

HYDROGEN-TRIGGERED X-RAY BURSTS ON SLOWLY ACCRETING NEUTRON STARS

By

Sierra Casten

A DISSERTATION

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

Physics and Astronomy—Doctor of Philosophy

2026

ABSTRACT

Many observed neutron stars are in binaries and accrete hydrogen-rich material from low-mass companions. The accumulating matter eventually triggers a thermonuclear runaway that results in an X-ray burst lasting 10–100s with a recurrence time of hours to days. Almost all observed thermonuclear X-ray bursts are thought to be triggered by the thermally unstable triple-alpha process, as most observed bursters are sufficiently hot that hydrogen burning is via the beta-limited, thermally stable hot CNO cycle. The conditions under which unstable hydrogen burning can produce observable bursts have remained uncertain. This dissertation investigates weak thermonuclear bursts from the accreting millisecond X-ray pulsar SAX J1808.4–3658 through a combined observational and theoretical study.

We analyzed two weak bursts detected with the Neutron Star Interior Composition Explorer and the Rossi X-ray Timing Explorer that occurred shortly after the onset of renewed accretion. The bursts exhibit significantly lower peak fluxes than the bright helium-triggered bursts typically observed from this source. They occur at low inferred accretion rates and modest accumulated columns. The burst properties, together with the elevated post-burst emission consistent with quasi-steady hot-CNO burning, strongly suggest ignition triggered by thermally unstable hydrogen burning.

To interpret these events, we used the stellar evolution code, MESA, to explore unstable hydrogen ignition on slowly accreting, cool neutron stars across a range of metallicities. We find that solar-metallicity accretion produces insufficient heating to drive convection or to generate an observable burst by several orders of magnitude. However, when the CNO abundance at the base of the accreted layer is enhanced, proton captures drive convection and produce a rapid burst rise consistent with the observed weak bursts from SAX J1808.4–3658. Following the burst, the models naturally settle into a prolonged phase of thermally stable hot-CNO burning that generates an extended tail similar to observations. The tail energetics and duration depend on the CNO abundance in the accreted matter, and thus open a new probe into mixing in the neutron star envelope.

Together, these observational and theoretical results strengthen the case that hydrogen-triggered

bursts occur on slowly accreting neutron stars and can reach observable detection thresholds. This work establishes a framework for identifying and interpreting hydrogen-triggered bursts and highlights the importance of metallicity in shaping burst behavior.

Copyright by
SIERRA CASTEN
2026

*Dedicated to my family and loved ones:
I love you to the farthest neutron star in the Universe and back.*

ACKNOWLEDGMENTS

I am deeply grateful to my advisor, Edward Brown, for his steady guidance and support throughout my graduate career. He consistently encouraged me to pursue my interests on the observational side even when they extended beyond his primary area of expertise. His openness to these directions was instrumental in shaping this work. He pushed me to find joy in stepping outside my comfort zone, often reminding me that “if we already knew the answer, that would be boring.” Beyond his scientific mentorship, Ed was patient and understanding during periods of burnout and always made time to offer thoughtful advice and perspective. I particularly valued our many science discussions over coffee and our many meetings where we often lost track of time, which challenged me to think more critically and helped foster my development into an independent researcher. I suspect his ulterior motive in nurturing that independence was to ensure he would always have someone available with whom to debate the finer points of neutron stars. It is a role I have thoroughly enjoyed and look forward to continuing.

I am also deeply grateful to Andrew Cumming and Duncan Galloway—both of whom generously took me under their wings during my graduate career. They consistently treated me with respect and kindness as a junior researcher. I especially appreciate Duncan’s initial skepticism prior to the publication of my first paper, followed by his willingness to be persuaded by the evidence—an experience that both strengthened the work and built my confidence as a scientist. Duncan’s continued encouragement and mentorship on the observational side were invaluable. Andrew truly lived up to his title as my academic uncle by consistently offering generous time, support, and mentorship. Andrew’s guidance was instrumental in developing physically rigorous MESA models for my second paper. I am particularly thankful to him for supporting my visit to McGill. I am also grateful to Simon Guichandut for his kindness and patience in answering my many questions about MESA and for his help troubleshooting modeling challenges along the way. The time and attention they all invested in me, well beyond any formal obligation, made an indelible impact on my development as a researcher, and I hold them in the highest respect and fondness.

Additionally, I am thankful to Tod Strohmayer, who took a chance on me as a recent undergrad-

uate and was my first mentor in neutron star research. He welcomed me into this field early in my career—I knew absolutely nothing about neutron stars or X-ray bursts—and he kindly guided me through my first project and continued to support our work as I began graduate school. I appreciate the opportunity he gave me to take my first steps into the study of neutron stars, and the role he played in launching my work in this field.

I am sincerely grateful to the graduate students, postdocs, and faculty at Michigan State University for fostering such an inclusive and supportive environment. The graduate students have become some of my closest friends and an incredible support system as we navigated coursework, impostor syndrome, and the many ups and downs of graduate school. Our time playing games and movie marathons are memories I especially cherish. I would also like to personally thank Jay Strader and Laura Chomiuk, whose encouragement and generosity meant a great deal to me. I always appreciated our many conversations—both scientific and otherwise—which made MSU feel all the more like home.

Last but certainly not least, I dedicate this work to my family and loved ones, whose support and love have shaped me into the person I am today. My parents, in particular, have shown me unconditional support at every stage of this journey. They sacrificed precious time together while my mom moved in with me during a difficult season of health and helped me recover from surgery. She took care of everything, so that I could focus my energy on finishing my coursework. I would likely not be completing my Ph.D. on this timeline without their sacrifices. To my brother Cole—who has always claimed he loved space first—while that may be true, I now have a doctorate in it, so I think it is fair to say I love it more. I am thankful for our sibling banter and over-the-phone hangouts, which always brightened my weeks. I am also grateful to my dog, Chico, for the constant companionship and (often persistent) reminders to step away from the computer and enjoy life beyond it. To my loved ones: thank you for supporting me, laughing with me, and reminding me to enjoy the journey along the way.

TABLE OF CONTENTS

LIST OF TABLES	ix
LIST OF FIGURES	x
CHAPTER 1 INTRODUCTION	1
1.1 Neutron Stars	1
1.2 Neutron Star Formation	2
1.3 Neutron Star Population	3
1.4 Accreting Neutron Stars	4
1.5 Dissertation Overview	10
CHAPTER 2 FIRST OBSERVATIONS OF H-TRIGGERED BURSTS	12
2.1 NICER Observations of SAX J1808.4–3658	15
2.2 Physical Scenario and Interpretation	22
2.3 Summary, Caveats, and Outlook	35
CHAPTER 3 EFFECTS OF METALLICITY ON H-TRIGGERED BURSTS	40
3.1 Burning via the Hot CNO cycle	41
3.2 Numerical models	43
3.3 Unstable hydrogen burning	45
3.4 Discussion	55
CHAPTER 4 CONCLUSION	58
4.1 Dissertation Summary	58
4.2 Future Modeling Outlook	59
4.3 Future Observational Outlook	60
BIBLIOGRAPHY	62

LIST OF TABLES

Table 2.1	Spectral Model Parameters for SAX J1808.4–3658: Persistent Emission	17
Table 2.2	Burst Parameters from NuSTAR and NICER Observations of the Outburst	35
Table 3.1	Isotopes in the reaction network <code>rp_157.net</code>	44

LIST OF FIGURES

Figure 1.1	Roche geometry for a neutron star (dot at left in the top panel) in a binary with a sun-like star (red sphere at right in the top panel) in a co-rotating reference frame. Equipotential contours are shown in the lower panel, along with the center of mass, “CM,” and the five Lagrange points L_i . The equipotential contour passing through L_1 is shown by the dark contour at bottom and the shaded surface at top. (Figure courtesy E. Brown, license CC BY-NC-SA 4.0).	4
Figure 1.2	Schematic of an accreting neutron star. Exterior to the core is a kilometer-thick crust composed of nuclei, electrons, and (at depth) free neutrons. Accreted material is deposited at the centimeter-deep <i>photosphere</i> , the visible “surface.” Below the photosphere is a meters-thick <i>envelope</i> , where the matter (mostly H and He) is ionized, non-degenerate, and opaque. As matter descends deeper into the star, nuclear reactions burn the H/He and other light elements to heavier nuclei. The electrons become degenerate and Coulomb interactions force the nuclei into a highly conductive, liquid state, or <i>ocean</i> , and eventually lock the nuclei into a rigid lattice—the <i>crust</i> . With increasing density, electron captures drive the composition more neutron-rich, and at $\sim 10^{-3}$ of nuclear density, neutrons “drip” from nuclei. When the spacing between nuclei becomes of order their size, nuclear interactions may distort the nuclei into sheets and rods, a phase known as <i>nuclear pasta</i> .	6
Figure 2.1	Light curve of the 2019 July–August outburst of in the 0.4–7 keV band observed with NICER. Count rates were computed in 2 s intervals. Note the logarithmic scale. The first, weaker burst is evident near day 3. The bright burst near day 14 was reported on by Bult et al. (2019).	16
Figure 2.2	Lightcurve of the first, weak burst from SAX J1808.4–3658 in the 0.4–7 keV band observed with NICER. Main panel: The count rates were computed in 1 s intervals, and the vertical dashed and dotted lines denote the intervals used to extract spectra for the pre- and post-burst spectral modeling, respectively. Inset panel: The same data are used, but the time bins are 16 s, and the logarithmic scale highlights the offset in count rate between the pre- and post-burst emission. The dashed red line is a constant value fit to the pre-burst level, and is meant as a guide to the eye.	19
Figure 2.3	Energy spectrum of the pre-burst interval, modeled with a phenomenological model similar to that employed by Patruno et al. (2009), which includes <code>diskbb</code> , <code>bbodyrad</code> and <code>powerlaw</code> components, in addition to the line at 1 keV. See the text in § 2.1.1 for further details.	20
Figure 2.4	Evolution of the weak X-ray burst derived from spectral modeling in the 0.5–10 keV band. We show from the top down: the bolometric flux, blackbody temperature, blackbody radius (at 3.5 kpc), and reduced χ^2 , respectively. The error bars indicate 1- σ confidence intervals.	23

Figure 2.5	Light curve from RXTE data (PCU 2, 3–30 keV) of the 2005 outburst from SAX J1808.4–3658. Note the logarithmic scale. A weak X-ray burst is seen early in this outburst. Much brighter and energetic bursts are seen near days 4 and 8. Note that the burst near day 8 was truncated by the RXTE exposure, and almost certainly the brightest part of this event was missed.	32
Figure 2.6	Light curves from NICER (black, left axis) and NuSTAR (red, right axis) around the time of the weak X-ray burst at $t = 0$. Both light curves are calculated using an 8 s time resolution.	34
Figure 3.1	Surface luminosity (<i>solid black</i>) and nuclear luminosity (<i>dashed red</i>) during the onset of unstable hydrogen burning and the evolution of the burning to a tail of quasi-steady burning. Models with $Z = 0.04$ (<i>top</i>) and 0.10 (<i>bottom</i>) are shown. The vertical dotted lines indicate where we sample the isotopic distribution in the accreted and burning neutron star envelope.	46
Figure 3.2	Mass fractions as functions of column depth y for accretion with $Z = 0.04$ (<i>top row</i>) and $Z = 0.10$ (<i>bottom row</i>). The columns, from left to right, are a time-ordered and correspond to the first three vertical lines depicted in Fig. 3.1. The lines terminate right at the base of the accreted layer. The isotope ^{56}Fe is not shown to avoid clutter.	47
Figure 3.3	Snapshot of the burst with $Z = 0.10$ just after the initial thermal instability as surface luminosity begins its rise. The main panel at left shows the mass fractions (omitting ^{56}Fe) in the accreted layer as a function of column. The top right panel shows the surface luminosity (<i>solid black curve</i>) and integrated nuclear heating (<i>dashed red curve</i>); the dots mark the current time (note that the nuclear heating extends beyond the plot limit, cf. Fig. 3.1, bottom). The bottom right panel shows the temperature as a function of column depth. The vertical dotted lines on the bottom panel indicate the range of column depth displayed in the main panel. The grey shaded region at right is the ^{56}Fe substrate of the initial, relaxed, NS envelope prior to accretion. The light tan region is where convection occurs. The temperature curve is turn orange-red where $\epsilon_{\text{nuc}} > 0.05\epsilon_{\text{CNO}}$ and red where $\epsilon_{\text{nuc}} > \epsilon_{\text{CNO}}$	49
Figure 3.4	Composition at the transition from burst peak to burst tail (the rightmost vertical dotted line in Fig. 3.1), in which H burns via the hot CNO cycle in quasi-steady-state, for $Z = 0.04$ (<i>top</i>) and $Z = 0.10$ (<i>bottom</i>). The isotope ^{56}Fe is not shown.	50
Figure 3.5	Ignition column depths (<i>dots</i>) as a function of metallicity in the accreted material. These decrease roughly as $y_{\text{ign}} \propto Z^{-0.4}$ (<i>dotted line</i>). For comparison, we also show the inferred burst ignition depth for SAX J1808.4–3658 (<i>shaded region</i>) which covers the range for $1.4 M_{\odot} < M < 2.0 M_{\odot}$. This range of inferred ignition depth is based on the observed fluence for the 2019 NICER burst with $d = 3.3$ kpc, $R = 12$ km, and anisotropy factor $\xi = 1$	51

Figure 3.6 Lightcurves for metallicities from $Z = 0.01$ to $Z = 0.3$, ordered lightest to darkest, respectively, and numbered by Z at the end of the tail. The lightcurves are aligned so that $L = 1.4 \times 10^{34} \text{ erg s}^{-1}$ at $t = 0$, and the time scale shifts from linear to logarithmic to follow the transition from peak to tail. Bursts for which no convection develops are indicated with dotted lines and are notable for their long rise. For comparison, we also plot (*gold*) the observations from the 2019 burst, including the measured tail luminosity (Casten et al., 2023), for $d = 3.3 \text{ kpc}$. Note that the observations have the pre-burst (accretion) emission subtracted. We also indicate (*horizontal dashed line*) the (Newtonian) accretion luminosity. 53

Figure 3.7 Duration t_{tail} (*top*) and total radiated energy E_{tail} (*bottom*) of the burst tail as functions of the metallicity Z . For comparison, we also show (*dashed lines*) the expected duration of hot CNO burning, t_d , Eq. (3.4), and $L_{\text{CNO}} \times t_d$, where we compute L_{CNO} using Eq. (3.5) and the fit for ignition column (Fig. 3.5). . . 54

Figure 3.8 The ratio of L_{peak} to L_{tail} , measured at the start of the tail. For comparison, we provide the metallicity corresponding to L_{tail} for our models along the top. . . 57

CHAPTER 1

INTRODUCTION

1.1 Neutron Stars

The concept of a neutron star emerged in the early 1930's as physicists started to speculate how matter might behave at extremely high densities. Even before the discovery of the neutron, Landau (1932) had speculated that stars might reach densities comparable to those of atomic nuclei, which form objects that resemble one "giant nucleus." This set the stage for Baade & Zwicky (1934) to propose that a supernova could mark the transition of a star into an extremely compact object composed primarily of neutrons. Although speculative at the time, their hypothesis introduced stellar densities exceeding those of atomic nuclei, which raised immediate questions about the structure and stability of such objects.

Early theoretical work quickly established that the equilibrium of these ultra-compact stars must be treated within the framework of general relativity. Tolman (1939) and Oppenheimer & Volkoff (1939) derived the relativistic equations governing hydrostatic equilibrium, now commonly known as the Tolman-Oppenheimer-Volkoff (TOV) equations. Analogous to the established treatment of electron-degenerate white dwarfs, Oppenheimer and Volkoff applied these equations to an idealized degenerate neutron gas (i.e., neglecting nucleon-nucleon interactions) and found a maximum stable mass of approximately $0.7M_{\odot}$.

Subsequent advances in nuclear physics significantly revised this early picture. The original TOV calculation neglected the effects of strong nuclear interactions between nucleons, which become increasingly important at supranuclear densities. The inclusion of nuclear interactions, particularly the short-range repulsive component of the strong force, stiffens the equation of state of dense matter and allows neutron stars to support substantially larger masses than predicted by the ideal neutron gas model. Early studies of dense nuclear matter, (e.g., Cameron, 1959; Wheeler, 1966), recognized that these interactions would play a crucial role in determining the maximum mass and internal structure of neutron stars. Current theoretical calculations and nuclear experiments place the maximum mass in the range $2.1\text{--}2.3M_{\odot}$ (e.g., Holt & Lim, 2019; Chatziioannou

et al., 2025). Observational constraints have yielded precise measurements of neutron star masses near or just above $2M_{\odot}$ (Demorest et al., 2010; Antoniadis et al., 2013; Linares et al., 2018).

The first observational confirmation that neutron stars existed, however, did not occur until more than three decades after their theoretical prediction. The earliest hint came from the discovery of galactic X-ray sources. Giacconi et al. (1962) reported the detection of Scorpius X-1, the first extrasolar X-ray source and the brightest persistent X-ray source in the sky. Although its nature was not immediately understood, Sco X-1 was later recognized as an accreting neutron star in a binary system.

In 1967, Jocelyn Bell Burnell and Anthony Hewish discovered rapidly pulsing radio sources, soon interpreted as rapidly rotating, magnetized neutron stars¹ (Hewish et al., 1968). Even before this interpretation was firmly established, Pacini (1967) suggested that a rapidly rotating neutron star with a strong magnetic field could power periodic emission through the loss of rotational energy. Shortly thereafter, Gold (1968) proposed that the observed pulses arise from radiation beamed from the magnetic poles of a rotating neutron star, producing a lighthouse-like effect as the star spins. Together, these ideas provided a compelling theoretical framework linking the newly discovered pulsars to evidence that neutron stars are a common outcome of stellar evolution.

1.2 Neutron Star Formation

After massive stars ($\gtrsim 8 M_{\odot}$) consume the hydrogen in their cores, they begin burning progressively heavier elements up to the iron peak. As the mass of the core grows through continued shell burning, its density increases, and the electrons become increasingly degenerate and eventually relativistic. As the inert core approaches the effective Chandrasekhar mass, the adiabatic index $\gamma \rightarrow 4/3$, at which point the core becomes mechanically unstable. As the collapse begins, the central density and temperature rise rapidly. Photodisintegrations and electron captures, e.g., $^{56}\text{Fe}(\gamma, \alpha)$ and $p + e^{-} \rightarrow n + \nu_e$, increase with density and temperature and further reduce the core's ability to support itself. Photodisintegration absorbs thermal energy from the collapsing core, effectively removing entropy from the star, while electron captures reduce the electron degeneracy

¹The first detected pulsar signal was initially labeled “LGM-1” for “Little Green Men” due to its remarkably regular periodicity, although a natural astrophysical origin (i.e., a rapidly spinning neutron star) was quickly established.

pressure that had previously helped stabilize the core. Together, these reactions further reduce γ below $4/3$, and the collapse runs away. The collapse proceeds on a dynamical (free-fall) timescale $(G\bar{\rho})^{-1/2} \lesssim 10^{-1}$ s and results in the rapid implosion of the core (see, e.g., Janka, 2012; Burrows, 2013).

The collapse continues until the central density exceeds that of nuclear matter, in excess of 10^{14} g cm $^{-3}$. At this density, the strong nuclear force stiffens the pressure, the collapse abruptly halts, and a proto-neutron star is born. The stiffening of the core and rebound of the core launches a shock outwards, which stalls as it moves outward through the iron-rich core, photodissociating nuclei along the way. If neutrinos liberated from the proto-neutron star can revive the shock, it can eventually unbind the star and produce a supernova. The current understanding is that some massive stars explode as core-collapse supernovae, with the collapsed core forming either a neutron star or a black hole, while others fail to produce a successful explosion and collapse directly into black holes. The precise progenitor mass ranges and the relative fractions leading to each outcome remain an active area of research (for a recent review, see Couch 2017).

1.3 Neutron Star Population

Following their birth, neutron stars evolve into a wide range of observable systems depending on their initial properties and subsequent environment. The majority of neutron stars in the Galaxy remain undetected. Young neutron stars are born extremely hot, but cool rapidly with their thermal surface emission declining below current detection thresholds on timescales of 10^5 – 10^6 yr (for a review, see, e.g., Yakovlev & Pethick, 2004; Page et al., 2009). Because older neutron stars are intrinsically faint and small, isolated systems without an additional energy source are difficult to observe. As a result, the detected population is strongly biased toward systems in which emission is sustained by ongoing energy input, such as rotation, magnetic field decay, or accretion. Of the observable neutron stars, several distinct classes have been identified (for a review, see, e.g., Harding, 2013). Pulsars are rapidly rotating neutron stars that emit radiation through the loss of rotational kinetic energy. They produce pulsed emission most prominently in radio but also at X-ray and γ -ray energies. Magnetars are powered by the decay and reconfiguration of their exceptionally strong

($\gtrsim 10^{14}$ G) magnetic fields, which lead to persistent high-energy emission and episodic bursts of X-rays and gamma-rays. In contrast, accretion-powered neutron stars are in binary systems and accrete matter from their companion star. The accretion produces bright X-ray emission both from the accretion itself and from nuclear reactions in the freshly accreted fuel.

1.4 Accreting Neutron Stars

Many observed neutron stars are in binaries and accrete from low-mass companions. Most of these are concentrated toward the center of the galaxy at distances of a few to 20 kpc (Sazonov et al., 2020). To understand accretion in such a system, consider a reference frame that rotates with the binary. In this co-rotating frame, the effective gravitational potential defines a set of equipotential surfaces. In between these surfaces is a saddle point, the inner Lagrange point (labeled L_1 in Fig. 1.1). This critical equipotential surface (Fig. 1.1, dark contour) forms two teardrop-shaped regions, called Roche lobes, which enclose each star and meet at L_1 .




Figure 1.1 Roche geometry for a neutron star (dot at left in the top panel) in a binary with a sun-like star (red sphere at right in the top panel) in a co-rotating reference frame. Equipotential contours are shown in the lower panel, along with the center of mass, “CM,” and the five Lagrange points L_i . The equipotential contour passing through L_1 is shown by the dark contour at bottom and the shaded surface at top. (Figure courtesy E. Brown, license CC BY-NC-SA 4.0).

When the companion overflows its Roche Lobe, matter crosses the inner Lagrange point and is

no longer gravitationally bound to the secondary but instead moves toward the primary. However, because the transferred material retains significant angular momentum, it does not fall directly onto the primary but instead forms an accretion disk. For accretion to proceed, angular momentum must be transported outward through the disk, allowing matter to spiral inward. This process requires dissipation within the disk, such as turbulent stresses driven by the magnetorotational instability, so that gravitational potential energy can be converted into heat and radiated away.

Accretion generates persistent emission as gravitational potential energy is released and radiated away, primarily in the X-ray:

$$L_{\text{acc}} \approx \frac{GM\dot{M}}{R} = \eta\dot{M}c^2, \quad (1.1)$$

where \dot{M} is the accretion rate and $\eta \approx 0.1$ is the accretion efficiency for a neutron star. The Eddington limit, defined by the balance between outward radiation pressure and inward gravitational force, sets an upper bound on the luminosity of a star,

$$L_{\text{Edd}} = \frac{4\pi GMm_{\text{u}}c}{Y_e\sigma_{\text{T}}} \approx 2 \times 10^{38} \left(\frac{M}{1.4 M_{\odot}} \right) \text{ erg s}^{-1}, \quad (1.2)$$

where σ_{T} is the Thomson cross section and Y_e is the electron abundance; $Y_e \approx 0.85$ for an ionized gas of solar composition.

As accretion proceeds, material accumulates on the neutron star surface, forming a layer that grows in mass until the conditions for thermonuclear ignition are reached. In hydrostatic equilibrium, the pressure gradient is determined by the TOV equation,

$$\frac{dP}{dr} = -\frac{Gm(r)}{r^2}\rho(r) \left[1 + \frac{P(r)}{\rho(r)c^2} \right] \left[1 + \frac{4\pi r^3 P(r)}{m(r)c^2} \right] \left[1 - \frac{2Gm(r)}{rc^2} \right]^{-1}. \quad (1.3)$$

In this form (Thorne, 1977), the TOV equation resembles the Newtonian expression for hydrostatic balance, with three relativistic correction factors in $[\cdot]$. Here $r = [\text{surface area}/(4\pi)]^{1/2}$ is the Schwarzschild radial coordinate, and ρ and P are the mass density and pressure. In the thin outer envelope, the radial extent of the accreted layer is much smaller than R , the neutron star radius, and the mass enclosed in r is $m(r) \approx M$, the total mass. Further, the first two relativistic correction factors are roughly unity; as a result, Eq. (1.3) can be cast into Newtonian form,

$$\frac{dP}{dr'} = -g\rho, \quad (1.4)$$




Figure 1.2 Schematic of an accreting neutron star. Exterior to the core is a kilometer-thick crust composed of nuclei, electrons, and (at depth) free neutrons. Accreted material is deposited at the centimeter-deep *photosphere*, the visible “surface.” Below the photosphere is a meters-thick *envelope*, where the matter (mostly H and He) is ionized, non-degenerate, and opaque. As matter descends deeper into the star, nuclear reactions burn the H/He and other light elements to heavier nuclei. The electrons become degenerate and Coulomb interactions force the nuclei into a highly conductive, liquid state, or *ocean*, and eventually lock the nuclei into a rigid lattice—the *crust*. With increasing density, electron captures drive the composition more neutron-rich, and at $\sim 10^{-3}$ of nuclear density, neutrons “drip” from nuclei. When the spacing between nuclei becomes of order their size, nuclear interactions may distort the nuclei into sheets and rods, a phase known as *nuclear pasta*.

where $dr' = dr [1 - 2Gm(r)/(rc^2)]^{-1/2}$ is the proper radial length and

$$g = \frac{GM}{R^2} \left[1 - \frac{2GM}{Rc^2} \right]^{-1/2} = \frac{GM}{R^2} (1+z) \quad (1.5)$$

is the surface gravitational acceleration with $1+z$ being the gravitational redshift at the neutron star surface. To check that this layer is indeed thin, we note that the pressure scale height is

$$H_P = \frac{P}{\rho g} \approx \frac{k_B T}{\mu m_u g} = 140 \text{ cm} \left(\frac{T}{10^8 \text{ K}} \right) \left(\frac{\mu}{0.6} \right)^{-1} \left(\frac{g}{10^{14} \text{ cm s}^{-2}} \right)^{-1} \ll R. \quad (1.6)$$

Here we've taken T and g at representative scales and set the mean molecular weight $\mu = 0.6$ (m_u is the atomic mass unit), as appropriate for a solar composition.

It is convenient to introduce the column depth y —roughly the mass per unit area exterior to r —as a radial coordinate:

$$y(r) = \int_r^\infty \rho(r) dr \approx \frac{P}{g}, \quad (1.7)$$

with the last term following from Eq. (1.4). At the photosphere, $y \approx \kappa^{-1} \sim 1 \text{ g cm}^{-2}$, where κ is the opacity (dominated by electron scattering at $T \gtrsim 10^7 \text{ K}$).

Continued accretion causes the density and temperature of the accreted matter to rise as it is buried deeper into the star. At only a few meters (cf. Eq. [1.6]) below the photosphere—at $y \sim 10^7 \text{ g cm}^{-2}$ for hydrogen and $y \sim 10^8 \text{ g cm}^{-2}$ for helium—conditions become favorable for thermonuclear fusion, and the envelope is locally heated by nuclear reactions (see Fig. 1.2).

Unstable thermonuclear ignition arises from a thin-shell instability (Hansen & Van Horn, 1975), in which fuel accumulates until an upward fluctuation in temperature causes the local heating rate from nuclear reactions to exceed the local cooling rate from thermal conduction. In a geometrically thin burning layer, thermal expansion does not lead to efficient cooling as in stellar cores. As a result, local heating drives a further rise in temperature, and the strong temperature dependence of unstable nuclear burning regimes leads to a thermonuclear runaway. This instability is further enhanced by partial electron degeneracy, which suppresses expansion and prevents the envelope from cooling efficiently. This causes a thermonuclear runaway called an X-ray burst. Typical burst energies are 10^{39} – 10^{40} erg and last for 10–100 s with a recurrence time of hours to days (Fujimoto

et al., 1981; Bildsten & Cumming, 1998; Strohmayer & Bildsten, 2006; Galloway & Keek, 2021). The two nuclear processes that are chiefly responsible for triggering the instability are the triple- α reaction and the CNO cycle (Fowler & Hoyle, 1965).

1.4.1 Observational Properties of X-ray Bursts

Type I X-ray bursts² were first identified in 1975 from the ultra-compact binary 3A 1820–30, independently detected by the *Small Astronomy Satellite 3 (SAS-3)* and the *Astronomical Netherlands Satellite (ANS)* (Grindlay et al., 1976; Belian et al., 1976). Within a year, additional sources were reported along the Galactic plane, many associated with globular clusters (Lewin, 1977). Today, type I bursts remain the most frequently observed thermonuclear explosions in the Galaxy, with over 120 confirmed Galactic sources and over 7,000 bursts (Galloway et al., 2020).

Most bursting systems have short orbital periods in the range 0.2–15 hr (Campana & Di Salvo, 2018). As a steady flow of material from an accretion disk falls on the neutron star surface, the persistent X-ray flux between bursts, F_{per} , provides an estimate of the accretion rate via

$$\dot{M} = 1.33 \times 10^{-11} M_{\odot} \text{ yr}^{-1} \left(\frac{F_{\text{per}} C_{\text{bol}}}{10^{-9} \text{ erg cm}^{-2} \text{ s}^{-1}} \right) \left(\frac{D}{10 \text{ kpc}} \right)^2 \left(\frac{M}{1.4 M_{\odot}} \right)^{-1} \left(\frac{1+z}{1.31} \right) \left(\frac{R}{10 \text{ km}} \right), \quad (1.8)$$

where C_{bol} is the bolometric correction to the measured flux and D is the distance (Galloway et al., 2008). The Eddington luminosity, Eq. (1.2), sets an upper bound on the accretion rate via Eq. (1.1), which corresponds to $\dot{M}_{\text{Edd}} \approx 2 \times 10^{-8} M_{\odot} \text{ yr}^{-1}$ for a $1.4 M_{\odot}$ neutron star.

In addition to the persistent flux, burst properties such as rise time, duration, fluence (total amount of energy or number of particles received per unit area over a period of time), and recurrence time provide key observational diagnostics of the underlying fuel composition and ignition conditions. He-triggered bursts that are triggered in a hydrogen-rich layer exhibit longer decay tails due to continued energy release from the rp-process, whereas He-dominated bursts are shorter and more symmetric.

²Observations of the “Rapid Burster” MXB 1730–335 soon revealed two distinct classes of bursts: thermonuclear type I bursts and accretion-driven type II bursts (Hoffman et al., 1978). Type I bursts are characterized by a rapid increase in X-ray luminosity followed by an exponential-like decay, reflecting the heating and cooling of the neutron star envelope, whereas type II bursts exhibit abrupt onsets and terminations without gradual decay, often recurring on minute timescales.

Comparing the persistent accretion luminosity radiated by the accumulated fuel between bursts to the energy released from burning the accumulated fuel,

$$\alpha = \frac{L_{\text{acc}}\Delta t}{E_{\text{burst}}} \approx \frac{GM}{RQ_{\text{nuc}}}, \quad (1.9)$$

where Δt is the recurrence time between bursts and Q_{nuc} is the heat released by nuclear reactions per unit mass, provides a useful diagnostic of nuclear burning. For pure helium bursts, $\alpha \approx 100\text{--}200$ and for mixed hydrogen/helium burning, $\alpha \approx 30\text{--}50$. These values can be understood by comparing the energy released per nucleon: accretion onto a neutron star releases ≈ 200 MeV per nucleon, while nuclear burning releases ≈ 1.6 MeV per nucleon for helium burning and ≈ 7 MeV per nucleon for hydrogen burning. As a result, α is smaller when hydrogen is present, since more nuclear energy is released per accreted nucleon.

Despite the large sample of known bursters, events occurring at the lowest accretion rates remain comparatively rare in observational catalogs because of their lower burst energetics. In fact, simulations have been unable to produce observable H-triggered bursts under solar composition accretion, and until recently, no such events had been detected observationally. This led to speculation that they might be unobservable, (see, e.g., Cooper & Narayan, 2007). As a result, the observational landscape of weak bursts—particularly those powered primarily by unstable hydrogen burning—remains incompletely explored and is the focus of this dissertation.

1.4.2 Burning Regimes of X-ray Bursts

There are a few main factors that determine what type of X-ray burst occurs: the accretion rate, the composition of the accreted material, and the neutron star surface temperature. The first of these factors, the accretion rate \dot{M} , has been studied extensively in the context of burst behavior (see, e.g., Taam & Picklum, 1979; Fujimoto et al., 1981; Cumming, 2004). The behavior has been mapped out using a mixture of semi-analytical calculations, in which the heating and cooling timescales were compared to determine stability, and single- and multi-zone models of the envelope. For a thorough review of these estimates, see Galloway & Keek (2021). Roughly speaking, for $\dot{M} \gtrsim 0.01 \dot{M}_{\text{Edd}}$, the β -limited hot CNO cycle stably burns H to He (Hoyle & Fowler, 1965).

In this cycle, the rate at which H is consumed is set by the temperature-insensitive β -decay lifetimes of ^{14}O and ^{15}O , and so the timescale over which H is depleted is solely set by the ratio of H to CNO. As a layer of He builds, it ignites via the triple- α process before H is exhausted. This results in a mixed H/He burst, in which the protons are captured onto seed nuclei formed from α -captures (Schatz et al., 2001). The shape of the lightcurve, characterized by a slower rise time ($\approx 5\text{--}10$ s), is dictated by the H/He ratio (Fisker et al., 2008), while the long decaying tail (\approx minutes) is sensitive to the efficiency of nuclear burning at various waiting points in the rp-process, as explored by Woosley et al. (2004) using multi-zone models. Additionally, the metallicity of the accreted material plays a key role in shaping burst properties, influencing recurrence times, peak luminosities, and nucleosynthesis pathways (José et al., 2010). When $\dot{M} \lesssim 0.1 \dot{M}_{\text{Edd}}$ and the surface temperature is $\gtrsim 8 \times 10^7$ K, there is sufficient time for the hot CNO cycle to completely consume the available H, so that the triple- α process ignites in a nearly pure He envelope. In contrast to the mixed H/He bursts, pure-He bursts are more intense, often Eddington-limited—as exhibited by photospheric radius expansion—and shorter in duration.

At $\dot{M} \lesssim 0.01 \dot{M}_{\text{Edd}}$, the envelope temperature may be cool enough, $\lesssim 8 \times 10^7$ K, that the CNO cycle is no longer β -limited and therefore becomes thermally unstable. If the unstable H burning heats the envelope sufficiently, He will also ignite and produce a mixed H/He burst. At even lower accretion rates, the heating from unstable H burning may not be sufficient to ignite He. Under these conditions, a weak burst powered solely by H burning would result. Although there have been several theoretical studies of burning at low \dot{M} (see, e.g., Cumming, 2004; Peng et al., 2007; Cooper & Narayan, 2007), observational examples of H-triggered bursts have been elusive. This dissertation presents two candidate H-triggered bursts and highlights remaining gaps in our understanding of the conditions required to produce observable H-triggered bursts in neutron star envelopes.

1.5 Dissertation Overview

Casten et al. (2023) reported a candidate H-triggered burst in 2019 from the accreting millisecond pulsar SAX J1808.4–3658 using observations from the *Neutron Star Interior Composition Explorer* (NICER). An archival search conducted in the same work identified a similar burst from

SAX J1808.4–3658 observed with the *Rossi X-ray Timing Explorer* (RXTE) in 2005. These two events are analyzed in detail in Chapter 2, where their observational properties and consistency with hydrogen-triggered ignition are examined.

At low accretion rates ($\dot{M} \lesssim 10^{-2} \dot{M}_{\text{Edd}}$), sedimentation of heavy elements in the neutron star envelope (Wallace et al., 1982; Bildsten et al., 1992; Peng et al., 2007) may enhance the CNO abundance at the base of the accreted layer and thereby strengthen hydrogen burning. Enhanced CNO abundances are known to drive vigorous hydrogen burning in classical nova models (Starfield, 1993; Denissenkov et al., 2013), although the mechanisms responsible for such enrichment in neutron star envelopes remain uncertain. To explore the impact of enhanced CNO abundances on hydrogen ignition, in Chapter 3 we use the open-source stellar evolution code MESA (Paxton et al., 2011; Jermyn et al., 2023) to investigate unstable hydrogen ignition at low accretion rates across a range of accreted metallicities. These models are used to assess whether CNO enhancement can produce bursts consistent with the observed candidate H-triggered events.

Results presented in Chapters 2 and 3 are based on work previously published in Casten et al. (2023, 2025), respectively. Figures in these chapters are taken directly from those works.

CHAPTER 2

FIRST OBSERVATIONS OF H-TRIGGERED BURSTS

In the case of the lowest accretion rates discussed previously (§1.4.2), temperatures in the accumulating shell may be low enough, $\lesssim 8 \times 10^7$ K, that steady hot-CNO hydrogen-burning essentially switches off, but instead can proceed in the unstable, temperature-sensitive cold-CNO regime. As material accumulates and is compressed, hydrogen will be consumed causing a rise in local temperature and driving a thermonuclear runaway once the local temperature exceeds the rate the envelope can cool. If the resulting increase in temperature is insufficient to ignite helium, the event produces a weak burst powered solely by hydrogen burning, which is referred to as a H-triggered burst.

Although there have been several theoretical studies of burning at low \dot{M} (see, e.g., Cumming, 2004; Peng et al., 2007; Cooper & Narayan, 2007), observational examples of H-triggered bursts have been elusive. There have been few published reports of X-ray bursts that can be clearly attributed to unstable hydrogen shell flashes. In one such case, Boirin et al. (2007) reported on the first observations of triple, short recurrence time (SRT) bursts from the high inclination, eclipsing source EXO 0748–676. They suggested that the initial bursts of singles, pairs, or triples (they call these the long waiting time, LWT, bursts), could be attributed to either helium-triggered, mixed H/He bursts at moderate accretion rates (10% of Eddington), or perhaps hydrogen-triggered bursts at lower accretion rates (1% of Eddington). Because the LWT bursts appeared somewhat underluminous compared with mixed H/He bursts in 1-D KEPLER models and the well-known example of such bursts from GS 1826–238, they suggested that this might be explained by the latter, hydrogen-triggered mechanism.

They further suggested that the SRT events, with waiting times close to 10–12 minutes, were likely caused by the re-ignition of unburned fuel, but they did not have a detailed explanation of how this occurs. More recently, Keek & Heger (2017) have outlined a theoretical mechanism to account for SRT bursts. Using detailed, 1-D KEPLER hydrodynamic simulations they showed that such events can be produced by opacity-driven convective mixing that transports fresh fuel to the ignition depth, and they also argued that this mechanism can produce simulated burst events that are

“strikingly similar” (in their words) to the SRT bursts seen from EXO 0748–676. If this mechanism is indeed at work, then it would further argue for the higher accretion rate (10% of Eddington), helium-triggered scenario in EXO 0748–676, as warmer envelopes, naturally produced by higher accretion rates, were required to produce the SRT events in their burst simulations. Moreover, they also showed that the fraction of fuel burned in the LWT events dropped as the envelope became hotter, and this relatively low fuel burning fraction could also naturally explain the apparently under luminous LWT bursts noted by Boirin et al. (2007). Thus, while Boirin et al. (2007) suggest that a hydrogen-triggered mechanism is possible for the LWT bursts from EXO 0748–676, we would characterize the current, overall evidence in support of that conclusion as tentative, particularly given the remaining uncertainties in the distance and anisotropy factors for this source. Indeed, in support of this we note that in their recent review of the field, Galloway & Keek (2021) also comment that, “No observations matching case I or case II bursting have been identified.” Here, cases I and II refer to the two hydrogen ignition paths at low accretion rates that were sketched above.

In this chapter we present a study of an apparently rarer class of weak X-ray bursts observed from SAX J1808.4–3658 that we argue show the hallmarks of being associated with the hydrogen ignition regime. This object was the first accreting millisecond X-ray pulsar discovered (Wijnands & van der Klis, 1998; Chakrabarty & Morgan, 1998), and hosts a neutron star in a 2.1 hr orbit with a low-mass brown dwarf (Bildsten & Chakrabarty, 2001). Its distance has been estimated at 3.5 ± 0.1 kpc (Galloway & Cumming, 2006), and it is likely that the donor provides a hydrogen-rich mix of matter to the neutron star during outbursts (Galloway & Cumming, 2006; Goodwin et al., 2019). To date, SAX J1808.4–3658 has been observed extensively during ten outbursts. While it is not our intention here to provide a broad observational overview of the source—for the purposes of this chapter we focus on issues relevant to its thermonuclear bursting behavior—readers can find elsewhere some recent studies on coherent pulse timing (Sanna et al., 2017; Bult et al., 2020; Illiano et al., 2022), X-ray spectral properties (Di Salvo et al., 2019), and aperiodic timing behavior (Bult & van der Klis, 2015; Sharma et al., 2022).

Observations of SAX J1808.4–3658 have revealed two types of thermonuclear bursts that show dramatically different peak fluxes and fluences. The bright photospheric radius expansion bursts show significantly higher total energy release and peak X-ray flux. The less frequently observed weak bursts produce much less energy and show peak fluxes about a factor of 25 less than the bright events; as such they are not Eddington-limited. When these weak bursts have been observed, they appear to be confined to earlier portions of the outbursts and occurred before the bright bursts were seen. This suggests there may be a window of occurrence for these bursts associated with the initial onset of accretion after a period of quiescence. This is particularly intriguing in the context of SAX J1808.4–3658 because it is known that the neutron star cools dramatically in quiescence (Heinke et al., 2009), and the unstable hydrogen-burning regime requires cooler temperatures in the accumulating layer. There has been more observational and theoretical research exploring the nature of the bright bursts than the weak class.

Here we present a detailed study of one of these weak bursts that was observed with NICER during the recent, 2019 August outburst from SAX J1808.4–3658. We also provide a briefer description of a similar burst observed with the RXTE in 2005 June. The chapter is organized as follows. In § 2.1 we introduce the NICER data and present lightcurves focusing on the initial part of the 2019 outburst, showing a single weak burst. We also present a spectral study of the persistent and burst emission (for the weak burst) in order to understand its energetics and to constrain the mass accretion rate and the likely accreted mass column at the time of its ignition. We present a discussion in § 2.2 of a likely physical scenario that results in the weak burst, arguing that the initial accretion onto a cool neutron star at the onset of the outburst naturally places the accumulating layer in the thermally unstable regime for CNO hydrogen ignition. Here, we also describe the 2005 June RXTE event, and we also report a brief summary of NuSTAR observations that began on 2019 August 10 and in which several brighter bursts were detected. We conclude in § 2.3 with a summary, a brief discussion of relevant uncertainties and other possible interpretations, and the outlook for future efforts.

2.1 NICER Observations of SAX J1808.4–3658

In late July 2019, it was reported that the optical flux from SAX J1808.4–3658 had increased, perhaps presaging a new X-ray outburst (Russell et al., 2019; Goodwin et al., 2020). This initiated an extensive monitoring campaign with NICER, which began on August 1, 2019 (Bult et al., 2020). NICER is an X-ray observatory that operates on the International Space Station (ISS). It observes across the 0.2–12 keV X-ray band and provides low-background, high-throughput ($\approx 1900 \text{ cm}^2$ at 1.5 keV), and high time resolution capabilities (Gendreau et al., 2012). The data obtained prior to the onset of the outburst, and up to and including the first observed X-ray burst are organized under observation IDs (ObsIDs), 205026010*mm*, and 25840101*nn*, where *mm* and *nn* run from 03–10 and 01–02, respectively. We used the standard screening criteria and NICERDAS version 8 to produce cleaned event lists. This means we retained only those epochs during which the pointing offset was $< 54''$, the Earth elevation angle was $> 15^\circ$, the elevation angle with respect to the bright Earth limb was $> 30^\circ$, and the instrument was not in the South Atlantic Anomaly (SAA). We used HEASOFT Version 6.29c to produce the light curves and spectra for the analyses reported here. The initial observations of the campaign did not reveal evidence of SAX J1808.4–3658 in X-ray outburst. The first indication that an accretion-driven flux was present occurred on August 6, 2019 at approximately 21:59 TT (Bult et al., 2019). Figure 2.1 shows the light curve (0.4–7 keV) of the outburst over approximately 20 days from the observed onset of significant X-ray activity. Time zero in the plot refers to the time of outburst onset, 58701.91597 MJD (TT). The two detected X-ray bursts are evident as “spikes” in the count rate near days 3 and 14, respectively. The much brighter second burst (near day 14) was reported on by Bult et al. (2019). Here we focus on a study of the much weaker first burst, which occurred at 58704.80764 MJD (TT), and is present in ObsID 2584010102.

2.1.1 Persistent Spectrum, Fluence, and Accreted Mass

To explore the weak burst energetics and ignition conditions we aim to constrain the total accreted mass from the beginning of the outburst up to the onset of the first burst. To do this we model the spectrum of the persistent emission to determine its flux and then integrate from the outburst




Figure 2.1 Light curve of the 2019 July–August outburst of in the 0.4–7 keV band observed with NICER. Count rates were computed in 2 s intervals. Note the logarithmic scale. The first, weaker burst is evident near day 3. The bright burst near day 14 was reported on by Bult et al. (2019).

onset to just prior to the burst. This integral provides an estimate of the energy fluence produced via accretion, which can then be converted to an accreted mass using standard assumptions for the accretion luminosity produced by the release of gravitational potential energy of the accreted matter.

In practice we find that the shape of the persistent spectrum gradually changes during this portion of the outburst, with the spectrum showing a modest hardening over time. We therefore measure the flux at a few intervals along the outburst rise, and use these measurements to estimate the flux per unit NICER count rate. We then use simple linear interpolation and the trapezoidal rule to

Table 2.1 Spectral Model Parameters for SAX J1808.4–3658: Persistent Emission

Parameter ^a	unit	ObsID		
		2584010102 (Pre-burst)	2050260110	2050260109 ^b
phabs				
n_H	10^{22} cm^{-2}	0.131 ± 0.016	0.160 ± 0.025	0.177 ± 0.028
diskbb				
$k_B T_{\text{in}}$	keV	0.849 ± 0.026	0.679 ± 0.024	0.528 ± 0.017
Norm ^c	$(R_{\text{km}}/D_{10})^2$	27.33 ± 5.36	25.19 ± 6.46	10.82 ± 1.79
bbodyrad				
$k_B T$	keV	2.03 ± 0.15	1.67 ± 0.15	–
Norm ^c	$(R_{\text{km}}/D_{10})^2$	0.98 ± 0.47	0.604 ± 0.336	–
powerlaw				
Index		1.96 ± 0.37	2.50 ± 1.05	2.01 ± 0.23
Norm	$10^{-3} \text{ keV}^{-1} \text{ cm}^{-2}$	26.6 ± 13.7	4.64	4.23 ± 1.6
gauss				
E	keV	0.990 ± 0.009	0.935 ± 0.014	–
σ_E	keV	0.015	0.015	–
Norm	$10^{-4} \text{ cm}^{-2} \text{ s}^{-1}$	8.44 ± 2.03	3.51 ± 1.00	–
$f_{0.1-20}$	$10^{-10} \text{ erg cm}^{-2} \text{ s}^{-1}$	7.06 ± 0.23	2.08 ± 0.40	0.535 ± 0.081
f_{bol}	$10^{-10} \text{ erg cm}^{-2} \text{ s}^{-1}$	7.9	2.5	0.64
χ^2/dof		117.9/112	98.9/97	94.5/106
Rate, 0.5–10 keV	s^{-1}	181.9 ± 0.4	58.70 ± 0.25	13.46 ± 0.07
Epoch ^d	d	2.89	1.45	0.59
Exposure	ks	0.740	0.927	2.807

^a Parameter uncertainties are 1σ .

^b For ObsID 2050260109, the bbodyrad and Gaussian components were not required in the fit (indicated with “–”).

^c Here R_{km}/D_{10} refers to the radius of the emitting area (or inner disk radius) in km divided by the source distance in units of 10 kpc.

^d Epoch refers to the center time of the spectral interval and corresponds to the time axis in Figure 2.1.

integrate the flux from outburst onset to the first burst to estimate the energy fluence.

The lightcurve in Figure 2.2 shows a close-up of the epoch around the first burst. We extracted a spectrum prior to the burst, the “pre-burst” interval (marked by the vertical dashed lines in Figure 2.2) and modeled its spectrum using XSPEC version 12.12.1 (Arnaud, 1996). We produced re-

sponse files with NICERDAS version 8, and we used the 3C50 background model, `nibackgen3c50` (Remillard et al., 2022), to produce a background spectrum appropriate for spectral modeling within XSPEC. We employed a phenomenological model similar to that discussed by Patruno et al. (2009), that includes thermal disk, power-law, and blackbody continuum components. In addition, and similarly to Bult et al. (2019), we find evidence for narrow-line emission near 1 keV, and we include a gaussian component to model this. In XSPEC notation the model has the form, `phabs*(diskbb + powerlaw + bbodyrad + gauss)`, where `phabs` represents the line of sight photoelectric absorption model parameterized by the column density of neutral hydrogen, n_H . This absorption model uses cross sections from Verner et al. (1996) and the chemical abundances from Anders & Grevesse (1989). We fit this model across the 0.5–10 keV bandpass and find that it provides an excellent fit, with a minimum $\chi^2 = 117.9$ for 112 degrees of freedom. The best-fitting model parameters are given in Table 2.1, and Figure 2.3 shows the unfolded photon spectrum (top), the observed count-rate spectrum and best-fitting model (middle), and the fit residuals in units of standard deviations (bottom). This model gives an unabsorbed flux (0.1–20 keV) of $7.06 \pm 0.23 \times 10^{-10} \text{ erg cm}^{-2} \text{ s}^{-1}$. If we extend the bandpass to estimate a bolometric flux we find a value (0.1–100 keV) of $7.9 \times 10^{-10} \text{ erg cm}^{-2} \text{ s}^{-1}$. The average count rate in the fitted energy band (0.5–10 keV) is $181.9 \pm 0.5 \text{ s}^{-1}$, so we estimate a flux per NICER count rate (0.5–10 keV) of $4.34 \times 10^{-12} \text{ erg cm}^{-2} \text{ s}^{-1} (\text{counts s}^{-1})^{-1}$ for this interval.

We extracted spectra from two other ObsIDs along the outburst rise, 2050260109 and 2050260110, and analyzed these spectra in the same manner as for the pre-burst interval just described. Results of these spectral fits are also reported in Table 2.1. For these intervals we estimate flux per NICER count rate values of 2.85×10^{-12} and $3.61 \times 10^{-12} \text{ erg cm}^{-2} \text{ s}^{-1} (\text{counts s}^{-1})^{-1}$, respectively.

For completeness we make a few additional comments regarding the 1 keV line component included in the spectral model. For the pre-burst interval (ObsID 2584010102), removing the gaussian line results in an increase in χ^2 of 31.3, and the ratio of the line normalization to its 1σ uncertainty is 4.2. The line is also evident in ObsID 2050260110, though at lower significance, with the ratio of the line normalization to its 1σ uncertainty now at 3.5. For ObsID 2050260109, the spectrum




Figure 2.2 Lightcurve of the first, weak burst from SAX J1808.4–3658 in the 0.4–7 keV band observed with NICER. Main panel: The count rates were computed in 1 s intervals, and the vertical dashed and dotted lines denote the intervals used to extract spectra for the pre- and post-burst spectral modeling, respectively. Inset panel: The same data are used, but the time bins are 16 s, and the logarithmic scale highlights the offset in count rate between the pre- and post-burst emission. The dashed red line is a constant value fit to the pre-burst level, and is meant as a guide to the eye.




Figure 2.3 Energy spectrum of the pre-burst interval, modeled with a phenomenological model similar to that employed by Patruno et al. (2009), which includes `diskbb`, `bbodyrad` and `powerlaw` components, in addition to the line at 1 keV. See the text in § 2.1.1 for further details.

extracted closest to the outburst onset and at the lowest observed flux, we no longer find evidence for the line. When detected the line is narrow in the sense that it is unresolved and we can only place an upper limit on its width of 0.09 keV (3σ). Finally, in this work our primary focus is to model the X-ray spectrum to infer the broadband flux. Excluding the 1 keV line from the spectral fits only changes the inferred flux at the few percent level, so including it, or not, does not significantly alter our inferences regarding the source flux. We elected to include it since doing so provides a better overall statistical description of the data.

To estimate the outburst fluence we use simple linear interpolation between data gaps, and we also apply linear interpolation of the flux per unit count rates, based on the spectral results discussed above. We employ the trapezoidal rule to integrate the total counts. We find a persistent

emission energy fluence of $E_{\text{bol}} = 7.92 \times 10^{-5} \text{ erg cm}^{-2}$, representing an estimate of the total energy associated with accretion from the outburst onset up to the initiation of the first observed burst.

Assuming the observed, accretion-driven luminosity for spherical accretion,

$$L_{\text{bol}} = 4\pi d^2 \xi_{\text{bol}} f_{\text{bol}} = \frac{z \dot{M} c^2}{(1+z)^3}, \quad (2.1)$$

where $(1+z) = (1 - 2GM/c^2R)^{-1/2}$ is the surface redshift, \dot{M} , ξ_{bol} is the persistent emission anisotropy factor, and f_{bol} is the observed bolometric flux. We can use this equation to estimate the total accreted mass required to produce the observed energy fluence (Johnston et al., 2020). We emphasize that L_{bol} and f_{bol} are the luminosity and flux as measured by an observer far from the neutron star. The anisotropy factor, ξ_{bol} , can be thought of as the solid angle into which the radiation is emitted, normalized by 4π , thus, isotropic emission is characterized by $\xi_{\text{bol}} = 1$. We write the accreted mass column in the local neutron star frame as,

$$y_a = \frac{M_a}{4\pi R^2} = \frac{\int \dot{M}(t') dt'}{4\pi R^2}, \quad (2.2)$$

where we use t' to emphasize that the \dot{M} integral is over the time as measured at the neutron star surface. With the use of equation (2.1) this becomes,

$$y_a = \frac{d^2 \xi_{\text{bol}} (1+z)^2}{z c^2 R^2} \int f_{\text{bol}}(t) (1+z) dt' = \frac{d^2 \xi_{\text{bol}} (1+z)^2}{z c^2 R^2} E_{\text{bol}}, \quad (2.3)$$

where t is the time measured in the observer's frame, and we note that $dt = (1+z) dt'$, and thus $\int f_{\text{bol}}(t) (1+z) dt' = \int f_{\text{bol}}(t) dt = E_{\text{bol}}$ is just the observed energy fluence.

If we assume $d = 3.5 \text{ kpc}$ (Galloway & Cumming, 2006), and use $E_{\text{bol}} = 7.92 \times 10^{-5} \text{ erg cm}^{-2}$, we can write the accreted column as,

$$y_a = 1.03 \times 10^7 \left[\frac{\xi_{\text{bol}} (1+z)^2}{z R_{10}^2} \right] \text{ g cm}^{-2}, \quad (2.4)$$

where R_{10} is the neutron star radius in units of 10 km. For $M = 1.4 M_{\odot}$, $R = 11 \text{ km}$, and adopting $\xi_{\text{bol}} = 1$ we find $y_a = 5.12 \times 10^7 \text{ g cm}^{-2}$. With $M = 2.0 M_{\odot}$ and $R = 11 \text{ km}$ we find that y_a decreases slightly to $3.91 \times 10^7 \text{ g cm}^{-2}$.

We can also use equation (2.1) to estimate the mass accretion rate, \dot{M} , at the time of the weak burst onset. Using the estimated pre-burst flux of $7.9 \times 10^{-10} \text{ erg cm}^{-2} \text{ s}^{-1}$, and a distance of 3.5 kpc, we find,

$$\dot{M} = 2.03 \times 10^{-11} \left[\frac{\xi_{\text{bol}}(1+z)^3}{z} \right] M_{\odot} \text{ yr}^{-1}. \quad (2.5)$$

Using the same parameter assumptions as above, we find estimates of $\dot{M} = 1.56 \times 10^{-10}$ and $1.38 \times 10^{-10} M_{\odot} \text{ yr}^{-1}$, respectively.

2.1.2 Burst Spectral Evolution: Peak Flux and Fluence

We first segmented the burst light curve into intervals of approximately 500 counts using a 1/8s time bin. We modeled the segmented spectra in the 0.5–10 keV band by adding a blackbody component to the pre-burst persistent emission model. The parameters of the persistent emission model were frozen to their best fit values, given in Table 2.1, only allowing the added blackbody component to vary, so that our model is `phabs(constant(diskbb + bbodyrad + powerlaw + gauss) + bbodyrad)`. We first tried multiplying the persistent emission model by a constant (Worpel et al., 2013), but found it was not statistically necessary, as it was possible to get a good fit with it left frozen at 1.0. The resulting evolution of the bolometric flux, the free parameters of the blackbody temperature and blackbody radius (at 3.5 kpc), along with the resulting χ^2 are shown in Figure 2.4. We found a peak bolometric burst flux of $f_p = 6.98 \pm 0.50 \times 10^{-9} \text{ erg cm}^{-2} \text{ s}^{-1}$. Using trapezoidal numerical integration of the flux, we calculated a bolometric fluence of $7.05 \pm 1.16 \times 10^{-8} \text{ erg cm}^{-2}$. Adopting $\xi_{\text{bol}} = 1$, and with $d = 3.5 \text{ kpc}$, we can then estimate (Eq. [2.1]) that the total energy released during the burst was $1.03 \times 10^{38} \text{ erg}$.

2.2 Physical Scenario and Interpretation

Transient systems like SAX J1808.4–3658 provide an interesting laboratory to explore the different predicted regimes of nuclear burning on neutron stars. Deep X-ray spectroscopic studies of the object in quiescence suggest rapid cooling of the neutron star core, perhaps by a form of enhanced neutrino emission such as Direct Urca (Heinke et al., 2009), which also provides some tentative evidence for a more massive neutron star ($\gtrsim 2 M_{\odot}$) in the system. Using the surface effective temperature constraints in quiescence from Heinke et al. (2009) and the theoretical results




Figure 2.4 Evolution of the weak X-ray burst derived from spectral modeling in the 0.5–10 keV band. We show from the top down: the bolometric flux, blackbody temperature, blackbody radius (at 3.5 kpc), and reduced χ^2 , respectively. The error bars indicate 1- σ confidence intervals.

of Potekhin et al. (1997), Mahmoodifar & Strohmayer (2013) estimated a core temperature for SAX J1808.4–3658 in the range from $7.2\text{--}7.7 \times 10^6$ K for neutron star masses between 1.4 and $2.0 M_{\odot}$. Due to the high thermal conductivity in the core and crust (Brown & Cumming, 2009), it is thus very likely that when accretion begins in SAX J1808.4–3658 after a period of quiescence, the accumulating layer starts out at temperatures $\lesssim 1 \times 10^7$ K, that is, well below the temperature at which CNO cycle hydrogen-burning becomes thermally stable (Fujimoto et al., 1981; Cumming, 2004; Galloway & Keek, 2021). In this temperature-sensitive regime hydrogen-burning proceeds at very low levels, and the thermal profile of the accumulating layer will be set principally by compressional heating (see Bildsten & Cutler, 1995; Brown & Bildsten, 1998, for an estimate of the heating rate). This is a much less efficient heat source than the energy released from hot-CNO cycle burning in the stable burning regime.

Fujimoto et al. (1981) estimate the accretion rate required to maintain a stable hydrogen-burning shell ($\dot{M}_{st}(B)$ in their Table 1) as $2.7 \times 10^{-10} M_{\odot} \text{ yr}^{-1}$, for a neutron star mass and radius of $1.41 M_{\odot}$ and 6.57 km. We note that the somewhat older neutron star models employed by Fujimoto et al. (1981) have rather small radii compared to that suggested by more recent observations (Miller et al., 2019; Riley et al., 2019). For a more typical radius of, say, 11 km (which we employed above), we would expect that the estimated rate would increase modestly by $\approx 10\%$, which would bring the value to $\approx 3 \times 10^{-10} M_{\odot} \text{ yr}^{-1}$. Expressed as a fraction of the Eddington accretion rate, \dot{M}_{Edd} , and adopting the value of $\dot{M}_{\text{Edd}} = 1.8 \times 10^{-8} M_{\odot} \text{ yr}^{-1}$ (Cumming, 2004), this is then equivalent to $\dot{M} = 0.0167 \dot{M}_{\text{Edd}}$. Note also that Cumming (2004) quotes a value of $\dot{M} \gtrsim 0.01 \dot{M}_{\text{Edd}}$ for stable, hot-CNO cycle hydrogen-burning. In addition, Cooper & Narayan (2007) used a two-zone model to carry out a linear stability analysis to specifically explore the conditions under which hydrogen-triggered bursts can occur at low accretion rates, and found that they occur for rates $\lesssim 0.003 \dot{M}_{\text{Edd}}$.

Above, we estimated an accretion rate at the time of the weak NICER burst in the range from $1.38\text{--}1.56 \times 10^{-10} M_{\odot} \text{ yr}^{-1}$ ($\dot{M} \approx 0.0077\text{--}0.0087 \dot{M}_{\text{Edd}}$). This is less than the required rates estimated by Cumming (2004) and Fujimoto et al. (1981) for stable CNO burning, but slightly higher than the accretion rate obtained by Cooper & Narayan (2007). These considerations provide strong

evidence that in the initial outburst stage, accretion onto SAX J1808.4–3658 proceeds in an \dot{M} range consistent with what Fujimoto et al. (1981) refer to as *case 3* shell flashes. In this regime the accumulating layer remains cool enough that CNO hydrogen-burning proceeds in the temperature-sensitive regime, that is, very little hydrogen is burned until the layer reaches the conditions for unstable ignition. Indeed, following Fujimoto et al. (1981, see their equation 11), we would estimate that only about 1–2% of the hydrogen would be burned prior to ignition. Further insights are provided by our estimates of the total column of matter accreted at the time of the burst, and its total energy fluence. In §2 above we estimated the accreted column to be in the range $3.91\text{--}5.12 \times 10^7 \text{ g cm}^{-2}$, and we measured an energy fluence in the burst of $\approx 1 \times 10^{38}$ erg (both at 3.5 kpc). For the following discussion we refer the reader to the illustrative hydrogen ignition curves presented by Cumming (2004, see their Fig. 1), Galloway & Keek (2021, see their Fig. 2), and Peng et al. (2007, see their Fig. 4). Based on these curves, we can estimate that a column of this size would be ignited at a temperature in the range from $4\text{--}5 \times 10^7$ K. What happens upon ignition of the hydrogen? The unstable burning will quickly heat the layer, raising the temperature to at least that at which the CNO energy generation rate saturates, but likely somewhat higher. Fujimoto et al. (1981) estimate that only a small fraction, ΔX , of the hydrogen needs to burn in order to raise the temperature. For a temperature change of 10^8 K, they estimate $\Delta X \approx 0.002$ (see their equation 12).

After ignition of the hydrogen, two subsequent paths have been described in the literature. First, if the ignition column is small enough, then an increase in its temperature may not cause it to cross the helium ignition curve, and additional accretion and/or an increase in the helium fraction is required before it will ignite. Alternatively, for deeper ignition columns, a temperature increase of a few 10^8 K would render the shell unstable to helium ignition, promptly producing a mixed H/He burst. We note that the work of Cooper & Narayan (2007) and Peng et al. (2007) also predict these two paths, and their calculations provide estimates of the hydrogen ignition columns that are broadly consistent with our estimate of the column accreted at the time of the weak burst. For example, at an accretion rate of $\dot{m} = 0.002 \dot{m}_{\text{Edd}}$ Cooper & Narayan (2007, see their Fig. 4, right column) find

behavior consistent with the first scenario, a sequence of weak hydrogen flashes occurs until the helium column grows sufficiently to reach ignition conditions. These calculations also provide an estimate of the temperature increase produced by the unstable hydrogen ignition, and suggest that changes of $\approx 2 \times 10^8$ K are likely. Peng et al. (2007, see their Fig. 7) also find a regime where hydrogen ignition does not lead to prompt ignition of a He burst. They also explore the effect of sedimentation on hydrogen-triggered bursts, which enhances the amount of CNO nuclei at the ignition depth and causes a sharper temperature rise. Sedimentation is likely to play an important role in setting the ignition conditions for the weak burst given the low estimated accretion rate.

Measurement of the burst fluence enables us to estimate the fraction, f_h , of accreted hydrogen needed to burn in order to produce that much energy. For an energy release (per gram) of $E_h = 6.4 \times 10^{18}$ erg g⁻¹ (Clayton, 1983), we would require $m_h = 1.6 \times 10^{19} (1+z)$ g of hydrogen to burn, where the factor of $(1+z)$ is included because we are interested in the energy released at the neutron star surface. Expressed as a column on the neutron star, $y_h = m_h/4\pi R^2$, and assuming $R = 11$ km, we find $y_h = 1.05 \times 10^6 (1+z)$ g cm⁻². The amount of hydrogen present in the accreted column is $y_a X$, where X is the mass fraction of hydrogen in the accreted material. We thus have,

$$f_h = \frac{y_h}{y_a X} = 0.105 \frac{(1+z)}{Y_a X}, \quad (2.6)$$

where Y_a is the estimated accreted column in units of 10^7 g cm⁻². Taking Y_a in the range from 3.9–5.1, a fractional hydrogen abundance in the accreted fuel of $X = 0.7$, and using the same M and R assumptions employed above to evaluate $(1+z)$, we find f_h in the range 0.04–0.06. Note that this value should be considered a lower limit, as it assumes that the estimated total accreted column produced only a single such burst, and the mass fraction would likely be reduced further if sedimentation is present (Peng et al., 2007, see below for additional discussion regarding potentially missed bursts). This value is larger than the estimate given by Fujimoto et al. (1981) for the fraction of hydrogen needed to burn to raise the temperature of the fuel by $\Delta T = 10^8$ K; we don't know the actual temperature increase however, and the estimate of Fujimoto et al. (1981) should be thought of as a lower limit to our estimate from the measured burst energy fluence, since burning will continue at the stable CNO burning rate.

Alternatively, we can ask the question, how much energy would we expect in the burst if the entire accreted column were to burn to heavy elements? The energy release per gram, Q_{nuc} , would depend on the details of the nuclear burning pathways, however, employing the value of $Q_{\text{nuc}} = (1.3 + 5.8X) \times 10^{18} \text{ erg g}^{-1}$ (Galloway & Cumming, 2006), and again adopting $X = 0.7$ we would expect $3.2\text{--}4.2 \times 10^{39}$ erg liberated at the neutron star surface by burning all the fuel. This estimate also assumes that the total accreted column produces a single burst. This is a factor of 30–40 larger than the observed fluence in the weak X-ray burst, and also argues that the weak burst is likely not a mixed H/He burst. Rather, our analysis suggests that it likely represents the unstable ignition of a modest fraction of the hydrogen in the accreted layer, which constitutes strong observational evidence for such a weak “hydrogen-only” flash.

Interestingly, in their two zone model Cooper & Narayan (2007) compute the peak energy fluxes produced during such weak hydrogen flashes. The range of fluxes that their model can produce is summarized in their Fig. 7 (bottom panel), where the peak flux is given as a fraction of the Eddington flux. Working backward, we measured a peak flux during the weak X-ray burst of $\approx 6.9 \times 10^{-9} \text{ erg cm}^{-2} \text{ s}^{-1}$. If we scale this by the peak flux ($2.3 \times 10^{-7} \text{ erg cm}^{-2} \text{ s}^{-1}$) of the Eddington-limited burst observed later in the outburst (Bult et al., 2019), we find a ratio of $\log(0.03) = -1.52$. Looking at Fig. 7 in Cooper & Narayan (2007) (bottom panel), we can find hydrogen flashes (the dotted-dashed lines in the figure) in a narrow range of accretion rate that reach this flux level.

The estimates in this discussion rely on analytic arguments and comparisons with existing ignition models, and therefore cannot fully capture the time-dependent thermal and compositional evolution of the accreted layer. In particular, the role of sedimentation, the detailed ignition conditions, and the subsequent burning pathways require a self-consistent, multi-zone treatment. Moreover, theoretical models such as Cooper & Narayan (2007) predict that hydrogen-triggered flashes at low accretion rates are intrinsically weak, with peak luminosities below the persistent accretion emission, which would make them difficult to detect observationally. If the observed weak burst is truly hydrogen-triggered, this would suggest that additional physical processes may be at play. In particular, enhancements to the CNO abundance—whether due to a higher metallicity donor, sedi-

mentation, diffusion, or other mixing processes—may alter the ignition conditions and increase the burst’s peak luminosity. These considerations motivate the more detailed numerical calculations presented in Chapter 3, where we use MESA to model hydrogen ignition at low accretion rates across a range of accreted metallicities.

2.2.1 Stable burning after the burst?

Previous theoretical studies concluded that the ignition of the hydrogen flash will raise the temperature in the layer to at least the stable burning regime, and likely higher. Thus, hot-CNO cycle burning would be expected to continue for some period of time after the unstable ignition. Can we see evidence for such stable burning in the NICER data? Interestingly, there is a clear “offset” between the pre- and post-burst flux levels. This offset can be seen in Figure 2.2. Note that the inset panel uses a larger time bin size and log scale to emphasize the persistent count rate levels, to more clearly highlight the offset. We also plot the average count rate value for the pre-burst level (red dashed line) as a guide to the eye. To explore this question further we used the same spectral model to characterize the post-burst data as we used for the pre-burst and other persistent emission intervals. The time interval used for the post-burst spectral extraction is marked by the vertical dotted lines in Figure 2.2 (main panel). We first tried to fit the post-burst spectrum using the same spectral shape as obtained from the pre-burst interval, allowing for the constant, f_a parameter to make up the flux difference. This did not provide an acceptable fit, and suggests the presence of an additional spectral component in the post-burst interval. To explore this further we subtracted the pre-burst spectrum from the post-burst and found that the remaining excess could be well fit by a soft thermal spectrum, characterized as a blackbody with $k_B T = 0.51 \pm 0.02$ keV, normalization of 82.5 ± 12.0 , and bolometric flux of $6.1 \pm 0.2 \times 10^{-11}$ erg cm⁻² s⁻¹. This is equivalent to a luminosity of 8.9×10^{34} erg s⁻¹ (at 3.5 kpc).

If the hydrogen burns stably at the same rate as it is accreted, then we would estimate a hydrogen-burning luminosity of $L_h = X\dot{m}E_h$, where X is the mass fraction of hydrogen in the accreted fuel, \dot{m} is the mass accretion rate at the burst onset, and E_h is the energy production per gram due to hydrogen-burning. With $\dot{m} = 1.4 \times 10^{-10} M_\odot \text{ yr}^{-1}$, $X = 0.7$, and $E_h = 6.4 \times 10^{18}$ erg g⁻¹, we would

predict a stable hydrogen-burning luminosity of $\approx 4 \times 10^{34} \text{ erg s}^{-1}$, which is a good fraction of the measured offset. Perhaps a better estimate can be obtained by evaluating the energy production rate associated with the saturated, hot-CNO burning rate as, $L_{\text{CNO}} = 4\pi R^2 y_a \epsilon_{\text{CNO}}$, where ϵ_{CNO} , y_a , and R , are the energy production rate due to hot-CNO burning, the accreted column depth, and the neutron star radius, respectively. With $\epsilon_{\text{CNO}} = 5.8 \times 10^{13} (Z_{\text{CNO}}/0.01) \text{ erg g}^{-1} \text{ s}^{-1}$ (Cumming & Bildsten, 2000a), $y_a = 4.5 \times 10^7 \text{ g cm}^{-2}$, and $R = 11 \text{ km}$, we find $L_{\text{CNO}} = 4 \times 10^{34} (Z_{\text{CNO}}/0.01) \text{ erg s}^{-1}$. Here, Z_{CNO} is the abundance of the CNO catalyzing elements. Employing the solar value $Z_{\text{CNO}} = 0.016$, we find $L_{\text{CNO}} = 6.4 \times 10^{34} \text{ erg s}^{-1}$; however, as noted above, at these low accretion rates sedimentation is very likely to be effective in enhancing the abundance of CNO elements near the base of the accreted fuel layer. For example, Peng et al. (2007, see their Figs. 2 & 3) report enhancements in CNO element abundances by factors of 2 to 5, depending on the accretion rate.

Based on these estimates it appears plausible that most or all of the observed flux offset can be accounted for by quasi-steady, hot-CNO burning of hydrogen. We note that the thermal nature of the spectral excess, and its $\approx 0.5 \text{ keV}$ temperature, similar to that at late times during the weak burst, is also consistent with this interpretation. This conclusion is also consistent with the hydrogen flash temperature and flux evolution calculations of Cooper & Narayan (2007). As an example, the hydrogen flashes shown in their Fig. 4 indicate that at ignition the flux rises abruptly, but then shows a “plateau-like” phase which decays over a timescale of several hours. The average flux levels near the beginning of these events are approximately consistent with the stable hydrogen-burning luminosity we estimated above. Once ignited, these flashes are burning hydrogen to helium in the fuel layer at essentially the saturated, hot-CNO cycle rate. We suggest that the two-zone model of Cooper & Narayan (2007) (with H and He zones) probably does not adequately track and resolve the fast, initial hydrogen-burning when the thermal instability is initiated, but better predicts the longer timescale, thermally stable burning. The hydrogen-only ignition modeled by Peng et al. (2007, see their Fig. 7) also appears at least approximately similar to what is observed for the weak NICER burst. Indeed, the ratio of the peak burst bolometric flux to the persistent, pre-burst flux is $7 \times 10^{-9} / 7.9 \times 10^{-10} \approx 8.8$, which is similar to the peak value of $F_{\text{cool}}/F_{\text{acc}}$ for the initial, burst-

like flux increase shown in their Figure 7 (middle panel), and the overall burst duration appears consistent with the observed burst as well.

More detailed, radially resolved, and perhaps multidimensional calculations will likely be needed to more accurately track the rapid hydrogen ignition phase and subsequent evolution of the accreted layer, which we suggest may account for the weak NICER burst. To briefly summarize, the weak NICER burst and post-burst flux offset appear to be consistent with the onset of a hydrogen-triggered shell flash in the cool, temperature-sensitive regime of the CNO cycle. The ignition column was likely shallow enough that the subsequent temperature increase was not sufficient to also promptly ignite a helium-burning instability. Existing models predict that hydrogen-triggered flashes at these accretion rates should be intrinsically weak and difficult to detect (Cooper & Narayan, 2007). The tension between theoretical expectations and the observed burst properties points to additional physics, such as enhanced CNO abundances from higher metallicity accretion or sedimentation that may alter the ignition conditions. Exploring these effects motivate the detailed numerical calculations presented in Chapter 3.

2.2.2 Missed bursts?

While NICER was able to begin observations quite close to the onset of accretion in the 2019 August outburst, the overall on-source coverage from onset to the time of the first observed burst was still rather modest, with a duty-cycle of about 4%. Thus, if other bursts occurred it is conceivable that the NICER observations simply missed them. However, based on our estimate of the size of the accreted column, as well as current theoretical estimates of the hydrogen ignition curve, we argue that likely only a few such bursts might have been missed. Firstly, while we don't know the precise temperature trajectory of the initial accumulating layer, it cannot plausibly be $\lesssim 2 \times 10^7$ K because at such low temperatures only columns much larger ($\gtrsim 2\text{--}3 \times 10^8 \text{ g cm}^{-2}$) than our estimate of the accreted column at the time of the weak burst ($3.9\text{--}5.1 \times 10^7 \text{ g cm}^{-2}$) would be needed to ignite unstable burning, and such an ignition would also very likely lead to a bright, mixed H/He burst, which was not observed, though could have perhaps been missed. Secondly, as the temperature of the fuel layer increases the size of the unstable column decreases, however, above temperatures

of about 8×10^7 K the hydrogen-burning will stabilize, precluding bursts. This sets a minimum combustible column for hydrogen ignition which is, using the ignition curve in Cumming (2004) as a guide, $\approx 1 \times 10^7$ g cm⁻². Based on our estimated accreted column this would set a limit of not more than about five such bursts potentially being produced, as that would just about exhaust the total column accreted at the time of the weak burst. Another benchmark can be set by the accretion rate. We estimated a value of $\dot{m} = 1.4\text{--}1.6 \times 10^{-10} M_{\odot} \text{ yr}^{-1}$ at the time of the weak burst. If we take half of this value as more representative of the mean rate during the 72 hr prior to the weak burst, we can estimate the time required to accrete the minimum unstable column of 1×10^7 g cm⁻². For $\dot{m} = 7 \times 10^{-11} M_{\odot} \text{ yr}^{-1}$, and assuming a radius $R = 11$ km, we find it would take 9.5 hr to accumulate such a column. Since the weak burst was observed after about 2.9 days, this also suggests the total number of such weak bursts is at most about 7. We suggest that the actual temperature trajectory is probably somewhere between the two extremes described above, perhaps consistent with an ignition column of $2\text{--}3 \times 10^7$ g cm⁻². If this is correct it would suggest that the NICER observations may have missed one or two such weak bursts.

2.2.3 Other examples: RXTE observations of the 2005 outburst

We searched the literature and previous observations of SAX J1808.4–3658 to try and identify similar examples of weak bursts. We found a quite similar event early in the 2005 June outburst that was observed with RXTE. We show in Figure 2.5 the light curve of this outburst as obtained from RXTE pointed observations. This burst occurred on 2 June at approximately 00:42:30 TT, and is evident near 0.25 days in the figure. We carried out a time resolved spectral analysis of this event, and found qualitatively similar properties for this burst as for the weak NICER burst. It reaches a peak bolometric flux of $1.54 \pm 0.11 \times 10^{-8}$ erg cm⁻² s⁻¹, about a factor of 2 greater than the NICER burst. It also had a peak blackbody temperature of 1.25 keV, which is about 25% larger than that of the NICER burst. We note that this burst appears in the Multi-Instrument Burst Archive (MINBAR) catalog, with a reported peak bolometric flux of 1.6×10^{-8} erg cm⁻² s⁻¹, and a fluence of 1.67×10^{-7} erg cm⁻² (Galloway et al., 2020).

The first evidence of active accretion for this outburst was provided by RXTE/PCA Galactic




Figure 2.5 Light curve from RXTE data (PCU 2, 3–30 keV) of the 2005 outburst from SAX J1808.4–3658. Note the logarithmic scale. A weak X-ray burst is seen early in this outburst. Much brighter and energetic bursts are seen near days 4 and 8. Note that the burst near day 8 was truncated by the RXTE exposure, and almost certainly the brightest part of this event was missed.

bulge scan observations on 31 May at 23:00:00 UTC, and indicated a persistent 2 - 10 keV X-ray flux level of ≈ 3 mCrab (Markwardt et al., 2005). This flux value is similar to that measured with NICER for ObsID 2050260109 during the 2019 outburst (see Table 2.1). The X-ray burst was observed approximately 25.7 hr later, and MINBAR reports a persistent flux at the time of the burst of 8.6×10^{-10} erg cm $^{-2}$ s $^{-1}$, just a bit larger than the value estimated prior to the 2019 NICER burst (again, see Table 2.1). We can use the pre-burst flux value reported by MINBAR and the earliest RXTE observations of the 2005 outburst reported by Markwardt et al. (2005) and Wijnands et al. (2005) to estimate the persistent, accretion-driven fluence prior to the weak 2005 burst. Evaluating a simple trapezoidal sum gives a value of 3.8×10^{-5} erg cm $^{-2}$ that is approximately half of the estimated fluence prior to the 2019 NICER event. This then suggests a total accreted column just prior to the 2005 RXTE event of about half that estimated for the 2019 NICER burst. Simply scaling our value estimated for the 2019 NICER burst suggests a range of $2.0\text{--}2.6 \times 10^7$ g cm $^{-2}$ for the total accreted column prior to the 2005 RXTE event.

2.2.4 Subsequent bursts detected with NuSTAR

Additional observations of SAX J1808.4–3658 were collected with NuSTAR between 2019 August 10 and 11 (MJD 58705.5–58706.5). While these data do not cover the time of the weak X-ray burst observed with NICER, NuSTAR did catch two subsequent bursts, providing some additional, interesting context to this early phase of the outburst. We processed the NuSTAR data (ObsID 90501335002) using NUSTARDAS version 2.1.2. Source data were extracted in the 3–79 keV energy range from a 40'' circular region centered on the source coordinates. The background was extracted using the same approach, but with the extraction region positioned in the background field. The NuSTAR light curve reveals two X-ray bursts, the first of which occurred 24.8 hours after the weak NICER burst, while the second occurred another 11 hours later. This light curve is shown in Figure 2.6. We emphasize that though some of the NICER exposure was simultaneous with NuSTAR, this did not include these two bursts, and they were only observed with NuSTAR.

We first investigate the persistent emission by extracting a spectrum from a 100 second window just prior to the first NuSTAR burst. As can be seen in Figure 2.6, this epoch was simultaneously

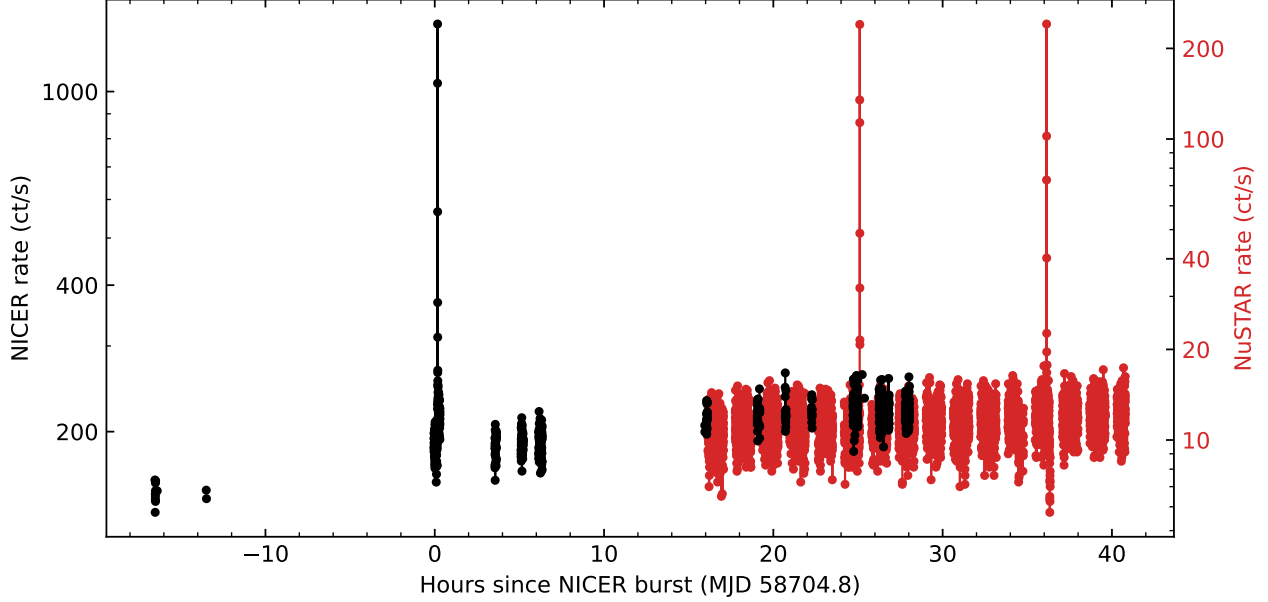


Figure 2.6 Light curves from NICER (black, left axis) and NuSTAR (red, right axis) around the time of the weak X-ray burst at $t = 0$. Both light curves are calculated using an 8 s time resolution.

observed with NICER, so we also extracted the contemporaneous NICER spectrum to obtain broad-band energy coverage. We model this spectrum using the same persistent emission used previously (see Table 2.1), allowing for a constant cross calibration factor between NICER and FPMA/B of NuSTAR. In keeping with the analysis of the NICER burst, we extrapolated the best spectral model over 0.1–100 keV to find a bolometric flux estimate of $1.47 \pm 0.05 \times 10^{-9} \text{ erg cm}^{-2} \text{ s}^{-1}$.

From the recurrence times between the observed bursts, we obtain an estimate of the fluence due to the accretion of $1.3 \times 10^{-4} \text{ erg cm}^{-2}$ and $5.8 \times 10^{-5} \text{ erg cm}^{-2}$ for the two bursts, respectively. Converting these measurements to column depths, we use equation 2.4 to find 8.4×10^7 and $3.7 \times 10^7 \text{ g cm}^{-2}$, respectively, where we again assumed a $1.4 M_{\odot}$ neutron star mass and an 11 km stellar radius. These column depths are of the same order as the one we calculated for the initial NICER burst. Indeed, given the observed 11 hr recurrence time between the two NuSTAR bursts, and the relatively constant persistent flux (and hence accretion rate), it is conceivable that a similar burst was missed in the gap between the weak NICER burst and the first NuSTAR burst. If so, then the accretion fluence for the two NuSTAR bursts would be essentially consistent with each other.

To explore the burst spectra, we proceeded by dividing the bursts into multiples of 1/8 seconds

Table 2.2 Burst Parameters from NuSTAR and NICER Observations of the Outburst

Dataset	Onset (MJD, TT)	Peak flux (10^{-9} erg cm $^{-2}$ s $^{-1}$)	Burst fluence (10^{-8} erg cm $^{-2}$)	Accretion fluence (10^{-4} erg cm $^{-2}$)	Peak $k_B T$ (keV)
NICER 1	58704.8068	7	7	0.8	1.0
NuSTAR 1	58705.8459	70	50	1.3	2.3
NuSTAR 2	58706.3058	40	30	0.58	1.7
NICER 2 ^a	58716.0861	300	200	–	2.5

^a The properties of the second NICER burst are taken from Bult et al. (2019) as an example of a bright Eddington-limited X-ray burst from SAX J1808.4–3658.

such that each bin contains at least 500 counts. We extract a spectrum for each bin and model it using an absorbed blackbody in addition to the fixed persistent emission. The inferred burst properties obtained from these fits are listed in Table 2.2. The two NuSTAR bursts had fluences of 5×10^{-7} erg cm $^{-2}$ and 3×10^{-7} erg cm $^{-2}$, respectively. This means that they are about a factor of 4–7 times more energetic than the weak X-ray burst observed with NICER. The first NuSTAR burst reached a peak flux ten times greater than that of the weak NICER burst, and it was also significantly “hotter,” reaching a peak blackbody temperature of 2.3 keV. At the same time, these bursts remain much fainter than the Eddington-limited bursts observed at later times in the outbursts of SAX J1808.4–3658, which typically have fluences of $2\text{--}4 \times 10^{-6}$ erg cm $^{-2}$ (Galloway et al., 2008; in’t Zand et al., 2013).

2.3 Summary, Caveats, and Outlook

Based on the considerations above we suggest a scenario similar to that discussed in the work of Cooper & Narayan (2007) and Peng et al. (2007) as a working hypothesis to account for the weak bursts observed by NICER and RXTE during the early days of the 2019 and 2005 outbursts of SAX J1808.4–3658. As accretion begins, the neutron star is cool enough and the accretion rate is low enough that CNO hydrogen-burning in the accumulating layer occurs in the temperature-sensitive regime. At these lower temperatures, $\lesssim 5 \times 10^7$ K, very little hydrogen is burned. Significant burning of hydrogen will only begin when the accumulated column reaches the conditions for the thermal instability to set in. For a temperature of $\approx 5 \times 10^7$ K this will occur at a column depth of about 3×10^7 g cm $^{-2}$. This value is not too dissimilar from the column estimated just

prior to the 2005 event. When the initial accumulating layer reaches ignition depth the hydrogen instability occurs, triggering a hydrogen flash. We suggest that the initial rapid increase in the nuclear energy generation rate ultimately results in the “heat pulse” that is observed as the weak X-ray burst, however, we think that more sophisticated, multi-dimensional theoretical calculations of the time-dependent nuclear energy generation coupled with the subsequent heat and radiation transport, will be needed to test the details of this hypothesis. After the initial hydrogen ignition, the burning layer will reach a high enough temperature that subsequent hydrogen-burning can proceed at the thermally stable level appropriate to the hot-CNO cycle. Above, we have argued that the observed offset between the pre- and post-burst flux levels of the 2019 event is consistent with this “quasi-steady” burning phase. This source of heat will keep the layer warm enough for burning to continue for a time, likely measured in hours if conditions are not too dissimilar from those modeled by Cooper & Narayan (2007). During this time the quasi-stable burning will increase the helium fraction of the layer. Given the gaps in NICER coverage after the weak burst, we cannot say how long this “quasi-steady” burning may have persisted, but we note that observations about 3.5 hrs after the burst show a count rate and flux approximately consistent with the pre-burst level. For the conditions described above, that is, a hydrogen ignition column of $\approx 3 \times 10^7 \text{ g cm}^{-2}$, such an initial hydrogen flash is unlikely to produce a prompt helium ignition, simply because at that column depth the helium will not be thermally unstable (Cumming, 2004).

As accretion continues, the hydrogen layer or layers that initially flashed will be pushed deeper, to higher column depths. The freshly accreted material above it will also reach the hydrogen ignition depth, and if so, produce another hydrogen flash, assuming its temperature is low enough. In this way, a sequence of hydrogen flashes could be produced. Eventually, the helium-enriched layers will likely reach column depths where the helium will ignite, producing more energetic, mixed H/He bursts. We suggest that the observed NuSTAR bursts are the result of this process. The steadily increasing accretion rate will also be an important variable, as this will tend to increase the temperature of the accreting layers. More complete theoretical modeling of this process will have to include the time-varying accretion rate (Johnston et al., 2018).

If the above scenario is approximately correct we can try to speculate further regarding a few other details of the observations. The 2005 event observed with RXTE was the earlier event in terms of the time since outburst onset, occurring approximately 1 day after onset. Other things being equal one would expect the accreting layer to be cooler than at later times, such as the 2.9 days post-outburst from the 2019 event. A cooler shell will have a larger unstable column, so that this could perhaps explain the fact that the RXTE event is the more energetic of the two weak bursts. This also provides some tentative evidence that the 2019 NICER event may have been preceded by at least one additional weak burst that was missed.

2.3.1 Remaining Uncertainties and Alternative Interpretations

In estimating accretion rates and accreted columns we allowed for variation in the neutron star mass, however, there are other uncertainties which complicate such estimates. These include the source distance, anisotropy factors, bolometric corrections, and the line-of-sight absorption. We note that the more recent work of Goodwin et al. (2019) reports a slightly closer distance of $3.3^{+0.3}_{-0.2}$ kpc for SAX J1808.4–3658. While their quoted uncertainty range includes the 3.5 kpc value we have adopted, a decrease from 3.5 to 3.3 kpc would reduce our estimates by a factor of 0.9. These authors also provide estimates of the anisotropy factors for both persistent and burst emission, finding $\xi_p = 0.87^{+0.12}_{-0.10}$ and $\xi_b = 0.74^{+0.10}_{-0.10}$. Applying these values would also reduce the estimated accretion rate and column, by a factor of 0.87, and decrease our estimate of the burst peak luminosity and fluence, by a factor of 0.74. Adopting the best values reported by Goodwin et al. (2019) for both d and ξ_p would reduce the estimated accretion rate and accreted column by a factor of 0.77.

We have argued that the weak bursts result from, principally, hydrogen-burning, but are there other possibilities involving the unstable burning of helium? One scenario that can produce weak or underluminous bursts is the phenomenon of short recurrence time (SRT) bursts (Keek et al., 2010). The idea behind SRT events is that they burn fuel remaining from a preceding burst. In the present case, a preceding, larger burst would have had to occur (and been missed) for this idea to be workable. In principle, this could account for the observed weak bursts, but there are some difficulties with this interpretation. First, the estimated accreted columns are uncomfortably low.

This scenario would require that a relatively bright, mixed H/He burst would have occurred prior to the observed weak events, and been missed in each case. As discussed above, this would require relatively large ignition columns, likely $\gtrsim 2 \times 10^8 \text{ g cm}^{-2}$, which is much larger than the estimated columns present just prior to each weak event. Our accretion column estimates would have to be underestimated by factors of 4–5 for this to be more plausible. Second, SAX J1808.4–3658 is not currently known to produce SRT events. There has been reasonably good coverage of past SAX J1808.4–3658 outbursts, and no SRT events have been definitively observed. For example, the compilation of SRT burst observations by Keek et al. (2010) does not include SAX J1808.4–3658, and we also note that the 401 Hz spin frequency for SAX J1808.4–3658 is less than the faster, $\gtrsim 500$ Hz, spins associated with some of the known SRT sources. Third, the flux offset between the pre- and post-burst emission seems to make more sense in the context of stable hydrogen-burning than what might be expected from an SRT event, for which one would not typically expect to find such a flux offset. We note also that the theoretical mechanism of opacity-driven convective mixing explored by Keek & Heger (2017) to account for SRT bursts occurs for ignition in relatively hot envelopes, which seems less applicable to the low accretion rate regime near burst onset that we have described above. It is difficult to completely rule out the SRT scenario, but we think the considerations above argue against it.

We have argued that the evolving accretion onto the “cool” neutron star in SAX J1808.4–3658 during the early part of the outburst provides a unique environment to explore the physics of nuclear burning on neutron stars, and most interestingly, the ignition of unstable hydrogen-burning in the temperature-sensitive regime of the CNO cycle. We suggest that the weak bursts seen by NICER and RXTE in the 2019 and 2005 outbursts, respectively, may result from this process. More complete, continuous observational coverage of the first 4–5 days of subsequent outbursts from SAX J1808.4–3658 could definitively test this hypothesis. Such data would also provide for detailed physical comparisons with new theoretical efforts to track the outcome of time-varying accretion onto neutron stars and the subsequent nuclear burning of the accreted matter. This could provide interesting constraints on such things as the accretion rate, the thermal profile of the accret-

ing matter and the nuclear energy generation and subsequent heat transport in the accreted layers.

CHAPTER 3

EFFECTS OF METALLICITY ON H-TRIGGERED BURSTS

In Chapter 2, we argued that the observed weak bursts were consistent with being solely due to unstable H burning. Specifically, the cool envelope temperature required for unstable H ignition was consistent with both observed weak bursts occurring 1–3 days after the onset of active accretion, and also that the existence of a long, steady tail of extended emission following the peak was consistent with stable CNO burning (see Peng et al., 2007). Adjusting for an updated distance of 3.3 kpc (Goodwin et al., 2019), Casten et al. found a peak luminosity of 9.1×10^{36} erg s⁻¹, a tail luminosity of $\approx 7.9 \times 10^{34}$ erg s⁻¹, and a column depth of $3.10\text{--}4.12 \times 10^7$ g cm⁻² for $1.4 M_{\odot} < M < 2.0 M_{\odot}$ and $R = 12$ km. These estimates provide a baseline for comparison from Chapter 2 with the detailed models presented in this chapter.

At low accretion rates, $\dot{M} \lesssim 10^{-2} \dot{M}_{\text{Edd}}$, sedimentation of heavy elements in the envelope (Bildsten et al., 1992; Wallace et al., 1982; Peng et al., 2007) could potentially enhance the CNO abundance at the base of the accreting envelope and lead to a more energetic burst. The role of enhanced CNO abundances in driving vigorous hydrogen burning is well established in models of classical novae (Starrfield 1993; see also Denissenkov et al. 2013 and references therein), though the exact mechanisms responsible for this enhancement is not well understood. Mixing processes, such as convective overshooting and Kelvin-Helmholtz instabilities, are thought to play a key role in enriching the accreted envelope with heavier elements from the underlying white dwarf (Casanova et al., 2018, 2011).

Peng et al. (2007) studied this effect in a simplified one-zone model, but it has not yet been included in multi-zone simulations. In this preliminary study, we do not include sedimentation directly, but instead study the effect of enhanced metallicity on the properties of the burst by considering a range of accreted metallicities. We use the open-source stellar evolution code MESA (Paxton et al. 2011, Jermyn et al. 2023 and references therein) to explore the unstable ignition of hydrogen at low mass accretion rates for a range of metallicities in the accreted fuel. As shown by the one-zone calculations of Peng et al. (2007) and two-zone model of Cooper & Narayan (2007),

after the initial peak in luminosity, there is a long tail of enhanced luminosity powered by stable hydrogen burning via the hot CNO cycle. This burning continues until it consumes the accumulated hydrogen, at which point the envelope cools and the luminosity drops.

We revisit quasi-steady-state H burning via the hot CNO cycle in § 3.1 and provide estimates for the luminosity and duration of the tail. We then describe our MESA models in § 3.2. In this preliminary study we do not include sedimentation but instead model a range of metallicities $Z = 0.01\text{--}0.30$, which is a larger range than expected, but allows us to fully explore the effect of enhanced CNO. For comparison, Peng et al. (2007) estimated that sedimentation in a hydrogen-carbon mixture could increase the abundance of ^{12}C at the base of the accreted layer from 0.02 to 0.23. Because the weak bursts were observed shortly after the onset of accretion, we do not run our simulations through several bursts but instead focus on the first burst following the onset of accretion. We show (§ 3.3) how the CNO abundance affects the ignition depth, peak luminosity, and burst lightcurve morphology, and we briefly compare our burst lightcurves and fluences with those of the 2019 burst observed from SAX J1808.4–3658. We conclude in § 3.4 with a summary, a brief discussion of uncertainties, and an outlook on future efforts.

3.1 Burning via the Hot CNO cycle

As shown by Peng et al. (2007), following the initial triggering of the burst by the cold CNO cycle reactions, the temperature increases enough that the burning enters the hot CNO cycle and becomes steady with the net rate set solely by the temperature-independent β -decays of ^{14}O and ^{15}O with half-lives $t_{1/2,014} = 70.64$ s and $t_{1/2,015} = 122.24$ s, respectively (JINA REACLIB; Cyburt et al., 2010). In this steady-state burning, all CNO nuclides are converted into ^{14}O and ^{15}O , with the number densities n_{014} and n_{015} being in the ratio $n_{014}/n_{015} = t_{1/2,014}/t_{1/2,015}$. In this regime, hydrogen is depleted at a rate

$$\frac{dn_{\text{H}}}{dt} = -4n_{\text{CNO}} \frac{\ln 2}{t_{1/2,014} + t_{1/2,015}}. \quad (3.1)$$

Here n_{CNO} denotes the number density of all CNO nuclides; during hot CNO burning, $n_{\text{CNO}} = n_{014} + n_{015}$. Thus, if H consumption proceeds via steady-state hot CNO burning in the burst tail,

the layer of hydrogen accumulated at the start of the burst will be depleted in a time

$$t_d = \frac{n_H}{n_{\text{CNO}}} \frac{t_{1/2, \text{O14}} + t_{1/2, \text{O15}}}{4 \ln 2}, \quad (3.2)$$

which depends solely on the ratio of hydrogen to CNO nuclides at the onset of burning.

The mean mass of the CNO nuclides in the accreted fuel is $\bar{A}_{\text{CNO}} = 14.56$ and the mass fraction of CNO elements for a given overall metallicity Z is $Z_{\text{CNO}} = 0.6677 Z$ (Asplund et al., 2005). As a consequence, the ratio of hydrogen (with mass fraction X) to CNO nuclides is

$$\frac{n_H}{n_{\text{CNO}}} = \frac{\bar{A}_{\text{CNO}} X}{Z_{\text{CNO}}} = 21.80 \frac{X}{Z}. \quad (3.3)$$

Substituting Eq. (3.3) into Eq. (3.2) along with $t_{1/2, \text{O14}}$ and $t_{1/2, \text{O15}}$ yields

$$t_d = 1.1 \times 10^5 \text{ s} \left(\frac{X}{0.7} \right) \left(\frac{0.01}{Z} \right). \quad (3.4)$$

Lampe et al. (2016) derived a similar expression; their prefactor differs from ours because they took the metallicity to entirely reside in CNO isotopes, i.e., $Z = Z_{\text{CNO}}$.

The production of each ${}^4\text{He}$ nucleus is accompanied by a release of heat, $Q = (26.73 - 3.22) \text{ MeV} = 24.7 \text{ MeV}$, in which the first term is the mass difference and the second term accounts for the neutrino losses¹. The luminosity from CNO burning in the burst tail is therefore

$$\begin{aligned} L_{\text{CNO}} &= \frac{4\pi R^2 y_{\text{ign}} X Q}{4m_u t_d} \\ &= 4.8 \times 10^{34} \text{ erg s}^{-1} \left(\frac{y_{\text{ign}}}{3 \times 10^7 \text{ g cm}^{-2}} \right) \left(\frac{Z}{0.02} \right). \end{aligned} \quad (3.5)$$

Here m_u is the atomic mass unit, and y_{ign} is the column depth of the base of the accreted H-rich layer, with column depth defined as $y = \int \rho dr \approx \Delta M_{\text{accr}} / (4\pi R^2)$. For $R = 12 \text{ km}$ the ignition mass is $\Delta M_{\text{ign}} = 2.7 \times 10^{-13} M_{\odot} (y_{\text{ign}} / 3 \times 10^7 \text{ g cm}^{-2})$ with a recurrence time $t_r = 24.0 \text{ hr} (y_{\text{ign}} / 3 \times 10^7 \text{ g cm}^{-2})$ for $\dot{M} = 10^{-10} M_{\odot} \text{ yr}^{-1}$. Note that L_{CNO} depends solely on y_{ign} and the metallicity Z of the accreted layer; in particular, it does not depend on the mass fraction of hydrogen X .

¹We adopt here the mean neutrino energy used by MESA, $\bar{E}_\nu = 3.22 \text{ MeV}$, which is larger than the value of 2.03 MeV used by Wallace & Woosley (1981).

For comparison, the luminosity produced by burning hydrogen at the rate at which it accretes is

$$\begin{aligned}
 L_{\text{SS}} &= \frac{\dot{M} X Q}{4m_{\text{u}}} \\
 &= 2.5 \times 10^{34} \text{ erg s}^{-1} \left(\frac{X}{0.7} \right) \left(\frac{\dot{M}}{10^{-10} M_{\odot} \text{ yr}^{-1}} \right). \tag{3.6}
 \end{aligned}$$

For $L_{\text{CNO}} > L_{\text{SS}}$, fuel cannot be replenished quickly enough to sustain L_{CNO} , and so, after a time t_d , the luminosity drops and the envelope cools on a thermal timescale. If the temperature is too cool, the proton captures in the CNO cycle cease and the envelope settles into a limit cycle.

3.2 Numerical models

Having described the basic physics of steady-state hot CNO burning following (unstable) H ignition, we now model the accretion and ignition. We do this using the open-source stellar evolution code MESA, version r23.05.01 (Paxton et al., 2011, 2013, 2015, 2018, 2019; Jermyn et al., 2023), and compiled with the MESA SDK, version x86-linux-23.7.3 (Townsend, 2024). The microphysics employed by MESA are as follows. The equation of state (EOS) is a blend of the OPAL (Rogers & Nayfonov, 2002), SCVH (Saumon et al., 1995), FreeEOS (Irwin, 2004), HELM (Timmes & Swesty, 2000), PC (Potekhin & Chabrier, 2010), and Skye (Jermyn et al., 2021) EOSs. Radiative opacities are primarily from OPAL (Iglesias & Rogers, 1993, 1996), with low-temperature data from Ferguson et al. (2005) and with the high-temperature, Compton-scattering dominated regime treated following Poutanen (2017). Electron conduction opacities are from Cassisi et al. (2007) and Blouin et al. (2020). Nuclear reaction rates are taken from JINA REACLIB (Cyburt et al., 2010) and NACRE (Angulo et al., 1999) and supplemented by additional tabulated weak reaction rates (Fuller et al., 1985; Oda et al., 1994; Langanke & Martínez-Pinedo, 2000). Screening is included via the prescription of Chugunov et al. (2007). Thermal neutrino loss rates are included using the prescription of Itoh et al. (1996).

We construct the envelope model using the MESA test suite case `make_env` with a neutron star mass $M = 1.4 M_{\odot}$, radius² $R = 12$ km, and luminosity $L_b = 10^{33} \text{ erg s}^{-1}$. We set the initial compo-

²We adopt here $R = 12$ km, which is consistent with recent NICER determinations of the radius of the high-mass

sition of the envelope to ^{56}Fe with a mass 1.8×10^{24} g, so that the base of the computational domain is initially at a column of 10^{11} g cm $^{-2}$. We then accrete H-rich matter at a rate $\dot{M} = 10^{-10} M_{\odot}$ yr $^{-1}$. The accretion rate is low in order to ensure a sufficiently cool envelope for the CNO cycle to be thermally unstable (Cumming, 2004; Peng et al., 2007; Casten et al., 2023). For this accretion rate, $L_b/\dot{M} = 0.16$ MeV u $^{-1}$ ($= 1.5 \times 10^{17}$ erg g $^{-1}$), which is consistent with the value needed for unstable H ignition (Peng et al., 2007).

The distribution of nuclides in the accreted matter follows the solar system abundances compiled by Asplund et al. (2005). To follow any potential breakout from the hot CNO cycle into an rp-process (Wallace & Woosley, 1981), we use a modified version of the truncated reaction network `rp_153.net`, which is a subset of the full rp-process network, as it only extends to ^{56}Ni . This is sufficient for modeling weak bursts, which do not develop a strong rp-process. We checked that the abundances of the heaviest isotopes did not appreciably change during the burst. To account for species that are not present in the network without renormalizing abundances, we use the default setting (`.true.`) for the input flag `accretion_dump_missing_metals_into_heaviest` and add ^{57}Co , ^{57}Ni , and ^{58}Ni (which is stable) to `rp_153.net`; otherwise there would be spurious heating from the decay of ^{56}Ni . In addition, we include ^{13}C to allow for decays of ^{13}N during the onset of the CNO cycle. Our modified reaction network, `rp_157.net` thus consists of the isotopes shown in Table 3.1.

Table 3.1 Isotopes in the reaction network `rp_157.net`.

Element ^a	A	Element	A	Element	A	Element	A
H	1–3	O	13–18	P	26–31	Ti	40–47
He	3,4	F	17–19	S	27–34	V	43–49
Li	7	Ne	18–21	Cl	30–35	Cr	44–52
Be	7,8	Na	20–23	Ar	31–38	Mn	47–53
B	8,11	Mg	21–25	K	35–39	Fe	48–56
C	9,11–13	Al	22–27	Ca	36–44	Co	51–57
N	12–15	Si	24–30	Sc	39–45	Ni	52–58

^a The reaction network also includes free neutrons.

pulsar PSR J0740+6620 (Dittmann et al., 2024; Salmi et al., 2024). For this initial, exploratory, study, we do not sample over the distribution of neutron star radii.

To better resolve the thermal instability, we reduce `time_delta_coeff`, the time resolution control, from 1.0 (the default) to 0.1. We model convection following the mixing length theory of Cox & Giuli (1968) with mixing length parameter $\alpha = 1.5$, and we use the (default) Schwarzschild criterion, rather than Ledoux, to determine convective stability. We find similar results using each criterion but using Schwarzschild improves numerical stability over all metallicities.

3.3 Unstable hydrogen burning

We model accreting neutron star envelopes with metallicities Z ranging from 0.01 to 0.30. In all cases, we set the hydrogen mass fraction $X = 0.70$ and the helium mass fraction $1 - X - Z$. We defer the exploration of more physically realistic conditions, such as would occur under sedimentation, to a subsequent paper. We first explore the onset of unstable hydrogen burning for two representative cases, $Z = 0.04$ and $Z = 0.10$. For $Z = 0.04$, the envelope does not become convective during the burst, and as a result the burst rise is lengthy; in contrast, for $Z = 0.10$ the burst rise is more violent and leads to the growth of a convective zone that rapidly transports heat to the surface. In both cases, following this peak the bursts exhibit a long tail of quasi-steady hydrogen burning via the hot CNO cycle. This continues until the hydrogen is depleted, at which point the envelope cools. After exploring these two cases in detail, we then examine features of the hydrogen burning over the range of sampled metallicities.

3.3.1 The onset of unstable hydrogen burning

Figure 3.1 illustrates the onset of the instability and evolution of the burst for $Z = 0.04$ and $Z = 0.10$. We show both the surface luminosity L (*solid black curves*) and the integrated nuclear heating $L_{\text{nuc}} = \int \epsilon_{\text{nuc}} dm$ (*dashed red curves*), which is stored in MESA as `log_power_nuc_burn`³. The thin vertical dashed lines denote times for which we shall show isotopes in the accreted envelope. A feature in the nuclear luminosity is the presence of an inflection, which steepens into a spike, followed by a rise to maximum. Similar behavior in the nuclear heating is observed in models of classical novae (Prialnik et al., 1979).

Figure 3.2 displays the evolution of the composition for these cases. The top row shows mass

³In a quirk of nomenclature, the parameter `log_Lnuc` subtracts off the heating from photodissociation reactions and so does not represent the full nuclear luminosity.




Figure 3.1 Surface luminosity (*solid black*) and nuclear luminosity (*dashed red*) during the onset of unstable hydrogen burning and the evolution of the burning to a tail of quasi-steady burning. Models with $Z = 0.04$ (*top*) and 0.10 (*bottom*) are shown. The vertical dotted lines indicate where we sample the isotopic distribution in the accreted and burning neutron star envelope.

fractions of isotopes as functions of column depth in the accreted envelope for $Z = 0.04$; the bottom, for $Z = 0.10$. Going from left to right, the plots represent snapshots corresponding to the first three thin vertical dotted lines in Fig. 3.1. The first snapshot (Fig. 3.2) occurs when the local nuclear heating timescale, $\tau_{\text{nuc}} = c_P T / \epsilon_{\text{nuc}}$, is comparable to the local thermal time (Heney & L'Ecuyer, 1969),

$$\tau_{\text{th}} = \frac{1}{4} \left[\int_0^{y_{\text{ign}}} \left(\frac{3c_P \kappa}{4acT^3} \right)^{1/2} dy \right]^2, \quad (3.7)$$

at the base of the accreted layer. Here c_P is the specific heat at constant pressure and κ is the opacity.

Before the onset of the burst, the reactions $^{12}\text{C}(p, \gamma)^{13}\text{N}(e^+ \nu_e)^{13}\text{C}(p, \gamma)^{14}\text{N}$ have stably depleted ^{12}C , and built up the abundance of ^{14}N . The onset of the burst is heralded by the rapidly increasing $^{14}\text{N}(p, \gamma)^{15}\text{O}$ as the envelope heats. The nuclear timescale τ_{nuc} does not, however, decrease faster than τ_{th} . As a consequence, the luminosity rises on a thermal timescale: the envelope

For $Z = 0.10$, the consumption of ^{12}C and buildup of ^{14}N (Fig. 3.2, *bottom left*) is also thermally stable; in this case, however, the ignition of ^{14}N is sufficiently energetic that τ_{nuc} decreases faster than τ_{th} at the base of the envelope. The spike in heating at $t = 21.1$ s (Fig. 3.1, *bottom*) rapidly heats the base of the accreted envelope and drives convection that mixes the deeper part of the envelope and transports heat outwards, which increases L . The buildup of ^{15}O with its relatively long half-life leads to a rapid, but temporary, fall in heating until $^{16}\text{O}(p, \gamma)^{17}\text{F}$ becomes established. With the fall in heating, convection turns off and does not resume. During the rise to the second peak (Fig. 3.2 at $t = 28.8$ s, *bottom right*) proton captures drive a flow towards ^{32}S through intermediate nuclides $^{23,25}\text{Mg}$, ^{25}Al , and $^{29,30}\text{P}$. This burst is sufficiently vigorous that the peak luminosity L_{peak} exceeds L_{acc} . A snapshot of the isotopic abundances, envelope thermal profile, and lightcurve is contained in Figure 3.3.

For both $Z = 0.04$ and $Z = 0.10$, the hot CNO cycle becomes established following the peak, and the envelope makes a transition from peak to tail, which we locate by finding where L_{nuc} is within 5% of L_{CNO} (Eq. 3.5) for 20 s. At this point (Fig. 3.4) the mass fractions of $^{14,15}\text{O}$ have settled into the ratio set by their half-lives.

3.3.2 Burst properties for different metallicities

After exploring the onset of unstable hydrogen burning for $Z = 0.04$, which produces a slowing rising luminosity, and for $Z = 0.10$, which launches convection and a strong peak in luminosity, we now present burst models over an ensemble of metallicities.

For $Z \geq 0.06$, we locate the time of burst onset, t_{ign} , to be when $(d \ln L_{\text{nuc}}/dt)^{-1} < 50$ s. For $Z = 0.04, 0.02$, and 0.01 , we increase this to 100, 500, and 600 s, respectively, since there isn't a strong thermal instability. Figure 3.5 shows the corresponding ignition column, $y_{\text{ign}} = t_{\text{ign}} \times \dot{M}/(4\pi R^2)$, for different Z . We can roughly (maximum error 6% at $Z = 0.30$) fit the ignition column as $y_{\text{ign}} \approx 8.0 \times 10^7 \text{ g cm}^{-2} (Z/0.01)^{-0.4}$. Comparing these ignition columns with those required for unstable triple- α ignition (Fujimoto et al., 1981), we see that the burst depth is sufficiently shallow to avoid helium ignition. To confirm this, we ran a case with $Z = 0.30$ and hydrogen and helium mass fractions scaled to 0.49 and 0.21, respectively; even with this enhanced helium abundance, there


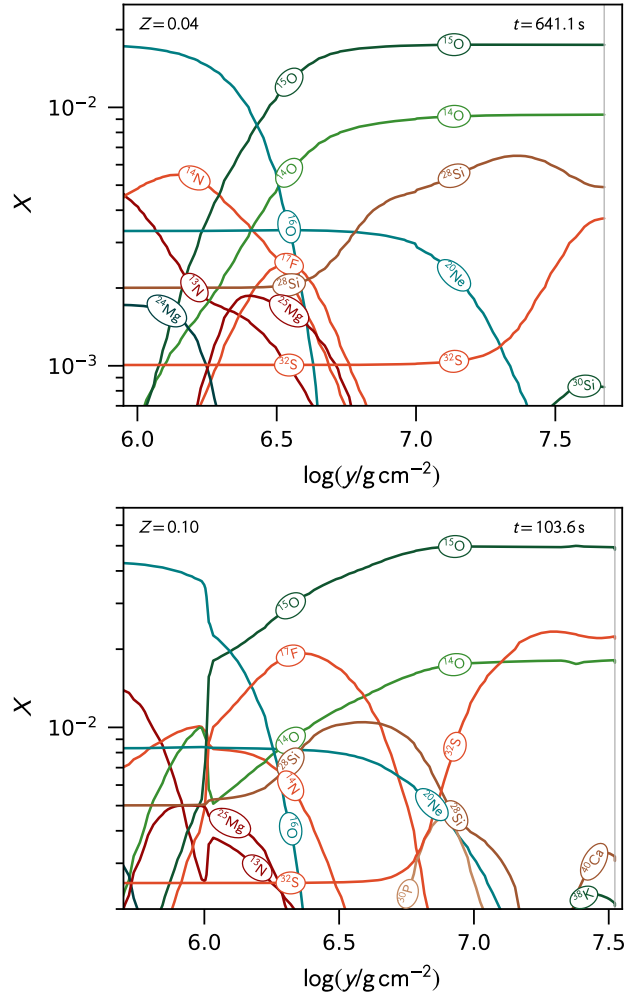


Figure 3.3 Snapshot of the burst with $Z = 0.10$ just after the initial thermal instability as surface luminosity begins its rise. The main panel at left shows the mass fractions (omitting ^{56}Fe) in the accreted layer as a function of column. The top right panel shows the surface luminosity (*solid black curve*) and integrated nuclear heating (*dashed red curve*); the dots mark the current time (note that the nuclear heating extends beyond the plot limit, cf. Fig. 3.1, bottom). The bottom right panel shows the temperature as a function of column depth. The vertical dotted lines on the bottom panel indicate the range of column depth displayed in the main panel. The grey shaded region at right is the ^{56}Fe substrate of the initial, relaxed, NS envelope prior to accretion. The light tan region is where convection occurs. The temperature curve is turn orange-red where $\epsilon_{\text{nuc}} > 0.05\epsilon_{\text{CNO}}$ and red where $\epsilon_{\text{nuc}} > \epsilon_{\text{CNO}}$.



of the radius of the high-mass pulsar PSR J0740+6620 (Dittmann et al., 2024; Salmi et al., 2024), and anisotropy factor $\xi = 1$. A higher mass lowers the inferred column, as does a lower distance estimate. Our models with $0.05 \lesssim Z \lesssim 0.1$ overlap the range inferred from the 2019 burst. Any residual hydrogen from the previous outburst would increase the ignition column, so our estimate is a lower limit on the accreted mass available for the 2019 burst. However, if any additional bursts were missed in the NICER observations prior to the weak burst, which seems plausible, then the estimate for the column would be lowered. In fact, the 2005 RXTE burst reported by Casten et al. (2023) in Chapter 2 ignited a couple days earlier relative to the NICER burst after the onset of accretion and the column was about a factor of 2 less than the column accumulated prior to the NICER burst. This further supports the need for elevated metallicities.

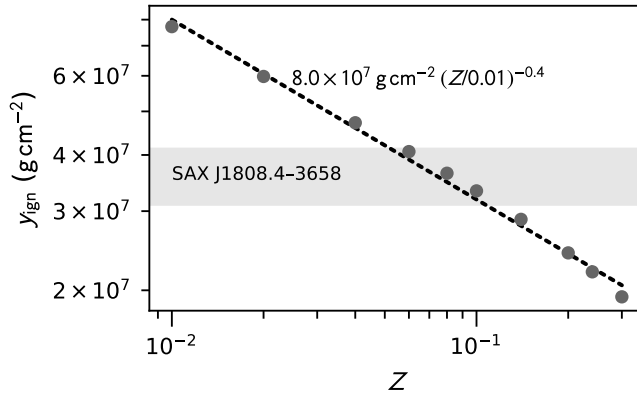


Figure 3.5 Ignition column depths (*dots*) as a function of metallicity in the accreted material. These decrease roughly as $y_{\text{ign}} \propto Z^{-0.4}$ (*dotted line*). For comparison, we also show the inferred burst ignition depth for SAX J1808.4–3658 (*shaded region*) which covers the range for $1.4 M_{\odot} < M < 2.0 M_{\odot}$. This range of inferred ignition depth is based on the observed fluence for the 2019 NICER burst with $d = 3.3$ kpc, $R = 12$ km, and anisotropy factor $\xi = 1$.

Figure 3.6 shows the burst lightcurves. The time axis is split, with linear scaling during the initial rise and logarithmic scaling during the long tail. The bursts segregate into two cases, *slow* and *fast*. For the slow bursts with $Z < 0.06$ (*dotted lines*), the burning is not sufficiently vigorous to launch convection, so the rise occurs on the thermal diffusion timescale. For $Z \geq 0.06$, a convective

a review, see Watts et al., 2016), so for simplicity we keep R fixed. Note, however, that the posterior credible interval for R from NICER observations (Dittmann et al., 2024; Salmi et al., 2024) spans over 2 km, which we do not factor in here.

zone briefly develops during the burst rise and produces a more pronounced peak. The contrast between peak and tail increases with Z as the initial spike in heating and resulting convection becomes stronger. At higher Z , there is sufficient ^{14}N to rapidly heat the envelope and drive convection. The resulting enhanced heat transport produces a sharp, rapid rise in the surface luminosity. Producing a strong peak thus requires a high concentration of CNO elements at the base of the accreted layer. Note, however, that only a small amount of hydrogen is consumed in the peak; for example, the equivalent α -value, defined as the ratio of the integrated accretion luminosity prior to the burst to the integrated luminosity of the peak, is $\alpha \gtrsim 10^3$ for $Z = 0.10$, far in excess of $\alpha \sim 100$ for typical X-ray bursts. Indeed, from observations discussed in Chapter 2, Casten et al. (2023) estimated the fraction of H consumed in the peak in the range 0.04–0.06 for a radius of 11 km.

For comparison, the observed 2019 NICER burst, with accretion luminosity (*horizontal dashed line*) subtracted, is overlaid (*gold points*). This observed burst has L_{peak} similar to that for our model $Z = 0.20$, albeit with a slower rise. This slower rise could be a consequence of the finite time for the burning front to propagate across the surface (see, e.g., Cavecchi et al., 2013; Eiden et al., 2020). In contrast, the measured tail luminosity (Casten et al., 2023) is consistent with $Z = 0.02$. Note that $\lesssim 500$ s of the tail were observed (Casten et al., 2023), however, which is far less than the expected tail duration.

For $Z \geq 0.02$, once the hydrogen in the accumulated layer is depleted, the burning shuts down and the envelope cools until sufficient hydrogen has again accreted to trigger an instability. The burning is then in a limit cycle that evolves on the secular timescale for the underlying neutron star ocean to warm. For $Z = 0.01$, the envelope mass increases during the long tail, which causes L to rise. As a result of this heating, the hot CNO burning does not shut down at the conclusion of the burst, but rather decreases to the lower steady-state burning rate L_{SS} (Eq. [3.6]).

The duration, and hence total energy radiated, of the tail depends on the hydrogen-to-CNO ratio (Eq. [3.2]). To check this scaling, we separate the lightcurves into peak and tail segments using the criteria described in § 3.3.1. We then integrate the luminosity to find the total energy radiated in the tail (E_{tail}). Fig. 3.7 displays t_{tail} (*top*) and E_{tail} (*bottom*). The burst models are in good agreement




Figure 3.6 Lightcurves for metallicities from $Z = 0.01$ to $Z = 0.3$, ordered lightest to darkest, respectively, and numbered by Z at the end of the tail. The lightcurves are aligned so that $L = 1.4 \times 10^{34} \text{ erg s}^{-1}$ at $t = 0$, and the time scale shifts from linear to logarithmic to follow the transition from peak to tail. Bursts for which no convection develops are indicated with dotted lines and are notable for their long rise. For comparison, we also plot (*gold*) the observations from the 2019 burst, including the measured tail luminosity (Casten et al., 2023), for $d = 3.3 \text{ kpc}$. Note that the observations have the pre-burst (accretion) emission subtracted. We also indicate (*horizontal dashed line*) the (Newtonian) accretion luminosity.

with our estimates (*dashed lines*) for t_d (Eq. [3.4]) and $L_{\text{CNO}} \times t_d$, where we compute L_{CNO} using Eq. ([3.5]) combined with the power-law fit (Fig. 3.5) for y_{ign} . Our estimate for E_{tail} differs from the MESA calculations at both low and high Z . At low Z , Eq. (3.4) leads to a slight underestimate (5%) of E_{tail} because the burst is sufficiently long-lived that the mass of the accreted layer significantly increases over the tail’s duration. As Z increases, t_{tail} decreases and the mass accumulated during t_{tail} becomes negligible. The estimate at high Z is 30% too high, however, because not all of the accumulated hydrogen is consumed. Locating the boundary (defined by where H has decreased to 50% of its surface abundance) that separates the unburnt hydrogen from the newly burnt ashes indicates that about 90% of the accumulated hydrogen is consumed during the burst.




Figure 3.7 Duration t_{tail} (*top*) and total radiated energy E_{tail} (*bottom*) of the burst tail as functions of the metallicity Z . For comparison, we also show (*dashed lines*) the expected duration of hot CNO burning, t_d , Eq. (3.4), and $L_{\text{CNO}} \times t_d$, where we compute L_{CNO} using Eq. (3.5) and the fit for ignition column (Fig. 3.5).

3.4 Discussion

Motivated by the recent discovery (Casten et al., 2023) of a weak flash during a recent outburst of SAX J1808.4–3658 discussed in Chapter 2, we have modeled the envelope of a cool, slowly accreting, neutron star, such as might occur just after the onset of an accretion outburst in a low-mass X-ray binary. Our models of hydrogen ignition span a range of envelope metallicities, and we follow the burst through the thermally unstable ignition of hydrogen and the transition to a steady-state burning via the hot CNO cycle. Our main conclusion is that there is indeed a regime of vigorous burning that launches convection and produces a sharp peak, in agreement with previous semi-analytical calculations (Peng et al., 2007; Cooper & Narayan, 2007). As with classical novae, the rapid growth of the instability cuts off after the initial flurry of proton captures locks the CNO nuclei into β^+ -decay bottlenecks. Producing a strong burst therefore requires enhanced CNO abundances. For our suite of models, we find that convection, and hence a well-defined, bright peak, only occurs for $Z \gtrsim 0.06$. At lower Z , the envelope heats on a thermal timescale and eases into a long plateau of hot CNO cycle burning. Without vigorous convection the flame front would not rapidly spread over the neutron star surface (Cavecchi et al., 2013) but perhaps would instead slowly propagate via a ring of fire (Bildsten, 1993). If so, the rise would be on the propagation time of the front and much longer than the rise in our one-dimensional simulations.

For higher- Z bursts, the initial, rapid peak in emission only consumes a small amount of H. Following the peak the burning settles into a quasi-steady-state with a luminosity set by the ignition column y_{ign} and metallicity, Eq. (3.5). This burning continues until the hydrogen is consumed, the timescale for which depends on the ratio X/Z in the accreted fuel (Eq. [3.4]). At the lowest metallicity ($Z = 0.01$) we explored, the burning did not extinguish at the end of the burst, but instead settled into burning hydrogen stably at the rate it accreted, Eq. (3.6).

Weak bursts are expected to occur primarily during the early stages of the accretion outburst, when the accretion rate is lower and the surface temperature cooler. In this study, we therefore focused on the first burst following the onset of accretion of hydrogen-rich material onto the neutron star envelope. The first burst is more energetic than subsequent bursts due to the initially cooler

temperature and the longer accumulation time required to reach ignition. As the envelope heats, the bursts are expected to evolve, and we see this in our simulations. For example, with $Z = 0.06$ y_{ign} decreases by 43% from the first to the second burst. The morphology of the lightcurve also evolves: during the first burst, the tail luminosity increases (see Fig. 3.6) as the nuclear heating warms the neutron star envelope. In subsequent bursts, the warmer envelope ensures a relatively constant tail luminosity.

We adopt a specific variation in composition: namely, holding X fixed while varying Z in order to isolate changes in metallicity. In reality, both hydrogen and helium abundances would vary. Models of helium bursts from SAX J1808.4–3658 at higher \dot{M} are consistent with solar metallicity. By tracking \dot{M} over an outburst and matching (He) burst recurrence times and fluences, Johnston et al. (2018) estimated $X = 0.44$ and $Z_{\text{CNO}} = 0.02$, for $d = 3.5$ kpc and $L_b/\dot{M} = 0.30$ MeV u⁻¹. By fitting a semi-analytical ignition model (Cumming & Bildsten, 2000b) to burst observations, Goodwin et al. (2019) inferred that $X = 0.57$ and $Z_{\text{CNO}} = 0.013$. These are consistent with the metallicity inferred from the tail luminosity (Casten et al., 2023), which aligns with our models having $Z = 0.02$. In this case, however, our models do not produce a strong peak; matching this peak with our models requires $Z \gtrsim 0.20$. If the weak burst observed by Casten et al. (2023) is indeed H-powered, then the base of the accreted layer must be greatly enriched relative to that in the accreted material.

As described in § 3.1, the luminosity and duration of the burst tail depend on the mass and abundance of hydrogen and CNO catalysts in the accreted fuel. Monitoring of the weak hydrogen flash, including the long tail, can therefore inform us not only about the composition of the accreted material, but also about the extent of sedimentation and mixing in the neutron star envelope. Figure 3.8 displays the distance-independent ratio $L_{\text{peak}}/L_{\text{tail}}$ over a range of L_{tail} , which depends on Z (indicated on top axis), both directly, Eq. (3.5), and via the ignition column (Fig. 3.5). Both L_{peak} and L_{tail} are sensitive to the CNO abundance, but at different times: L_{peak} is set by the amount of CNO available at the base of the accreted layer during the initial instability; whereas L_{tail} depends on the available CNO after the envelope has been mixed. Observations of these quantities

can therefore inform us not only about the metallicity of the accreted material, but also about the degree of stratification, such as arising from sedimentation or entrainment of previously burned material. Measurements of L_{peak} and L_{tail} from future observations that fall above the points in Fig. 3.8 would suggest an enhancement of CNO at the base of the layer.




Figure 3.8 The ratio of L_{peak} to L_{tail} , measured at the start of the tail. For comparison, we provide the metallicity corresponding to L_{tail} for our models along the top.

The onset of accretion in an X-ray transient is difficult to predict, and hence the first days of an accretion outburst tend to be sparsely observed. If H-triggered bursts are in fact confined to the few first days of accretion, then it is likely that our sample is quite incomplete. Recently it has become feasible for optical monitoring, by the X-ray Binary New Early Warning System (XB-NEWS), to provide the necessary trigger to catch the X-ray rise of a new outburst (Russell et al., 2019; Goodwin et al., 2020). This approach in fact enabled the detection of the 2019 NICER burst and has the potential to aid in observing future additional weak bursts. Obtaining more observations with continuous and complete coverage over the first few days of outbursts from SAX J1808.4–3658 and other X-ray transients could significantly increase the number of detected weak bursts and provide a more detailed picture of the physics in the neutron star envelope.

CHAPTER 4

CONCLUSION

4.1 Dissertation Summary

Hydrogen-triggered bursts have long been predicted at low accretion rates but had been evasive to detection. This dissertation addressed this elusive regime of unstable nuclear burning through a combined observational and theoretical investigation of weak bursts from the accreting millisecond pulsar SAX J1808.4–3658 with the goal of clarifying the conditions under which unstable hydrogen burning can produce observable X-ray bursts. These results demonstrate that H-triggered bursts can produce observable X-ray signatures under realistic conditions.

In Chapter 2, we presented an observational study of two weak bursts from the accreting millisecond pulsar SAX J1808.4–3658. Using NICER observations of a burst detected shortly after the 2019 outburst onset, together with an archival RXTE event from 2005, we showed that these bursts occur at low accretion rates and modest accumulated columns. The observed peak fluxes are factors of 15–30 below those of the bright He-powered bursts from SAX J1808.4–3658, and the post-burst emission exhibits behavior consistent with quasi-steady hot-CNO burning. Taken together, these properties provide strong observational evidence that these events are triggered by thermally unstable hydrogen burning rather than helium ignition. More broadly, these results suggest that weak, H-triggered bursts may represent a population of thermonuclear bursts that have been systematically missed due to observational biases against the earliest phases of the outburst.

In Chapter 3, we investigated the conditions required for unstable hydrogen ignition using the stellar evolution code MESA. Our models of the accreting neutron star envelope demonstrate that for solar metallicity hydrogen burning is too weak to drive convection and produces only a slow luminosity rise with a peak luminosity several orders of magnitude below the accretion luminosity—far below our ability of detection and does not match the recent observations. When the CNO abundance is enhanced ($Z \gtrsim 0.06$) however, proton captures onto CNO nuclei rapidly heat the envelope sufficiently to launch convection and produce a sharp luminosity peak comparable to the observed 2019 SAX J1808.4–3658 weak burst. Following the peak, the burning naturally settles into a quasi-

steady hot-CNO-powered tail whose luminosity and duration depend primarily on the ignition depth and CNO abundance. This tail is also seen in the observed bursts.

A similar sensitivity to CNO abundance is well established in classical novae, where enhanced metallicity leads to more rapid energy generation and more explosive ignition. In the neutron star envelope, processes such as sedimentation can enhance the abundance of CNO nuclei at the ignition depth, increasing the nuclear heating rate and modifying both the peak luminosity and the subsequent stable burning phase. The dependence of the burst morphology on CNO abundance therefore provides a potential diagnostic of mixing and diffusion processes in the neutron star envelope. These results in Chapter 3 show that weak, H-triggered bursts can reproduce the observed burst provided the accreted layer is sufficiently enriched in CNO isotopes, although some differences between the peak luminosity and tail luminosity remain in comparison with the observed burst. This indicates there may be additional physics may be in play, such as sedimentation, to enhance metallicity in the burst peak, rather than enhancing the metallicity over the entire burst.

In conclusion, the observational and theoretical results presented in this dissertation strengthen the case that H-triggered bursts not only occur on slowly accreting neutron stars but can also reach observable detection thresholds. The theoretical insights presented provide an initial framework for interpreting the weak bursts from SAX J1808.4–3658 and highlight the sensitivity of the burst properties to the metallicity and thermal state of the accreted envelope. These results establish a foundation for future observational searches and more detailed modeling of H-triggered bursts.

4.2 Future Modeling Outlook

While the results presented in this dissertation strengthen the case for observable H-triggered bursts on slowly accreting neutron stars, several important physical processes remain to be explored. Incorporating these effects will be essential for developing a more complete theoretical framework and for improving comparisons with observations.

One important next step is the inclusion of elemental diffusion and sedimentation in the accreted envelope. These processes are expected to operate in neutron star envelopes and may significantly modify the CNO abundance at the base of the accreted layer. Incorporating diffusion and sedi-

mentation into the MESA burst models will help determine whether the metallicity enhancements required in the present models can arise naturally during quiescence and early accretion within the first few days after the onset.

In addition, the models presented here are one-dimensional and therefore neglect the inherently multidimensional and turbulent nature of mixing and rotationally induced transport (e.g., meridional circulation and shear instabilities). Extending these calculations to multi-dimensional simulations would allow for a more realistic treatment of mixing in the accreted envelope and may alter the ignition conditions and burst morphology. In particular, mixing plays an important role in redistributing fuel and heat on the neutron star surface.

Another key limitation is the neglect of flame spreading. The present models assume instantaneous global ignition, whereas in reality thermonuclear burning is expected to begin locally and propagate across the surface. Incorporating flame spreading will be important for connecting ignition models to observed burst rise times, lightcurves, and burst oscillation behavior. Such work will also help determine whether H-triggered bursts exhibit distinct spreading signatures compared to He-triggered bursts.

4.3 Future Observational Outlook

A challenge to interpreting weak bursts is that they have only recently been detected despite decades of observing X-ray bursts. The unpredictable onset of accretion in X-ray transients often results in sparse coverage during the first days of an outburst. Given that the two weak bursts occurred 1–3 days after the onset of accretion, and cooler envelope temperatures are required for thermally unstable H ignition, it is plausible that weak bursts may be confined to the early stages of the outburst. If so, weak bursts would easily be missed without a reliable trigger for the onset of accretion. With the success of XB-NEWS’s optical-triggering campaign (Russell et al., 2019), we now have, for the first time, a promising way to search for weak bursts. Subsequent work correlating the optical and X-ray observations of the August 2019 outburst demonstrated a 4 day lag between the optical and X-ray rises (X-ray lagging the optical; Goodwin et al., 2020). These results provide a strong demonstration that optical monitoring can provide the necessary trigger for X-ray monitoring

to catch the X-ray rise of a new outburst and lead to the discovery of more weak bursts.

Observations of the onset of active accretion in LMXB transients, especially with coordinated multi-wavelength coverage, are rare because the rise to outburst typically occurs over only a few days. In the standard picture, transient outbursts are often associated with a thermal-viscous instability in the accretion disk, although irradiation, inner-disk truncation, and possibly additional source-dependent physics can significantly modify the evolution. Optical monitoring is particularly valuable because it can trace changes in the outer disk before the X-ray flux from the inner flow fully rises. The delay between the optical and X-ray brightening therefore probes the timescale on which the outburst propagates inward through the accretion flow and can constrain disk-evolution models, including the roles of heating-front propagation, truncation, and irradiation. Coordinated optical and X-ray campaigns therefore offer a rare opportunity to connect the development of the accretion flow during outburst onset with the earliest phases of thermonuclear burning on accreting neutron stars.

A systematic search for weak bursts in other slowly accreting neutron star systems will be essential for determining how common H-triggered bursts are across the bursting population. If weak bursts are indeed restricted to these early phases, they would provide a new probe of the accreted envelope, with their properties offering constraints on composition and mixing through their sensitivity to the CNO abundance. High-cadence X-ray coverage during the first few days of an outburst will be particularly important for testing whether weak bursts are confined to the earliest phases of renewed accretion. Although such coverage has historically been rare due to the rapid rise times of LMXB outbursts, ongoing and upcoming wide-field optical surveys (e.g., ZTF, GOTO, and Rubin) are expected to enable earlier identification of outbursts. In some cases, optical emission precedes the X-ray rise by several days, which allows for rapid X-ray follow-up during the onset of accretion.

BIBLIOGRAPHY

- Anders, E., & Grevesse, N. 1989, *Geochim. Cosmochim. Acta*, 53, 197, doi: 10.1016/0016-7037(89)90286-X
- Angulo, C., Arnould, M., Rayet, M., et al. 1999, *Nucl. Phys. A*, 656, 3, doi: 10.1016/S0375-9474(99)00030-5
- Antoniadis, J., Freire, P. C. C., Wex, N., et al. 2013, *Science*, 340, 448, doi: 10.1126/science.1233232
- Arnaud, K. A. 1996, in *Astronomical Society of the Pacific Conference Series*, Vol. 101, *Astronomical Data Analysis Software and Systems V*, ed. G. H. Jacoby & J. Barnes, 17
- Asplund, M., Grevesse, N., & Sauval, A. J. 2005, in *Astronomical Society of the Pacific Conference Series*, Vol. 336, *Cosmic Abundances as Records of Stellar Evolution and Nucleosynthesis*, ed. T. G. Barnes, III & F. N. Bash, 25
- Baade, W., & Zwicky, F. 1934, *Phys. Rev.*, 45, 138
- Belian, R. D., Conner, J. P., & Evans, W. D. 1976, *ApJ*, 206, L135
- Bildsten, L. 1993, *ApJ*, 418, L21, doi: 10.1086/187106
- Bildsten, L., & Chakrabarty, D. 2001, *ApJ*, 557, 292
- Bildsten, L., & Cumming, A. 1998, *ApJ*, 506, 842
- Bildsten, L., & Cutler, C. 1995, *ApJ*, 449, 800
- Bildsten, L., Salpeter, E. E., & Wasserman, I. 1992, *ApJ*, 384, 143
- Blouin, S., Shaffer, N. R., Saumon, D., & Starrett, C. E. 2020, *ApJ*, 899, 46, doi: 10.3847/1538-4357/ab9e75
- Boirin, L., Keek, L., Méndez, M., et al. 2007, *A&A*, 465, 559
- Brown, E. F., & Bildsten, L. 1998, *ApJ*, 496, 915
- Brown, E. F., & Cumming, A. 2009, *ApJ*, 698, 1020
- Bult, P., Chakrabarty, D., Arzoumanian, Z., et al. 2020, *ApJ*, 898, 38, doi: 10.3847/1538-4357/ab9827
- Bult, P., & van der Klis, M. 2015, *ApJ*, 806, 90, doi: 10.1088/0004-637X/806/1/90

- Bult, P. M., Gendreau, K. C., Arzoumanian, Z., et al. 2019, *The Astronomer's Telegram*, 13001, 1
- Burrows, A. 2013, *Reviews of Modern Physics*, 85, 245, doi: 10.1103/RevModPhys.85.245
- Cameron, A. G. 1959, *ApJ*, 130, 884, doi: 10.1086/146780
- Campana, S., & Di Salvo, T. 2018, in *Astrophysics and Space Science Library*, Vol. 457, *Astrophysics and Space Science Library*, ed. L. Rezzolla, P. Pizzochero, D. I. Jones, N. Rea, & I. Vidaña, 149, doi: 10.1007/978-3-319-97616-7_4
- Casanova, J., José, J., García-Berro, E., Shore, S. N., & Calder, A. C. 2011, *Nature*, 478, 490, doi: 10.1038/nature10520
- Casanova, J., José, J., & Shore, S. N. 2018, *A&A*, 619, A121, doi: 10.1051/0004-6361/201833422
- Cassisi, S., Potekhin, A. Y., Pietrinferni, A., Catelan, M., & Salaris, M. 2007, *ApJ*, 661, 1094, doi: 10.1086/516819
- Casten, S., Guichandut, S., Cumming, A., & Brown, E. F. 2025, *ApJ*, 989, 195, doi: 10.3847/1538-4357/adf106
- Casten, S., Strohmayer, T. E., & Bult, P. 2023, *ApJ*, 948, 117, doi: 10.3847/1538-4357/acc24f
- Cavecchi, Y., Watts, A. L., Braithwaite, J., & Levin, Y. 2013, *MNRAS*, 434, 3526, doi: 10.1093/mnras/stt1273
- Chakrabarty, D., & Morgan, E. H. 1998, *Nature*, 394, 346
- Chatziioannou, K., Cromartie, H. T., Gandolfi, S., et al. 2025, *Reviews of Modern Physics*, 97, 045007, doi: 10.1103/ymsq-cfcw
- Chugunov, A. I., Dewitt, H. E., & Yakovlev, D. G. 2007, *Phys. Rev. D*, 76, 025028, doi: 10.1103/PhysRevD.76.025028
- Clayton, D. D. 1983, *Principles of stellar evolution and nucleosynthesis* (University of Chicago Press)
- Cooper, R. L., & Narayan, R. 2007, *ApJ*, 661, 468, doi: 10.1086/513461
- Couch, S. M. 2017, *Philosophical Transactions of the Royal Society of London Series A*, 375, 20160271, doi: 10.1098/rsta.2016.0271
- Cox, J. P., & Giuli, R. T. 1968, *Principles of Stellar Structure*, Vol. 1 (New York: Gordon and Breach)

- Cumming, A. 2004, *Nuclear Physics B Proceedings Supplements*, 132, 435, doi: 10.1016/j.nuclphysbps.2004.04.078
- Cumming, A., & Bildsten, L. 2000a, *ApJ*, 544, 453, doi: 10.1086/317191
- . 2000b, *ApJ*, 544, 453
- Cyburt, R. H., Amthor, A. M., Ferguson, R., et al. 2010, *ApJS*, 189, 240, doi: 10.1088/0067-0049/189/1/240
- Demorest, P. B., Pennucci, T., Ransom, S. M., Roberts, M. S. E., & Hessels, J. W. T. 2010, *Nature*, 467, 1081
- Denissenkov, P. A., Herwig, F., Bildsten, L., & Paxton, B. 2013, *ApJ*, 762, 8, doi: 10.1088/0004-637X/762/1/8
- Di Salvo, T., Sanna, A., Burderi, L., et al. 2019, *MNRAS*, 483, 767, doi: 10.1093/mnras/sty2974
- Dittmann, A. J., Miller, M. C., Lamb, F. K., et al. 2024, *ApJ*, 974, 295, doi: 10.3847/1538-4357/ad5f1e
- Eiden, K., Zingale, M., Harpole, A., et al. 2020, *ApJ*, 894, 6, doi: 10.3847/1538-4357/ab80bc
- Ferguson, J. W., Alexander, D. R., Allard, F., et al. 2005, *ApJ*, 623, 585, doi: 10.1086/428642
- Fisker, J. L., Schatz, H., & Thielemann, F.-K. 2008, *ApJS*, 174, 261, doi: 10.1086/521104
- Fowler, W. A., & Hoyle, F. 1965, *Nucleosynthesis in massive stars and supernovae* (University of Chicago Press)
- Fujimoto, M. Y., Hanawa, T., & Miyaji, S. 1981, *ApJ*, 247, 267
- Fuller, G. M., Fowler, W. A., & Newman, M. J. 1985, *ApJ*, 293, 1, doi: 10.1086/163208
- Galloway, D. K., & Cumming, A. 2006, *ApJ*, 652, 559
- Galloway, D. K., & Keek, L. 2021, in *Astrophysics and Space Science Library*, Vol. 461, *Timing Neutron Stars: Pulsations, Oscillations and Explosions*, ed. T. M. Belloni, M. Méndez, & C. Zhang, 209–262, doi: 10.1007/978-3-662-62110-3_5
- Galloway, D. K., Muno, M. P., Hartman, J. M., Psaltis, D., & Chakrabarty, D. 2008, *ApJS*, 179, 360
- Galloway, D. K., in 't Zand, J., Chenevez, J., et al. 2020, *ApJS*, 249, 32, doi: 10.3847/1538-4365/ab9f2e

- Gendreau, K. C., Arzoumanian, Z., & Okajima, T. 2012, in Proc. SPIE, Vol. 8443, Space Telescopes and Instrumentation 2012: Ultraviolet to Gamma Ray, 844313, doi: 10.1117/12.926396
- Giacconi, R., Gursky, H., Paolini, F. R., & Rossi, B. B. 1962, Phys. Rev. Lett., 9, 439
- Gold, T. 1968, Nature, 218, 731, doi: 10.1038/218731a0
- Goodwin, A. J., Galloway, D. K., Heger, A., Cumming, A., & Johnston, Z. 2019, MNRAS, 490, 2228, doi: 10.1093/mnras/stz2638
- Goodwin, A. J., Russell, D. M., Galloway, D. K., et al. 2020, MNRAS, 498, 3429, doi: 10.1093/mnras/staa2588
- Grindlay, J., Gursky, H., Schnopper, H., et al. 1976, ApJ, 205, L127, doi: 10.1086/182105
- Hansen, C. J., & Van Horn, H. M. 1975, ApJ, 195, 735
- Harding, A. K. 2013, Frontiers of Physics, 8, 679, doi: 10.1007/s11467-013-0285-0
- Heinke, C. O., Jonker, P. G., Wijnands, R., Deloye, C. J., & Taam, R. E. 2009, ApJ, 691, 1035, doi: 10.1088/0004-637X/691/2/1035
- Henry, L., & L'Ecuyer, J. L. 1969, ApJ, 156, 549
- Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., & Collins, R. A. 1968, Nature, 217, 709
- Hoffman, J. A., Lewin, W. H. G., Doty, J., et al. 1978, ApJ, 221, L57
- Holt, J. W., & Lim, Y. 2019, in American Institute of Physics Conference Series, Vol. 2127, Xiamen-CUSTIPEN Workshop on the Equation of State of Dense Neutron-Rich Matter in the Era of Gravitational Wave Astronomy (AIP), 020019, doi: 10.1063/1.5117809
- Hoyle, F., & Fowler, W. A. 1965, in Quasi-Stellar Sources and Gravitational Collapse, ed. I. Robinson, A. Schild, & E. L. Schucking, 17
- Iglesias, C. A., & Rogers, F. J. 1993, ApJ, 412, 752, doi: 10.1086/172958
- . 1996, ApJ, 464, 943, doi: 10.1086/177381
- Illiano, G., Papitto, A., Sanna, A., et al. 2022, arXiv e-prints, arXiv:2212.09778. <https://arxiv.org/abs/2212.09778>
- in't Zand, J. J. M., Galloway, D. K., Marshall, H. L., et al. 2013, A&A, 553, A83, doi: 10.1051/0004-6361/201321056

- Irwin, A. W. 2004, The FreeEOS Code for Calculating the Equation of State for Stellar Interiors.
<http://freeeos.sourceforge.net/>
- Itoh, N., Hayashi, H., Nishikawa, A., & Kohyama, Y. 1996, *ApJS*, 102, 411, doi: 10.1086/192264
- Janka, H.-T. 2012, *Annual Review of Nuclear and Particle Science*, 62, 407, doi: 10.1146/annurev-nucl-102711-094901
- Jermyn, A. S., Schwab, J., Bauer, E., Timmes, F. X., & Potekhin, A. Y. 2021, *ApJ*, 913, 72, doi: 10.3847/1538-4357/abf48e
- Jermyn, A. S., Bauer, E. B., Schwab, J., et al. 2023, *ApJS*, 265, 15, doi: 10.3847/1538-4365/acae8d
- Johnston, Z., Heger, A., & Galloway, D. K. 2018, *MNRAS*, 477, 2112, doi: 10.1093/mnras/sty757
- . 2020, *MNRAS*, 494, 4576, doi: 10.1093/mnras/staa1054
- José, J., Moreno, F., Parikh, A., & Iliadis, C. 2010, *ApJS*, 189, 204, doi: 10.1088/0067-0049/189/1/204
- Keek, L., Galloway, D. K., in't Zand, J. J. M., & Heger, A. 2010, *ApJ*, 718, 292, doi: 10.1088/0004-637X/718/1/292
- Keek, L., & Heger, A. 2017, *ApJ*, 842, 113, doi: 10.3847/1538-4357/aa7748
- Lampe, N., Heger, A., & Galloway, D. K. 2016, *The Astrophysical Journal*, 819, 46, doi: 10.3847/0004-637X/819/1/46
- Landau, L. D. 1932, *Phys. Zs. Sowjet*, 1, 285
- Langanke, K., & Martínez-Pinedo, G. 2000, *Nuclear Physics A*, 673, 481, doi: 10.1016/S0375-9474(00)00131-7
- Lewin, W. H. G. 1977, in *Eighth Texas Symposium on Relativistic Astrophysics*, ed. M. D. Papa-
giannis, Vol. 302, 210, doi: 10.1111/j.1749-6632.1977.tb37050.x
- Linares, M., Shahbaz, T., & Casares, J. 2018, *ApJ*, 859, 54, doi: 10.3847/1538-4357/aabde6
- Mahmoodifar, S., & Strohmayer, T. 2013, *ApJ*, 773, 140, doi: 10.1088/0004-637X/773/2/140
- Markwardt, C. B., Swank, J., Wijnands, R., & in't Zand, J. 2005, *The Astronomer's Telegram*, 505,
1
- Miller, M. C., Lamb, F. K., Dittmann, A. J., et al. 2019, *ApJ*, 887, L24, doi: 10.3847/2041-8213/ab50c5

- Oda, T., Hino, M., Muto, K., Takahara, M., & Sato, K. 1994, *Atomic Data and Nuclear Data Tables*, 56, 231, doi: 10.1006/adnd.1994.1007
- Oppenheimer, J. R., & Volkoff, G. M. 1939, *Phys. Rev.*, 55, 374
- Pacini, F. 1967, *Nature*, 216, 567, doi: 10.1038/216567a0
- Page, D., Lattimer, J. M., Prakash, M., & Steiner, A. W. 2009, *ApJ*, 707, 1131
- Patruno, A., Rea, N., Altamirano, D., et al. 2009, *MNRAS*, 396, L51, doi: 10.1111/j.1745-3933.2009.00660.x
- Paxton, B., Bildsten, L., Dotter, A., et al. 2011, *ApJS*, 192, 3, doi: 10.1088/0067-0049/192/1/3
- Paxton, B., Cantiello, M., Arras, P., et al. 2013, *ApJS*, 208, 4, doi: 10.1088/0067-0049/208/1/4
- Paxton, B., Marchant, P., Schwab, J., et al. 2015, *ApJS*, 220, 15, doi: 10.1088/0067-0049/220/1/15
- Paxton, B., Schwab, J., Bauer, E. B., et al. 2018, *ApJS*, 234, 34, doi: 10.3847/1538-4365/aaa5a8
- Paxton, B., Smolec, R., Schwab, J., et al. 2019, *ApJS*, 243, 10, doi: 10.3847/1538-4365/ab2241
- Peng, F., Brown, E. F., & Truran, J. W. 2007, *ApJ*, 654, 1022, doi: 10.1086/509628
- Potekhin, A. Y., & Chabrier, G. 2010, *Contributions to Plasma Physics*, 50, 82, doi: 10.1002/ctpp.201010017
- Potekhin, A. Y., Chabrier, G., & Yakovlev, D. G. 1997, *A&A*, 323, 415
- Poutanen, J. 2017, *ApJ*, 835, 119, doi: 10.3847/1538-4357/835/2/119
- Prialnik, D., Shara, M. M., & Shaviv, G. 1979, *A&A*, 72, 192
- Remillard, R. A., Loewenstein, M., Steiner, J. F., et al. 2022, *AJ*, 163, 130, doi: 10.3847/1538-3881/ac4ae6
- Riley, T. E., Watts, A. L., Bogdanov, S., et al. 2019, *ApJ*, 887, L21, doi: 10.3847/2041-8213/ab481c
- Rogers, F. J., & Nayfonov, A. 2002, *ApJ*, 576, 1064, doi: 10.1086/341894
- Russell, D. M., Bramich, D. M., Lewis, F., et al. 2019, *Astronomische Nachrichten*, 340, 278, doi: 10.1002/asna.201913610
- Salmi, T., Choudhury, D., Kini, Y., et al. 2024, *ApJ*, 974, 294, doi: 10.3847/1538-4357/ad5f1f

Sanna, A., Di Salvo, T., Burderi, L., et al. 2017, MNRAS, 471, 463, doi: 10.1093/mnras/stx1588

Saumon, D., Chabrier, G., & van Horn, H. M. 1995, ApJS, 99, 713, doi: 10.1086/192204

Sazonov, S., Paizis, A., Bazzano, A., et al. 2020, New Astronomy Reviews, 88, 101536, doi: 10.1016/j.newar.2020.101536

Schatz, H., Aprahamian, A., Barnard, V., et al. 2001, Phys. Rev. Lett., 86, 3471

Sharma, R., Sanna, A., & Beri, A. 2022, MNRAS, doi: 10.1093/mnras/stac3779

Starrfield, S. 1993, in Astrophysics and Space Science Library, Vol. 177, Astrophysics and Space Science Library, ed. J. Sahade, G. E. McCluskey, & Y. Kondo, 209, doi: 10.1007/978-94-011-2416-4_12

Strohmayer, T., & Bildsten, L. 2006, in Compact stellar X-ray sources, Vol. 39 (Cambridge University Press), 113–156

Taam, R. E., & Picklum, R. E. 1979, ApJ, 233, 327

Thorne, K. S. 1977, ApJ, 212, 825

Timmes, F. X., & Swesty, F. D. 2000, ApJS, 126, 501, doi: 10.1086/313304

Tolman, R. C. 1939, Physical Review, 55, 364

Townsend, R. 2024, MESA SDK for Linux, 23.7.3, Zenodo, doi: 10.5281/zenodo.10624843

Verner, D. A., Ferland, G. J., Korista, K. T., & Yakovlev, D. G. 1996, ApJ, 465, 487, doi: 10.1086/177435

Wallace, R. K., & Woosley, S. E. 1981, ApJS, 45, 389

Wallace, R. K., Woosley, S. E., & Weaver, T. A. 1982, ApJ, 258, 696

Watts, A. L., Andersson, N., Chakrabarty, D., et al. 2016, Reviews of Modern Physics, 88, 021001, doi: 10.1103/RevModPhys.88.021001

Wheeler, J. A. 1966, ARA&A, 4, 393, doi: 10.1146/annurev.aa.04.090166.002141

Wijnands, R., & van der Klis, M. 1998, Nature, 394, 344

Wijnands, R., Wolt, M. K., Linares, M., et al. 2005, The Astronomer's Telegram, 507, 1

Woosley, S. E., Heger, A., Cumming, A., et al. 2004, ApJS, 151, 75

Worpel, H., Galloway, D. K., & Price, D. J. 2013, *ApJ*, 772, 94, doi: 10.1088/0004-637X/772/2/94

Yakovlev, D. G., & Pethick, C. J. 2004, *ARA&A*, 42, 169