "A method for directional detection of dark matter using spectroscopy of crystal defects," Phys. Rev. D **96**, 035009 (2017), arXiv:1705.09760 [hep-ph].

[132] M. C. Marshall, M. J. Turner, M. J. H. Ku, D. F. Phillips, and R. L. Walsworth, "Directional detection of dark matter with diamond," Quantum Sci. Technol. **6**, 024011 (2021), arXiv:2009.01028 [physics.ins-det].

[133] R. Ebadi *et al.*, "Directional detection of dark matter using solid-state quantum sensing," in *2022 Snowmass Summer Study* (2022) arXiv:2203.06037 [physics.ins-det].

[134] N. Agafonova *et al.* (NEWSdm Collaboration), "Discovery potential for directional dark matter detection with nuclear emulsions," Eur. Phys. J. C **78**, 578 (2018), arXiv:1705.00613 [astro-ph.CO].

[135] A. Golovatiuk, A. Ustyuzhanin, A. Alexandrov, and G. De Lellis, "Deep learning for direct dark matter search with nuclear emulsions," Comput. Phys. Commun. **275**, 108312 (2022), arXiv:2106.11995 [hep-ex].

[136] R. Bernabei *et al.*, "Further results from DAMA/LIBRA-phase 2 and perspectives," Nucl. Phys. At. Energy **22**, 329 (2021).

[137] J. L. Feng, J. Kumar, D. Marfatia, and D. Sanford, "Isospin-violating dark matter," Phys. Lett. B **703**, 124 (2011), arXiv:1102.4331 [hep-ph].

[138] G. Adhikari *et al.* (COSINE-100 Collaboration), "Three-year annual modulation search with COSINE-100," Phys. Rev. D **106**, 052005 (2022), arXiv:2111.08863 [hep-ex].

[139] J. Amare *et al.*, "Annual modulation results from three-year exposure of ANAIS-112," Phys. Rev. D **103**, 102005 (2021), arXiv:2103.01175 [astro-ph.IM].

[140] D. Buttazzo, P. Panci, N. Rossi, and A. Strumia, "Annual modulations from secular variations: Relaxing DAMA?" J. High Energy Phys. **04**, 137 (2020), arXiv:2002.00459 [hep-ex].

[141] G. Adhikari *et al.* (COSINE-100 Collaboration), "An induced annual modulation signature in COSINE-100 data by DAMA/LIBRA's analysis method," Sci. Rep. **13**, 4676 (2023), arXiv:2208.05158 [hep-ex].

[142] M. Antonello *et al.* (SABRE Collaboration), "The SABRE project and the SABRE Proof-of-Principle," Eur. Phys. J. C **79**, 363 (2019), arXiv:1806.09340 [physics.ins-det].

[143] V. Zema (COSINUS Collaboration), "COSINUS: A NaI-based cryogenic calorimeter for direct dark matter search," Nuovo Cimento C **42**, 228 (2020).

[144] J. M. Gaskins, "A review of indirect searches for particle dark matter," Contemp. Phys. **57**, 496 (2016), arXiv:1604.00014 [astro-ph.HE].

[145] C. Pérez de los Heros, “Status, challenges and directions in indirect dark matter searches,” Symmetry **12**, 1648 (2020), arXiv:2008.11561 [astro-ph.HE].

[146] M. Cirelli, G. Corcella, A. Hektor, G. Hutsi, M. Kadastik, P. Panci, M. Raidal, F. Sala, and A. Strumia, “PPPC 4 DM ID: A Poor Particle Physicist Cookbook for Dark Matter Indirect Detection,” J. Cosmol. Astropart. Phys. **03**, 051 (2011), [Erratum: JCAP **10**, E01 (2012)], arXiv:1012.4515 [hep-ph].

[147] J. Buch, M. Cirelli, G. Giesen, and M. Taoso, “PPPC 4 DM Secondary: A Poor Particle Physicist Cookbook for secondary radiation from dark matter,” J. Cosmol. Astropart. Phys. **09**, 037 (2015), arXiv:1505.01049 [hep-ph].

[148] C. A. Argüelles, A. Diaz, A. Kheirandish, A. Olivares-Del-Campo, I. Safa, and A. C. Vincent, “Dark matter annihilation to neutrinos,” Rev. Mod. Phys. **93**, 035007 (2021), arXiv:1912.09486 [hep-ph].

[149] C. A. Argüelles, D. Delgado, A. Friedlander, A. Kheirandish, I. Safa, A. C. Vincent, and H. White, “Dark matter decay to neutrinos,” (2022), arXiv:2210.01303 [hep-ph].

[150] J. Conrad and O. Reimer, “Indirect dark matter searches in gamma and cosmic rays,” Nat. Phys. **13**, 224 (2017), arXiv:1705.11165 [astro-ph.HE].

[151] S. Ando *et al.*, “Snowmass2021 Cosmic Frontier: Synergies between dark matter searches and multiwavelength/multimessenger astrophysics,” in *Snowmass 2021* (2022) arXiv:2203.06781 [hep-ph].

[152] M. Baryakhtar *et al.*, “Dark matter In extreme astrophysical environments,” in *Snowmass 2021* (2022) arXiv:2203.07984 [hep-ph].

[153] K. Bechtol *et al.*, “Snowmass2021 Cosmic Frontier white paper: Dark matter physics from halo measurements,” in *2022 Snowmass Summer Study* (2022) arXiv:2203.07354 [hep-ph].

[154] S. Chakrabarti *et al.*, “Snowmass2021 Cosmic Frontier white paper: Observational facilities to study dark matter,” in *Snowmass 2021* (2022) arXiv:2203.06200 [astro-ph.CO].

[155] R. Brito *et al.*, “Snowmass2021 Cosmic Frontier white paper: Probing dark matter with small-scale astrophysical observations,” in *Snowmass 2021* (2022) arXiv:2203.15954 [hep-ph].

[156] A. Drlica-Wagner *et al.*, “Report of the Topical Group on Cosmic Probes of dark matter for Snowmass 2021,” in *Snowmass 2021* (2022) arXiv:2209.08215 [hep-ph].

[157] B. Carr and F. Kuhnel, “Primordial black holes as dark matter: Recent developments,” Annu. Rev. Nucl. Part. Sci. **70**, 355 (2020), arXiv:2006.02838 [astro-ph.CO].

[158] A. M. Green and B. J. Kavanagh, “Primordial black holes as a dark matter candidate,” J. Phys. G: Nucl. Part. Phys. **48**, 043001 (2021), arXiv:2007.10722 [astro-ph.CO].

[159] B. J. Kavanagh, “bradkav/pbhbounds: Release version,” Zenodo (2019), 10.5281/zenodo.3538999.

[160] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), “Observation of gravitational waves from a binary black hole merger,” Phys. Rev. Lett. **116**, 061102 (2016).

[161] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), “GW170817: Observation of gravitational waves from a binary neutron star inspiral,” Phys. Rev. Lett. **119**, 161101 (2017).

[162] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, Fermi GBM Collaboration, INTEGRAL, IceCube Collaboration, AstroSat Cadmium Zinc Telluride Imager Team, IPN Collaboration, Insight-HMXT Collaboration, ANTARES Collaboration, Swift Collaboration, AGILE Team, 1M2H Team, Dark Energy Camera GW-EM Collaboration, DES Collaboration, DLT40 Collaboration, GRAWITA (GRAvitational Wave Inaf TeAm), Fermi Large Area Telescope Collaboration, ATCA (Australian Telescope Compact Array), ASKAP (Australian SKA Pathfinder), Las Cumbres Observatory Group, OzGrav Collaboration, DWF (Deeper, Wider, Faster program) Collaboration, AST3 Collaboration, CAASTRO Collaboration, VINROUGE Collaboration, MASTER Collaboration, J-GEM, GROWTH Collaboration, JAGWAR Collaboration, Caltech-NRAO, TTU-NRAO, NuSTAR Collaboration, Pan-STARRS, MAXI Team, TZAC Consortium, KU Collaboration, Nordic Optical Telescope, ePESSTO, GROND, Texas Tech University, SALT Group, TOROS (Transient Robotic Telescope of the South) Collaboration, BOOTES Collaboration, MWA (Murchison Widefield Array), CALET Collaboration, IKI-GW Follow-up Collaboration, H.E.S.S. Collaboration, LOFAR Collaboration, LWA (Long Wavelength Array), HAWC Collaboration, Pierre Auger Collaboration, ALMA Collaboration, EURO VLBI Team, Pi of the Sky Collaboration, Chandra Team at McGill University, DFN (Desert Fireball Newtork), ATLAS, High Time Resolution Universe Survey, RIMAS, RATIR, and SKA South Africa/MeerKAT), “Multi-messenger observations of a binary neutron star merger,” Astrophys. J. Lett. **848**, L12 (2017), arXiv:1710.05833 [astro-ph.HE].

[163] R. Abbott *et al.* (LIGO Scientific Collaboration, KAGRA Collaboration, and Virgo Collaboration), “Constraints on dark photon dark matter using data from LIGO’s and Virgo’s third observing run,” Phys. Rev. D **105**, 063030 (2022), arXiv:2105.13085 [astro-ph.CO].

[164] K. Nagano, H. Nakatsuka, S. Morisaki, T. Fujita, Y. Michimura, and I. Obata, “Axion dark matter search using arm cavity transmitted beams of gravita-

tional wave detectors," Phys. Rev. D **104**, 062008 (2021), arXiv:2106.06800 [hep-ph].

[165] L. Sagunski, J. Zhang, M. C. Johnson, L. Lehner, M. Sakellariadou, S. L. Liebling, C. Palenzuela, and D. Neilsen, "Neutron star mergers as a probe of modifications of general relativity with finite-range scalar forces," Phys. Rev. D **97**, 064016 (2018), arXiv:1709.06634 [gr-qc].

[166] J. Huang, M. C. Johnson, L. Sagunski, M. Sakellariadou, and J. Zhang, "Prospects for axion searches with Advanced LIGO through binary mergers," Phys. Rev. D **99**, 063013 (2019), arXiv:1807.02133 [hep-ph].

[167] J. Zhang, Z. Lyu, J. Huang, M. C. Johnson, L. Sagunski, M. Sakellariadou, and H. Yang, "First constraints on nuclear coupling of axionlike particles from the binary neutron star gravitational wave event GW170817," Phys. Rev. Lett. **127**, 161101 (2021), arXiv:2105.13963 [hep-ph].

[168] R. Brito, V. Cardoso, and P. Pani, "Superradiance: New frontiers in black hole physics," Lect. Notes Phys. **906**, 1 (2015), arXiv:1501.06570 [gr-qc].

[169] V. Cardoso, Ó. Dias, G. S. Hartnett, M. Middleton, P. Pani, and J. E. Santos, "Constraining the mass of dark photons and axion-like particles through black-hole superradiance," J. Cosmol. Astropart. Phys. **03**, 043 (2018), arXiv:1801.01420 [gr-qc].

[170] M. Baryakhtar, M. Galanis, R. Lasenby, and O. Simon, "Black hole superradiance of self-interacting scalar fields," Phys. Rev. D **103**, 095019 (2021), arXiv:2011.11646 [hep-ph].

[171] K. K. Y. Ng, S. Vitale, O. A. Hannuksela, and T. G. F. Li, "Constraints on ultralight scalar bosons within black hole spin measurements from the LIGO-Virgo GWTC-2," Phys. Rev. Lett. **126**, 151102 (2021), arXiv:2011.06010 [gr-qc].

[172] R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs, "The Australia Telescope National Facility pulsar catalogue," Astron. J. **129**, 1993 (2005), arXiv:astro-ph/0412641.

[173] M. Maiorano, F. De Paolis, and A. A. Nucita, "Principles of gravitational-wave detection with pulsar timing arrays," Symmetry **13**, 2418 (2021), arXiv:2112.08064 [astro-ph.GA].

[174] A. N. Lommen, "Pulsar timing arrays: The promise of gravitational wave detection," Rep. Prog. Phys. **78**, 124901 (2015).

[175] G. Hobbs and S. Dai, "Gravitational wave research using pulsar timing arrays," Natl. Sci. Rev. **4**, 707 (2017), arXiv:1707.01615 [astro-ph.IM].

[176] N. K. Porayko *et al.* (PPTA Collaboration), “Parkes Pulsar Timing Array constraints on ultralight scalar-field dark matter,” Phys. Rev. D **98**, 102002 (2018), arXiv:1810.03227 [astro-ph.CO].

[177] R. Kato and J. Soda, “Search for ultralight scalar dark matter with NANOGrav pulsar timing arrays,” J. Cosmol. Astropart. Phys. **09**, 036 (2020), arXiv:1904.09143 [astro-ph.HE].

[178] E. R. Siegel, M. P. Hertzberg, and J. N. Fry, “Probing dark matter substructure with pulsar timing,” Mon. Not. R. Astron. Soc. **382**, 879 (2007), arXiv:astro-ph/0702546.

[179] H. Ramani, T. Trickle, and K. M. Zurek, “Observability of dark matter substructure with pulsar timing correlations,” J. Cosmol. Astropart. Phys. **12**, 033 (2020), arXiv:2005.03030 [astro-ph.CO].

[180] M. S. Delos and T. Linden, “Dark matter microhalos in the solar neighborhood: Pulsar timing signatures of early matter domination,” Phys. Rev. D **105**, 123514 (2022), arXiv:2109.03240 [astro-ph.CO].

[181] S. Vegetti, L. V. E. Koopmans, A. Bolton, T. Treu, and R. Gavazzi, “Detection of a dark substructure through gravitational imaging,” Mon. Not. R. Astron. Soc. **408**, 1969 (2010), arXiv:0910.0760 [astro-ph.CO].

[182] S. Vegetti, D. Lagattuta, J. McKean, M. W. Auger, C. D. Fassnacht, and L. V. E. Koopmans, “Gravitational detection of a low-mass dark satellite galaxy at cosmological distance,” Nature **481**, 341 (2012), arXiv:1201.3643 [astro-ph.CO].

[183] Y. D. Hezaveh *et al.*, “Detection of lensing substructure using ALMA observations of the dusty galaxy SDP.81,” Astrophys. J. **823**, 37 (2016), arXiv:1601.01388 [astro-ph.CO].

[184] P. Natarajan *et al.*, “Mapping substructure in the HST Frontier Fields cluster lenses and in cosmological simulations,” Mon. Not. R. Astron. Soc. **468**, 1962 (2017), arXiv:1702.04348 [astro-ph.GA].

[185] A. M. Nierenberg, T. Treu, S. A. Wright, C. D. Fassnacht, and M. W. Auger, “Detection of substructure with adaptive optics integral field spectroscopy of the gravitational lens B1422+231,” Mon. Not. R. Astron. Soc. **442**, 2434 (2014), arXiv:1402.1496 [astro-ph.GA].

[186] A. M. Nierenberg *et al.*, “Probing dark matter substructure in the gravitational lens HE 0435−1223 with the WFC3 grism,” Mon. Not. R. Astron. Soc. **471**, 2224 (2017), arXiv:1701.05188 [astro-ph.CO].

[187] D. J. Lagattuta, J. Richard, H. Ebeling, Q. Basto, C. Cerny, A. Edge, M. Jauzac, G. Mahler, and R. Massey, “RXJ0437+00: Constraining dark matter with

exotic gravitational lenses," Mon. Not. R. Astron. Soc. **522**, 1091 (2023), arXiv:2303.09568 [astro-ph.CO].

[188] J. W. Hsueh, C. D. Fassnacht, S. Vegetti, J. P. McKean, C. Spingola, M. W. Auger, L. V. E. Koopmans, and D. J. Lagattuta, "SHARP — II. Mass structure in strong lenses is not necessarily dark matter substructure: A flux ratio anomaly from an edge-on disc in B1555+375," Mon. Not. R. Astron. Soc. **463**, L51 (2016), arXiv:1601.01671 [astro-ph.CO].

[189] J.-W. Hsueh, G. Despali, S. Vegetti, D. Xu, C. D. Fassnacht, and R. B. Metcalf, "Flux-ratio anomalies from discs and other baryonic structures in the Illustris simulation," Mon. Not. Roy. Astron. Soc. **475**, 2438 (2018), arXiv:1707.07680 [astro-ph.GA].

[190] D. Gilman, S. Birrer, T. Treu, C. R. Keeton, and A. Nierenberg, "Probing the nature of dark matter by forward modelling flux ratios in strong gravitational lenses," Mon. Not. Roy. Astron. Soc. **481**, 819 (2018), arXiv:1712.04945 [astro-ph.CO].

[191] W. Enzi, S. Vegetti, G. Despali, J.-W. Hsueh, and R. B. Metcalf, "Systematic errors in strong gravitational lensing reconstructions, a numerical simulation perspective," Mon. Not. R. Astron. Soc. **496**, 1718 (2020), arXiv:1911.02581 [astro-ph.GA].

[192] Ž. Ivezić *et al.* (LSST Collaboration), "LSST: From science drivers to reference design and anticipated data products," Astrophys. J. **873**, 111 (2019), arXiv:0805.2366 [astro-ph].

[193] I. Tutusaus *et al.* (EUCLID Consortium), "Euclid: The importance of galaxy clustering and weak lensing cross-correlations within the photometric Euclid survey," Astron. Astrophys. **643**, A70 (2020), arXiv:2005.00055 [astro-ph.CO].

[194] C. Weiner, S. Serjeant, and C. Sedgwick, "Predictions for strong-lens detections with the Nancy Grace Roman Space Telescope," Res. Notes of the Am. Astron. Soc. **4**, 190 (2020), arXiv:2010.15173 [astro-ph.GA].

[195] A. Lewis and A. Challinor, "Weak gravitational lensing of the CMB," Phys. Rep. **429**, 1 (2006), arXiv:astro-ph/0601594.

[196] H. N. Nguyễn, N. Sehgal, and M. Madhavacheril, "Measuring the small-scale matter power spectrum with high-resolution CMB lensing," Phys. Rev. D **99**, 023502 (2019), arXiv:1710.03747 [astro-ph.CO].

[197] A. Krolewski, S. Ferraro, and M. White, "Cosmological constraints from unWISE and Planck CMB lensing tomography," J. Cosmol. Astropart. Phys. **12**, 028 (2021), arXiv:2105.03421 [astro-ph.CO].

[198] H. Miyatake, Y. Harikane, M. Ouchi, Y. Ono, N. Yamamoto, A. J. Nishizawa, N. Bahcall, S. Miyazaki, and A. A. P. Malagón, "First identification of a CMB lensing signal produced by 1.5 million galaxies at $z \sim 4$: Constraints on matter density fluctuations at high redshift," Phys. Rev. Lett. **129**, 061301 (2022), arXiv:2103.15862 [astro-ph.CO].

[199] A. Kovács *et al.* (DES Collaboration), "Dark Energy Survey Year 3 results: Imprints of cosmic voids and superclusters in the Planck CMB lensing map," Mon. Not. Roy. Astron. Soc. **515**, 4417 (2022), arXiv:2203.11306 [astro-ph.CO].

[200] K. Abazajian *et al.* (CMB-S4 Collaboration), "CMB-S4 science case, reference design, and project plan," (2019), arXiv:1907.04473 [astro-ph.IM].

[201] P. Ade *et al.* (Simons Observatory Collaboration), "The Simons Observatory: Science goals and forecasts," J. Cosmol. Astropart. Phys. **02**, 056 (2019), arXiv:1808.07445 [astro-ph.CO].

[202] N. Sehgal *et al.*, "CMB-HD: An ultra-deep, high-resolution millimeter-wave survey over half the sky," in *Bulletin of the American Astronomical Society*, Vol. 51 (2019) p. 6, arXiv:1906.10134 [astro-ph.CO].

[203] Z. Li, V. Gluscevic, K. K. Boddy, and M. S. Madhavacheril, "Disentangling dark physics with cosmic microwave background experiments," Phys. Rev. D **98**, 123524 (2018), arXiv:1806.10165 [astro-ph.CO].

[204] C. Dvorkin *et al.*, "Dark matter physics from the CMB-S4 experiment," in *Snowmass 2021* (2022) arXiv:2203.07064 [hep-ph].

[205] M. Milgrom, "A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis." Astrophys. J. **270**, 365 (1983).

[206] R. B. Tully and J. R. Fisher, "A new method of determining distances to galaxies," Astron. Astrophys. **54**, 661 (1977).

[207] S. S. McGaugh, J. M. Schombert, G. D. Bothun, and W. J. G. de Blok, "The baryonic Tully-Fisher relation," Astrophys. J. Lett. **533**, L99 (2000), arXiv:astro-ph/0003001.

[208] M. Milgrom, "The modified dynamics as a vacuum effect," Phys. Lett. A **253**, 273 (1999), arXiv:astro-ph/9805346.

[209] K.-H. Chae, F. Lelli, H. Desmond, S. S. McGaugh, P. Li, and J. M. Schombert, "Testing the strong equivalence principle: Detection of the external field effect in rotationally supported galaxies," Astrophys. J. **904**, 51 (2020), [Erratum: ApJ 910, 81 (2021)], arXiv:2009.11525 [astro-ph.GA].

[210] K. H. Chae, H. Desmond, F. Lelli, S. S. McGaugh, and J. M. Schombert, "Testing the strong equivalence principle. II. Relating the external field effect in galaxy rotation curves to the large-scale structure of the Universe," Astrophys. J. **921**, 104 (2021), arXiv:2109.04745 [astro-ph.GA].

[211] C. Skordis and T. Złośnik, "New relativistic theory for modified Newtonian dynamics," Phys. Rev. Lett. **127**, 161302 (2021), arXiv:2007.00082 [astro-ph.CO].

[212] S. S. McGaugh, "Testing galaxy formation and dark matter with low surface brightness galaxies," Stud. Hist. Philos. Sci. **88**, 220 (2021), arXiv:2103.05003 [astro-ph.GA].

[213] S. Schael *et al.* (ALEPH Collaboration, DELPHI Collaboration, L3 Collaboration, OPAL Collaboration, SLD Collaboration, LEP Electroweak Working Group, SLD Electroweak Group, and SLD Heavy Flavour Group), "Precision electroweak measurements on the Z resonance," Phys. Rep. **427**, 257 (2006), arXiv:hep-ex/0509008.

[214] I. J. R. Aitchison and A. J. G. Hey, *Gauge Theories in Particle Physics: A Practical Introduction*, 4th ed. (Taylor & Francis, Boca Raton, 2013).

[215] M. Thomson, *Modern Particle Physics* (Cambridge Univ. Press, Cambridge, 2013).

[216] M. D. Schwartz, *Quantum Field Theory and the Standard Model* (Cambridge Univ. Press, Cambridge, 2014).

[217] T. D. Lee and C.-N. Yang, "Question of parity conservation in weak interactions," Phys. Rev. **104**, 254 (1956).

[218] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson, "Experimental test of parity conservation in beta decay," Phys. Rev. **105**, 1413 (1957).

[219] M. Goldhaber, L. Grodzins, and A. W. Sunyar, "Helicity of neutrinos," Phys. Rev. **109**, 1015 (1958).

[220] F. Reines and C. L. Cowan, "Detection of the free neutrino," Phys. Rev. **92**, 830 (1953).

[221] G. Danby, J. M. Gaillard, K. A. Goulianos, L. M. Lederman, N. B. Mistry, M. Schwartz, and J. Steinberger, "Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos," Phys. Rev. Lett. **9**, 36 (1962).

[222] K. Kodama *et al.* (DONuT Collaboration), "Final tau-neutrino results from the DONuT experiment," Phys. Rev. D **78**, 052002 (2008), arXiv:0711.0728 [hep-ex].

[223] P. Janot and S. Jadach, "Improved Bhabha cross section at LEP and the number of light neutrino species," Phys. Lett. B **803**, 135319 (2020), arXiv:1912.02067 [hep-ph].

[224] G. 't Hooft, "Symmetry breaking through Bell-Jackiw anomalies," Phys. Rev. Lett. **37**, 8 (1976).

[225] G. 't Hooft, "Computation of the quantum effects due to a four-dimensional pseudoparticle," Phys. Rev. D **14**, 3432 (1976), [Erratum: PRD **18**, 2199 (1978)].

[226] S. Weinberg, "A model of leptons," Phys. Rev. Lett. **19**, 1264 (1967).

[227] S. F. King, "Models of neutrino mass, mixing and CP violation," J. Phys. G: Nucl. Part. Phys. **42**, 123001 (2015), arXiv:1510.02091 [hep-ph].

[228] A. de Gouvêa, "Neutrino mass models," Annu. Rev. Nucl. Part. Sci. **66**, 197 (2016).

[229] E. Majorana, "Teoria simmetrica dell'elettrone e del positrone," Nuovo Cim. **14**, 171 (1937).

[230] S. Weinberg, "Baryon and lepton nonconserving processes," Phys. Rev. Lett. **43**, 1566 (1979).

[231] F. Bonnet, M. Hirsch, T. Ota, and W. Winter, "Systematic study of the $d = 5$ Weinberg operator at one-loop order," J. High Energy Phys. **07**, 153 (2012), arXiv:1204.5862 [hep-ph].

[232] A. Abada, C. Biggio, F. Bonnet, M. B. Gavela, and T. Hambye, "Low energy effects of neutrino masses," J. High Energy Phys. **12**, 061 (2007), arXiv:0707.4058 [hep-ph].

[233] J. Schechter and J. W. F. Valle, "Neutrinoless double beta decay in SU(2) $\times$ U(1) theories," Phys. Rev. D **25**, 2951 (1982).

[234] M. J. Dolinski, A. W. P. Poon, and W. Rodejohann, "Neutrinoless double-beta decay: Status and prospects," Annu. Rev. Nucl. Part. Sci. **69**, 219 (2019), arXiv:1902.04097 [nucl-ex].

[235] J. Engel and J. Menéndez, "Status and future of nuclear matrix elements for neutrinoless double-beta decay: A review," Rep. Prog. Phys. **80**, 046301 (2017).

[236] J. Heeck and W. Rodejohann, "Neutrinoless quadruple beta decay," Europhys. Lett. **103**, 32001 (2013), arXiv:1306.0580 [hep-ph].

[237] M. Hirsch, R. Srivastava, and J. W. F. Valle, "Can one ever prove that neutrinos are Dirac particles?" Phys. Lett. B **781**, 302 (2018), arXiv:1711.06181 [hep-ph].

[238] S. Bilenky, "Neutrino oscillations: From a historical perspective to the present status," Nucl. Phys. B **908**, 2 (2016), arXiv:1602.00170 [hep-ph].

[239] C. Giganti, S. Lavignac, and M. Zito, "Neutrino oscillations: The rise of the PMNS paradigm," Prog. Part. Nucl. Phys. **98**, 1 (2018), arXiv:1710.00715 [hep-ex].

[240] R. Davis Jr., D. S. Harmer, and K. C. Hoffman, "Search for neutrinos from the Sun," Phys. Rev. Lett. **20**, 1205 (1968).

[241] W. C. Haxton, R. G. Hamish Robertson, and A. M. Serenelli, "Solar neutrinos: Status and prospects," Annu. Rev. Astron. Astrophys. **51**, 21 (2013), arXiv:1208.5723 [astro-ph.SR].

[242] C. K. Jung, C. McGrew, T. Kajita, and T. Mann, "Oscillations of atmospheric neutrinos," Annu. Rev. Nucl. Part. Sci. **51**, 451 (2001).

[243] T. Kajita, "The measurement of neutrino properties with atmospheric neutrinos," Annu. Rev. Nucl. Part. Sci. **64**, 343 (2014).

[244] D. Casper *et al.*, "Measurement of atmospheric neutrino composition with the IMB-3 detector," Phys. Rev. Lett. **66**, 2561 (1991).

[245] R. Becker-Szendy *et al.*, "Neutrino measurements with the IMB detector," Nucl. Phys. B Proc. Suppl. **38**, 331 (1995).

[246] K. S. Hirata *et al.* (Kamiokande-II Collaboration), "Observation of a small atmospheric ν_μ/ν_e ratio in Kamiokande," Phys. Lett. B **280**, 146 (1992).

[247] W. W. M. Allison *et al.* (Soudan 2 Collaboration), "The atmospheric neutrino flavor ratio from a 3.9 fiducial kiloton-year exposure of Soudan 2," Phys. Lett. B **449**, 137 (1999), arXiv:hep-ex/9901024.

[248] Y. Fukuda *et al.* (Super-Kamiokande Collaboration), "Measurement of a small atmospheric ν_μ/ν_e ratio," Phys. Lett. B **433**, 9 (1998), arXiv:hep-ex/9803006.

[249] Y. Fukuda *et al.* (Super-Kamiokande Collaboration), "Study of the atmospheric neutrino flux in the multi-GeV energy range," Phys. Lett. B **436**, 33 (1998), arXiv:hep-ex/9805006.

[250] A. Gando *et al.* (KamLAND Collaboration), "Reactor on-off antineutrino measurement with KamLAND," Phys. Rev. D **88**, 033001 (2013), arXiv:1303.4667 [hep-ex].

[251] A. Bellerive *et al.* (SNO Collaboration), "The Sudbury Neutrino Observatory," Nucl. Phys. B **908**, 30 (2016), arXiv:1602.02469 [nucl-ex].

[252] K. Eguchi *et al.* (KamLAND Collaboration), "First results from KamLAND: Evidence for reactor anti-neutrino disappearance," Phys. Rev. Lett. **90**, 021802 (2003), arXiv:hep-ex/0212021.

[253] M. H. Ahn *et al.* (K2K Collaboration), "Measurement of neutrino oscillation by the K2K experiment," Phys. Rev. D **74**, 072003 (2006), arXiv:hep-ex/0606032.

[254] P. Adamson *et al.* (MINOS Collaboration), "Combined analysis of ν_μ disappearance and $\nu_\mu \to \nu_e$ appearance in MINOS using accelerator and atmospheric neutrinos," Phys. Rev. Lett. **112**, 191801 (2014), arXiv:1403.0867 [hep-ex].

[255] B. Pontecorvo, "Neutrino experiments and the problem of conservation of leptonic charge," Zh. Eksp. Teor. Fiz. **53**, 1717 (1967).

[256] F. F. Deppisch, *A Modern Introduction to Neutrino Physics* (Morgan & Claypool, San Rafael, 2019).

[257] Z. Xing and S. Zhou, *Neutrinos in Particle Physics, Astronomy and Cosmology* (Springer Berlin, Heidelberg, 2011).

[258] S. P. Mikheyev and A. Yu. Smirnov, "Resonant enhancement of ν oscillations in matter and solar-neutrino spectroscopy," Nuovo Cimento C **9**, 17 (1986), [Yad. Fiz. **42**, 1441 (1985)].

[259] L. Wolfenstein, "Neutrino oscillations in matter," Phys. Rev. D **17**, 2369 (1978).

[260] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, T. Schwetz, and A. Zhou, "The fate of hints: Updated global analysis of three-flavor neutrino oscillations," J. High Energy Phys. **09**, 178 (2020), arXiv:2007.14792 [hep-ph].

[261] M. Aker *et al.* (KATRIN Collaboration), "Direct neutrino-mass measurement with sub-electronvolt sensitivity," Nat. Phys. **18**, 160 (2022), arXiv:2105.08533 [hep-ex].

[262] A. Ashtari Esfahani *et al.* (Project 8 Collaboration), "Determining the neutrino mass with cyclotron radiation emission spectroscopy—Project 8," J. Phys. G: Nucl. Part. Phys. **44**, 054004 (2017), arXiv:1703.02037 [physics.ins-det].

[263] A. Ashtari Esfahani *et al.*, "Bayesian analysis of a future β decay experiment's sensitivity to neutrino mass scale and ordering," Phys. Rev. C **103**, 065501 (2021), arXiv:2012.14341 [physics.data-an].

[264] F. An *et al.* (JUNO Collaboration), "Neutrino physics with JUNO," J. Phys. G: Nucl. Part. Phys. **43**, 030401 (2016), arXiv:1507.05613 [physics.ins-det].

[265] B. Abi *et al.* (DUNE Collaboration), "Long-baseline neutrino oscillation physics potential of the DUNE experiment," Eur. Phys. J. C **80**, 978 (2020), arXiv:2006.16043 [hep-ex].

[266] P. Minkowski, “$\mu \to e\gamma$ at a rate of one out of 10^9 muon decays?” Phys. Lett. B **67**, 421 (1977).

[267] M. Gell-Mann, P. Ramond, and R. Slansky, “Complex spinors and unified theories,” Conf. Proc. C **790927**, 315 (1979), arXiv:1306.4669 [hep-th].

[268] R. N. Mohapatra and G. Senjanovic, “Neutrino mass and spontaneous parity nonconservation,” Phys. Rev. Lett. **44**, 912 (1980).

[269] T. Yanagida, “Horizontal symmetry and masses of neutrinos,” Prog. Theor. Phys. **64**, 1103 (1980).

[270] A. M. Abdullahi *et al.*, “The present and future status of heavy neutral leptons,” in *2022 Snowmass Summer Study* (2022) arXiv:2203.08039 [hep-ph].

[271] M. Drewes, “The phenomenology of right handed neutrinos,” Int. J. Mod. Phys. E **22**, 1330019 (2013), arXiv:1303.6912 [hep-ph].

[272] A. Boyarsky, M. Drewes, T. Lasserre, S. Mertens, and O. Ruchayskiy, “Sterile neutrino dark matter,” Prog. Part. Nucl. Phys. **104**, 1 (2019).

[273] T. Asaka, S. Blanchet, and M. Shaposhnikov, “The νMSM, dark matter and neutrino masses,” Phys. Lett. B **631**, 151 (2005), arXiv:hep-ph/0503065.

[274] T. Asaka and M. Shaposhnikov, “The νMSM, dark matter and baryon asymmetry of the universe,” Phys. Lett. B **620**, 17 (2005), arXiv:hep-ph/0505013.

[275] A. D. Sakharov, “Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe,” Sov. Phys. Usp. **34**, 392 (1991), [Pisma Zh. Eksp. Teor. Fiz. **5**, 32 (1967)].

[276] A. I. Vainshtein, V. I. Zakharov, V. A. Novikov, and M. A. Shifman, “ABC’s of instantons,” Sov. Phys. Usp. **25**, 195 (1982).

[277] V. A. Kuzmin, V. A. Rubakov, and M. E. Shaposhnikov, “On anomalous electroweak baryon-number non-conservation in the early universe,” Phys. Lett. B **155**, 36 (1985).

[278] N. S. Manton, “Topology in the Weinberg-Salam theory,” Phys. Rev. D **28**, 2019 (1983).

[279] F. R. Klinkhamer and N. S. Manton, “A saddle point solution in the Weinberg-Salam theory,” Phys. Rev. D **30**, 2212 (1984).

[280] V. A. Kuzmin, V. A. Rubakov, and M. E. Shaposhnikov, “Anomalous electroweak baryon number nonconservation and GUT mechanism for baryogenesis,” Phys. Lett. B **191**, 171 (1987).

[281] S. Yu. Khlebnikov and M. E. Shaposhnikov, “The statistical theory of anomalous fermion number nonconservation,” Nucl. Phys. B **308**, 885 (1988).

[282] M. E. Shaposhnikov, "Sphalerons and baryogenesis," Class. Quant. Grav. **10**, S147 (1993).

[283] A. Riotto and M. Trodden, "Recent progress in baryogenesis," Annu. Rev. Nucl. Part. Sci. **49**, 35 (1999), arXiv:hep-ph/9901362.

[284] D. Bodeker and W. Buchmuller, "Baryogenesis from the weak scale to the grand unification scale," Rev. Mod. Phys. **93**, 035004 (2021), arXiv:2009.07294 [hep-ph].

[285] K. Kajantie, M. Laine, K. Rummukainen, and M. Shaposhnikov, "Is there a hot electroweak phase transition at $m_H \gtrsim m_W$?" Phys. Rev. Lett. **77**, 2887 (1996).

[286] K. Rummukainen, M. Tsypin, K. Kajantie, M. Laine, and M. E. Shaposhnikov, "The universality class of the electroweak theory," Nucl. Phys. B **532**, 283 (1998), arXiv:hep-lat/9805013.

[287] D. E. Morrissey and M. J. Ramsey-Musolf, "Electroweak baryogenesis," New J. Phys. **14**, 125003 (2012), arXiv:1206.2942 [hep-ph].

[288] M. B. Gavela, P. Hernandez, J. Orloff, and O. Pene, "Standard model CP violation and baryon asymmetry," Mod. Phys. Lett. A **9**, 795 (1994), arXiv:hep-ph/9312215.

[289] E. K. Akhmedov, V. A. Rubakov, and A. Yu. Smirnov, "Baryogenesis via neutrino oscillations," Phys. Rev. Lett. **81**, 1359 (1998), arXiv:hep-ph/9803255.

[290] M. Shaposhnikov, "The νMSM, leptonic asymmetries, and properties of singlet fermions," J. High Energy Phys. **08**, 008 (2008), arXiv:0804.4542 [hep-ph].

[291] A. Boyarsky, O. Ruchayskiy, and M. Shaposhnikov, "The role of sterile neutrinos in cosmology and astrophysics," Annu. Rev. Nucl. Part. Sci. **59**, 191 (2009), arXiv:0901.0011 [hep-ph].

[292] L. Canetti, M. Drewes, and M. Shaposhnikov, "Sterile neutrinos as the origin of dark and baryonic matter," Phys. Rev. Lett. **110**, 061801 (2013), arXiv:1204.3902 [hep-ph].

[293] L. Canetti, M. Drewes, T. Frossard, and M. Shaposhnikov, "Dark matter, baryogenesis and neutrino oscillations from right handed neutrinos," Phys. Rev. D **87**, 093006 (2013), arXiv:1208.4607 [hep-ph].

[294] J. Klarić, M. Shaposhnikov, and I. Timiryasov, "Uniting low-scale leptogenesis mechanisms," Phys. Rev. Lett. **127**, 111802 (2021), arXiv:2008.13771 [hep-ph].

[295] M. Drewes and B. Garbrecht, "Leptogenesis from a GeV seesaw without mass degeneracy," J. High Energy Phys. **03**, 096 (2013), arXiv:1206.5537 [hep-ph].

[296] T. Venumadhav, F.-Y. Cyr-Racine, K. N. Abazajian, and C. M. Hirata, "Sterile neutrino dark matter: Weak interactions in the strong coupling epoch," Phys. Rev. D **94**, 043515 (2016), arXiv:1507.06655 [astro-ph.CO].

[297] K. N. Abazajian, "Sterile neutrinos in cosmology," Phys. Rep. **711-712**, 1 (2017), arXiv:1705.01837 [hep-ph].

[298] E. W. Kolb and M. S. Turner, *The Early Universe*, Frontiers in Physics, Vol. 69 (CRC Press, Boca Raton, 1990).

[299] X.-D. Shi and G. M. Fuller, "A new dark matter candidate: Nonthermal sterile neutrinos," Phys. Rev. Lett. **82**, 2832 (1999), arXiv:astro-ph/9810076 [astro-ph].

[300] S. Dodelson and L. M. Widrow, "Sterile neutrinos as dark matter," Phys. Rev. Lett. **72**, 17 (1994), arXiv:hep-ph/9303287 [hep-ph].

[301] F. Bezrukov, "Neutrino minimal standard model predictions for neutrinoless double beta decay," Phys. Rev. D **72**, 071303 (2005).

[302] K. G. Leach and S. Friedrich (BeEST Collaboration), "The BeEST experiment: Searching for beyond Standard Model neutrinos using ^{7}Be decay in STJs," J. Low Temp. Phys. **209**, 796 (2022), arXiv:2112.02029 [nucl-ex].

[303] S. Mertens, T. Lasserre, S. Groh, G. Drexlin, F. Glueck, A. Huber, A. W. P. Poon, M. Steidl, N. Steinbrink, and C. Weinheimer, "Sensitivity of next-generation tritium beta-decay experiments for keV-scale sterile neutrinos," J. Cosmol. Astropart. Phys. **02**, 020 (2015), arXiv:1409.0920 [physics.ins-det].

[304] J. N. Abdurashitov *et al.*, "First measurements in search for keV sterile neutrino in tritium beta-decay by Troitsk nu-mass experiment," Pisma Zh. Eksp. Teor. Fiz. [JETP Lett.] **105**, 723 (2017), arXiv:1703.10779 [hep-ex].

[305] J. A. L. Canning, F. F. Deppisch, and W. Pei, "Sensitivity of future tritium decay experiments to new physics," J. High Energy Phys. **03**, 144 (2023), arXiv:2212.06106 [hep-ph].

[306] M. Aker *et al.* (KATRIN Collaboration), "Search for keV-scale sterile neutrinos with first KATRIN data," (2022), arXiv:2207.06337 [nucl-ex].

[307] K. Bondarenko, A. Boyarsky, D. Gorbunov, and O. Ruchayskiy, "Phenomenology of GeV-scale heavy neutral leptons," J. High Energy Phys. **11**, 032 (2018), arXiv:1805.08567 [hep-ph].

[308] S. Ito *et al.*, "Search for massive neutrinos in $\pi^+ \to e^+\nu_e$ decay," Hyperfine Interact. **238**, 1 (2017).

[309] A. Aguilar-Arevalo *et al.* (PIENU Collaboration), “Search for heavy neutrinos in $\pi \to \mu\nu$ decay,” Phys. Lett. B **798**, 134980 (2019), arXiv:1904.03269 [hep-ex].

[310] E. Cortina Gil *et al.* (NA62 Collaboration), “Search for heavy neutral lepton production in K^+ decays to positrons,” Phys. Lett. B **807**, 135599 (2020), arXiv:2005.09575 [hep-ex].

[311] C. J. Parkinson *et al.* (NA62 Collaboration), “Search for heavy neutral lepton production at the NA62 experiment,” PoS Proc. Sci. **EPS-HEP2021**, 686 (2022).

[312] P. Abreu *et al.* (DELPHI Collaboration), “Search for neutral heavy leptons produced in Z decays,” Z. Phys. C **74**, 57 (1997), [Erratum: Z. Phys. C **75**, 580 (1997)].

[313] D. Liventsev *et al.* (Belle Collaboration), “Search for heavy neutrinos at Belle,” Phys. Rev. D **87**, 071102 (2013), [Erratum: PRD **95**, 099903 (2017)], arXiv:1301.1105 [hep-ex].

[314] D. Liventsev *et al.* (Belle Collaboration), “Search for a heavy neutrino in tau decays at Belle,” (2022), arXiv:2212.10095 [hep-ex].

[315] J. P. Lees *et al.* (BaBar Collaboration), “Search for heavy neutral leptons using tau lepton decays at BaBaR,” Phys. Rev. D **107**, 052009 (2023), arXiv:2207.09575 [hep-ex].

[316] C. S. Kim, Y. Kwon, D. Lee, S. Oh, and D. Sahoo, “Probing sterile neutrinos in B (D) meson decays at Belle II (BESIII),” Eur. Phys. J. C **80**, 730 (2020), arXiv:1908.00376 [hep-ph].

[317] P. Abratenko *et al.* (MicroBooNE Collaboration), “Search for heavy neutral leptons decaying into muon-pion pairs in the MicroBooNE detector,” Phys. Rev. D **101**, 052001 (2020), arXiv:1911.10545 [hep-ex].

[318] K. Abe *et al.* (T2K), “Search for heavy neutrinos with the T2K near detector ND280,” Phys. Rev. D **100**, 052006 (2019), arXiv:1902.07598 [hep-ex].

[319] F. Bergsma *et al.* (CHARM Collaboration), “A search for decays of heavy neutrinos,” Phys. Lett. B **128**, 361 (1983).

[320] J. Orloff, A. N. Rozanov, and C. Santoni, “Limits on the mixing of tau neutrino to heavy neutrinos,” Phys. Lett. B **550**, 8 (2002), arXiv:hep-ph/0208075.

[321] I. Boiarska, A. Boyarsky, O. Mikulenko, and M. Ovchynnikov, “Constraints from the CHARM experiment on heavy neutral leptons with tau mixing,” Phys. Rev. D **104**, 095019 (2021), arXiv:2107.14685 [hep-ph].

[322] G. Aad *et al.* (ATLAS Collaboration), "Search for heavy Majorana neutrinos with the ATLAS detector in pp collisions at $\sqrt{s} = 8$ TeV," J. High Energy Phys. **07**, 162 (2015), arXiv:1506.06020 [hep-ex].

[323] G. Aad *et al.* (ATLAS Collaboration), "Search for heavy neutral leptons in decays of W bosons produced in 13 TeV pp collisions using prompt and displaced signatures with the ATLAS detector," J. High Energy Phys. **10**, 265 (2019), arXiv:1905.09787 [hep-ex].

[324] A. M. Sirunyan *et al.* (CMS Collaboration), "Search for heavy neutral leptons in events with three charged leptons in proton-proton collisions at $\sqrt{s} =$ 13 TeV," Phys. Rev. Lett. **120**, 221801 (2018), arXiv:1802.02965 [hep-ex].

[325] A. Tumasyan *et al.* (CMS Collaboration), "Search for long-lived heavy neutral leptons with displaced vertices in proton-proton collisions at $\sqrt{s} = 13$ TeV," J. High Energy Phys. **07**, 081 (2022), arXiv:2201.05578 [hep-ex].

[326] R. Aaij *et al.* (LHCb Collaboration), "Search for heavy neutral leptons in $W^+ \to \mu^+\mu^\pm$jet decays," Eur. Phys. J. C **81**, 248 (2021), arXiv:2011.05263 [hep-ex].

[327] P. Giffin, S. Gori, Y.-D. Tsai, and D. Tuckler, "Heavy neutral leptons at beam dump experiments of future lepton colliders," J. High Energy Phys. **04**, 046 (2023), arXiv:2206.13745 [hep-ph].

[328] K. Abazajian, "Production and evolution of perturbations of sterile neutrino dark matter," Phys. Rev. D **73**, 063506 (2006), arXiv:astro-ph/0511630.

[329] M. Laine and M. Shaposhnikov, "Sterile neutrino dark matter as a consequence of νMSM-induced lepton asymmetry," J. Cosmol. Astropart. Phys. **806**, 031 (2008), arXiv:0804.4543 [hep-ph].

[330] P. D. Serpico and G. G. Raffelt, "Lepton asymmetry and primordial nucleosynthesis in the era of precision cosmology," Phys. Rev. D **71**, 127301 (2005), arXiv:astro-ph/0506162 [astro-ph].

[331] A. Matsumoto *et al.*, "EMPRESS. VIII. A new determination of primordial He abundance with extremely metal-poor galaxies: A suggestion of the lepton asymmetry and implications for the Hubble tension," Astrophys. J. **941**, 167 (2022), arXiv:2203.09617 [astro-ph.CO].

[332] M. Escudero, A. Ibarra, and V. Maura, "Primordial lepton asymmetries in the precision cosmology era: Current status and future sensitivities from BBN and the CMB," Phys. Rev. D **107**, 035024 (2023), arXiv:2208.03201 [hep-ph].

[333] J. F. Cherry and S. Horiuchi, "Closing in on resonantly produced sterile neutrino dark matter," Phys. Rev. D **95**, 083015 (2017), arXiv:1701.07874 [hep-ph].

[334] D. Blas, J. Lesgourgues, and T. Tram, "The Cosmic Linear Anisotropy Solving System (CLASS). Part II: Approximation schemes," J. Cosmol. Astropart. Phys. **2011**, 034 (2011), arXiv:1104.2933 [astro-ph.CO].

[335] C. Howlett, A. Lewis, A. Hall, and A. Challinor, "CMB power spectrum parameter degeneracies in the era of precision cosmology," J. Cosmol. Astropart. Phys. **2012**, 027 (2012), arXiv:1201.3654 [astro-ph.CO].

[336] E. O. Nadler *et al.* (DES Collaboration), "Milky Way satellite census. III. Constraints on dark matter properties from observations of Milky Way satellite galaxies," Phys. Rev. Lett. **126**, 091101 (2021), arXiv:2008.00022 [astro-ph.CO].

[337] J. S. Bullock and M. Boylan-Kolchin, "Small-scale challenges to the ΛCDM paradigm," Annu. Rev. Astron. Astrophys. **55**, 343 (2017), arXiv:1707.04256 [astro-ph.CO].

[338] S. Y. Kim, A. H. G. Peter, and J. R. Hargis, "Missing satellites problem: Completeness corrections to the number of satellite galaxies in the Milky Way are consistent with cold dark matter predictions," Phys. Rev. Lett. **121**, 211302 (2018), arXiv:1711.06267 [astro-ph.CO].

[339] S. Y. Kim and A. H. G. Peter, "The Milky Way satellite velocity function is a sharp probe of small-scale structure problems," (2021), arXiv:2106.09050 [astro-ph.GA].

[340] A. Dekker, S. Ando, C. A. Correa, and K. C. Y. Ng, "Warm dark matter constraints using Milky Way satellite observations and subhalo evolution modeling," Phys. Rev. D **106**, 123026 (2022), arXiv:2111.13137 [astro-ph.CO].

[341] S. Tremaine and J. E. Gunn, "Dynamical role of light neutral leptons in cosmology," Phys. Rev. Lett. **42**, 407 (1979).

[342] A. Boyarsky, O. Ruchayskiy, and D. Iakubovskyi, "A lower bound on the mass of dark matter particles," J. Cosmol. Astropart. Phys. **03**, 005 (2009), arXiv:0808.3902 [hep-ph].

[343] G. W. Angus, "A lower limit on the dark particle mass from dSphs," J. Cosmol. Astropart. Phys. **2010**, 026 (2010), arXiv:0907.1526 [astro-ph.CO].

[344] M. Rauch, "The Lyman alpha forest in the spectra of quasistellar objects," Annu. Rev. Astron. Astrophys. **36**, 267 (1998), arXiv:astro-ph/9806286.

[345] M. Ouchi, Y. Ono, and T. Shibuya, "Observations of the Lyman-α Universe," Annu. Rev. Astron. Astrophys. **58**, 617 (2020), arXiv:2012.07960 [astro-ph.GA].

[346] I. A. Zelko, T. Treu, K. N. Abazajian, D. Gilman, A. J. Benson, S. Birrer, A. M. Nierenberg, and A. Kusenko, "Constraints on sterile neutrino models from strong gravitational lensing, Milky Way satellites, and the Lyman-α forest," Phys. Rev. Lett. **129**, 191301 (2022).

[347] S. Vegetti, L. V. E. Koopmans, M. W. Auger, T. Treu, and A. S. Bolton, "Inference of the cold dark matter substructure mass function at $z = 0.2$ using strong gravitational lenses," Mon. Not. R. Astron. Soc. **442**, 2017 (2014), arXiv:1405.3666 [astro-ph.GA].

[348] Y. Hezaveh, N. Dalal, G. Holder, T. Kisner, M. Kuhlen, and L. Perreault Levasseur, "Measuring the power spectrum of dark matter substructure using strong gravitational lensing," J. Cosmol. Astropart. Phys. **11**, 048 (2016), arXiv:1403.2720 [astro-ph.CO].

[349] J.-W. Hsueh, W. Enzi, S. Vegetti, M. Auger, C. D. Fassnacht, G. Despali, L. V. E. Koopmans, and J. P. McKean, "SHARP – VII. New constraints on the dark matter free-streaming properties and substructure abundance from gravitationally lensed quasars," Mon. Not. R. Astron. Soc. **492**, 3047 (2020), arXiv:1905.04182 [astro-ph.CO].

[350] A. Schneider, "Astrophysical constraints on resonantly produced sterile neutrino dark matter," J. Cosmol. Astropart. Phys. **1604**, 059 (2016), arXiv:1601.07553 [astro-ph.CO].

[351] R. Shrock, "Decay $L^0 \to \nu_l \gamma$ in gauge theories of weak and electromagnetic interactions," Phys. Rev. D **9**, 743 (1974).

[352] P. B. Pal and L. Wolfenstein, "Radiative decays of massive neutrinos," Phys. Rev. D **25**, 766 (1982).

[353] J. Ellis, "TikZ-Feynman: Feynman diagrams with TikZ," Comput. Phys. Commun. **210**, 103 (2017), arXiv:1601.05437 [hep-ph].

[354] K. Arnaud, R. Smith, and A. Siemiginowska, eds., *Handbook of X-ray Astronomy*, Cambridge Observing Handbook for Research Astronomers, Vol. 7 (Cambridge Univ. Press, Cambridge, 2011).

[355] J. Neyman and E. S. Pearson, "IX. On the problem of the most efficient tests of statistical hypotheses," Philos. Trans. R. Soc. London, Ser. A **231**, 289 (1933).

[356] S. S. Wilks, "The large-sample distribution of the likelihood ratio for testing composite hypotheses," Annu. Math. Stat. **9**, 60 (1938).

[357] A. Wald, "Tests of statistical hypotheses concerning several parameters when the number of observations is large," Trans. Am. Math. Soc. **54**, 426 (1943).

[358] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, "Asymptotic formulae for likelihood-based tests of new physics," Eur. Phys. J. C **71**, 1554 (2011), [Erratum: EPJ C **73**, 2501 (2013)], arXiv:1007.1727 [physics.data-an].

[359] H. Chernoff, "On the distribution of the likelihood ratio," Ann. Math. Stat. **25**, 573 (1954).

[360] S. Algeri, J. Aalbers, K. D. Morå, and J. Conrad, "Searching for new phenomena with profile likelihood ratio tests," Nat. Rev. Phys. **2**, 245 (2020).

[361] A. Boyarsky, J. W. den Herder, A. Neronov, and O. Ruchayskiy, "Search for the light dark matter with an x-ray spectrometer," Astropart. Phys. **28**, 303 (2007), arXiv:astro-ph/0612219.

[362] H. V. D. Bradt, T. Ohashi, and K. A. Pounds, "X-ray astronomy missions." Annu. Rev. Astron. Astrophys. **30**, 391 (1992).

[363] E. Kalemci, "Summary of the past, present and future of the x-ray astronomy," Eur. Phys. J. Plus **133**, 407 (2018).

[364] C. Bambi, ed., *Tutorial Guide to X-ray and Gamma-ray Astronomy: Data Reduction and Analysis* (Springer Nature, Singapore, 2020).

[365] G. Vedrenne *et al.*, "SPI: The spectrometer aboard INTEGRAL," Astron. Astrophys. **411**, L63 (2003).

[366] A. Boyarsky, D. Malyshev, A. Neronov, and O. Ruchayskiy, "Constraining dark matter properties with SPI," Mon. Not. R. Astron. Soc. **387**, 1345 (2008), arXiv:0710.4922 [astro-ph].

[367] F. Calore, A. Dekker, P. D. Serpico, and T. Siegert, "Constraints on light decaying dark matter candidates from 16 years of INTEGRAL/SPI observations," Mon. Not. R. Astron. Soc. **520**, 4167 (2023), arXiv:2209.06299 [hep-ph].

[368] P. Ubertini *et al.*, "IBIS: The imager on-board INTEGRAL," Astron. Astrophys. **411**, L131 (2003).

[369] N. Lund *et al.*, "JEM-X: The x-ray monitor aboard INTEGRAL," Astron. Astrophys. **411**, L231 (2003).

[370] W. B. Atwood *et al.* (Fermi-LAT Collaboration), "The Large Area Telescope on the Fermi Gamma-ray Space Telescope mission," Astrophys. J. **697**, 1071 (2009), arXiv:0902.1089 [astro-ph.IM].

[371] M. Di Mauro, X. Hou, C. Eckner, G. Zaharijas, and E. Charles, "Search for γ-ray emission from dark matter particle interactions from Andromeda and Triangulum galaxies with the Fermi Large Area Telescope," Phys. Rev. D **99**, 123027 (2019), arXiv:1904.10977 [astro-ph.HE].

[372] S. Murgia, "The Fermi–LAT Galactic Center excess: Evidence of annihilating dark matter?" Annu. Rev. Nucl. Part. Sci. **70**, 455 (2020).

[373] M. N. Mazziotta, F. Loparco, D. Serini, A. Cuoco, P. De La Torre Luque, F. Gargano, and M. Gustafsson, "Search for dark matter signatures in the gamma-ray emission towards the Sun with the Fermi Large Area Telescope," Phys. Rev. D **102**, 022003 (2020), arXiv:2006.04114 [astro-ph.HE].

[374] M. Di Mauro, M. Stref, and F. Calore, "Investigating the effect of Milky Way dwarf spheroidal galaxies extension on dark matter searches with Fermi-LAT data," Phys. Rev. D **106**, 123032 (2022), arXiv:2212.06850 [astro-ph.HE].

[375] C. Meegan *et al.*, "The Fermi Gamma-ray Burst Monitor," Astrophys. J. **702**, 791 (2009), arXiv:0908.0450 [astro-ph.IM].

[376] K. C. Y. Ng, S. Horiuchi, J. M. Gaskins, M. Smith, and R. Preece, "Improved limits on sterile neutrino dark matter using full-sky Fermi Gamma-ray Burst Monitor data," Phys. Rev. D **92**, 043503 (2015), arXiv:1504.04027 [astro-ph.CO].

[377] P. Kaaret *et al.*, "HaloSat: A CubeSat to Study the hot Galactic halo," Astrophys. J. **884**, 162 (2019), arXiv:1909.13822 [astro-ph.IM].

[378] E. M. Silich, K. Jahoda, L. Angelini, P. Kaaret, A. Zajczyk, D. M. LaRocca, R. Ringuette, and J. Richardson, "A search for the 3.5 keV line from the Milky Way's dark matter halo with HaloSat," Astrophys. J. **916**, 2 (2021), arXiv:2105.12252 [astro-ph.HE].

[379] J. S. Adams *et al.*, "Micro-X sounding rocket: Transitioning from first flight to a dark matter configuration," J. Low Temp. Phys. **199**, 1072 (2020), arXiv:1908.09010 [astro-ph.IM].

[380] A. J. F. Hubbard *et al.*, "Design and status of the Micro-X microcalorimeter sounding rocket," J. Phys. Conf. Ser. **1342**, 012096 (2020).

[381] J. S. Adams *et al.*, "Micro-X sounding rocket payload re-flight progress," J. Low Temp. Phys. **209**, 832 (2022), arXiv:2111.06952 [astro-ph.IM].

[382] J. W. den Herder *et al.*, "The Reflection Grating Spectrometer on board XMM-Newton," Astron. Astrophys. **365**, L7 (2001).

[383] M. J. L. Turner *et al.*, "The European Photon Imaging Camera on XMM-Newton: The MOS cameras," Astron. Astrophys. **365**, L27 (2001), arXiv:astro-ph/0011498 [astro-ph].

[384] L. Strüder *et al.*, "The European Photon Imaging Camera on XMM-Newton: The pn-CCD camera," Astron. Astrophys. **365**, L18 (2001).

[385] A. Boyarsky, D. Iakubovskyi, O. Ruchayskiy, and V. Savchenko, "Constraints on decaying dark matter from XMM-Newton observations of M31," Mon. Not. R. Astron. Soc. **387**, 1361 (2008), arXiv:0709.2301 [astro-ph].

[386] E. Bulbul, M. Markevitch, A. Foster, R. K. Smith, M. Loewenstein, and S. W. Randall, "Detection of an unidentified emission line in the stacked x-ray spectrum of galaxy clusters," Astrophys. J. **789**, 13 (2014), arXiv:1402.2301 [astro-ph.CO].

[387] A. Boyarsky, O. Ruchayskiy, D. Iakubovskyi, and J. Franse, "Unidentified line in x-ray spectra of the Andromeda galaxy and Perseus galaxy cluster," Phys. Rev. Lett. **113**, 251301 (2014), arXiv:1402.4119 [astro-ph.CO].

[388] C. Dessert, N. L. Rodd, and B. R. Safdi, "The dark matter interpretation of the 3.5-keV line is inconsistent with blank-sky observations," Science **367**, 1465 (2020), arXiv:1812.06976 [astro-ph.CO].

[389] S. Bhargava *et al.*, "The XMM Cluster Survey: New evidence for the 3.5-keV feature in clusters is inconsistent with a dark matter origin," Mon. Not. R. Astron. Soc. **497**, 656 (2020), arXiv:2006.13955 [astro-ph.CO].

[390] J. W. Foster, M. Kongsore, C. Dessert, Y. Park, N. L. Rodd, K. Cranmer, and B. R. Safdi, "Deep search for decaying dark matter with XMM-Newton blank-sky observations," Phys. Rev. Lett. **127**, 051101 (2021), arXiv:2102.02207 [astro-ph.CO].

[391] S. S. Murray *et al.*, "In-flight performance of the Chandra High-Resolution Camera," in *X-Ray Optics, Instruments, and Missions III*, SPIE Conf. Ser., Vol. 4012, edited by J. E. Truemper and B. Aschenbach (2000) p. 68.

[392] A. C. Brinkman *et al.*, "Description and performance of the Low-Energy Transmission Grating Spectrometer on board Chandra," in *X-Ray Optics, Instruments, and Missions III*, SPIE Conf. Ser., Vol. 4012, edited by J. E. Truemper and B. Aschenbach (2000) p. 81.

[393] C. Canizares *et al.*, "The Chandra High-Energy Transmission Grating: Design, fabrication, ground calibration, and 5 years in flight," Publ. Astron. Soc. Pac. **117**, 1144 (2005).

[394] 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 *X-Ray and Gamma-Ray Telescopes and Instruments for Astronomy.*, SPIE Conf. Ser., Vol. 4851, edited by J. E. Truemper and H. D. Tananbaum (2003) p. 28.

[395] S. Horiuchi, P. J. Humphrey, J. Onorbe, K. N. Abazajian, M. Kaplinghat, and S. Garrison-Kimmel, "Sterile neutrino dark matter bounds from galaxies of the Local Group," Phys. Rev. D **89**, 025017 (2014), arXiv:1311.0282 [astro-ph.CO].

[396] N. Cappelluti, E. Bulbul, A. Foster, P. Natarajan, M. C. Urry, M. W. Bautz, F. Civano, E. Miller, and R. K. Smith, “Searching for the 3.5 keV line in the Deep Fields with Chandra: The 10 Ms observations,” Astrophys. J. **854**, 179 (2018), arXiv:1701.07932 [astro-ph.CO].

[397] D. Sicilian, N. Cappelluti, E. Bulbul, F. Civano, M. Moscetti, and C. S. Reynolds, “Probing the Milky Way’s dark matter halo for the 3.5 keV line,” Astrophys. J. **905**, 146 (2020), arXiv:2008.02283 [astro-ph.HE].

[398] R. L. Kelley *et al.*, “The Suzaku high resolution X-Ray Spectrometer,” Publ. Astron. Soc. Jpn. **59**, 77 (2007).

[399] K. Koyama *et al.*, “X-ray Imaging Spectrometer (XIS) on board Suzaku,” Publ. Astron. Soc. Jpn. **59**, S23 (2007).

[400] T. Takahashi *et al.*, “Hard X-Ray Detector (HXD) on board Suzaku,” Publ. Astron. Soc. Jpn. **59**, S35 (2007).

[401] M. Loewenstein, A. Kusenko, and P. L. Biermann, “New limits on sterile neutrinos from Suzaku observations of the Ursa Minor dwarf spheroidal galaxy,” Astrophys. J. **700**, 426 (2009), arXiv:0812.2710 [astro-ph].

[402] T. Tamura, R. Iizuka, Y. Maeda, K. Mitsuda, and N. Y. Yamasaki, “An x-ray spectroscopic search for dark matter in the Perseus cluster with Suzaku,” Publ. Astron. Soc. Jpn. **67**, 23 (2015), arXiv:1412.1869 [astro-ph.HE].

[403] O. Urban, N. Werner, S. W. Allen, A. Simionescu, J. S. Kaastra, and L. E. Strigari, “A Suzaku search for dark matter emission lines in the x-ray brightest galaxy glusters,” Mon. Not. R. Astron. Soc. **451**, 2447 (2015), arXiv:1411.0050 [astro-ph.CO].

[404] N. Sekiya, N. Y. Yamasaki, and K. Mitsuda, “A search for a keV signature of radiatively decaying dark matter with Suzaku XIS observations of the x-ray diffuse background,” Publ. Astron. Soc. Jpn. **68**, S31 (2016), arXiv:1504.02826 [astro-ph.HE].

[405] E. Bulbul, M. Markevitch, A. Foster, E. Miller, M. Bautz, M. Loewenstein, S. W. Randall, and R. K. Smith, “Searching for the 3.5 keV line in the stacked Suzaku observations of galaxy clusters,” Astrophys. J. **831**, 55 (2016), arXiv:1605.02034 [astro-ph.HE].

[406] N. Gehrels *et al.*, “The Swift gamma-ray burst mission,” Astrophys. J. **611**, 1005 (2004), arXiv:astro-ph/0405233 [astro-ph].

[407] D. N. Burrows *et al.*, “The Swift X-Ray Telescope,” Space Sci. Rev. **120**, 165 (2005), arXiv:astro-ph/0508071 [astro-ph].

[408] S. D. Barthelmy *et al.*, “The Burst Alert Telescope (BAT) on the Swift MIDEX mission,” Space Sci. Rev. **120**, 143 (2005), arXiv:astro-ph/0507410.

[409] P. Predehl *et al.* (eROSITA Collaboration), "The eROSITA x-ray telescope on SRG," Astron. Astrophys. **647**, A1 (2021), arXiv:2010.03477 [astro-ph.HE].

[410] M. Pavlinsky *et al.*, "The ART-XC telescope on board the SRG observatory," Astron. Astrophys. **650**, A42 (2021), arXiv:2103.12479 [astro-ph.HE].

[411] V. V. Barinov, R. A. Burenin, D. S. Gorbunov, and R. A. Krivonos, "Towards testing sterile neutrino dark matter with the Spectrum-Roentgen-Gamma mission," Phys. Rev. D **103**, 063512 (2021), arXiv:2007.07969 [astro-ph.CO].

[412] A. Dekker, E. Peerbooms, F. Zimmer, K. C. Y. Ng, and S. Ando, "Searches for sterile neutrinos and axionlike particles from the Galactic halo with eROSITA," Phys. Rev. D **104**, 023021 (2021), arXiv:2103.13241 [astro-ph.HE].

[413] E. I. Zakharov *et al.*, "All-sky limits on sterile neutrino Galactic dark matter obtained with SRG/ART-XC after two years of operations," (2023), arXiv:2303.12673 [astro-ph.HE].

[414] C. A. Kilbourne *et al.*, "Design, implementation, and performance of the Astro-H SXS calorimeter array and anticoincidence detector," J. Astron. Telesc. Instrum. Syst. **4**, 011214 (2018).

[415] T. Tanaka *et al.*, "Soft x-ray imager aboard Hitomi (ASTRO-H)," J. Astron. Telesc. Instrum. Syst. **4**, 011211 (2018), arXiv:1801.06932 [astro-ph.IM].

[416] K. Nakazawa *et al.*, "Hard X-ray Imager onboard Hitomi (ASTRO-H)," J. Astron. Telesc. Instrum. Syst. **4**, 021410 (2018).

[417] F. A. Aharonian *et al.* (Hitomi Collaboration), "Hitomi constraints on the 3.5 keV line in the Perseus galaxy gluster," Astrophys. J. Lett. **837**, L15 (2017), arXiv:1607.07420 [astro-ph.HE].

[418] T. Tamura *et al.*, "An x-ray spectroscopic search for dark matter and unidentified line signatures in the Perseus cluster with Hitomi," Publ. Astron. Soc. Jpn. **71** (2019), arXiv:1811.05767 [astro-ph.HE].

[419] M. Tashiro *et al.*, "Concept of the X-ray Astronomy Recovery Mission," in *Space Telescopes and Instrumentation 2018: Ultraviolet to Gamma Ray*, SPIE Conf. Ser., Vol. 10699, edited by J.-W. A. den Herder, S. Nikzad, and K. Nakazawa (2018) p. 1069922.

[420] L. Angelini *et al.* (XRISM Science Team), "Science with the X-Ray Imaging and Spectroscopy Mission (XRISM)," (2020), arXiv:2003.04962 [astro-ph.HE].

[421] Y. Ishisaki *et al.*, "Status of Resolve instrument onboard X-Ray Imaging and Spectroscopy Mission (XRISM)," in *Space Telescopes and Instrumentation 2022: Ultraviolet to Gamma Ray*, SPIE Conf. Proc., Vol. 12181, edited by J.-W. A. den Herder, S. Nikzad, and K. Nakazawa (2022) p. 121811S.

[422] H. Nakajima *et al.*, "Soft X-ray Imager (SXI) for Xtend onboard X-Ray Imaging and Spectroscopy Mission (XRISM)," in *Space Telescopes and Instrumentation 2020: Ultraviolet to Gamma Ray*, SPIE Conf. Proc., Vol. 11444, edited by J.-W. A. den Herder, S. Nikzad, and K. Nakazawa (2020) p. 1144423.

[423] A. Boyarsky, J. Franse, D. Iakubovskyi, and O. Ruchayskiy, "Checking the dark matter origin of a 3.53 keV line with the Milky Way center," Phys. Rev. Lett. **115**, 161301 (2015), arXiv:1408.2503 [astro-ph.CO].

[424] A. Neronov and D. Malyshev, "Toward a full test of the νMSM sterile neutrino dark matter model with Athena," Phys. Rev. D **93**, 063518 (2016), arXiv:1509.02758 [astro-ph.HE].

[425] K. Perez, K. C. Y. Ng, J. F. Beacom, C. Hersh, S. Horiuchi, and R. Krivonos, "Almost closing the νMSM sterile neutrino dark matter window with NuSTAR," Phys. Rev. D **95**, 123002 (2017), arXiv:1609.00667 [astro-ph.HE].

[426] D. R. Wik *et al.*, "NuSTAR observations of the Bullet cluster: Constraints on inverse Compton emission," Astrophys. J. **792**, 48 (2014), arXiv:1403.2722 [astro-ph.HE].

[427] M. Drewes *et al.*, "A white paper on keV sterile neutrino dark matter," J. Cosmol. Astropart. Phys. **1701**, 025 (2017), arXiv:1602.04816 [hep-ph].

[428] T. Aramaki *et al.*, "Snowmass2021 Cosmic Frontier: The landscape of cosmic-ray and high-energy photon probes of particle dark matter," *2022 Snowmass Summer Study*, (2022), arXiv:2203.06894 [hep-ex].

[429] A. Boyarsky, D. Malyshev, O. Ruchayskiy, and D. Savchenko, "Technical comment on the paper of Dessert *et al*. 'The dark matter interpretation of the 3.5 keV line is inconsistent with blank-sky observations'," (2020), arXiv:2004.06601 [astro-ph.CO].

[430] K. N. Abazajian, "Technical comment on 'The dark matter interpretation of the 3.5-keV line is inconsistent with blank-sky observations'," (2020), arXiv:2004.06170 [astro-ph.HE].

[431] C. Dessert, N. L. Rodd, and B. R. Safdi, "Response to a comment on Dessert *et al*. 'The dark matter interpretation of the 3.5 keV line is inconsistent with blank-bky observations'," Phys. Dark Univ. **30**, 100656 (2020), arXiv:2006.03974 [astro-ph.CO].

[432] F. A. Harrison *et al.* (NuSTAR Collaboration), "The Nuclear Spectroscopic Telescope Array (NuSTAR) high-energy x-ray mission," Astrophys. J. **770**, 103 (2013), arXiv:1301.7307 [astro-ph.IM].

[433] J. E. Koglin *et al.*, "NuSTAR hard x-ray optics design and performance," in *Optics for EUV, X-Ray, and Gamma-Ray Astronomy IV*, SPIE Conf. Proc., Vol. 7437, edited by S. L. O'Dell and G. Pareschi (2009) p. 74370C.

[434] T. Kitaguchi *et al.*, "Spectral calibration and modeling of the NuSTAR CdZnTe pixel detectors," in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, SPIE Conf. Proc., Vol. 8145, edited by O. H. Siegmund (2011) p. 814507, arXiv:1109.0378 [astro-ph.IM].

[435] T. Kitaguchi *et al.*, "Inflight performance and calibration of the NuSTAR CdZnTe pixel detectors," in *Space Telescopes and Instrumentation 2014: Ultraviolet to Gamma Ray*, SPIE Conf. Proc., Vol. 9144, edited by T. Takahashi, J.-W. A. den Herder, and M. Bautz (2014) p. 91441R.

[436] E. Schönfeld and H. Janßen, "Calculation of emission probabilities of x-rays and Auger electrons emitted in radioactive disintegration processes," Appl. Radiat. Isot. **52**, 595 (2000).

[437] N. Nica, "Nuclear data sheets for $A = 155$," Nucl. Data Sheets **160**, 1 (2019).

[438] B. Grefenstette, M. Brightman, H. P. Earnshaw, K. Forster, K. K. Madsen, and H. Miyasaka, "Measuring the evolution of the NuSTAR detector gains," (2022), arXiv:2206.04058 [astro-ph.IM].

[439] O. Ruchayskiy, A. Boyarsky, D. Iakubovskyi, E. Bulbul, D. Eckert, J. Franse, D. Malyshev, M. Markevitch, and A. Neronov, "Searching for decaying dark matter in deep XMM–Newton observation of the Draco dwarf spheroidal," Mon. Not. R. Astron. Soc. **460**, 1390 (2016), arXiv:1512.07217 [astro-ph.HE].

[440] S. Riemer-Sørensen *et al.*, "Dark matter line emission constraints from NuSTAR observations of the Bullet Cluster," Astrophys. J. **810**, 48 (2015), arXiv:1507.01378 [astro-ph.CO].

[441] A. Boyarsky, A. Neronov, O. Ruchayskiy, and M. Shaposhnikov, "Constraints on sterile neutrinos as dark matter candidates from the diffuse x-ray background," Mon. Not. R. Astron. Soc. **370**, 213 (2006), arXiv:astro-ph/0512509.

[442] A. Neronov, D. Malyshev, and D. Eckert, "Decaying dark matter search with NuSTAR deep sky observations," Phys. Rev. D **94**, 123504 (2016), arXiv:1607.07328 [astro-ph.HE].

[443] H. Kaneda, K. Makishima, S. Yamauchi, K. Koyama, K. Matsuzaki, and N. Y. Yamasaki, "Complex spectra of the Galactic ridge x-rays observed with ASCA," Astrophys. J. **491**, 638 (1997).

[444] M. Revnivtsev, S. Sazonov, E. Churazov, W. Forman, A. Vikhlinin, and R. Sunyaev, "Discrete sources as the origin of the Galactic x-ray ridge emission," Nature **458**, 1142 (2009), arXiv:0904.4649 [astro-ph.GA].

[445] T. Yuasa, K. Makishima, and K. Nakazawa, "Broadband spectral analysis of the Galactic ridge x-ray emission," Astrophys. J. **753**, 129 (2012), arXiv:1205.1574 [astro-ph.GA].

[446] K. Perez, R. Krivonos, and D. R. Wik, "The Galactic bulge diffuse emission in broadband x-rays with NuSTAR," Astrophys. J. **884**, 153 (2019), arXiv:1909.05916 [astro-ph.HE].

[447] H. Stiele, W. Pietsch, F. Haberl, D. Hatzidimitriou, R. Barnard, B. F. Williams, A. K. H. Kong, and U. Kolb, "The deep XMM-Newton survey of M31," Astron. Astrophys. **534**, A55 (2011), arXiv:1106.4755 [astro-ph.HE].

[448] "XMM-Newton science data images – M31 galaxy," `https://heasarc.gsfc.nasa.gov/docs/xmm/gallery/esas-gallery/xmm_gal_science_m31.html`, [Last accessed 2023-03-15].

[449] J. Wilms, A. Allen, and R. McCray, "On the Absorption of X-Rays in the Interstellar Medium," Astrophys. J. **542**, 914 (2000), arXiv:astro-ph/0008425 [astro-ph].

[450] A. Tamm, E. Tempel, P. Tenjes, O. Tihhonova, and T. Tuvikene, "Stellar mass map and dark matter distribution in M31," Astron. Astrophys. **546**, A4 (2012), arXiv:1208.5712.

[451] R. Launhardt, R. Zylka, and P. G. Mezger, "The nuclear bulge of the Galaxy. III. Large scale physical characteristics of stars and interstellar matter," Astron. Astrophys. **384**, 112 (2002), arXiv:astro-ph/0201294 [astro-ph].

[452] M. Revnivtsev, S. Molkov, and S. Sazonov, "Map of the Galaxy in the 6.7-keV emission line," Mon. Not. R. Astron. Soc. Lett. **373**, L11 (2006), arXiv:astro-ph/0605693 [astro-ph].

[453] R. A. Krivonos, S. S. Tsygankov, I. A. Mereminskiy, A. A. Lutovinov, S. Yu. Sazonov, and R. A. Sunyaev, "New hard x-ray sources discovered in the ongoing INTEGRAL Galactic plane survey after 14 yr of observations," Mon. Not. R. Astron. Soc. **470**, 512 (2017), arXiv:1704.03364 [astro-ph.HE].

[454] D. E. Gruber, J. L. Matteson, L. E. Peterson, and G. V. Jung, "The spectrum of diffuse cosmic hard x-rays measured with HEAO 1," Astrophys. J. **520**, 124 (1999), arXiv:astro-ph/9903492 [astro-ph].

[455] E. Churazov *et al.*, "INTEGRAL observations of the cosmic x-ray background in the 5–100 keV range via occultation by the Earth," Astron. Astrophys. **467**, 529 (2006), arXiv:astro-ph/0608250 [astro-ph].

[456] R. Krivonos, D. Wik, B. Grefenstette, K. Madsen, K. Perez, S. Rossland, S. Sazonov, and A. Zoglauer, "NuSTAR measurement of the cosmic x-ray background in the 3–20 keV energy band," Mon. Not. R. Astron. Soc. **502**, 3966 (2021), arXiv:2011.11469 [astro-ph.HE].

[457] J. W. Foster, N. L. Rodd, and B. R. Safdi, "Revealing the dark matter halo with axion direct detection," Phys. Rev. D **97**, 123006 (2018), arXiv:1711.10489 [astro-ph.CO].

[458] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, "Power-constrained limits," (2011), arXiv:1105.3166 [physics.data-an].

[459] M. Berg, J. P. Conlon, F. Day, N. Jennings, S. Krippendorf, A. J. Powell, and M. Rummel, "Constraints on axion-like particles from x-ray observations of NGC1275," Astrophys. J. **847**, 101 (2017), arXiv:1605.01043 [astro-ph.HE].

[460] M. C. D. Marsh, H. R. Russell, A. C. Fabian, B. P. McNamara, P. Nulsen, and C. S. Reynolds, "A new bound on axion-like particles," J. Cosmol. Astropart. Phys. **12**, 036 (2017), arXiv:1703.07354 [hep-ph].

[461] C. S. Reynolds, M. C. D. Marsh, H. R. Russell, A. C. Fabian, R. Smith, F. Tombesi, and S. Veilleux, "Astrophysical limits on very light axion-like particles from Chandra grating spectroscopy of NGC 1275," Astrophys. J. **890**, 59 (2020), arXiv:1907.05475 [hep-ph].

[462] C. Dessert, A. J. Long, and B. R. Safdi, "No evidence for axions from Chandra observation of the magnetic white dwarf RE J0317-853," Phys. Rev. Lett. **128**, 071102 (2022), arXiv:2104.12772 [hep-ph].

[463] J. S. Reynés, J. H. Matthews, C. S. Reynolds, H. R. Russell, R. N. Smith, and M. C. D. Marsh, "New constraints on light axion-like particles using Chandra transmission grating spectroscopy of the powerful cluster-hosted quasar H1821+643," Mon. Not. R. Astron. Soc. **510**, 1264 (2021), arXiv:2109.03261 [astro-ph.HE].

[464] M. Buschmann, R. T. Co, C. Dessert, and B. R. Safdi, "Axion emission can explain a new hard x-ray excess from nearby isolated neutron stars," Phys. Rev. Lett. **126**, 021102 (2021).

[465] C. Dessert, J. W. Foster, and B. R. Safdi, "X-ray searches for axions from super star clusters," Phys. Rev. Lett. **125**, 261102 (2020), arXiv:2008.03305 [hep-ph].

[466] M. Xiao, K. M. Perez, M. Giannotti, O. Straniero, A. Mirizzi, B. W. Grefenstette, B. M. Roach, and M. Nynka, "Constraints on axionlike particles from a hard x-ray observation of Betelgeuse," Phys. Rev. Lett. **126**, 031101 (2021), arXiv:2009.09059 [astro-ph.HE].

[467] W. DeRocco, S. Wegsman, B. Grefenstette, J. Huang, and K. Van Tilburg, "First indirect detection constraints on axions in the solar basin," Phys. Rev. Lett. **129**, 101101 (2022), arXiv:2205.05700 [hep-ph].

[468] C. O'Hare, "cajohare/AxionLimits: AxionLimits v1.0," Zenodo (2020), 10.5281/zenodo.3932430.

[469] M. Cirelli, N. Fornengo, B. J. Kavanagh, and E. Pinetti, "INTEGRAL x-ray constraints on sub-GeV dark matter," Phys. Rev. D **103**, 063022 (2021), arXiv:2007.11493 [hep-ph].

[470] M. Cirelli, N. Fornengo, J. Koechler, E. Pinetti, and B. M. Roach, "Putting all the X in one basket: Updated x-ray constraints on sub-GeV dark matter," (2023), arXiv:2303.08854 [hep-ph].

[471] D. Maurin, H. P. Dembinski, J. Gonzalez, I. C. Mariş, and F. Melot, "Cosmic-Ray Database update: Ultra-high energy, ultra-heavy, and antinuclei cosmic-ray data (`CRDB v4.0`)," Universe **6**, 102 (2020), arXiv:2005.14663 [astro-ph.HE].

[472] M. S. Longair, *High Energy Astrophysics*, 3rd ed. (Cambridge Univ. Press, Cambridge, 2011).

[473] M. Aguilar *et al.* (AMS Collaboration), "Properties of neon, magnesium, and silicon primary cosmic rays results from the Alpha Magnetic Spectrometer," Phys. Rev. Lett. **124**, 211102 (2020).

[474] M. Aguilar *et al.* (AMS Collaboration), "Properties of a new group of cosmic nuclei: Results from the Alpha Magnetic Spectrometer on sodium, aluminum, and nitrogen," Phys. Rev. Lett. **127**, 02101 (2021), [Erratum: PRL 127, 159901 (2021)].

[475] M. Aguilar *et al.* (AMS Collaboration), "The Alpha Magnetic Spectrometer (AMS) on the international space station: Part II — Results from the first seven years," Phys. Rep. **894**, 1 (2021).

[476] M. Aguilar *et al.* (AMS Collaboration), "Properties of iron primary cosmic rays: Results from the Alpha Magnetic Spectrometer," Phys. Rev. Lett. **126**, 041104 (2021).

[477] A. D. Panov *et al.*, "Energy spectra of abundant nuclei of primary cosmic rays from the data of ATIC-2 experiment: Final results," Bull. Russ. Acad. Sci. Phys. **73**, 564 (2009), arXiv:1101.3246 [astro-ph.HE].

[478] Y. Shikaze *et al.*, "Measurements of 0.2 to 20-GeV/n cosmic-ray proton and helium spectra from 1997 through 2002 with the BESS spectrometer," Astropart. Phys. **28**, 154 (2007), arXiv:astro-ph/0611388.

[479] K. Abe *et al.*, "Measurement of the cosmic-ray antiproton spectrum at solar minimum with a long-duration balloon flight over Antarctica," Phys. Rev. Lett. **108**, 051102 (2012), arXiv:1107.6000 [astro-ph.HE].

[480] J. J. Engelmann, P. Ferrando, A. Soutoul, P. Goret, and E. Juliusson, "Charge composition and energy spectra of cosmic-ray for elements from Be to Ni—Results from HEAO-3-C2," Astron. Astrophys. **233**, 96 (1990).

[481] K. Asakimori *et al.*, "Cosmic-ray proton and helium spectra: Results from the JACEE experiment," Astrophys. J. **502**, 278 (1998).

[482] V. Grebenyuk *et al.*, "Energy spectra of abundant cosmic-ray nuclei in the NUCLEON experiment," Adv. Space Res. **64**, 2546 (2019).

[483] V. Grebenyuk *et al.*, "Secondary cosmic rays in the NUCLEON space experiment," Adv. Space Res. **64**, 2559 (2019), arXiv:1809.09665 [astro-ph.HE].

[484] V. A. Derbina *et al.*, "Cosmic-ray spectra and composition in the energy range of 10–1000 TeV per particle obtained by the RUNJOB experiment," Astrophys. J. Lett. **628**, L41 (2005).

[485] F. Gahbauer, G. Hermann, J. R. Hörandel, D. Müller, and A. A. Radu, "A new Measurement of the intensities of the heavy primary cosmic-Ray nuclei around 1 TeV amu^{-1}," Astrophys. J. **607**, 333 (2004).

[486] M. Ave, P. J. Boyle, F. Gahbauer, C. Höppner, J. R. Hörandel, M. Ichimura, D. Müller, and A. Romero-Wolf, "Composition of primary cosmic-ray nuclei at high energies," Astrophys. J. **678**, 262 (2008), arXiv:0801.0582 [astro-ph].

[487] A. Obermeier, M. Ave, P. Boyle, C. Hoppner, J. Horandel, and D. Muller, "Energy spectra of primary and secondary cosmic-ray nuclei measured with TRACER," Astrophys. J. **742**, 14 (2011), arXiv:1108.4838 [astro-ph.HE].

[488] A. W. Strong, I. V. Moskalenko, and V. S. Ptuskin, "Cosmic-ray propagation and interactions in the Galaxy," Annu. Rev. Nucl. Part. Sci. **57**, 285 (2007), arXiv:astro-ph/0701517.

[489] 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 (2012), arXiv:1206.1593 [astro-ph.HE].

[490] I. A. Grenier, J. H. Black, and A. W. Strong, "The nine lives of cosmic rays in galaxies," Annu. Rev. Astron. Astrophys. **53**, 199 (2015).

[491] E. Fermi, "On the origin of the cosmic radiation," Phys. Rev. **75**, 1169 (1949).

[492] G. A. Tammann, W. Loeffler, and A. Schroeder, "The Galactic supernova rate," Astrophys. J. Suppl. Ser. **92**, 487 (1994).

[493] A. M. Hillas, "Topical Review: Can diffusive shock acceleration in supernova remnants account for high-energy Galactic cosmic rays?" J. Phys. G: Nucl. Part. Phys. **31**, R95 (2005).

[494] A. M. Bykov, D. C. Ellison, A. Marcowith, and S. M. Osipov, "Cosmic ray production in supernovae," Space Sci. Rev. **214**, 41 (2018), arXiv:1801.08890 [astro-ph.HE].

[495] P. Cristofari, "The hunt for pevatrons: The case of supernova remnants," Universe **7**, 324 (2021), arXiv:2110.07956 [astro-ph.HE].

[496] R. Schlickeiser, *Cosmic Ray Astrophysics* (Springer Berlin, Heidelberg, 2002).

[497] T. A. Porter, G. Johannesson, and I. V. Moskalenko, "The GALPROP cosmic-ray propagation and nonthermal emissions framework: Release v57," Astrophys. J. Suppl. Ser. **262**, 30 (2022), arXiv:2112.12745 [astro-ph.HE].

[498] C. Evoli, D. Gaggero, A. Vittino, G. Di Bernardo, M. Di Mauro, A. Ligorini, P. Ullio, and D. Grasso, "Cosmic-ray propagation with `DRAGON2`: I. Numerical solver and astrophysical ingredients," J. Cosmol. Astropart. Phys. **02**, 015 (2017), arXiv:1607.07886 [astro-ph.HE].

[499] E. M. Burbidge, G. R. Burbidge, W. A. Fowler, and F. Hoyle, "Synthesis of the elements in stars," Rev. Mod. Phys. **29**, 547 (1957).

[500] M. Aguilar *et al.* (AMS Collaboration), "Observation of new properties of secondary cosmic rays lithium, beryllium, and boron by the Alpha Magnetic Spectrometer on the International Space Station," Phys. Rev. Lett. **120**, 021101 (2018).

[501] M. S. Potgieter, "Solar modulation of cosmic rays," Living Rev. Sol. Phys. **10**, 3 (2013), arXiv:1306.4421 [physics.space-ph].

[502] L. J. Gleeson and W. I. Axford, "Solar modulation of Galactic cosmic rays," Astrophys. J. **154**, 1011 (1968).

[503] C. Corti, V. Bindi, C. Consolandi, C. Freeman, A. Kuhlman, C. Light, M. Palermo, and S. Wang, "Test of validity of the force-field approximation with AMS-02 and PAMELA monthly fluxes," PoS Proc. Sci. **ICRC2019**, 1070 (2020), arXiv:1910.00027 [astro-ph.HE].

[504] C. Corti, V. Bindi, C. Consolandi, and K. Whitman, "Solar modulation of the local interstellar spectrum with Voyager 1, AMS-02, PAMELA, and BESS," Astrophys. J. **829**, 8 (2016), arXiv:1511.08790 [astro-ph.HE].

[505] D. Bisschoff, M. S. Potgieter, and O. P. M. Aslam, "New very local interstellar spectra for electrons, positrons, protons and light cosmic ray nuclei," Astrophys. J. **878**, 59 (2019), arXiv:1902.10438 [astro-ph.HE].

[506] D. F. Smart and M. A. Shea, "A review of geomagnetic cutoff rigidities for earth-orbiting spacecraft," Adv. Space Res. **36**, 2012 (2005).

[507] O. Chamberlain, E. Segrè, C. Wiegand, and T. Ypsilantis, "Observation of antiprotons," Phys. Rev. **100**, 947 (1955).

[508] B. Cork, G. R. Lambertson, O. Piccioni, and W. A. Wenzel, "Antineutrons produced from antiprotons in charge-exchange collisions," Phys. Rev. **104**, 1193 (1956).

[509] D. E. Dorfan, J. Eades, L. M. Lederman, W. Lee, and C. C. Ting, "Observation of antideuterons," Phys. Rev. Lett. **14**, 1003 (1965).

[510] T. Massam, T. Muller, B. Righini, M. Schneegans, and A. Zichichi, "Experimental observation of antideuteron production," Nuovo Cimento **39**, 10 (1965).

[511] N. K. Vishnevsky *et al.*, "Observation of antitritium," Yad. Fiz. [Sov. J. Nucl. Phys.] **20**, 694 (1974).

[512] Yu. M. Antipov *et al.*, "Observation of antihelium-3," Nucl. Phys. B **31**, 235 (1971).

[513] H. Agakishiev *et al.* (STAR Collaboration), "Observation of the antimatter helium-4 nucleus," Nature **473**, 353 (2011), [Erratum: Nature **475**, 412 (2011)], arXiv:1103.3312 [nucl-ex].

[514] L. P. Csernai and J. I. Kapusta, "Entropy and cluster production in nuclear collisions," Phys. Rep. **131**, 223 (1986).

[515] S. T. Butler and C. A. Pearson, "Deuterons from high-energy proton bombardment of matter," Phys. Rev. **129**, 836 (1963).

[516] H. Sato and K. Yazaki, "On the coalescence model for high energy nuclear reactions," Physics Lett. B **98**, 153 (1981).

[517] D.-M. Gomez-Coral, A. Menchaca Rocha, V. Grabski, A. Datta, P. von Doetinchem, and A. Shukla, "Deuteron and antideuteron production simulation in cosmic-ray interactions," Phys. Rev. D **98**, 023012 (2018), arXiv:1806.09303 [astro-ph.HE].

[518] M. Kachelrieß, S. Ostapchenko, and J. Tjemsland, "Alternative coalescence model for deuteron, tritium, helium-3 and their antinuclei," Eur. Phys. J. A **56**, 4 (2020), arXiv:1905.01192 [hep-ph].

[519] T. Aramaki *et al.*, "Review of the theoretical and experimental status of dark matter identification with cosmic-ray antideuterons," Phys. Rep. **618**, 1 (2016), arXiv:1505.07785 [hep-ph].

[520] P. von Doetinchem *et al.*, "Cosmic-ray antinuclei as messengers of new physics: Status and outlook for the new decade," J. Cosmol. Astropart. Phys. **08**, 035 (2020), arXiv:2002.04163 [astro-ph.HE].

[521] M. Kachelrieß, S. Ostapchenko, and J. Tjemsland, "On nuclear coalescence in small interacting systems," Eur. Phys. J. A **57**, 167 (2021), arXiv:2012.04352 [hep-ph].

[522] T. Shao, J. Chen, Y.-G. Ma, and Z. Xu, "Production of light antinuclei in pp collisions by dynamical coalescence and their fluxes in cosmic rays near earth," Phys. Rev. C **105**, 065801 (2022).

[523] P. Braun-Munzinger, K. Redlich, and J. Stachel, "Particle production in heavy ion collisions," in *Quark-Gluon Plasma*, edited by R. C. Hwa and X.-N. Wang (2003) p. 491, arXiv:nucl-th/0304013.

[524] R. Pasechnik and M. Šumbera, "Phenomenological review on quark–gluon plasma: Concepts vs. observations," Universe **3**, 7 (2017), arXiv:1611.01533 [hep-ph].

[525] J. Chen, D. Keane, Y.-G. Ma, A. Tang, and Z. Xu, "Antinuclei in heavy-ion collisions," Phys. Rep. **760**, 1 (2018), arXiv:1808.09619 [nucl-ex].

[526] P. Braun-Munzinger, J. Stachel, and C. Wetterich, "Chemical freezeout and the QCD phase transition temperature," Phys. Lett. B **596**, 61 (2004), arXiv:nucl-th/0311005.

[527] L. Šerkšnytė *et al.*, "Reevaluation of the cosmic antideuteron flux from cosmic-ray interactions and from exotic sources," Phys. Rev. D **105**, 083021 (2022), arXiv:2201.00925 [astro-ph.HE].

[528] B. V. Jacak and B. Müller, "The exploration of hot nuclear matter," Science **337**, 310 (2012).

[529] A. Andronic, P. Braun-Munzinger, K. Redlich, and J. Stachel, "Decoding the phase structure of QCD via particle production at high energy," Nature **561**, 321 (2018), arXiv:1710.09425 [nucl-th].

[530] J. Herms, A. Ibarra, A. Vittino, and S. Wild, "Antideuterons in cosmic rays: Sources and discovery potential," J. Cosmol. Astropart. Phys. **02**, 018 (2017), arXiv:1610.00699 [astro-ph.HE].

[531] N. Abgrall *et al.* (NA61/SHINE Collaboration), "NA61/SHINE facility at the CERN SPS: beams and detector system," J. Instrum. **9**, P06005 (2014), arXiv:1401.4699 [physics.ins-det].

[532] D. Larsen (NA61/SHINE Collaboration), "Upgrade of the NA61/SHINE facility beyond 2020 for an expanded physics programme," Universe **5**, 24 (2019).

[533] A. Aduszkiewicz *et al.* (NA61/SHINE Collaboration), "Measurements of $\pi^{\pm}$, $\mathrm{K}^{\pm}$, p and $\bar{\mathrm{p}}$ spectra in proton-proton interactions at 20, 31, 40, 80 and 158 GeV/c with the NA61/SHINE spectrometer at the CERN SPS," Eur. Phys. J. C **77**, 671 (2017), arXiv:1705.02467 [nucl-ex].

[534] M. Korsmeier, F. Donato, and M. Di Mauro, "Production cross sections of cosmic antiprotons in the light of new data from the NA61 and LHCb experiments," Phys. Rev. D **97**, 103019 (2018), arXiv:1802.03030 [astro-ph.HE].

[535] R. Aaij *et al.* (LHCb Collaboration), "Measurement of antiproton production in pHe collisions at $\sqrt{s_{NN}} = 110$ GeV," Phys. Rev. Lett. **121**, 222001 (2018), arXiv:1808.06127 [hep-ex].

[536] G. Barucca *et al.* (PANDA Collaboration), "PANDA Phase One," Eur. Phys. J. A **57**, 184 (2021), arXiv:2101.11877 [hep-ex].

[537] H. Fuke *et al.*, "Search for cosmic-ray antideuterons," Phys. Rev. Lett. **95**, 081101 (2005), arXiv:astro-ph/0504361.

[538] V. Choutko and F. Giovacchini, "Cosmic rays antideuteron sensitivity for AMS-02 experiment," in *International Cosmic Ray Conference*, International Cosmic Ray Conference, Vol. 4 (2008) pp. 765–768.

[539] K. Perez *et al.*, "Astro 2020 science white paper: Cosmic-ray antinuclei as messengers for dark matter," (2019), arXiv:1904.05938 [astro-ph.HE].

[540] F. Donato, N. Fornengo, and P. Salati, "Anti-deuterons as a signature of supersymmetric dark matter," Phys. Rev. D **62**, 043003 (2000), arXiv:hep-ph/9904481.

[541] A. Barrau, G. Boudoul, F. Donato, D. Maurin, P. Salati, I. Stefanon, and R. Taillet, "Antideuterons as a probe of primordial black holes," Astron. Astrophys. **398**, 403 (2003), arXiv:astro-ph/0207395.

[542] H. Baer and S. Profumo, "Low energy antideuterons: Shedding light on dark matter," J. Cosmol. Astropart. Phys. **12**, 008 (2005), arXiv:astro-ph/0510722.

[543] F. Donato, N. Fornengo, and D. Maurin, "Antideuteron fluxes from dark matter annihilation in diffusion models," Phys. Rev. D **78**, 043506 (2008), arXiv:0803.2640 [hep-ph].

[544] L. A. Dal and A. R. Raklev, "Antideuteron limits on decaying dark matter with a tuned formation model," Phys. Rev. D **89**, 103504 (2014), arXiv:1402.6259 [hep-ph].

[545] I. Cholis, T. Linden, and D. Hooper, "Antideuterons and antihelium nuclei from annihilating dark matter," Phys. Rev. D **102**, 103019 (2020), arXiv:2001.08749 [astro-ph.HE].

[546] R. L. Golden, S. Horan, B. G. Mauger, G. D. Badhwar, J. L. Lacy, S. A. Stephens, R. R. Daniel, and J. E. Zipse, "Evidence for the existence of cosmic-ray antiprotons," Phys. Rev. Lett. **43**, 1196 (1979).

[547] T. Hams *et al.*, "BESS-Polar II experiment," in *30th International Cosmic Ray Conference*, Vol. 2 (2007) pp. 91–94.

[548] K. Abe *et al.*, "Search for antihelium with the BESS-Polar spectrometer," Phys. Rev. Lett. **108**, 131301 (2012), arXiv:1201.2967 [astro-ph.CO].

[549] N. Tomassetti (AMS Collaboration), "AMS-02 in space: Physics results, overview, and challenges," Nucl. Part. Phys. Proc. **265-266**, 245 (2015), arXiv:1511.00052 [astro-ph.HE].

[550] V. Bindi *et al.*, "Calibration and performance of the AMS-02 time of flight detector in space," Nucl. Instrum. Methods Phys. Res. A **743**, 22 (2014).

[551] F. Giovacchini, J. Casaus, and A. Oliva (AMS-02 RICH Collaboration), "The AMS-02 RICH detector: Status and physics results," Nucl. Instrum. Methods Phys. Res. A **952**, 161797 (2020).

[552] S. Rosier-Lees (AMS Collaboration), "Performance of the AMS-02 electromagnetic calorimeter in space," J. Phys. Conf. Ser. **404**, 012034 (2012).

[553] J. Heisig, "Cosmic-ray antiprotons in the AMS-02 era: A sensitive probe of dark matter," Mod. Phys. Lett. A **36**, 2130003 (2021), arXiv:2012.03956 [astro-ph.HE].

[554] J. Heisig, M. Korsmeier, and M. W. Winkler, "Dark matter or correlated errors: Systematics of the AMS-02 antiproton excess," Phys. Rev. Res. **2**, 043017 (2020).

[555] V. A. Choutko, AMS days at La Palma, Spain (2018, unpublished).

[556] S. Ting, "The first five years of the Alpha Magnetic Spectrometer on the International Space Station," CERN Colloquium (2016), last accessed 2023-04-01.

[557] V. Poulin, P. Salati, I. Cholis, M. Kamionkowski, and J. Silk, "Where do the AMS-02 antihelium events come from?" Phys. Rev. D **99**, 023016 (2019), arXiv:1808.08961 [astro-ph.HE].

[558] F. Rogers *et al.* (GAPS Collaboration), "Sensitivity of the GAPS experiment to low-energy cosmic-ray antiprotons," Astropart. Phys. **145**, 102791 (2023), arXiv:2206.12991 [astro-ph.HE].

[559] N. Saffold *et al.*, "Passivation of Si(Li) detectors operated above cryogenic temperatures for space-based applications," Nucl. Instrum. Methods Phys. Res. A **997**, 165015 (2021), arXiv:2102.06168 [astro-ph.IM].

[560] N. Saffold *et al.* (GAPS Collaboration), "Cosmic antihelium-3 nuclei sensitivity of the GAPS experiment," Astropart. Phys. **130**, 102580 (2021), arXiv:2012.05834 [hep-ph].

[561] R. Seki and C. E. Wiegand, "Kaonic and other exotic atoms," Annu. Rev. Nucl. Part. Sci. **25**, 241 (1975).

[562] E. Borie and M. Leon, "X-ray yields in protonium and mesic hydrogen," Phys. Rev. A **21**, 1460 (1980).

[563] T. P. Terada and R. S. Hayano, "Cascade calculation of hadronic hydrogen atoms," Phys. Rev. C **55**, 73 (1997).

[564] V. Markushin and T. Jensen, "Atomic cascade and x-ray yields in light exotic atoms," Nucl. Phys. A **691**, 318 (2001).

[565] D. Gotta, "Precision spectroscopy of light exotic atoms," Prog. Part. Nucl. Phys. **52**, 133 (2004).

[566] T. Aramaki *et al.*, "A measurement of atomic x-ray yields in exotic atoms and implications for an antideuteron-based dark matter search," Astropart. Phys. **49**, 52 (2013), arXiv:1303.3871 [astro-ph.IM].

[567] G. Backenstoss, "Pionic atoms," Annu. Rev. Nucl. Part. Sci. **20**, 467 (1970).

[568] E. Borie and G. A. Rinker, "The energy levels of muonic atoms," Rev. Mod. Phys. **54**, 67 (1982).

[569] C. J. Batty, "Anti-protonic hydrogen atoms," Rep. Prog. Phys. **52**, 1165 (1989).

[570] C. J. Batty, E. Friedman, and A. Gal, "Strong interaction physics from hadronic atoms," Phys. Rep. **287**, 385 (1997).

[571] L. I. Menshikov and M. K. Eseev, "Some problems of the physics of exotic atoms," Phys. Usp. **44**, 135 (2001).

[572] D. Horváth, "Exotic atoms and muonium," in *Handbook of Nuclear Chemistry*, edited by A. Vértes, S. Nagy, Z. Klencsár, R. G. Lovas, and F. Rösch (Springer U.S., Boston, 2011) p. 1485.

[573] C. Curceanu *et al.*, "The modern era of light kaonic atom experiments," Rev. Mod. Phys. **91**, 025006 (2019).

[574] M. Doser, "Antiprotonic bound systems," Prog. Part. Nucl. Phys. **125**, 103964 (2022).

[575] C. Amsler, "Proton-antiproton annihilation and meson spectroscopy with the Crystal Barrel," Rev. Mod. Phys. **70**, 1293 (1998).

[576] J. Cugnon and J. Vandermeulen, "Antiproton-nucleus interaction," Ann. Phys. Fr. **14**, 49 (1989).

[577] J. Cugnon, "Antideuteron annihilation on nuclei," Nucl. Phys. A **542**, 559 (1992).

[578] A. S. Sudov *et al.*, "Production of light particles after anti-proton nucleus annihilation and their interpretation with statistical models," Nucl. Phys. A **554**, 223 (1993).

[579] D. Polster *et al.*, "Light particle emission induced by stopped antiprotons in nuclei: Energy dissipation and neutron-to-proton ratio," Phys. Rev. C **51**, 1167 (1995).

[580] M. Manghisoni, V. Re, E. Riceputi, M. Sonzogni, L. Ratti, and L. Fabris, "Low-noise analog channel for the readout of the Si(Li) detector of the GAPS experiment," IEEE Trans. Nucl. Sci. **68**, 2661 (2021).

[581] V. Re, L. Ghislotti, P. Lazzaroni, M. Manghisoni, E. Riceputi, L. Ratti, M. Boezio, G. Zampa, and L. Fabris, "A mixed-signal processor for x-ray spectrometry and tracking in the GAPS experiment," Nucl. Instrum. Methods Phys. Res. A **1045**, 167617 (2023).

[582] E. Riceputi, M. Boezio, L. Fabris, L. Ghislotti, P. Lazzaroni, M. Manghisoni, L. Ratti, V. Re, and G. Zampa, "The 32 analog channels readout for the long-flight gaps balloon experiment tracking system," in *Proceedings of SIE 2022*, edited by G. Cocorullo, F. Crupi, and E. Limiti (Springer Nature Switzerland, Cham, 2023) pp. 27–32.

[583] M. Manghisoni, D. Comotti, L. Gaioni, L. Ratti, and V. Re, "Dynamic compression of the signal in a charge sensitive amplifier: From concept to design," IEEE Trans. Nucl. Sci. **62**, 2318 (2015).

[584] S. Quinn (GAPS Collaboration), "The GAPS instrument: A large area time of flight and high resolution exotic atom spectrometer for cosmic antinuclei," PoS Proc. Sci. **ICRC2021**, 079 (2021).

[585] J. L. Ryan, *Dark Matter Searches with Gamma Rays from the Galactic Center Halo and Cosmic-Ray Antimatter*, Ph.D. thesis, University of California, Los Angeles (2022), `https://escholarship.org/uc/item/78d2f4gc`.

[586] S. Ritt, "The DRS chip: Cheap waveform digitizing in the GHz range," Nucl. Instrum. Methods Phys. Res. A **518**, 470 (2004).

[587] H. Fuke *et al.*, "Development of a cooling system for GAPS using oscillating heat pipe," Trans. Jpn. Soc. Aeronaut. Space. Sci., Aerosp. Tech. Jpn. **14**, 17 (2016).

[588] K. Mori, C. J. Hailey, E. A. Baltz, W. W. Craig, M. Kamionkowski, W. T. Serber, and P. Ullio, "A novel antimatter detector based on x-ray deexcitation of exotic atoms," Astrophys. J. **566**, 604 (2002), arXiv:astro-ph/0109463.

[589] C. J. Hailey *et al.*, "Accelerator testing of the general antiparticle spectrometer; a novel approach to indirect dark matter detection," J. Cosmol. Astropart. Phys. **01**, 007 (2006), arXiv:astro-ph/0509587.

[590] H. Fuke *et al.*, "The pGAPS experiment: An engineering balloon flight of prototype GAPS," Adv. Space Res. **53**, 1432 (2014), arXiv:1303.0380 [astro-ph.IM].

[591] W. R. Leo, *Techniques for Nuclear and Particle Physics Experiments: A How to Approach* (Springer-Verlag Berlin, Heidelberg, 1987).

[592] Yu. P. Raizer, *Gas Discharge Physics*, edited by J. E. Allen (Springer Berlin, Heidelberg, 1991).

[593] R. K. Crouch and T. Gilmer, "Thermal ionization energy of lithium and lithium-oxygen complexes in single-crystal silicon," J. Phys. Chem. Solids **30**, 2037 (1969).

[594] E. M. Pell, "Ion drift in an *n–p* junction," J. Appl. Phys. **31**, 291 (1960).

[595] H. Reiss, C. S. Fuller, and F. J. Morin, "Chemical interactions among defects in germanium and silicon," Bell Sys. Tech. J. **35**, 535 (1956).

[596] D. Freck and J. Wakefield, "Gamma-ray spectrum obtained with a lithium-drifted *p–i–n* junction in germanium," Nature **193**, 669 (1962).

[597] A. Tavendale and G. Ewan, "A high resolution lithium-drift germanium gamma-ray spectrometer," Nucl. Instrum. Methods **25**, 185 (1963).

[598] A. Lauber, "The theory of compensation in lithium drifted semiconductor detectors," Nucl. Instrum. Methods **75**, 297 (1969).

[599] F. Goulding, "Semiconductor detectors for nuclear spectrometry, I," Nucl. Instrum. Methods **43**, 1 (1966), proceedings of the Tenth Summer Meeting of Nuclear Physicists.

[600] I. Ahmad and F. Wagner, "A simple cooled Si(Li) electron spectrometer," Nucl. Instrum. Methods **116**, 465 (1974).

[601] S. M. Krimigis, T. P. Armstrong, W. I. Axford, C. O. Bostrom, C. Y. Fan, G. Gloeckler, and L. J. Lanzerotti, "The Low Energy Charged Particle (LECP) experiment on the Voyager spacecraft," Space Sci. Rev. **21**, 329 (1977).

[602] R. Giacconi *et al.*, "The Einstein (HEAO 2) x-ray observatory," Astrophys. J. **230**, 540 (1979).

[603] G. Gerbier *et al.*, "Measurement of the ionization of slow silicon nuclei in silicon for the calibration of a silicon dark-matter detector," Phys. Rev. D **42**, 3211 (1990).

[604] M. L. Roberts, G. S. Bench, D. W. Heikkinen, D. H. Morse, P. R. Bach, and A. E. Pontau, "The new nuclear microprobe at Livermore," Nucl. Instrum. Methods Phys. Res., Sect. B **104**, 13 (1995).

[605] E. C. Stone *et al.*, "The Cosmic-Ray Isotope Spectrometer for the Advanced Composition Explorer," Space Sci. Rev. **86**, 285 (1998).

[606] L. A. Popeko, I. M. Kotina, G. A. Shishkina, L. A. Grigorieva, A. G. Artyukh, Y. G. Teterev, and Y. M. Sereda, “Thick Si(Li) coaxial detectors for registration of intermediate energy heavy ions,” Nucl. Instrum. Methods Phys. Res. A **596**, 235 (2008).

[607] Y. B. Gurov *et al.*, “Planar Si(Li) detectors with a large active volume,” Instrum. Exp. Tech. **53**, 36 (2010).

[608] A. Simon, A. Warczak, T. ElKafrawy, and J. A. Tanis, “Radiative double electron capture in collisions of O^{8+} ions with carbon,” Phys. Rev. Lett. **104**, 123001 (2010).

[609] V. N. Aseev *et al.*, “Upper limit on the electron antineutrino mass from the Troitsk experiment,” Phys. Rev. D **84**, 112003 (2011).

[610] A. V. Derbin, I. S. Drachnev, A. S. Kayunov, and V. N. Muratova, “Constraints on the axion-electron coupling constant for solar axions appearing owing to bremsstrahlung and the Compton process,” JETP Lett. **95**, 339 (2012), 10.1134/S002136401207003X.

[611] Y. Zheng *et al.*, “Mini-Orange spectrometer at CIAE,” Chinese Phys. C **40**, 086002 (2016), arXiv:1601.01078 [physics.ins-det].

[612] T. K. Eriksen *et al.*, “Improved precision on the experimental $E0$ decay branching ratio of the Hoyle state,” Phys. Rev. C **102**, 024320 (2020), arXiv:2007.15374 [nucl-ex].

[613] J. A. Briz *et al.*, “Clarifying the structure of low-lying states in ^{72}Br,” Phys. Rev. C **105**, 014323 (2022).

[614] K. Perez *et al.*, “Fabrication of low-cost, large-area prototype Si(Li) detectors for the GAPS experiment,” Nucl. Instrum. Methods Phys. Res. A **905**, 12 (2018), arXiv:1807.07912 [astro-ph.IM].

[615] M. Kozai *et al.*, “Developing a mass-production model of large-area Si(Li) detectors with high operating temperatures,” Nucl. Instrum. Methods Phys. Res. A **947**, 162695 (2019), arXiv:1906.05577 [physics.ins-det].

[616] M. Kozai *et al.*, “Statistical investigation of the large-area Si(Li) detectors mass-produced for the GAPS experiment,” Nucl. Instrum. Methods Phys. Res. A **1034**, 166820 (2022).

[617] F. Rogers *et al.*, “Large-area Si(Li) detectors for X-ray spectrometry and particle tracking in the GAPS experiment,” J. Instrum. **14**, P10009 (2019), arXiv:1906.00054 [physics.ins-det].

[618] J. Llacer, “Study of surface effects in thick lithium drifted silicon radiation detectors,” IEEE Trans. Nucl. Sci. **11**, 221 (1964).

[619] F. R. Rogers, *Applications of X-Ray Instrumentation for Dark Matter Searches with Cosmic Antiparticles*, Ph.D. thesis, Massachusetts Institute of Technology (2022).

[620] H. Spieler, *Semiconductor Detector Systems*, Semiconductor Science and Technology, Vol. 12 (Oxford Univ. Press, Oxford, 2005).

[621] L. Fabris, N. W. Madden, and H. Yaver, "A fast, compact solution for low noise charge preamplifiers," Nucl. Instrum. Methods Phys. Res. A **424**, 545 (1999).

[622] S. Kumar, J. Chen, and F. Kondev, "Nuclear data sheets for $A = 109$," Nucl. Data Sheets **137**, 1 (2016).

[623] A. Stoessl, "`dactylos`—Python interface to CAEN N6725 digitizer," https://github.com/achim1/dactylos [Last accessed 2023-04-01].

[624] S. Ohkawa, M. Yoshizawa, and K. Husimi, "Direct synthesis of the Gaussian filter for nuclear pulse amplifiers," Nucl. Instrum. Methods **138**, 85 (1976).

[625] J. B. Johnson, "Thermal agitation of electricity in conductors," Phys. Rev. **32**, 97 (1928).

[626] H. Nyquist, "Thermal agitation of electric charge in conductors," Phys. Rev. **32**, 110 (1928).

[627] R. Voss, "$1/f$ (flicker) noise: A brief review," in *33rd Annual Symposium on Frequency Control* (1979) pp. 40–46.

[628] P. Dutta and P. M. Horn, "Low-frequency fluctuations in solids: $1/f$ noise," Rev. Mod. Phys. **53**, 497 (1981).

[629] M. B. Weissman, "$1/f$ noise and other slow, nonexponential kinetics in condensed matter," Rev. Mod. Phys. **60**, 537 (1988).

[630] Sh. Kogan, *Electronic Noise and Fluctuations in Solids* (Cambridge Univ. Press, Cambridge, 1996).

[631] E. Milotti, "$1/f$ noise: A pedagogical review," , physics/0204033 (2002), arXiv:physics/0204033 [physics.class-ph].

[632] E. Paladino, Y. M. Galperin, G. Falci, and B. L. Altshuler, "$1/f$ noise: Implications for solid-state quantum information," Rev. Mod. Phys. **86**, 361 (2014).

[633] W. Schottky, "Small-shot effect and flicker effect," Phys. Rev. **28**, 74 (1926).

[634] J. B. Johnson, "The Schottky effect in low frequency circuits," Phys. Rev. **26**, 71 (1925).

[635] J. Bernamont, "Fluctuations de potentiel aux bornes d'un conducteur métallique de faible volume parcouru par un courant," Ann. Phys. **11**, 71 (1937).

[636] A. L. McWhorter, *1/f Noise and Related Surface Effects in Germanium*, Ph.D. thesis, Massachusetts Institute of Technology (1955), `https://dspace.mit.edu/handle/1721.1/12061`.

[637] G. B. Rybicki and A. P. Lightman, *Radiative Processes in Astrophysics* (Wiley, 1986).

[638] F. S. Goulding, "Pulse-shaping in low-noise nuclear amplifiers—A physical approach to noise analysis," Nucl. Instrum. Methods **100**, 493 (1972).

[639] V. Radeka, "Low noise techniques in detectors," Annu. Rev. Nucl. Part. Sci. **38**, 217 (1988).

[640] B. G. Lowe and R. A. Sareen, "A measurement of the electron-hole pair creation energy and the Fano factor in silicon for 5.9 keV x-rays and their temperature dependence in the range 80–270 K," Nucl. Instrum. Methods Phys. Res. A **576**, 367 (2007).

[641] D. Rodrigues *et al.*, "Absolute measurement of the Fano factor using a Skipper-CCD," Nucl. Instrum. Methods Phys. Res. A **1010**, 165511 (2021), arXiv:2004.11499 [physics.ins-det].

[642] A. Chilingarov, "Temperature dependence of the current generated in Si bulk," J. Instrum. **8**, P10003 (2013).

[643] W. Bludau, A. Onton, and W. Heinke, "Temperature dependence of the band gap of silicon," J. Appl. Phys. **45**, 1846 (1974).