MONASH UNIVERSITY

DOCTORAL THESIS

# Modelling Thermonuclear X-ray Bursts on Accreting Neutron Stars

Author:
Zac JOHNSTON

Supervisors:
Prof. Alexander HEGER
Assoc. Prof. Duncan GALLOWAY

A thesis submitted in fulfilment of the requirements
for the degree of Doctor of Philosophy

Faculty of Science
School of Physics and Astronomy

16 February 2020

Declaration of Authorship

I, Zac Johnston, hereby declare that this thesis contains no material which has been accepted for the award of any other degree or diploma at any university or equivalent institution and that, to the best of my knowledge and belief, this thesis contains no material previously published or written by another person, except where due reference is made in the text of the thesis.

This thesis includes one (1) original paper published in a peer reviewed journal and one (1) unpublished publication. My contribution to these works is listed on the following page. The core theme of the thesis is the computational modelling of type-I X-ray bursts. The ideas, development and writing up of the papers in the thesis were the principal responsibility of myself, the student, working within the School of Physics and Astronomy, Monash University, under the supervision of Prof. Alexander Heger and Assoc. Prof. Duncan K. Galloway. The inclusion of co-authors reflects the fact that the work came from active collaboration between researchers and acknowledges input into team-based research.

Signed,

Zac Johnston   10 October 2019

The undersigned hereby certify that the above declaration correctly reflects the nature and extent of the student's and co-authors' contributions to this work.

Signed,

Prof. Alexander Heger   10 October 2019

Chapter	Publication Title	Status	Student's Contribution	Co-authors' Contribution	Monash Student Co-authors
5	Simulating X-ray bursts during a transient accretion event	Published	80% Computational models, data analysis, writing manuscript.	Alexander Heger: 10% Feedback and discussion Duncan Galloway: 10% Feedback and discussion	No
6	Multi-epoch X-ray burst modelling: MCMC with large grids of 1D simulations	Submitted	80% Computational models, data analysis, writing manuscript.	Alexander Heger: 10% Feedback and discussion. Duncan Galloway: 10% Feedback and discussion.	No

MONASH UNIVERSITY

Abstract

Faculty of Science
School of Physics and Astronomy

Doctor of Philosophy

Modelling Thermonuclear X-ray Bursts on Accreting Neutron Stars

by Zac JOHNSTON

In low-mass X-ray binaries, the accretion of stellar material onto a neutron star can fuel unstable thermonuclear flashes known as Type I X-ray bursts. In a matter of seconds, the thin shell of hydrogen and/or helium is converted into heavier elements through nuclear fusion, heating the envelope to $\sim 10^9$ K. The burst of thermal emission, dominated by X-rays, lasts ≈ 10 $- 100$ s and is observable with satellite-based X-ray telescopes. The properties of the burst reflect the local conditions of the neutron star surface. Simulating these events using computational models can provide valuable information about the nature of the accreting system. Measuring neutron star properties, especially the mass and radius, has been a longstanding objective in astrophysics because it can constrain the equation of state of dense nuclear matter. One-dimensional (1D) astrophysics codes with large nuclear reaction networks are the current state-of-the-art for simulating X-ray bursts. These codes can track the evolution of isotopes through thousands of nuclear reaction pathways, to predict the released nuclear energy and final composition of the ashes. In this thesis, I make extensive use of KEPLER, a 1D code at the forefront of these efforts. I first present improvements to the setup and analysis of KEPLER burst models. By accounting for nuclear heating in the initial conditions, I shorten the thermal burn-in time, thereby reducing computational expense and producing more consistent burst trains. To model bursts fueled by transient accretion events, I perform the first such simulations with fully time-dependent accretion rates. Building upon previous efforts to model the "Clocked Burster", GS 1826−238, I precompute a grid of 3840 simulations and sample the interpolated results using Markov Chain Monte Carlo (MCMC) methods. By comparing the predictions to multi-epoch observations, I obtain posterior probability distributions for the system parameters. I then extend these MCMC methods to the pure-helium burster, 4U 1820−30, using a grid of 168 simulations. Finally, I discuss potential improvements for future studies, to further develop the computational modelling of accreting neutron stars.

Acknowledgements

I would like to thank my supervisors, Alexander Heger and Duncan Galloway, for their immeasurable feedback, support, and encouragement.

For useful discussions, I would like to thank Adam Jacobs, Laurens Keek, Hendrik Schatz, Ed Brown, Frank Chambers, Zach Meisel, Andrew Casey, Adelle Goodwin, and countless others.

I thank my office mates and close friends, Conrad Chan, David Liptai, and Hayley Macpherson, for keeping me sane.

I thank Belinda Jude for her crucial support over the years.

I thank my parents, Carolyn and Peter Johnston, for their love and support since the very beginning, and for their unwavering encouragement throughout my studies.

I thank Sara Hugentobler, for her love and support. I couldn't have done it without you.

Contents

List of Abbreviations

1D	1-Dimensional
2D	2-Dimensional
3D	3-Dimensional
AMXP	Accreting Millisecond X-ray Pulsar
GR	General Relativity
LMXB	Low Mass X-ray Binary
MCMC	Markov Chain Monte Carlo
NICER	Neutron Star Interior Composition Explorer
NS	Neutron Star
PNS	Proto-Neutron Star
PRE	Photospheric Radius Expansion
RXTE	Rossi X-ray Timing Explorer
XRB	X-Ray Burst

List of Symbols

d	Distance	cm
E_b	Burst energy	erg
f_b	Burst fluence	$\mathrm{erg\,cm^{-2}}$
f_p	Persistent fluence	$\mathrm{erg\,cm^{-2}}$
F_b	Burst flux	$\mathrm{erg\,s^{-1}\,cm^{-2}}$
F_{Edd}	Eddington flux	$\mathrm{erg\,s^{-1}\,cm^{-2}}$
F_p	Persistent flux	$\mathrm{erg\,s^{-1}\,cm^{-2}}$
F_{peak}	Peak burst flux	$\mathrm{erg\,s^{-1}\,cm^{-2}}$
g	Gravitational Acceleration	$\mathrm{cm\,s^{-2}}$
i	Inclination	deg
L_{acc}	Accretion luminosity	$\mathrm{erg\,s^{-1}}$
L_b	Burst luminosity	$\mathrm{erg\,s^{-1}}$
L_{Edd}	Eddington luminosity	$\mathrm{erg\,s^{-1}}$
L_p	Persistent luminosity	$\mathrm{erg\,s^{-1}}$
L_{peak}	Peak burst luminosity	$\mathrm{erg\,s^{-1}}$
M	Stellar mass	$M_\odot$
$\dot{M}$	Global accretion rate	$\dot{M}_{Edd}$
$\dot{M}_{Edd}$	Global Eddington accretion rate	$M_\odot\,\mathrm{yr^{-1}}$
$\dot{m}$	Local accretion rate	$\dot{m}_{Edd}$
$\dot{m}_{Edd}$	Local Eddington accretion rate	$\mathrm{g\,cm^{-2}\,s^{-1}}$
Q_b	Crustal heating rate	$\mathrm{MeV\,nucleon^{-1}}$
Q_{nuc}	Nuclear heating rate	$\mathrm{MeV\,nucleon^{-1}}$
R	Stellar radius	km
t_{pre}	Pre-burst time	s
t_{start}	Burst start time	s
t_{peak}	Burst peak time	s
t_{end}	Burst end time	s
X_0	Hydrogen mass fraction	–
Y_0	Helium mass fraction	–
y	Column depth	$\mathrm{g\,cm^{-2}}$
y_{ig}	Ignition column depth	$\mathrm{g\,cm^{-2}}$
Z_{CNO}	CNO mass fraction	–
z	Gravitational redshift	–
α	Alpha ratio	–
Δt	Recurrence time	s
ν	Burst rate	$\mathrm{s^{-1}}$
κ	Opacity	$\mathrm{cm^2\,g^{-1}}$
ξ	GR-corrected radius ratio	–
ξ_b	Burst anisotropy	–
ξ_p	Persistent anisotropy	–
τ	Autocorrelation time	–
φ	GR-corrected mass ratio	–

Chapter 1

Introduction and Background

Type I X-ray bursts are recurring thermonuclear flashes on accreting neutron stars. They are distinct from Type II X-ray bursts, which are caused by sporadic accretion (Hoffman, Marshall, and Lewin, 1978).

Throughout this work, we will simply use "bursts" to refer to Type I X-ray bursts. Previous reviews have been provided by Lewin, Paradijs, and Taam (1993), Bildsten (1997), Strohmayer and Bildsten (2006), and Galloway and Keek (2017).

In this chapter, we provide a brief overview of Type I X-ray bursts. We describe the first detections and modelling efforts (Section 1.1), the catalogues of burst observations (Section 1.2), the mechanisms of burst ignition (Section 1.3), and previous works with computational burst modelling (Section 1.4).

In Chapter 2, we describe the astrophysical code used throughout this thesis, KEPLER, and the process of extracting burst properties from the models and predicting observable quantities. In Chapter 3, we present improvements to the setup of KEPLER burst models, and direct comparisons between KEPLER and MESA burst models. In Chapter 4, we present the first burst simulations with time-dependent accretion rates. In Chapter 5, we present the application of Markov chain Monte Carlo methods to large grids of KEPLER models. In Chapter 6, we present the extension of these MCMC methods to a hydrogen-poor system, 4U 1820−30. In Chapter 7, we summarise the work presented in this thesis, and discuss potential improvements for future work.

1.1 Early History

The thin shell instability was discovered by Schwarzschild and Härm (1965), in which thermonuclear burning restricted to a thin shell ($\Delta R \ll R$) can undergo a thermal runaway due to its inability to expand and cool. Hansen and Horn (1975) modelled the thermonuclear stability of accreted hydrogen and helium on neutron stars, and found that most configurations were subject to the thin shell instability. Concurrent to this work, Type I X-ray bursts were discovered independently by Babushkina et al. (1975), Grindlay et al. (1976), and Belian, Conner, and Evans (1976), shortly followed with further detections by Lewin et al. (1976) and Clark et al. (1976).

Woosley and Taam (1976) and Maraschi and Cavaliere (1977) independently attributed the newly-discovered bursts to a thermonuclear origin – the unstable burning regime first uncovered by Hansen and Horn (1975). The thermonuclear model was further developed by Joss (1977), Lamb and

Lamb (1978), and Taam and Picklum (1978). Alongside these efforts, Sugimoto and Fujimoto (1978) and Fujimoto and Sugimoto (1979) examined the analogous phenomenon of helium shell flashes on accreting white dwarfs. Shortly thereafter, Fujimoto, Hanawa, and Miyaji (1981) presented a foundational work placing the various ignition conditions of hydrogen and helium under a consistent framework.

The first detailed numerical models were produced by Joss (1978), who simulated helium bursts using a modified one-dimensional stellar evolution code, ASTRA (Rakavy, Shaviv, and Zinamon, 1967). The characteristic properties of the observed bursts were successfully reproduced, including the onset and decay timescales, overall luminosities, and recurrence times. These models marked the beginning of the computational modelling of X-ray bursts (Section 1.4).

1.2 Observational Catalogues

New generations of X-ray telescopes were launched in the following decades, and the growing "zoo" of bursting behaviour offered multiple pathways into the study of thermonuclear burning on neutron stars. Large catalogues of these observations enable population studies of bursting systems, and the identification of global patterns in bursting behaviour. In particular, the collections of burst data serve as test beds for computational models.

An early collection was compiled from the literature by Paradijs, Penninx, and Lewin (1988), containing 45 bursts from ten systems. They found that the burst duration was anti-correlated with the persistent flux, hinting at a common relationship between systems.

Cornelisse et al. (2003) compiled 1823 bursts from nine systems observed with *BeppoSAX*, covering the years 1996 to 2002. Their analysis found global transitions between bursting regimes. When the persistent luminosity increased to $L_{\mathrm{p}} \approx 2 \times 10^{37} \, \mathrm{erg \, s^{-1}}$, the sources transitioned from long, frequent, and quasi-periodic bursts, to short and infrequent bursts. Above this limit, the bursts grew more irregular until ceasing altogether above $L_{\mathrm{p}} \approx 5.5 \times 10^{37} \, \mathrm{erg \, s^{-1}}$.

Galloway et al. (2008) provided another extensive catalogue, with observations from the X-ray telescope, *RXTE*. Bringing together 1187 individual bursts from 48 accreting neutron stars, this catalogue further enabled the study of diverse bursting patterns. The phenomena included burst oscillations, short-waiting time bursts, photospheric radius-expansion bursts, and the unexpected "turn over" in burst rate at high accretion rates. A successor to this catalogue, the *Multi-INstrument Burst ARchive* (MINBAR)[1], extends the library to over 7000 bursts from 85 source, with data from multiple instruments, including *RXTE*, *BeppoSAX*, and *INTEGRAL*. Its unprecedented size makes MINBAR the most comprehensive burst catalogue to date, and we make use of its data throughout this thesis.

---

[1]http://burst.sci.monash.edu/minbar

1.3 The Burst Mechanism

Thermonuclear X-ray bursts are the result of stellar material accreting onto neutron stars. In the idealised picture, the accreted material, rich in hydrogen and/or helium, spreads evenly over the surface and forms a shell only $\sim 1\,\mathrm{m}$ thick. The base of this envelope is steadily buried deeper as new material is accreted on top, and is compressed to higher pressures and temperatures. Under hydrostatic equilibrium, the pressure at the base is given by the weight of the overlying fuel,

$$P = yg, \tag{1.1}$$

where $y \sim 10^7$–$10^8\,\mathrm{g\,cm^{-2}}$ is the column depth and $g \sim 10^{14}\,\mathrm{cm\,s^{-2}}$ is the neutron star surface gravity.

When the base of the envelope reaches temperatures of $T \sim 10^8\,\mathrm{K}$, nuclear burning can become thermally unstable. An increase in the rate of nuclear burning increases the temperature, in turn accelerating fusion in an unstable feedback loop. This thermal runaway occurs when the temperature sensitivity of nuclear heating exceeds that of radiative cooling. For a fixed pressure, this threshold is given by

$$\frac{d\epsilon_{\mathrm{nuc}}}{dT} = \frac{d\epsilon_{\mathrm{cool}}}{dT}, \tag{1.2}$$

where ϵ_{nuc} is the specific nuclear heating rate, and ϵ_{cool} is the specific radiative cooling rate.

Thermonuclear burning spreads throughout the envelope as a convective "flame", consuming most of the available fuel in $\lesssim 1\,\mathrm{s}$. The nuclear flash releases $\sim 10^{39}\,\mathrm{erg}$, heating the envelope to $\sim 10^9\,\mathrm{K}$. Due to the extreme gravitational potential of $\approx 200\,\mathrm{MeV\,nucleon^{-1}}$, in contrast to $\approx 5\,\mathrm{MeV\,nucleon^{-1}}$ from nuclear burning, the vast bulk of the accreted material remains bound to the surface. As the outflowing thermal energy reaches the neutron star photosphere, it radiates strongly in X-rays. The surface luminosity peaks within ≈ 1–$10\,\mathrm{s}$, and then decays to pre-burst levels as envelope cools over the following ≈ 10–$100\,\mathrm{s}$. Examples of observed lightcurves are shown in Figure 1.1.

After a burst, the accreted envelope has been processed into "ashes" – the leftover products of nuclear fusion. Fresh material is accreted on top, burying the ashes and eventually incorporating them into the neutron star crust. The composition of these ashes impacts the thermal properties of the crust that is formed (e.g., Brown and Bildsten, 1998; Gupta et al., 2007).

1.3.1 Nuclear pathways

The dominant nuclear pathways on accreting neutron stars include the hot (β-limited) CNO cycle, the 3α (triple-α) process, the α-process, the αp-process, and the rp-process. Detailed descriptions of these pathways and their role in bursts can be found in Lewin, Paradijs, and Taam (1993), Woosley et al. (2004), Fisker, Schatz, and Thielemann (2008), and José et al. (2010).

For temperatures of $T > 8 \times 10^7\,\mathrm{K}$, hydrogen is converted to helium through the hot CNO cycle (e.g., Strohmayer and Bildsten, 2006; Wiescher
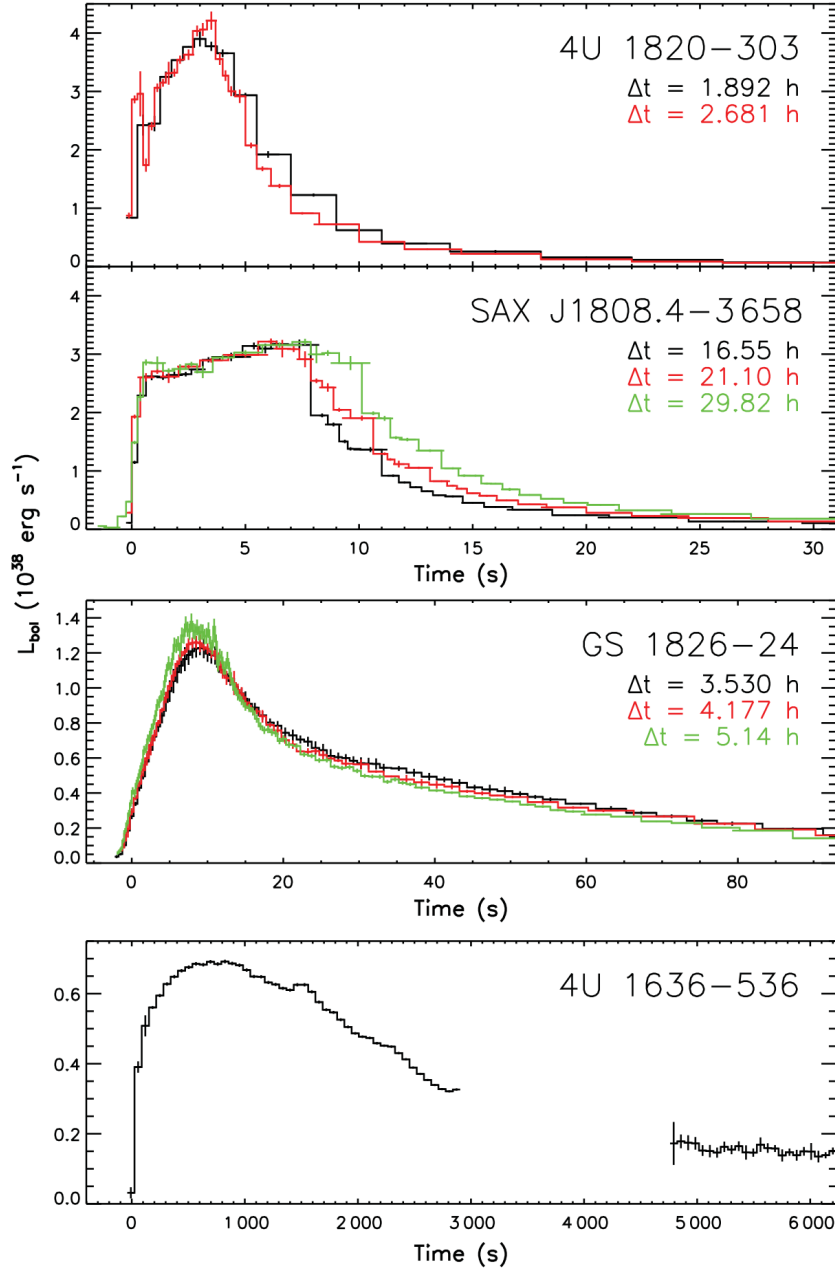
Figure 1.1: The variety of lightcurves observed from four bursting sources (figure from Galloway, Goodwin, and Keek, 2017). 4U 1820−30 is an ultra-compact binary accreting almost pure helium, and exhibits short PRE bursts at relatively high accretion rates ($\dot{m} \approx 0.2\,\dot{m}_{\mathrm{Edd}}$). In SAX J1808.4−3658, the low accretion rates ($\dot{m} \lesssim 0.05\,\dot{m}_{\mathrm{Edd}}$) allow hydrogen to be depleted through hot CNO burning, leading to powerful PRE bursts with a characteristic "plateau" at the Eddington luminosity. GS 1826−238 accretes roughly solar material, $X_0 \approx 0.7$, where the hydrogen fuels *rp*-process burning and leads to long burst tails. The persistent accretor, 4U 1636−536, in addition to hydrogen-rich bursts similar to GS 1826−238, has exhibited four superbursts – extremely long bursts with recurrence times of $\Delta t \sim 1\,\mathrm{yr}$.

et al., 2010),

$$^{12}\text{C}(p, \gamma)^{13}\text{N}(p, \gamma)^{14}\text{O}(\beta^+)^{14}\text{N}(p, \gamma)^{15}\text{O}(\beta^+)^{15}\text{N}(p, \alpha)^{12}\text{C}, \qquad (1.3)$$

where the net result is $4p \rightarrow \alpha$. In this catalytic cycle, the rate is limited by slow β-decays, and the burning of hydrogen becomes independent of temperature. When active, steady CNO burning between bursts dominates the heating of the accreted layers. For a given parcel of accreted material, the time to deplete hydrogen (e.g., as derived in Lampe, Heger, and Galloway, 2016) is given by

$$t_{\text{CNO}} \approx 20\,\text{h}(\frac{X_0}{0.7})(\frac{Z_{\text{CNO}}}{0.01})^{-1}, \qquad (1.4)$$

where X_0 is the initial hydrogen mass fraction and cno is the mass fraction of CNO isotopes. Equation 1.4 is only approximate because X_0 and Z_{CNO} are mass fractions, not number fractions, and depend on the chosen distribution of CNO isotopes.

The 3α-process converts helium into carbon with a net reaction of $3\alpha \rightarrow^{12}\text{C}$. Due to its strong temperature sensitivity, the 3α-process is thermally unstable across typical accretion rates, and serves as the ignition reaction for most burst regimes (described further in § 1.3.2).

For temperatures above $T \gtrsim 1 \times 10^9\,\text{K}$, successive $alpha$ captures overtake 3α as the main source of energy (Fujimoto, Hanawa, and Miyaji, 1981), building progressively heavier elements,

$$^{12}\text{C}(\alpha, \gamma)^{16}\text{O}(\alpha, \gamma)^{20}\text{Ne}(\alpha, \gamma)^{24}\text{Mg}(\alpha, \gamma)^{28}\text{Si} \dots \qquad (1.5)$$

If hydrogen is present in the accreted fuel, a complex network of reactions can build heavy isotopes to the iron-group and beyond (Wallace and Woosley, 1981; Hanawa and Fujimoto, 1984). For temperatures of $T \gtrsim 4 \times 10^8\,\text{K}$, the break-out reactions of $^{14}\text{O}(\alpha, p)^{17}\text{F}$ and $^{15}\text{O}(\alpha, \gamma)^{19}\text{Ne}$ destroy catalysts of the CNO cycle, and interrupt the burning of hydrogen to helium (Lewin, Paradijs, and Taam, 1993). This breakout paves the way for the αp-process, whereby sequences of (α, p) and (p, γ) reactions proceed to iron-group nuclei.

The presence of hydrogen at these high temperatures also allows the rapid proton (rp) process to occur. Successive proton-captures and β^+-decays proceed along the proton drip-line far from the valley of stability, and produce heavy nuclei with atomic mass numbers of $A \approx 60$–100 (Wallace and Woosley, 1981; Schatz et al., 1998; Koike et al., 1999; Schatz et al., 2001; Koike et al., 2004) The rp-process is relatively slow in comparison to the α reactions of the burst ignition. This delayed energy release results in extended burst tails which are characteristic of mixed hydrogen/helium bursts (e.g., GS 1826$-$238 in Figure 1.1).

The nuclear pathways described above span thousands of reactions, the vast majority of which have not been measured experimentally. When nuclear reaction rates are manually varied within their uncertainty, burst models predict altered lightcurves and ashes (Koike et al., 1999; Parikh et al., 2008; Parikh et al., 2009). Existing uncertainties in reaction rates thus contribute to burst model uncertainties, and limit their ability to constrain observed systems. Sensitivity studies are used to determine which reactions have the strongest influence on burst properties, and thus which nuclear

experiments to prioritise in future (Cyburt et al., 2016; Meisel, Merz, and Medvid, 2019).

1.3.2   Bursting regimes

Different regimes of bursting are predicted to occur depending on the local conditions. The composition of the accreted fuel, the accretion rate, and the thermal structure of the envelope all influence the thermonuclear path to ignition. The transitions between these regimes were first laid out by Fujimoto, Hanawa, and Miyaji (1981), and summarised again by Bildsten (1997). Most recently, Galloway, Goodwin, and Keek (2017) revised the list to a total of seven classifications.

   We provide here a brief summary of these regimes as predicted by models, adapted from Galloway, Goodwin, and Keek (2017). The transitions assume an accreted composition of $X_0 = 0.73$ and $Z_{\mathrm{CNO}} = 0.02$, and a crustal heating strength of $Q_{\mathrm{b}} = 0.1\,\mathrm{MeV\,nucleon^{-1}}$. The accretion rates are given as a fraction of the canonical Eddington rate for solar composition, $\dot{m}_{\mathrm{Edd}} = 8.775 \times 10^4\,\mathrm{g\,cm^{-2}\,s^{-1}}$ (equivalent to a global rate of $\dot{M}_{\mathrm{Edd}} = 1.75 \times 10^{-8}\,\mathrm{M_\odot\,yr^{-1}}$ for a neutron star mass and radius of $M = 1.4\,\mathrm{M_\odot}$ and $R = 10\,\mathrm{km}$, and assuming spherically-symmetric accretion). Note that Cases I and II have not yet been observed, and that the transition to stable hydrogen/helium burning is observed to occur at roughly half the predicted accretion rate listed here.

   I   For $\dot{m} \lesssim 0.001\,\dot{m}_{\mathrm{Edd}}$ ($T \lesssim 7 \times 10^8\,\mathrm{K}$), the CNO cycle is thermally unstable, and hydrogen ignition occurs. The hydrogen burst also triggers unstable helium burning.

  II   For $0.001 \lesssim \dot{m} \lesssim 0.004\,\dot{m}_{\mathrm{Edd}}$, the hydrogen burst of Case I is too shallow to ignite helium, due to sedimentation of helium and the CNO isotopes (Peng, Brown, and Truran, 2007). The accreted helium continues to build a column of fuel below the hydrogen bursts, eventually reaching unstable helium ignition on its own.

 III   For $0.004 \lesssim \dot{m} \lesssim 0.08\,\dot{m}_{\mathrm{Edd}}$ ($T \gtrsim 7 \times 10^8\,\mathrm{K}$), hydrogen burning is stable, and burns completely to helium via the hot CNO cycle. A pure helium layer accumulates below the hydrogen-burning region, and eventually reaches unstable helium ignition.

  IV   For $\dot{m} \approx 0.1\,\dot{m}_{\mathrm{Edd}}$, no bursts occur because helium is stably burned to carbon before it can ignite, although this regime may lead to carbon-fuelled superbursts (Keek and Heger, 2016).

   V   For $0.1 \lesssim \dot{m} \lesssim 1.0\,\dot{m}_{\mathrm{Edd}}$, hydrogen is steadily burned to helium, as in Case III, but the helium ignites before the hydrogen is depleted. The mixed hydrogen/helium burst can produce heavy ashes through a combination of 3α, αp and rp-processes (Wallace and Woosley, 1981).

  VI   For $\dot{m} \approx 1.0\,\dot{m}_{\mathrm{Edd}}$, the accreted fuel undergoes oscillatory burning, producing a "marginally stable" burning regime (Heger, Cumming, and Woosley, 2007).

VII For $\dot{m} \gtrsim 1.0\,\dot{m}_{\mathrm{Edd}}$, both hydrogen and helium burning are stable, and no bursts occur. This model-predicted transition to stable hydrogen/helium burning occurs much later than inferred from observations, which stabilise around $0.1 \lesssim \dot{m} \lesssim 0.3\,\dot{m}_{\mathrm{Edd}}$ (Paradijs, Penninx, and Lewin, 1988; Cornelisse et al., 2003).

These regimes produce two broad categories of observed bursts: helium (He), and hydrogen/helium (H/He) bursts.

Helium bursts are typically characterised by short rise times ($< 1\,\mathrm{s}$), short overall durations ($\lesssim 10\,\mathrm{s}$), and photospheric radius-expansion (PRE; Tawara et al., 1984; Lewin, Vacca, and Basinska, 1984). During PRE, spectral fits to the bursts yield an expanded blackbody radius, a reduced effective temperature, and approximately constant bolometric luminosity. These features are thought to indicate that the bursts reach the Eddington luminosity – the limit at which radiation pressure balances gravitational pressure. Example lightcurves of observed PRE bursts are shown in the upper two panels of Figure 1.1.

In contrast, mixed hydrogen/helium bursts are characterised by relatively long rise times ($\approx$ 1–5 s), broad sub-Eddington peaks, and long tails ($\sim 10^2\,\mathrm{s}$). Their burst profiles are understood to result from a more extended nuclear energy release than He bursts, leading to longer tails. Example lightcurves of H/He bursts are shown in the third panel of Figure 1.1, observed from the famous "clocked burster", GS 1826$-$238.

1.4 Computational Models

Much of our current understanding of the bursting mechanism can be attributed to computational models. In the decades since the first efforts (Section 1.1), models of X-ray bursts have progressed from relatively simple sets of analytic equations to state-of-the-art astrophysics codes. These codes are now capable of simulating dozens of sequential bursts, while tracking the nuclear reactions of thousands of isotopes. We briefly summarise here previous modelling efforts, with a focus on one-dimensional burst codes, in particular KEPLER, which we make extensive use of in this thesis.

One-dimensional (1D) burst codes approximate the neutron star envelope as a spherically symmetric shell. While 1D models lack the inherently multi-dimensional (multi-D) effects of convection and flame spreading (e.g., Zingale et al., 2015; Cavecchi et al., 2016), multi-D models are constrained by computational expense to a simulation time of $\lesssim 1\,\mathrm{s}$. Computationally cheaper methods, such as semi-analytic models (e.g., Cumming, 2003) and one-zone models (e.g., Schatz et al., 2001), allow for extensive parameter explorations, but lack feedback from nuclear burning on the radial temperature and density profiles, and ultimately their influence on the burst lightcurves. 1D burst codes, therefore, remain the best tools currently available for performing parameter studies which self-consistently follow the mixing and burning of accreted material.

AGILE is an implicit 1D general relativistic code (Liebendörfer, Rosswog, and Thielemann, 2002). The code was used by Fisker et al. (2006) to model X-ray bursts and test the lower limit of the $^{15}\mathrm{O}(\alpha, \gamma)^{19}\mathrm{Ne}$ reaction rate, using a reaction network of 298 isotopes. AGILE was subsequently

used to model the extent of the *rp*-process in mixed H/He bursts, using a network with 304 isotopes (Fisker, Schatz, and Thielemann, 2008).

SHIVA is a 1D astrophysics code that was originally applied to classical novae on accreting white dwarfs (José and Hernanz, 1998). The code was used by to model X-ray bursts with a nuclear network of 324 isotopes (José et al., 2010), computing a simulation with a setup similar to the "ZM" KEPLER model from Woosley et al. (2004). SHIVA predicted recurrence times that were twice as long as KEPLER, which is likely at least partly due to the incorrect opacities previously used in the latter (see Section 3.1).

MESA is another 1D astrophysics code capable of simulating bursts (Paxton et al., 2015). Compositions of burst ashes predicted using MESA were used in Meisel and Deibel (2017) to examine their impact on crust cooling. MESA was not used for large-scale burst modelling until the multi-epoch models of GS 1826−238 by Meisel (2018, in Chapter 5 we present multi-epoch KEPLER models using the same reference dataset). A follow-up study also investigated the influence of reaction rate uncertainties on the inferred neutron star properties (Meisel, Merz, and Medvid, 2019).

1.4.1 KEPLER

KEPLER (Weaver, Zimmerman, and Woosley, 1978) is a 1D stellar astrophysics code which has been used to model multiple aspects of the bursting process. In Chapter 2, we discuss the setup and execution of these models.

KEPLER was first applied to bursts by Wallace, Woosley, and Weaver (1982), who produced four simulations of hydrogen/helium fuel, using only a 19-isotope nuclear network. The first use of KEPLER in its modern burst configuration was performed by Woosley et al. (2004). They introduced a fully adaptive nuclear network, which could simulate the extended nuclear reactions of the rp-process up to the proton drip line, for the first time in 1D.

Expanding on this work, Heger et al. (2007) performed seven models and compared the predictions with the system GS 1826−238 – the Clocked Burster. One of these models, labelled *A3*, matched the observed lightcurve morphology with surprising accuracy, and is now a common reference point for burst models.

Keek and Heger (2011) applied KEPLER to the superburst regime, modelling the ignition of deep carbon oceans and the resulting hours-long bursts. In a follow-up study, Keek, Heger, and Zand (2012) accreted helium-rich fuel onto a carbon ocean already close to ignition, to simulate a superburst during regular bursts. Keek and Heger (2016) discovered a stable helium burning regime (Case IV in Section 1.3.2), which only occurred for a narrow range of accretion rates around $\dot{m} \approx 0.1\,\dot{m}_{\mathrm{Edd}}$ and which could explain the production of carbon oceans as superburst fuel.

Lampe, Heger, and Galloway (2016) presented the largest grid of 1D models to date. They explored the dependence of burst properties on accretion rate and metallicity, and compared the results to observed trends.

Cyburt et al. (2016) used KEPLER to explore the sensitivity of burst models to uncertainties in nuclear reaction rates. By varying key reaction rates in the nuclear network, they ranked rates by their influence on the burst lightcurves and properties. They found that the CNO breakout reaction $^{15}\mathrm{O}(\alpha, \gamma)^{19}\mathrm{Ne}$ had the strongest sensitivity.

Chapter 2

Methods: Modelling X-ray Bursts

For the research presented in this thesis, we use the KEPLER code to simulate X-ray bursts. In this chapter, we provide an overview of the KEPLER model for bursting (Section 2.1), the methods used to extract burst properties from the model output (Section 2.2), and the general relativity (GR) corrections applied to the Newtonian KEPLER quantities (Section 2.3).

2.1   KEPLER: A 1D Hydrodynamic Burst Code

KEPLER was developed in the 1970s for modelling the pre-supernova evolution of massive stars (Weaver, Zimmerman, and Woosley, 1978). In the decades since, the code has been applied to regimes of stellar evolution and explosive nucleosynthesis (e.g., Woosley, Heger, and Weaver, 2002; Rauscher et al., 2002; Heger et al., 2003; Woosley et al., 2004). In this thesis, we will focus on the application of KEPLER to thermonuclear X-ray bursts. A selection of previous burst studies are described in Section 1.4.1.

Previous descriptions of KEPLER for burst modelling can be found in Woosley et al. (2004) and Keek and Heger (2011) A KEPLER burst model consists of a Lagrangian grid of zones in the radial direction. The zones span a thin shell of material at the neutron star surface, extending down from column depths of $y \sim 10^3 \, \mathrm{g \, cm^{-2}}$ to the deep ocean at $y \sim 10^{12} \, \mathrm{g \, cm^{-2}}$. The neutron star crust is located below this lower boundary. Zones are added and removed according to zoning parameters, which can be tuned to control the grid resolution. Convection is parametrised in 1D using mixing length theory, where the diffusion coefficient is set by the estimated convective velocity. KEPLER uses an adaptive nuclear network which can simulate all the nuclear processes described in Section 1.3.1. Isotopes are dynamically added and removed from the network during the simulation as they are created and destroyed. A public reaction rate library has been maintained at REACLIB[1] (Cyburt et al., 2010), and we use this library for the KEPLER models presented in Chapter 5 and Chapter 6.

2.1.1   Input Parameters

Parameters can be used to adjust the properties and behaviour of the modelled envelope. Resolution and zoning parameters control the numerical structure, and are tuned to ensure convergence of the results. The particular system being modelled is represented by physical parameters, such as

---

[1] https://reaclib.jinaweb.org/

the accretion rate and fuel composition. These variables can be explored and modified to determine the likely values which best reproduce the observed data. We can thus use these parameters to understand and constrain the physical characteristics of observed bursting systems.

Mass, Radius, and Gravity

The parameter for the neutron star mass, M, sets the gravitational mass interior to the base of the model. Along with the radius, R, this determines the gravitational acceleration, g, experience by each zone. The thin shell of the envelope, $\Delta R \ll R$, results in a gravitational acceleration that differs by $\lesssim 1\,\%$ across the model domain.

Chemical Composition

The accreted composition is set with the parameters of the hydrogen mass fraction, X_0, and the CNO mass fraction ("metallicity"), Z_{CNO}. For simplicity, the CNO metallicity is accreted into KEPLER as ^{14}N, which is the dominant isotope from stellar CNO burning. The remainder of the fuel is put into helium (^{4}He): $Y_0 = 1 - X_0 - Z_{\mathrm{CNO}}$.

Accretion Rate

The mass accretion rate, $\dot{M}$, is the dominant parameter for modelling X-ray bursts. It sets the rate at which nuclear fuel is added to the envelope, and thus how quickly an explosive layer can be built up. It also determines the total rate of crustal heating, Q_{b}. The accretion rate of a KEPLER model can be constant to represent persistent accretors, or time-varying to represent transient accretors (see Chapter 4). The accretion rate is used during both model initialisation (Section 2.1.2) and execution (Section 2.1.3). It is typically expressed as a fraction of the Eddington-limited accretion rate for solar composition, $\dot{M}_{\mathrm{Edd}} = 1.75 \times 10^{-8}\,\mathrm{M_\odot\,yr^{-1}}$. Although the actual $\dot{M}_{\mathrm{Edd}}$ depends on the composition and neutron star gravity, the canonical value serves as a common reference point between models and codes.

Crustal Heating

The crustal heating parameter, Q_{b}, controls the heat flowing into the envelope from the crust below. It is effectively a lower boundary condition, setting the heat flux at the innermost zone. During accretion, the weight of new material compresses the neutron star crust, inducing electron captures and pycnonuclear (density-driven) reactions (Haensel and Zdunik, 2008). The total energy yield from this process is 1–2 $\mathrm{MeV\,nucleon^{-1}}$ (Haensel and Zdunik, 2008), but the net flux reaching the envelope is typically assumed to be $\approx 0.1\,\mathrm{MeV\,nucleon^{-1}}$ (e.g., Heger et al., 2007).

Assumptions

No model is without assumptions and limitations. KEPLER, being a 1D code, inherently assumes spherical symmetry. The only spatial degree of freedom is in the radial direction, meaning that energy, heat, and chemical composition can only move "up" or "down", but not across the surface. To

a certain degree this is an adequate approximation on the global scale, but it cannot truly capture the multidimensional nature of accretion, convection, turbulence, and the lateral spreading of the nuclear flame across the neutron star surface (e.g., Shara, 1982; Zingale et al., 2015; Cavecchi et al., 2013). These effects likely play important roles on the ignition conditions, the rise of the burst lightcurve, and quasi-periodic oscillations (Watts, 2012).

KEPLER burst models do not include rotation or magnetic fields, which can also influence the dynamics of flame spreading (Cavecchi et al., 2015; Cavecchi et al., 2016).

2.1.2 Setup and Initialisation

The setup phase of a model involves multiple steps. First, an approximate thermal and chemical profile is constructed. The composition profile consists of a heavy substrate lying below a lighter envelope. The substrate, representing the heavy ashes of previous burning, is composed of ^{54}Fe and acts as an inert thermal buffer for the heating generated during the model. New material is then accreted, without nuclear burning, to build the outer layers to the thin surface. With accretion and nuclear reactions disabled, the thermal profile is relaxed, with the base heat flux at the inner boundary determined by the Q_b parameter. The envelope then relaxes into thermal equilibrium, where the flux leaving the surface is equal to the flux entering the base. In Section 3.2, we present an improvement to this method, where we include a heat source representing nuclear burning, which improves the model equilibrium and reduces burn-in. The model is now initialised and ready for the full simulation to begin.

2.1.3 Execution and Output

The full simulation is started by switching on accretion and the nuclear reaction network. The model accretes material at the specified rate and composition, and nuclear reactions begin processing the fuel during steady burning until the unstable ignition of a burst occurs (Section 1.3). The bursting process repeats, producing a "train" of sequential bursts. We typically choose the total simulation time to generate the desired number of bursts, using a prediction of the recurrence time. The primary model output we use in this thesis is the lightcurve – the bolometric surface luminosity as a function of time. A model lightcurve with a train of bursts is shown in Figure 2.1. The individual bursts can then be sliced out (as described below), and the burst properties calculated.

2.2 Extracting Bursts from Models

For the extraction of burst properties from KEPLER models, we have developed software tools over the course of this project, which we have collected under PYBURST, a PYTHON library available on the open-source platform github[2]. We describe here the general procedure of the model analysis.

Given a model lightcurve (Figure 2.1), we identify each burst and extract its properties. We then average the properties over the burst sequence

---

[2]https://github.com/zacjohnston/pyburst

Figure 2.1: Example lightcurves from a KEPLER simulation. An entire model is plotted in the upper panel, and the individual extracted bursts stacked in the lower panel. This figure also appears in Johnston, Heger, and Galloway (2019) (Chapter 5).

(excluding some number of initial bursts, Section 2.2.2) to obtain summary properties for the model. The standard deviation of these quantities indicates the inherent burst-to-burst variation, which we take as the model uncertainty (note that this does not include systematic uncertainties, for example due to reaction rate uncertainties).

To analyse a model, PYBURST follows a pipeline that proceeds roughly as follows:

    I Identify local maxima in the model lightcurve.

    II Discard non-burst maxima, such as numerical spikes or bumps.

    III Identify the start and end of each burst lightcurve.

    IV Calculate the properties of each burst.

    V Average the properties across all bursts.

From this procedure, we obtain a collection of bursts with their individual and averaged properties, along with associated uncertainties.

2.2.1 Calculating burst properties

The most common burst properties of interest are the recurrence time, Δt, the peak luminosity, L_{peak}, the burst energy (sometimes called fluence, see below), E_{b}, and the alpha ratio, α.

Identifying bursts

Firstly, the maxima in the model lightcurve are found. We save computational time, and avoid false-positives from small bumps, by imposing a minimum luminosity threshold. If we are confident that the modelled bursts will always peak above $10^{37}\,\mathrm{erg\,s^{-1}}$, for example, we need only consider these sections of the lightcurve. Once we have identified the maxima in a lightcurve, we filter out any that are not deemed to be bursts. Once the candidate burst peaks, L_{peak}, have been verified, PYBURST proceeds to calculate the remaining burst properties.

Recurrence time and burst rate

The recurrence time, Δt, is the elapsed time since the previous burst. Specifically, PYBURST defines it as the time between burst peaks. Note that the recurrence time is (by definition) undefined for the first burst in the sequence. The burst rate is simply $\nu = 1/\Delta t$.

**Ignition column**

A burst ignites when the base of the accreted layer reaches unstable thermonuclear conditions. Given a burst with a recurrence time of Δt and a constant accretion rate of $\dot{m}$, the mass of the accreted layer is then simply

$$M_{\mathrm{acc}} = 4\pi R^2 \dot{m}\Delta t, \tag{2.1}$$

and the ignition depth is

$$y_{\mathrm{ig}} = \dot{m}\Delta t. \tag{2.2}$$

Figure 2.2: An example KEPLER burst lightcurve with key points as identified by PYBURST. The green circles indicate, from left to right, t_{pre}, t_{start}, and t_{end}. The orange circle indicates the burst peak, $(t_{\mathrm{peak}}, L_{\mathrm{peak}})$.

Lightcurve points

To extract the burst energetics, we determine the start and end points of the burst lightcurve (Figure 2.2). We define a reference point t_{pre}, that is a set time interval prior to the peak, chosen to ensure the entire burst rise is captured, $30\,\mathrm{s}$, for example. The start of the burst rise, t_{start}, is defined as the point where the lightcurve has reached some fraction of the peak luminosity, $25\,\%$, for example. The burst end, t_{end}, is defined as the point where the luminosity has decayed to a given fraction of the peak luminosity, $1\,\%$, for example.

Burst energetics

The burst energetics are calculated using the above lightcurve reference points. The burst energy, E_{b}, is obtained by integrating the luminosity between t_{pre} and t_{end}, after subtracting a baseline luminosity, L_{pre}, which is taken at t_{pre}. Some emission prior to the burst may be included in this integration, but the contribution is negligible ($< 0.1\,\%$) if L_{pre} is subtracted.

Note that the burst energy here is often called the burst fluence, which is used interchangeably with the observed quantity of time-integrated burst flux ($\int F_{\mathrm{b}}\,\mathrm{d}t$). Although these quantities are indeed related (Section 2.4.4), the term "fluence" specifically refers to time-integrated flux, and to avoid confusion we will here maintain the distinction between the burst energy, E_{b}, and the fluence, f_{b}.

α (alpha) ratio

We then obtain α, the ratio between the persistent energy and burst energy:

$$\alpha = \frac{-\phi \Delta t \dot{M}}{E_{\mathrm{b}}}, \tag{2.3}$$

where $\phi = -zc^2/(1+z)$ is the gravitational potential at the neutron star surface. This ratio represents the relative efficiency of accretion compared to thermonuclear burning, and can be used to infer the dominant composition of the accreted fuel. A small α ratio of ≈ 40 corresponds to strong nuclear energy efficiency, indicating hydrogen-rich fuel. A large α ratio of $\gtrsim 100$ corresponds to relatively inefficient nuclear energy generation, indicating fuel that is hydrogen-deficient (i.e., helium-rich).

2.2.2 Average Model Properties

Having calculated the above properties for each burst in the sequence, we then calculate summary quantities for the model as a whole. Because the initial bursts are often much more energetic than the subsequent train (Woosley et al., 2004), they're usually treated as a "burn-in" phase of the simulation, and excluded from analysis. In previous studies, only the first one or two bursts were excluded (e.g., Heger et al., 2007; Cyburt et al., 2010; Lampe, Heger, and Galloway, 2016). We discovered the presence of extended burn-in (Section 3.2).

We can also calculate average burst lightcurves for comparison to observations. The most common approach is to stack the burst lightcurves, aligned by peak, for example, and calculate the average and standard deviation of the luminosity along the curve. With this method, a mean burst lightcurve can be obtained and compared with observations.

2.3 Correcting for General Relativity

Due to its origin as a stellar evolution code, KEPLER uses Newtonian gravity, which is sufficiently accurate for regular stars. X-ray bursts, however, occur in a highly-relativistic environment[3]. To accurately model bursts, it is crucial to account for the effects of General Relativity (GR) when comparing models with observations.

Fortunately, the results of Newtonian KEPLER models require only straightforward corrections, due to the thin-shell nature of the envelope. These corrections have been described in detail previously (Woosley et al., 2004; Keek and Heger, 2011; Lampe, Heger, and Galloway, 2016), and we will provide here a brief summary. This section largely follows the notation and conventions used in Appendix B of Keek and Heger (2011).

In this section we will signify Newtonian KEPLERquantities with the subscript ′k′, GR quantities with the subscript ′g′, and quantities in the frame of a distant observer with the subscript ′∞′. These corrections rely on the assumption of a thin shell at the surface of a neutron star.

2.3.1 Definitions

To resolve the discrepancy between the Newtonian gravity of KEPLER and the GR gravity of actual neutron stars, we first note that gravitational acceleration is approximately constant throughout the thin shell of an accreted

---

[3]A canonical neutron star, with mass $M = 1.4\,M_\odot$ and radius $R = 10\,\mathrm{km}$, has a gravitational acceleration at the surface of $g \approx 2.4 \times 10^{14}\,\mathrm{cm\,s^{-2}}$ (i.e. $\approx$8000 times the speed of light per second).

envelope. For a 10 m thick envelope on the surface of a canonical neutron star, the acceleration differs by only $\approx 2\,\%$ from top to bottom.

Under the approximation that gravity is constant, a Newtonian KEPLER model with a mass and radius $(M_{\mathrm{k}}, R_{\mathrm{k}})$ is equivalent to an "actual" neutron star with a different mass and radius $(M_{\mathrm{g}}, R_{\mathrm{g}})$ if it has the same acceleration under GR.

The acceleration under Newtonian gravity is given by

$$g_{\mathrm{k}} = \frac{GM}{R^2},\tag{2.4}$$

where G is the gravitational constant. The acceleration under GR is instead given by

$$g_{\mathrm{g}} = \frac{GM(1+z)}{R^2},\tag{2.5}$$

where the gravitational redshift is given by

$$1 + z = \frac{1}{\sqrt{1 - \frac{2GM}{c^2 R}}}.\tag{2.6}$$

In other words, $g_{\mathrm{g}} = (1+z)g_{\mathrm{k}}$ for a given mass and radius[4]. If we impose the requirement that the acceleration is equal under both the Newtonian and GR regimes, we obtain

$$
\begin{aligned}
g_{\mathrm{k}} &= g_{\mathrm{g}}, \\
\frac{GM_{\mathrm{k}}}{R_{\mathrm{k}}^2} &= \frac{GM_{\mathrm{g}}(1+z)}{R_{\mathrm{g}}^2}, \\
\left(\frac{R_{\mathrm{g}}}{R_{\mathrm{k}}}\right)^2 &= \frac{M_{\mathrm{g}}}{M_{\mathrm{k}}}(1+z).
\end{aligned}
\tag{2.7}
$$

Note that from here on, $(1+z)$ from Equation (2.6) is always calculated with the GR mass and radius. If we define the ratios[5]

$$\varphi = \frac{M_{\mathrm{g}}}{M_{\mathrm{k}}}, \quad \xi = \frac{R_{\mathrm{g}}}{R_{\mathrm{k}}},\tag{2.8}$$

we can write Equation (2.7) as

$$\xi^2 = \varphi(1+z).\tag{2.9}$$

If we define the gravitational radius for the Newtonian case

$$\zeta = \frac{GM_{\mathrm{k}}}{c^2 R_{\mathrm{k}}},\tag{2.10}$$

we can also write Equation (2.6) as

$$1 + z = \frac{1}{\sqrt{1 - \frac{2\varphi\zeta}{\xi}}}.\tag{2.11}$$

---

[4] For a canonical neutron star, $1 + z \approx 1.26$.
[5] Note that the radius ratio ξ is unrelated to the anisotropy factors ξ_{b} and ξ_{p}, which we will introduce later.

Substituting Equation (2.11) into Equation (2.9) we obtain

$$\varphi^2 + 2\varphi\xi^3\zeta - \xi^4 = 0. \tag{2.12}$$

Thus, a Newtonian KEPLER model with a given mass and radius (M_k, R_k) corresponds to any pair (M_g, R_g) which satisfies Equation (2.12). Note that a given gravitational acceleration corresponds to a contour of mass-radius pairs. Depending on which quantities are already known, we can solve the above problem for the remaining variables.

Solving for Mass

If the actual neutron star radius (R_g) is known, we can determine the mass (M_g) by solving Equation (2.12) for the mass ratio:

$$\varphi = \zeta\xi^3 \left(\sqrt{1 + \zeta^{-2}\xi^{-2}} - 1 \right), \tag{2.13}$$

and then using $M_\mathrm{g} = \varphi M_\mathrm{k}$.

Solving for Radius

If instead the actual neutron star mass (M_g) is known, we can find the radius (R_g) by solving Equation (2.12) for the radius ratio:

$$\xi = \frac{\zeta\varphi}{2} \left(1 + \sqrt{1 - A} + \sqrt{2 + A + \frac{2}{\sqrt{1 - A}}} \right), \tag{2.14}$$

where we have defined

$$A = \sqrt[3]{\frac{2}{9}} \frac{\left(\frac{B^2}{\varphi^2} - 2\sqrt[3]{6} \right)}{B\zeta^2},$$

$$\tag{2.15}$$

$$B = \sqrt[3]{9\zeta^2\varphi^4 + \sqrt{3}\varphi^3 \sqrt{16 + 27\zeta^4\varphi^2}}.$$

We can then simply use $R_\mathrm{g} = \xi R_\mathrm{k}$

2.3.2 Correcting local Newtonian quantities

In the previous section, we solved for the GR neutron star masses and radii (M_g, R_g) which are equivalent to a Newtonian KEPLER model. The physical quantities predicted by the model are then corrected to be consistent with the actual mass and radius.

Accretion rate

The Newtonian KEPLER model is equivalent to a GR-corrected neutron star with the same local accretion rate, $\dot{m}$. From Equation (2.8), the ratio of the GR-corrected neutron star surface area to the Newtonian surface area is equal to ξ^2. The global accretion rate, $\dot{M} = 4\pi R^2 \dot{m}$, is thus scaled according

to

$$\begin{aligned}
\dot{M}_g &= \xi^2 \dot{M}_k \\
&= \varphi(1+z)\dot{M}_k,
\end{aligned} \tag{2.16}$$

where we have used the relation $\xi^2 = \varphi(1+z)$ from Equation (2.9).

Luminosity

Similarly, the Newtonian luminosities are scaled by the surface area,

$$\begin{aligned}
L_g &= \xi^2 L_k \\
&= \varphi(1+z)L_k.
\end{aligned} \tag{2.17}$$

Accretion luminosity

The accretion luminosity (L_{acc}) originates from material accreting at a certain rate ($\dot{M}$) through a gravitational potential (ϕ),

$$L_{acc} = -\dot{M}\phi. \tag{2.18}$$

The potentials for the Newtonian and GR regimes are given by

$$\phi_k = -\frac{GM_k}{R_k}, \quad \phi_g = -\frac{c^2 z}{1+z}. \tag{2.19}$$

Using Equations (2.16) and (2.19), we obtain

$$\begin{aligned}
L_{acc,g} &= -\dot{M}_g \phi_g \\
&= -\varphi(1+z)\dot{M}_k \phi_g \\
&= \varphi c^2 z \dot{M}_k.
\end{aligned} \tag{2.20}$$

Burst energy (fluence)

The burst energy (E_b) is the total energy radiated during the burst, corresponding to the time-integrated burst luminosity ($\int L_b \, dt$). Using the area ratio (ξ^2), we obtain

$$\begin{aligned}
E_{b,g} &= \xi^2 E_{b,k} \\
&= \varphi(1+z)E_{b,k}.
\end{aligned} \tag{2.21}$$

Eddington luminosity

The Eddington luminosity (for a spherically-symmetric object) is found by balancing the gravitational force with the radiation pressure, and is given by

$$L_{Edd} = \frac{4\pi R^2 g c}{\kappa}, \tag{2.22}$$

where κ is the opacity. Substituting g from Equations (2.4) and (2.5), we have

$$L_{Edd,g} = \frac{4\pi G c(1+z)M_g}{\kappa}, \quad L_{Edd,k} = \frac{4\pi G c M_k}{\kappa}. \tag{2.23}$$

Noting that $M_g = \varphi M_k$ and $1 + z = \xi^2/\varphi$, we obtain

$$
\begin{aligned}
L_{\mathrm{Edd,g}} &= \frac{4\pi G c \xi^2 M_k}{\kappa} \\
&= \xi^2 L_{\mathrm{Edd,k}}
\end{aligned}
\tag{2.24}
$$

We could also have noted that g is equal in both cases (by definition), leaving a factor of $R_g^2/R_k^2 = \xi^2$.

2.3.3 Transforming to an observer frame

The GR-corrected quantities are then converted from the local reference frame of the neutron star surface to the frame of a distant observer, which we signify with the subscript "∞".

Timescales

Timescale quantities are time-dilated by

$$
t_\infty = (1 + z)t_g = (1 + z)t_k.
\tag{2.25}
$$

Luminosity

Photons are redshifted to lower energies upon leaving the gravitational potential of the neutron star, in addition to their rate of arrival becoming time-dilated. Combined with Equation (2.17), the luminosity for an observer is given by

$$
L_\infty = \frac{L_g}{(1 + z)^2} = \frac{\varphi L_k}{1 + z}.
\tag{2.26}
$$

Burst energy

Because the burst energy (E_b) is integrated over time, the time-dilation from Equation (2.25) is removed, giving

$$
E_{b,\infty} = \frac{E_{b,g}}{1 + z} = \varphi E_{b,k},
\tag{2.27}
$$

where we have used Equation (2.21), and $\xi^2 = \varphi(1 + z)$ from Equation (2.9).

2.4 Predicting Observable Burst Properties

The extracted burst properties (Section 2.2), once corrected to an equivalent neutron star surface under GR (Section 2.3), can then be used to calculate "observables" as measured with Earth-based instruments.

2.4.1 Burst rate and recurrence time

The burst rate (ν) is time-dilated according to Equation (2.25),

$$
\nu_\infty = \frac{\nu_g}{1 + z} = \frac{\nu_k}{1 + z}.
\tag{2.28}
$$

Figure 2.3: Anisotropy factors versus inclination angle, as predicted by He and Keek (2016) with their *disc a* model for a thin, flat accretion disc. The factors alter the apparent isotropic flux given by Equation (2.31).

The recurrence time is simply the inverse,

$$\Delta t_\infty = (1 + z)\Delta t_{\mathrm{g}} = (1 + z)\Delta t_{\mathrm{k}}. \tag{2.29}$$

2.4.2   Flux – accounting for anisotropy

X-ray telescopes count photons within a given energy band, from which the incident bolometric flux (F) can be inferred. Assuming a uniform (i.e., isotropic) angular distribution for the source of radiation, the flux is given by

$$F = \frac{L_\infty}{4\pi d^2}, \tag{2.30}$$

where d is the distance to the source. Note that for visual clarity, from here on we assume that fluxes are always in the reference frame of an observer, and omit the subscript '∞'.

KEPLER models predict the local bolometric luminosity (L_{k}), which is GR-corrected (L_{g}) and shifted to the frame of a distant observer (L_∞) using Equation (2.26). If we assume the radiation is isotropic, this luminosity corresponds to the luminosity in Equation (2.30), and the observed flux can be directly predicted.

In reality, however, the neutron star is surrounded by an accretion disc, which can intercept, scatter, and obscure photons, resulting in a non-uniform (i.e., anisotropic) angular distribution of radiation (e.g. Lapidus and Sunyaev, 1985; Sztajno et al., 1987; Fujimoto, 1988; He and Keek, 2016). The apparent luminosity for an observer depends on the inclination, i, of the line of sight as measured from the rotation axis of the disc.

The angular distribution of radiation is dependent on the morphology and radiative properties of the accretion disc. These properties, including

the inclination itself, are often highly uncertain. Anisotropic effects are thus typically represented with generalised factors, given by ξ (introduced by Sztajno et al., 1987):

$$F_{\text{b}} = \frac{L_{\text{b},\infty}}{4\pi d^2 \xi_{\text{b}}}, \quad F_{\text{p}} = \frac{L_{\text{p},\infty}}{4\pi d^2 \xi_{\text{p}}}, \tag{2.31}$$

where the the burst ('b') and persistent ('p') emission are treated separately. Because ξ_{b} and ξ_{p} are degenerate with d, inferred distances are typically reported in the form $d\sqrt{\xi}$. Models which predict the dependence of ξ_{b} and ξ_{p} on the inclination (e.g. He and Keek, 2016, Figure 2.3) can then be used to constrain the absolute distance, d.

2.4.3 Persistent accretion flux

The persistent flux (F_{p}) is the steady emission observed between bursts. This quantity is thought to originate primarily from the luminosity generated by accretion (L_{acc}) given in Equation (2.18). Assuming that the contribution of steady nuclear burning is negligible (i.e., $L_{\text{p}} = L_{\text{acc}}$), we have, from Equation (2.31),

$$F_{\text{p}} = \frac{L_{\text{acc},\infty}}{4\pi d^2 \xi_{\text{p}}}. \tag{2.32}$$

From Equations (2.20) and (2.26), we obtain

$$L_{\text{acc},\infty} = \frac{c^2 z \dot{M}_{\text{g}}}{(1+z)^3} = \frac{c^2 z \varphi \dot{M}_{\text{k}}}{(1+z)^2}. \tag{2.33}$$

Equation (2.32) then becomes

$$F_{\text{p}} = \frac{c^2 z \dot{M}_{\text{g}}}{4\pi d^2 \xi_{\text{p}} (1+z)^3} = \frac{c^2 z \varphi \dot{M}_{\text{k}}}{4\pi d^2 \xi_{\text{p}} (1+z)^2}. \tag{2.34}$$

Thus, the persistent accretion flux can be predicted directly with the model parameters.

2.4.4 Burst fluence

The burst fluence (f_{b}) is the time-integrated burst flux ($\int F_{\text{b}} \, dt$), and is the observable equivalent of the burst energy (E_{b}). From Equation (2.31) we obtain

$$f_{\text{b}} = \frac{E_{\text{b},\infty}}{4\pi d^2 \xi_{\text{b}}}. \tag{2.35}$$

Substituting Equation (2.27) gives

$$f_{\text{b}} = \frac{E_{\text{b,g}}}{4\pi d^2 \xi_{\text{b}} (1+z)} = \frac{\varphi E_{\text{b,k}}}{4\pi d^2 \xi_{\text{b}}}. \tag{2.36}$$

2.4.5 Persistent fluence

The persistent fluence (f_{p}) is the time-integrated persistent flux ($\int F_{\text{p}} \, dt$) since the previous burst. Assuming that L_{acc} is constant between bursts,

the persistent fluence is given by

$$f_{\rm p} = \frac{L_{\rm acc,\infty}\Delta t_\infty}{4\pi d^2 \xi_{\rm p}}.$$ (2.37)

From Equations (2.28) and (2.34), we obtain

$$f_{\rm p} = \frac{c^2 z \dot{M}_{\rm g}\Delta t_{\rm g}}{4\pi d^2 \xi_{\rm p}(1+z)^2} = \frac{c^2 z \varphi \dot{M}_{\rm k}\Delta t_{\rm k}}{4\pi d^2 \xi_{\rm p}(1+z)}$$ (2.38)

2.4.6 Alpha ratio

A common measure of burst energetics is the ratio of the persistent to burst fluence (Section 2.2.1),

$$\alpha = \frac{f_{\rm p}}{f_{\rm b}}.$$ (2.39)

From Equations (2.35) and (2.37), we obtain

$$\alpha = \frac{\xi_{\rm b}}{\xi_{\rm p}} \cdot \frac{L_{\rm acc,\infty}\Delta t_\infty}{E_{\rm b,\infty}}.$$ (2.40)

Thus, α is independent of distance, and can also be used to infer the anisotropy ratio $\xi_{\rm b}/\xi_{\rm p}$. Using Equations (2.29) and (2.33), we obtain

$$\alpha = \frac{\xi_{\rm b} c^2 z \dot{M}_{\rm g}\Delta t_{\rm g}}{\xi_{\rm p}(1+z)E_{\rm b,g}} = \frac{\xi_{\rm b} c^2 z \dot{M}_{\rm k}\Delta t_{\rm k}}{\xi_{\rm p}\varphi(1+z)E_{\rm b,k}}.$$ (2.41)

2.4.7 Eddington flux

The Eddington-limited flux ($F_{\rm Edd}$) is typically inferred from the peak of PRE bursts, which are thought to reach the local Eddington luminosity ($L_{\rm Edd}$). Once again, from Equation (2.31) we have

$$F_{\rm Edd} = \frac{L_{\rm Edd,\infty}}{4\pi d^2 \xi_{\rm b}}.$$ (2.42)

Using Equations (2.23) and (2.26), we obtain

$$L_{\rm Edd,\infty} = \frac{4\pi G c M_{\rm g}}{\kappa(1+z)} = \frac{4\pi G c \varphi M_{\rm k}}{\kappa(1+z)}.$$ (2.43)

Equation (2.42) then becomes

$$F_{\rm Edd} = \frac{G c M_{\rm g}}{\kappa d^2 \xi_{\rm b}(1+z)} = \frac{G c \varphi M_{\rm k}}{\kappa d^2 \xi_{\rm b}(1+z)}.$$ (2.44)

For ionised material, the radiation pressure is exerted on electrons via Thomson scattering, whereas the mass is dominated by nucleons. For hydrogen, we can make the approximation that the opacity is given by $\kappa = \sigma_{\rm T}/m_{\rm p}$, where $\sigma_{\rm T}$ is the Thomson scattering cross section, and $m_{\rm p}$ is the proton mass.

Additionally, if we assume that the accreted material is a mixture of hydrogen and helium, we can introduce a factor to account for the composition, $2/(1+X_0)$, where X_0 is the hydrogen mass fraction. Compared to

pure hydrogen ($X_0 = 1.0$), pure helium ($X_0 = 0.0$) has double the mass per charge, resulting in double the Eddington limit.

Applying both of these values, Equation (2.44) becomes

$$F_{\mathrm{Edd}} = \frac{2Gm_{\mathrm{p}}cM_{\mathrm{g}}}{\sigma_{\mathrm{T}}d^2\xi_{\mathrm{b}}(1+z)(1+X_0)} = \frac{2Gm_{\mathrm{p}}c\varphi M_{\mathrm{k}}}{\sigma_{\mathrm{T}}d^2\xi_{\mathrm{b}}(1+z)(1+X_0)}. \tag{2.45}$$

Chapter 3

Results: Improvements to KEPLER Burst Models

In this chapter we outline the main improvements to KEPLER burst models used for this research.

In Section 3.1, we discuss a previous error in the opacities, which caused an artificially hot envelope and increased burst rate compared to other codes – a discrepancy already noticed by others in the literature. This correction is applied to all burst models presented in this thesis.

In Section 3.2, we describe the inclusion of a mock nuclear heat source during the model setup, which improved issues with thermal "burn-in" of the envelope. This treatment was applied to the models presented in Chapters 5 and 6, but was developed after the publication of Chapter 4.

Finally, in Section 3.3, we compare KEPLER models to another one-dimensional (1D) burst code, MESA, after including the above improvements. This test is the closest direct comparison of 1D burst codes to date, and we demonstrate that the improvements made to KEPLER reduce the discrepancy between the models.

3.1 Corrected Opacities

During the preparation of models for Johnston, Heger, and Galloway (2018) (Chapter 4), it was discovered[1] that the KEPLER burst models mistakenly included an opacity multiplication factor of ≈ 1.5.

The opacity modification originated from tests for an alternative to the GR-corrections described in Section 2.3. The boosted opacity was intended to slow heat transport such that the time dilation effects of GR would be replicated. The approach, however, was abandoned in favour of applying GR-corrections in post-processing. The multiplication factor, however, was mistakenly left in the model setup files, transmitting the error to all subsequent studies, including Woosley et al. (2004), Heger et al. (2007), Cyburt et al. (2016), Lampe, Heger, and Galloway (2016), and Galloway, Goodwin, and Keek (2017).

To test the effect of the corrected opacity, we computed two KEPLER models for comparison. We used the parameters from the popular reference model *A3* from Heger et al. (2007). The parameters used were a hydrogen mass fraction of $X_0 = 0.7048$, CNO mass fraction of $Z_{\mathrm{CNO}} = 0.02$, crustal heating rate of $Q_{\mathrm{b}} = 0.1\,\mathrm{MeV\,nucleon^{-1}}$, surface gravity of $g = 1.858 \times 10^{14}\,\mathrm{cm\,s^{-2}}$, and accretion rate of $\dot{m} = 0.0903\,\dot{m}_{\mathrm{Edd}}$, where

---

[1]by Adam Jacobs, Michigan State University, pers. comm.
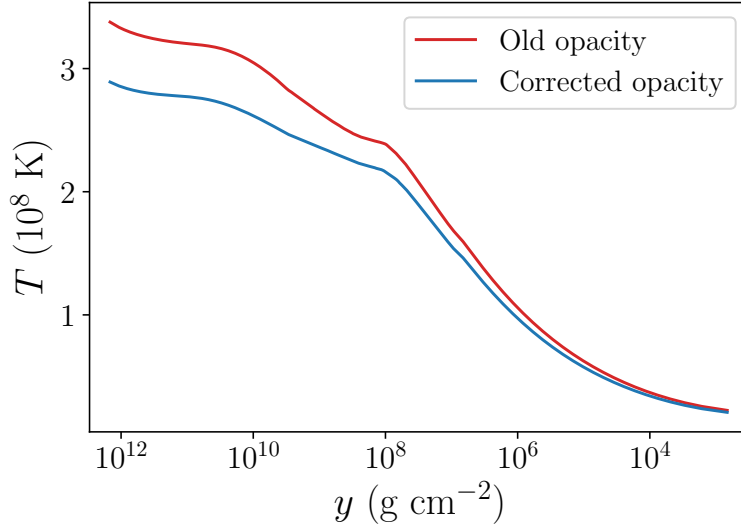
Figure 3.1: KEPLER temperature profiles, for a model with the erroneous opacity multiplier of ≈ 1.5 (red curve), and with the multiplier removed (blue curve). The larger opacities reduce the efficiency of thermal transport, producing a hotter envelope and shorter burst recurrence times (see also Table 3.1). This error was corrected for the models in Johnston, Heger, and Galloway (2018, (Chapter 4)), and is present in all previously published KEPLER burst models.

$\dot{m}_{\mathrm{Edd}} = 8.775 \times 10^4 \, \mathrm{g \, cm^{-2} \, s^{-1}}$. The opacity multiplier was left in place for the first model, and removed from the second. The models were run for $\approx 100 \, \mathrm{h}$ in the model rest frame, producing ≈ 30 bursts each.

The temperature profiles are compared in Figure 3.1. The snapshots are taken near the end of each model, prior to the ignition of the next burst, to allow the envelopes sufficient time to settle into a steady bursting state. The larger opacity produces systematically higher temperatures, which alters the conditions for burst ignition. The average burst properties for each model, after excluding the first 10 bursts, are listed in Table 3.1. The predictions of the two models are inconsistent, and the original opacity multiplier results in recurrence times of $\Delta t = 2.71 \pm 0.06 \, \mathrm{h}$, compared to $3.38 \pm 0.11 \, \mathrm{h}$ when it is removed.

Because the increased opacity produces an artificially hotter envelope, the conditions for burst ignition are reached at shallower depths, and thus earlier in time for a given accretion rate. The result is larger burst rates, ν, and reduced energetics, such as the peak luminosity and burst energy.

Indeed, discrepancies between the burst codes had already been noticed. Using the 1D code SHIVA, José et al. (2010) found that their burst models predicted recurrence times a factor of ≈ 2 longer than KEPLER models from Woosley et al. (2004). Similarly, with the 1D code MESA, Paxton et al. (2015) required accretion rates a factor of ≈ 2 larger to achieve a similar Δt to Heger et al. (2007). This inconsistency with MESA was reproduced by Meisel (2018), with comparisons to KEPLER models from Lampe, Heger, and Galloway (2016).

This systematic error in previous KEPLER models should be taken into account when comparing to the previous studies. Without directly recomputing the original models, however, it is difficult to apply a straightforward correction to the previous results. As a rule of thumb, the original

Table 3.1: Burst properties from models with and without the erroneous multiplier on opacity. Each value is an average over the burst sequence, excluding the initial 10 bursts, where the uncertainties are 1σ standard deviations. The artificially-large opacity in the original model results in systematic differences in the burst predictions. All values are in the local Newtonian frame of the KEPLER model, i.e., *not* corrected for GR (see Section 2.3).

		Opacity Multiplier	
		Old	Corrected
N bursts	–	38	30
Δt	(h)	2.71 ± 0.06	3.38 ± 0.11
ν	(day^{-1})	8.9 ± 0.2	7.1 ± 0.2
L_{peak}	$(10^{38}\,\text{erg}\,\text{s}^{-1})$	1.25 ± 0.08	2.04 ± 0.17
E_{b}	$(10^{39}\,\text{erg})$	4.08 ± 0.10	4.92 ± 0.11

models can be considered equivalent to models with larger accretion rates or base heating.

We perform a more direct comparison of the updated KEPLER models with MESA in Section 3.3.

3.2 Nuclear Preheating and Model Burn-in

The thermal history of the neutron star envelope can shape its bursting behaviour (Taam, 1980). Bursts produce nuclear heating and leftover ashes, which determine the thermal and compositional state of the envelope for subsequent bursts. This thermal and compositional "inertia" necessitates the simulation of many consecutive bursts, in order to reach a quasi-periodic limit cycle (Woosley et al., 2004). Time-dependent burst models, therefore, are subject to an initial "burn-in" phase, which is then excluded from our analysis.

In previous KEPLER studies, typically only the first 1–2 bursts were discarded (e.g., Heger et al., 2007; Cyburt et al., 2010; Lampe, Heger, and Galloway, 2016). The models were assumed to have reached a steady state by that point, particularly because the first burst is such an energetic outlier by comparison (Woosley et al., 2004). During our model tests, however, we discovered systematic trends in the burst properties which can persist for tens of bursts. Because 10–30 bursts are usually simulated per model (e.g., Heger et al., 2007; Lampe, Heger, and Galloway, 2016), this extended burn-in can potentially impact the entire model sequence.

A possible contribution to model burn-in is that nuclear heating, Q_{nuc}, is not accounted for when setting up the envelope. The base flux from crustal heating, Q_{b}, is used as a lower boundary condition at $y \approx 10^{12}\,\text{g}\,\text{cm}^{-2}$, and the envelope is relaxed to thermal equilibrium. Mass accretion and nuclear burning are then switched on, and the full simulation begins. Because the thermal profile is dominated by heat from the crust, only the crustal heating rate, $Q_{\text{b}} \approx 0.1\,\text{MeV}\,\text{nucleon}^{-1}$, is included. Nuclear heating in the shallower layers of $y \sim 10^7$–$10^8\,\text{g}\,\text{cm}^{-2}$ was assumed to stabilise within the first few

bursts, and have little impact on the overall thermal profile. If nuclear heating does significantly contribute to the long-term thermal structure of the envelope, then the existing models are out-of-equilibrium.

We tested the influence of nuclear heating on the model setup and subsequent simulation. In addition to Q_{b}, we include a heat source of $Q_{\mathrm{nuc}} = 5\,\mathrm{MeV\,nucleon^{-1}}$ at a depth of $y = 8 \times 10^7\,\mathrm{g\,cm^{-2}}$, distributed with a Gaussian width of $\sigma = 8 \times 10^6\,\mathrm{g\,cm^{-2}}$. The envelope is relaxed to equilibrium, and when the simulation begins, the Q_{nuc} source is switched off and the full nuclear network calculations are enabled.

We tested this implementation with three sets of model parameters. The first set was for mixed hydrogen/helium (H/He) bursts, such as those observed from GS 1826$-$238 (as modelled by Heger et al., 2007). The second and third sets were for "pure" helium (He) bursts, such as those observed from 4U 1820$-$30 (as modelled by Cumming, 2003), which exhibit photospheric-radius expansion (PRE). For the H/He set, we used an accreted hydrogen mass fraction of $X_0 = 0.73$, a CNO mass fraction of $Z_{\mathrm{CNO}} = 0.005$, a crustal heating of $Q_{\mathrm{b}} = 0.05\,\mathrm{MeV\,nucleon^{-1}}$, a surface gravity of $g = 2.654 \times 10^{14}\,\mathrm{cm\,s^{-2}}$, and an accretion rate of $\dot{m} = 0.2\,\dot{m}_{\mathrm{Edd}}$, where $\dot{m}_{\mathrm{Edd}} = 8.775 \times 10^4\,\mathrm{g\,cm^{-2}\,s^{-1}}$. For the pure He sets, we used $X_0 = 0.0$, $Z_{\mathrm{CNO}} = 0.015$, $Q_{\mathrm{b}} = 0.1\,\mathrm{MeV\,nucleon^{-1}}$, $g = 1.858 \times 10^{14}\,\mathrm{cm\,s^{-2}}$, and accretion rates of $\dot{m} = 0.2$ and $0.4\,\dot{m}_{\mathrm{Edd}}$, respectively. For each set, we computed one model with the original Q_{b}-only setup, and another with the nuclear "preheating" setup.

3.2.1   Mixed Hydrogen/Helium Models

In Figure 3.2, the resulting burst sequences (upper panel) and thermal profiles (lower panel) for the H/He models are shown. A clear trend of decreasing recurrence time, Δt, can be seen for the original setup without preheating (red points). This trend continues even up to ≈ 60 bursts – roughly double the length of typical KEPLER simulations (e.g., Heger et al., 2007; Lampe, Heger, and Galloway, 2016). In contrast, when nuclear preheating is included, the simulation has reached a steady bursting state within the first few bursts. As indicated by the dashed lines, an average taken from the original setup would overestimate Δt by $\approx 20\,\%$.

The burn-in is further illustrated by the temperature profiles (Figure 3.2, lower panel). Without accounting for nuclear heating, the envelope begins $\approx 50\,\%$ colder than equilibrium, and steadily heats up once the full nuclear calculations of bursts are included (red curves). After 58 bursts, the temperature of the deeper layers ($y \gtrsim 10^{10}\,\mathrm{g\,cm^{-2}}$) has increased by $\approx 50\,\%$. With nuclear preheating, the simulation begins much closer to equilibrium, and the temperature of the deeper layers has only increased by $\approx 2\,\%$ after 58 bursts. This small increase in temperature suggests that equilibrium has not yet been achieved, although a slight offset is expected given the relatively crude treatment for nuclear heating. Nevertheless, the effect of this leftover burn-in on Δt appears to be smaller than the inherent burst-to-burst variation.
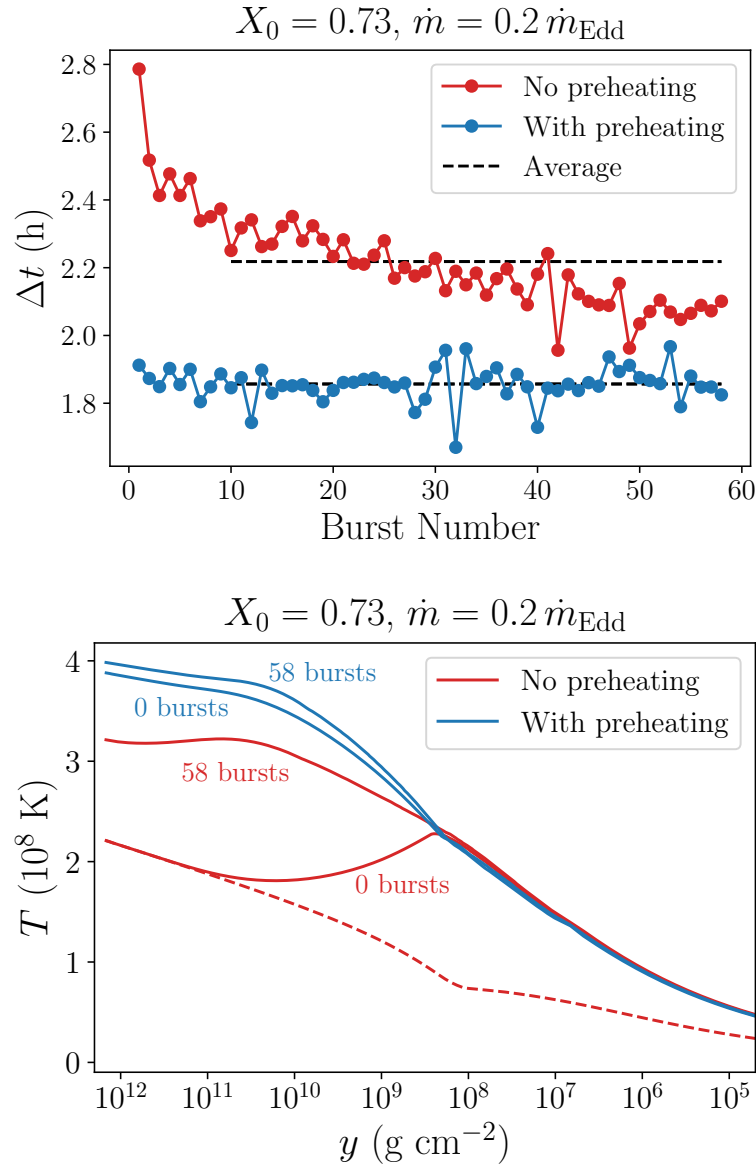
Figure 3.2: The effect of including mock nuclear heating during the setup of mixed hydrogen/helium models. (*Upper panel*): The recurrence times, Δt, of model burst sequences both with and without nuclear preheating. The horizontal dashed lines are the average values, after excluding the first 10 bursts. (*Lower panel*): The corresponding temperature profiles for the two models, at selected points along the burst sequence. The profiles are taken shortly before the following burst ignites. The coloured text associated with each curve indicates the number of bursts that have elapsed. The red dashed curve is at $t = 0$, and the adjacent red "0 bursts" curve is taken before the first burst ignites. A version of this figure appears in Johnston, Heger, and Galloway (2019, Chapter 5).
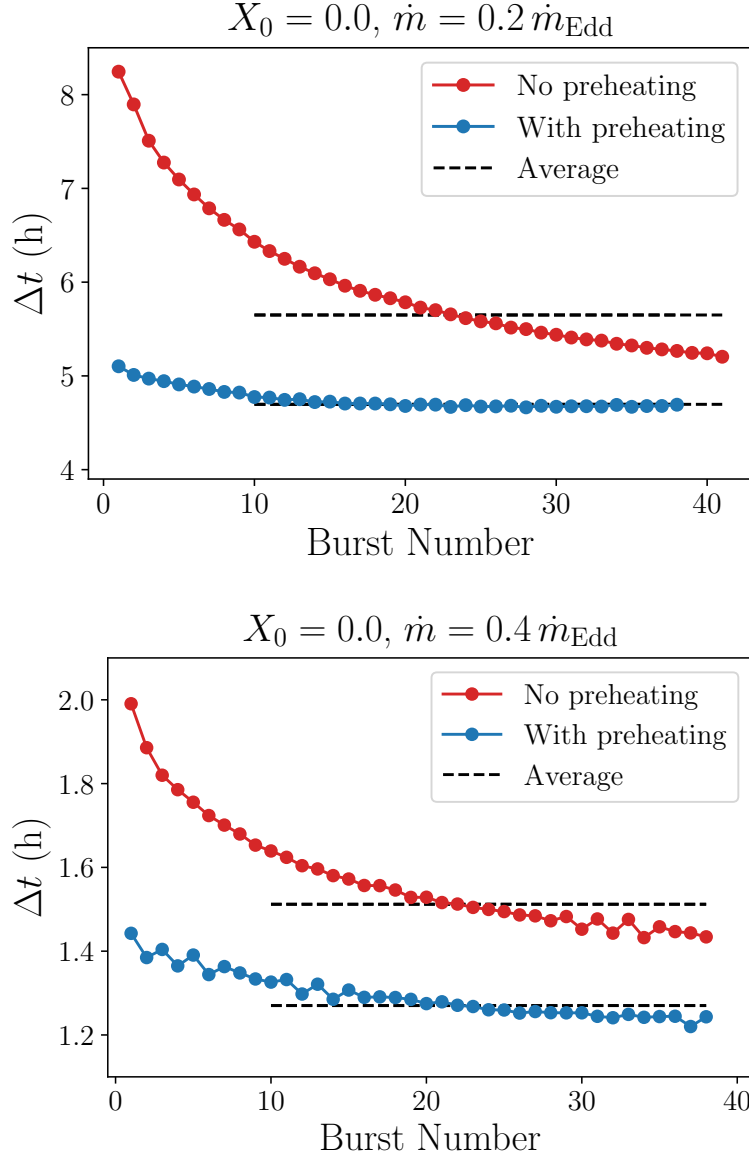
Figure 3.3: The same as the upper panel of Figure 3.2, but for two pure helium models with $\dot{m} = 0.2$ (*upper panel*) and $0.4\,\dot{m}_{\mathrm{Edd}}$ (*lower panel*). The preheating implementation is identical to the H/He models, with a heat source of $Q_{\mathrm{nuc}} = 5\,\mathrm{MeV\,nucleon^{-1}}$ centred at a depth of $y = 8 \times 10^7\,\mathrm{g\,cm^{-2}}$.

3.2.2 Pure Helium Models

The burst sequences for the pure He models are shown in Figure 3.3, and the corresponding temperature profiles are shown in Figure 3.4. Similar to the mixed H/He case, the models without nuclear preheating exhibit systematic trends in Δt. The addition of preheating suppresses the time required for burn-in for $\dot{m} = 0.2\,\dot{m}_{\mathrm{Edd}}$, and the bursts stabilise after 10–20 bursts. For $\dot{m} = 0.4\,\dot{m}_{\mathrm{Edd}}$, however, the burn-in appears to continue even after 38 bursts, although at a reduced rate. The burst-to-burst variation in Δt for the last 10 bursts is $< 1\,\%$.

The temperature profiles for both $\dot{m} = 0.2$ and $0.4\,\dot{m}_{\mathrm{Edd}}$ begin colder without preheating, and systematically heat up over the course of 38 bursts. The models with preheating, however, behave differently to the H/He case. Whereas the shallow layers ($y \lesssim 10^9\,\mathrm{g\,cm^{-2}}$) start hotter than the original setup and slightly heat up, the deeper layers ($y \gtrsim 10^9\,\mathrm{g\,cm^{-2}}$) actually cool down after 38 bursts. The "kink" in each profile corresponds to the transition from the accreted material to the inert iron substrate, which represents the deep ocean of previous burst ashes (see Section 2.1.2). In the H/He models (Figure 3.2), this interface is smoothed out as the simulated ashes become closer in composition to the substrate. In the pure He models, however, this interface persists, suggesting that the iron substrate is not a good representation for the ashes of helium bursts.

3.2.3 Discussion

The addition of a nuclear heat source during the setup of KEPLER models significantly shortened the systematic model burn-in. For the mixed hydrogen/helium case tested here, the number of bursts required was reduced from $\gtrsim 60$ to $\lesssim 10$ (Figure 3.2). A small amount of heating persisted in the deeper envelope, although any slight trend in Δt appears to be hidden by the burst-to-burst variation. For the two sets of pure helium bursts tested, some burn-in remained even after 30–40 bursts, particularly for the higher accretion rate of $\dot{m} = 0.4\,\dot{m}_{\mathrm{Edd}}$ (Figure 3.3). Nevertheless, the burst-to-burst variation dropped to $< 1\,\%$, but longer simulations are still required to determine how long the downward trend in Δt continues. Longer simulations could also test whether both treatments – with and without preheating – do eventually converge to the same burst properties.

Future work is still needed to investigate the optimal strength, location, and distribution of nuclear preheating, depending on the model parameters. In these tests, we used a heating strength of $Q_{\mathrm{nuc}} = 5\,\mathrm{MeV\,nucleon^{-1}}$ located at $y = 8 \times 10^7\,\mathrm{g\,cm^{-2}}$, with a Gaussian distribution of width $\sigma = 8 \times 10^6\,\mathrm{g\,cm^{-2}}$. Whereas the model burn-in was essentially eliminated for H/He models, the ignition depth and total energy release could be refined for pure He models. Additionally, the persistence of a kink in the pure He temperature profiles (Figure 3.4) suggests that a lighter composition may be needed than the existing iron substrate. Despite its limitations, our preheating treatment can potentially save days of computation time, by reducing the number of bursts required to obtain reliable predictions.
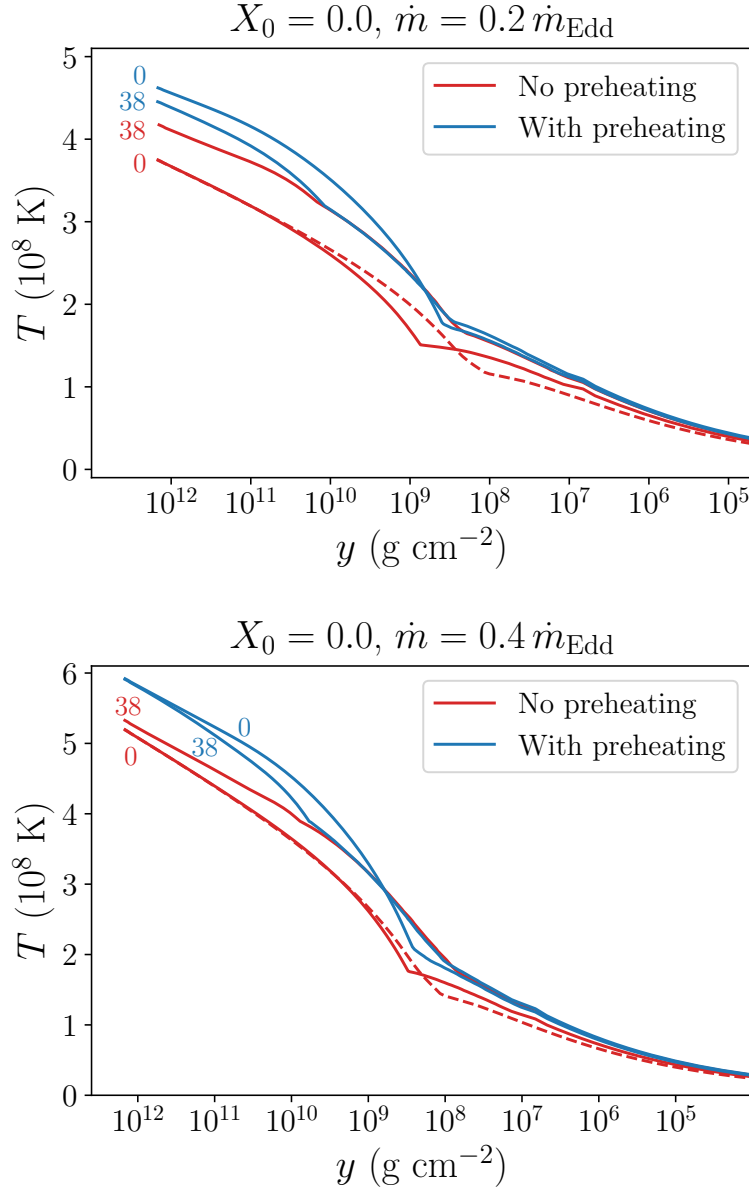
Figure 3.4: The same as the lower panel of Figure 3.2, but for the corresponding pure He models in Figure 3.3. The coloured numbers next to each curve signify the number of bursts elapsed. The red dashed curve is $t = 0$. Note the inversion of heating/cooling in the deeper layers, at $y \gtrsim 10^9 \, \mathrm{g \, cm^{-2}}$. The kink in each curve corresponds to the transition in composition from the accreted material to the inert iron substrate.

3.3   Updated Comparison to MESA Burst Models

Comparing the predictions of different codes is an important test of code verification and reproducibility. Although comparisons have been made in previous works (e.g., José et al., 2010; Paxton et al., 2015; Meisel, 2018), the model input parameters are typically slightly different, partly due to the limited number of published models available. No comparison yet exists for two burst codes using the same set of X_0, Z_{CNO}, Q_b, $\dot{m}$, and g.

Discrepancies between the predictions of KEPLER and other 1D burst codes have been noted in the literature (see Section 3.1). During the course of this research, we discovered that incorrect opacities were being used in KEPLER, which likely contributed to this inconsistency (Section 3.1). To test our updated model setup, and whether the corrected opacity improves code agreement, we have computed a set of simulations for direct comparison with existing MESA models.

3.3.1   Model setup

For our comparison, we used MESA burst models of the mixed hydrogen/helium burster, GS 1826−238, produced by Meisel (2018). We chose a subset of five models which were among the best fits to the observations, labelled as *ma*1–*ma*6 in the dataset[2] (excluding *ma*3, which did not produce bursts).

All five models used a hydrogen mass fraction of $X_0 = 0.7$, a CNO mass fraction of $Z_{CNO} = 0.02$, a crustal heating of $Q_b = 0.1\,\mathrm{MeV\,nucleon^{-1}}$, and a surface gravity of $g = 1.858 \times 10^{14}\,\mathrm{cm\,s^{-2}}$, corresponding to a gravitational mass of $M = 1.4\,\mathrm{M_\odot}$, and a radius of $R = 11.2\,\mathrm{km}$. The models differed only by accretion rate, for $\dot{M} = 0.061$, 0.079, 0.123, 0.164, and $0.185\,\dot{M}_{Edd}$, where $\dot{M}_{Edd} = 1.75 \times 10^{-8}\,\mathrm{M_\odot\,yr^{-1}}$. These values differ slightly from those reported in Meisel (2018), because MESA "settles" into a target $\dot{M}$, and we have taken the accretion rate averaged over the whole model.

We computed a set of five KEPLER models using these same parameters, with some modifications to ensure consistency between the codes. Because in KEPLER Newtonian gravity, we used $M = 1.4\,\mathrm{M_\odot}$ and $R = 10\,\mathrm{km}$ to reproduce the same g. With appropriate corrections for general relativity (GR), the KEPLER models are equivalent to a neutron star of $M = 1.4\,\mathrm{M_\odot}$ and $R = 11.2\,\mathrm{km}$ (see Section 2.3). The radius ratio is then $\xi = R_g/R_k$, where the subscripts "g" and "k" correspond to MESA and KEPLER, respectively. For a given accretion rate in MESA, the equivalent accretion rate used in KEPLER is $\dot{M}_k = \xi^{-2}\dot{M}_g$. We used V2.2 of the nuclear reaction rate library, REACLIB (Cyburt et al., 2010), as used in the MESA models. To our knowledge, the nuclear preheating treatment we describe in Section 3.2 is not implemented in MESA, and so we disabled it for this comparison.

Sequences of ≈ 30 bursts were computed for each model, and the average burst properties calculated using the methods described in Section 2.2. The 1σ standard deviations were taken as the model uncertainties. To avoid potential differences in analysis techniques, we used these same routines to extract the burst properties from MESA. The Newtonian quantities predicted by KEPLER were corrected to the equivalent GR neutron star frame of MESA, using the procedure described in Section 2.3.2.

---

[2]available at https://inpp.ohio.edu/~meisel/MESA/mesaresults.html
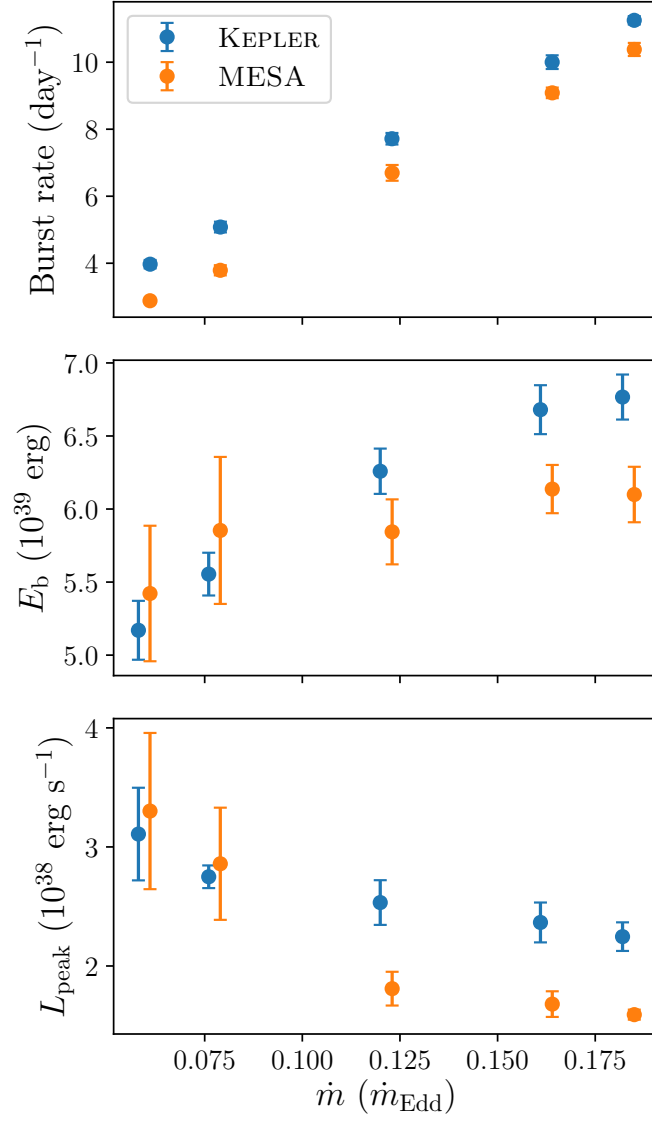
Figure 3.5: Comparison of burst properties predicted by KEPLER and MESA models, using the same input parameters. The KEPLER values, including $\dot{M}$, have been GR-corrected to the equivalent local frame of the MESA models. The points have been slightly offset horizontally for clarity.
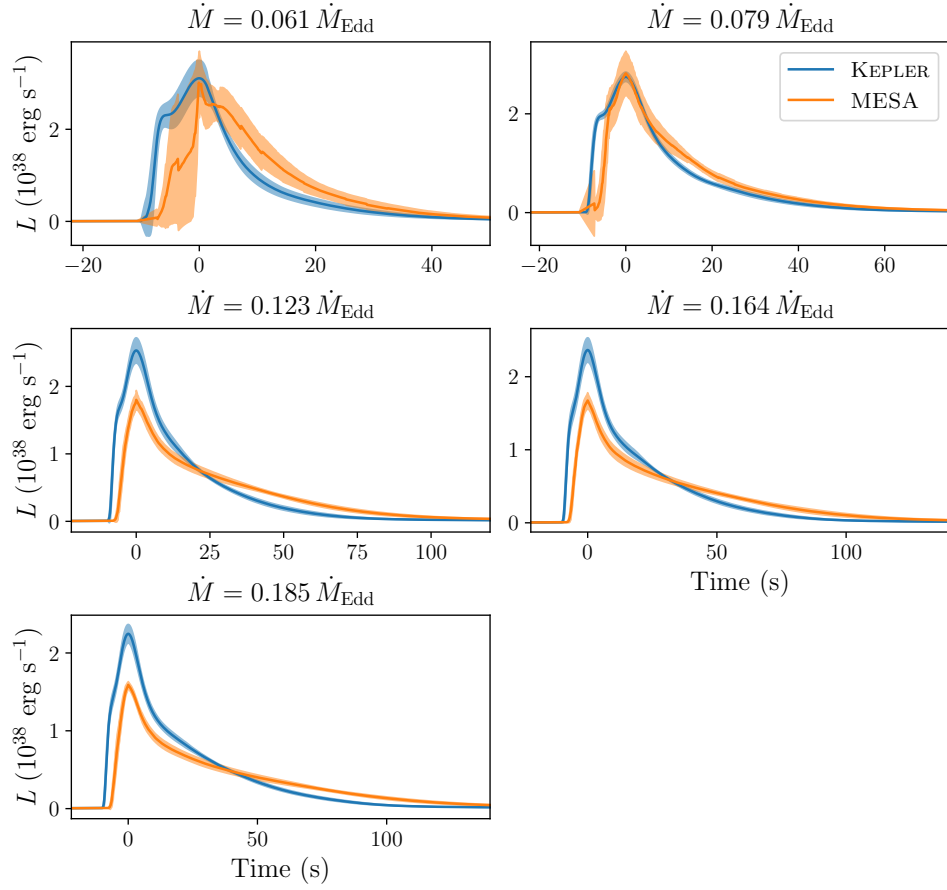
Figure 3.6: Comparison of the average burst lightcurves predicted by KEPLER and MESA, for five accretion rates, aligned by peak. The shaded regions are the 1σ standard deviations in L. The KEPLER luminosities have been corrected for the area ratio, ξ^2. Quantities are in the local frame of the neutron star surface, and are not redshifted.

3.3.2   Results

The predicted burst rate, ν, burst energy, E_b, and peak luminosity, L_{peak}, are plotted for each model in Figure 3.5. Despite broad similarities for the predicted trends and values, disagreement remains between the codes.

For the burst rates, there is a persistent offset of ≈ 1 burst per day, although the linear relationship with $\dot{M}$ itself is consistent. This relatively uniform offset suggests there may still be systematic issues affecting the codes.

The burst energies agree within uncertainties for the lowest accretion rates of $\dot{M} = 0.061$ and $0.079\,\dot{M}_{Edd}$, but diverge for the higher rates of 0.123–0.185 $\dot{M}_{Edd}$, for which KEPLER produces $\approx 10\,\%$ larger E_b. Similar to E_b, the peak luminosities agree for the two lowest $\dot{M}$ values, but the KEPLER values are $\approx 40\,\%$ larger.

The average burst lightcurves are compared in Figure 3.6, and further illustrate the pattern noted above for E_b and L_{peak}. The two lowest $\dot{M}$ broadly agree, whereas above $0.123\,\dot{M}_{Edd}$ the KEPLER models predict systematically stronger bursts. On the other hand, MESA yields longer lightcurve tails, which would normally indicate a larger hydrogen fraction and stronger *rp*-process burning. The longer recurrence times of the MESA bursts, however, should instead result in less hydrogen at ignition due to hot CNO burning. The MESA lightcurve for $\dot{M} = 0.061\,\dot{M}$ exhibits erratic behaviour, suggesting an issue with the average lightcurve, perhaps from misaligned individual lightcurves.

This limited study represents the first direct comparison between 1D burst codes, using matched input parameters for $\dot{M}$, X_0, Z_{CNO}, Q_b, and g. Whereas the models generally agree within uncertainties at low accretion rates, KEPLER consistently predicts stronger bursts at higher $\dot{M}$. KEPLER also produces larger burst rates than MESA for all $\dot{M}$ considered here, although this result appears to be at odds with the expected behaviour for shorter recurrence times. Future studies are required to quantify these differences in detail, and their possible dependence on burst regime. Additional sets of composition, crustal heating, and gravity should be explored for parameter sensitivities. Closer inspection of the model profiles, thermal structure, and reaction networks will also help to understand the discrepancies between the codes.

Chapter 4

Simulating X-ray Bursts During a Transient Accretion Event

Simulating X-ray bursts during a transient accretion event

Zac Johnston,[1,2,3]★ Alexander Heger[1,2,3,4,5] and Duncan K. Galloway[1,2,3]

[1]*School of Physics and Astronomy, Monash University, Victoria 3800, Australia*
[2]*Monash Centre for Astrophysics (MoCA), Monash University, Victoria 3800, Australia*
[3]*Joint Institute for Nuclear Astrophysics – Center for the Evolution of the Elements, East Lansing, MI 48824, USA*
[4]*Department of Astronomy, Shanghai Jiao Tong University, Shanghai 200240, China*
[5]*School of Physics and Astronomy, University of Minnesota, Minneapolis, MN 55455, USA*

ABSTRACT

Modelling of thermonuclear X-ray bursts on accreting neutron stars has to date focused on stable accretion rates. However, bursts are also observed during episodes of transient accretion. During such events, the accretion rate can evolve significantly between bursts, and this regime provides a unique test for burst models. The accretion-powered millisecond pulsar SAX J1808.4-3658 exhibits accretion outbursts every 2–3 yr. During the well-sampled month-long outburst of 2002 October, four helium-rich X-ray bursts were observed. Using this event as a test case, we present the first multizone simulations of X-ray bursts under a time-dependent accretion rate. We investigate the effect of using a time-dependent accretion rate in comparison to constant, averaged rates. Initial results suggest that using a constant, average accretion rate between bursts may underestimate the recurrence time when the accretion rate is decreasing, and overestimate it when the accretion rate is increasing. Our model, with an accreted hydrogen fraction of $X = 0.44$ and a CNO metallicity of $Z_{CNO} = 0.02$, reproduces the observed burst arrival times and fluences with root mean square (rms) errors of 2.8 h, and 0.11×10^{-6} erg cm^{-2}, respectively. Our results support previous modelling that predicted two unobserved bursts and indicate that additional bursts were also missed by observations.

Key words: methods: numerical – stars: neutron – pulsars: individual: (SAX J1808.4-3658) – X-rays: bursts.

1 INTRODUCTION

Type I X-ray bursts are thermonuclear flashes in the accreted envelopes of neutron stars (Belian, Conner & Evans 1976; Grindlay et al. 1976). In low-mass X-ray binaries (LMXBs), a mix of hydrogen and helium is transferred from a low-mass companion ($\lesssim 1\,M_\odot$) to a neutron star via Roche lobe overflow, forming an accretion disc that feeds nuclear fuel to the neutron star surface. The base of the accreted layer is buried deeper under new material and heated to the point of thermonuclear runaway (Woosley & Taam 1976; Joss 1977). The heat released by the rapid fusion of the accreted layer is observable as a burst of X-rays, lasting approximately 10–100 s. Fresh fuel is then accreted on to the ashes, and bursts recur within hours to days (for reviews, refer to Lewin, Paradijs & Taam 1993; Strohmayer & Bildsten 2003; Galloway et al. 2008).

As each new layer of fuel is buried deeper, its hydrogen is steadily converted to helium via the beta-limited (hot) CNO cycle. If the burst recurrence time is longer than the time to deplete hydrogen, the burst will ignite in a deep helium layer (Case 2; Fujimoto, Hanawa

& Miyaji 1981). This class of helium bursts reach Eddington luminosity (L_{Edd}) and exhibit photospheric radius expansion (PRE; Lewin, Vacca & Basinska 1984; Tawara et al. 1984). Their light-curves feature rapid onsets ($\lesssim 1$ s), broad plateaus (≈ 10 s), and short tails ($\lesssim 30$ s) due to the absence of extended *rp*-process burning.

Some X-ray burst systems, such as GS 1826-24 (Ubertini et al. 1999), accrete and produce bursts at a consistent rate. X-ray burst modelling to date has focused on stable accretion rates (e.g. Woosley et al. 2004; Heger et al. 2007; Keek & Heger 2011), and the dependence of burst properties on those rates (Lampe, Heger & Galloway 2016). Transient X-ray binaries, on the other hand, can remain dormant for years at a time and experience unstable surges of accretion known as outbursts. The addition of fresh material to the neutron star surface, as with stable accretors, can trigger series of X-ray bursts.

SAX J1808.4-3658 was the first accreting millisecond pulsar (AMXP) to be observed (Chakrabarty & Morgan 1998; in 't Zand et al. 1998; Wijnands & van der Klis 1998) and undergoes month-long outbursts every 2–3 yr (e.g. Wijnands 2004; Galloway 2006; Hartman et al. 2008; Hartman, Watts & Chakrabarty 2009; Patruno & Watts 2012; Patruno et al. 2012, 2017). During the well-sampled outburst of 2002 October, four thermonuclear X-ray bursts were

★ E-mail: zac.johnston@monash.edu

observed (Chakrabarty et al. 2003). Subsequent modelling determined these to be helium bursts (Galloway & Cumming 2006, hereafter G06). The bursts from this event have been proposed as a standard test case for numerical modelling (Case 2, Galloway, Goodwin & Keek 2017).

Using a semi-analytic model with a one-zone ignition criterion (described in Cumming & Bildsten 2000), G06 matched models to the observed burst properties and constrained the distance of the system to 3.5 ± 0.1 kpc. In order to solve for the ignition conditions, the accretion rate was averaged between bursts. Because these models were computationally inexpensive, many thousands could efficiently explore parameter space. Despite these advantages, the semi-analytic nature of the model lacks chemical and thermal inertia from one burst to the next.

To improve upon this modelling of the 2002 October outburst, we present a multizone simulation produced with the KEPLER code, which tracks the composition and thermal state of the envelope as the outburst evolves (Woosley et al. 2004). This is the first application of time-dependent accretion rates to multizone burst simulations, allowing us to capture the thermal and chemical inertia of the envelope throughout an accretion episode.

In Section 2, we describe the X-ray data from the 2002 October outburst, the code we used to simulate the bursts, the system parameters chosen for the simulation, how we constructed the accretion rate curve, and how burst properties were extracted for comparison with observations. In Section 3, we examine the effect that a time-varying accretion rate has compared to a constant rate and compare the simulated burst properties with those observed, including times of arrival, fluences, and light curves. We summarize our results in Section 4 and discuss planned improvements to the model.

2 METHOD

2.1 Observational data

We used data from the 2002 October outburst taken with the Proportional Counter Array (PCA; Jahoda et al. 1996) and the All Sky Monitor (ASM; Levine et al. 1996) of the *Rossi X-ray Timing Explorer* (*RXTE*). The PCA instrument is composed of five proportional counter units (PCUs) sensitive to photon energies of 2–60 keV. Two components of the PCA data were used: the persistent accretion flux, F_p (Section 2.5), and the time-resolved burst light curves, F_b (Section 3.2). The purpose of the ASM instrument is primarily to trigger alerts of transient events, and produces only low signal-to-noise data. It scans the sky every 90-min orbit, and scans a given object roughly 5–10 times each day in 90-s exposures, from which 1-d average count rates are calculated.

The burst data were taken from the catalogue of Galloway et al. (2008), which was since re-analysed for the Multi-INstrument Burst ARchive (MINBAR, in-development).[1] This re-analysis consisted of fitting absorbed blackbody models over the range of 2.5–20 keV, with the neutral column density fixed at $n_H = 1.2 \times 10^{21}$ cm^{-2} (Wang et al. 2001). Updated PCA response matrices (v11.7[2]) were used, and the recommended systematic error of 0.5 per cent was adopted. Churazov weighting was employed to address the issue of low-count spectra in XSPEC (Dorman & Arnaud 2001). Deadtime was

estimated using the Standard-1 mode data,[3] and the exposure time was then reduced by the deadtime correction factor, contributing an additional 25–30 per cent to the photon flux at the burst peaks.

For the persistent flux (F_p), n_H was fixed, and updated response matrices and deadtime correction used, as with the burst data. Additionally, spectra were averaged over each observation separately for each PCU, excluding the burst times. Each spectrum was then fitted with one of a family of models, including blackbody+power law or Comptonization, often adopting a Gaussian component to model Fe Kα emission around 6.4 keV. The models were then integrated over the range of 3–25 keV, and the bolometric flux estimated.

2.2 Numerical method

To simulate the neutron star envelope during an accretion outburst, we used the 1D stellar hydrodynamics code KEPLER, which models a grid of Lagrangian zones in the radial direction (Weaver, Zimmerman & Woosley 1978; Woosley et al. 2004). Each zone, representing a spherically symmetric shell of stellar material, has its own isotopic abundances and thermal properties. Convection of heat and nuclei between zones is modelled using mixing length theory, where a time-dependent diffusion coefficient is set by the convective velocity (implementation described in Heger, Langer & Woosley 2000).

KEPLER uses an adaptive nuclear network that can track the reactions between more than 1000 isotopes up to the proton drip line (Rauscher et al. 2002). Isotopes are automatically added and removed from the network as needed (Woosley et al. 2004). This allows us to efficiently model the β-limited CNO cycle, the 3α-process, the αp-process, and the rp-process (Cyburt et al. 2010, 2016).

KEPLER also uses an adaptive spatial grid, in which zones are actively split or combined at each time-step in order to maintain resolution of thermodynamic gradients. Multiple criteria govern this rezoning; we impose a minimum zone thickness of 10 cm, a surface zone mass of $\sim 10^{18}$ g, and the above-mentioned accretion depth of 10^{19} g. These were chosen to avoid needlessly creating large numbers of zones, while maintaining consistency of the resulting burst properties. The model described in Section 3 has 71 initial zones, growing as mass is accreted to 122 zones at the time of the first burst, and to 174 zones by the end of the simulation.

The simulation domain extends from the neutron star (NS) photosphere to the deep ocean near the crust, covering column depths of $10^4 \lesssim y \lesssim 10^{12}$ g cm^{-2} ($10^4 \lesssim \rho \lesssim 10^9$ g cm^{-3}). Accretion-driven heating in the crust from electron captures and pycnonuclear reactions is included as a luminosity at the base of the grid, L_{crust}. To set up the initial thermal state, the envelope is relaxed until thermal equilibrium is met between L_{crust} and the luminosity at the surface.

Nuclear reactions in the envelope are then switched on, and accretion is simulated by adding mass to the zone with an exterior mass coordinate of 10^{19} g ($y \approx 8 \times 10^5$ g cm^{-2}) at the rate $\dot{M}$, which may evolve with time. This depth is chosen to avoid unnecessary re-zoning at the surface. Above this, the accreted composition is advected, and heating from accretion and compression is included (Keek & Heger 2011).

The accreted fuel is composed of ^{1}H, ^{4}He, and ^{14}N, given by their mass fractions, X_0, Y_0, and Z_{CNO}, respectively. We choose ^{14}N for simplicity, because it is the most abundant CNO metal in solar

Table 1. Summary of parameters used in the model. Please note that these should not yet be considered best-fitting values to the system.

Quantity		Units	Description
g	1.86	$10^{14}\,\mathrm{cm\,s^{-2}}$	Surface gravity
M	1.4^a	$\mathrm{M_\odot}$	NS mass
R	11.2^a	km	NS radius
Q_b	0.3	MeV nucleon^{-1}	Crustal heating
X_0	0.44	Mass fraction	Accreted hydrogen
Z_{CNO}	0.02	Mass fraction	Accreted metallicity
d	3.5^b	kpc	Distance
ξ_p	1.1^b	–	Persistent anisotropy

Notes. ^{a}Effective GR-corrected values. Other combinations are still valid if they preserve g, but will alter $(1 + z)$ and the conversion of model results to observable values (Section 2.3).
^{b}Assumed values for inferring the accretion rate from persistent flux. Other combinations are still valid (Section 2.5).

material, and the hot-CNO cycle rapidly stabilizes to equilibrium values of ^{14}O and ^{15}O.

2.3 General relativistic corrections

KEPLER performs calculations in the local NS frame using Newtonian gravity. This approximation is acceptable because the grid spans only a thin surface shell of the neutron star, over which the gravitational acceleration, g, varies by $\lesssim 2$ per cent. However, in order to compare results with observations, general relativity (GR) must be accounted for. Given a model Newtonian NS mass and radius (M, R), there are combinations of GR mass and radius (M_{GR}, R_{GR}) such that the surface gravity is equal under both regimes (refer to appendix B of Keek & Heger 2011). In other words, the model can be considered equivalent to a neutron star with an 'actual' mass and radius of M_{GR} and R_{GR}. This is satisfied when

$$\frac{GM}{R^2} = \frac{GM_{GR}}{R_{GR}^2}(1 + z), \tag{1}$$

where the gravitational redshift factor is

$$1 + z = \frac{1}{\sqrt{1 - 2GM_{GR}/(c^2 R_{GR})}}. \tag{2}$$

If we choose $M_{GR} = M$, the solutions are simplified, and the luminosity and time can be converted from the model's Newtonian frame (L, t) to a distant observer frame (L_∞, t_∞) with

$$L_\infty = \frac{L}{1 + z}, \qquad t_\infty = t\,(1 + z). \tag{3}$$

2.4 Model parameters

The model parameters are summarized in Table 1.

We adopt a gravitational mass of $M = M_{GR} = 1.4\,\mathrm{M_\odot}$ and a Newtonian model radius of $R = 10$ km, equivalent to $R_{GR} \approx 11.2$ km (equation 1). This gives a surface gravity of $g \approx 1.858 \times 10^{14}\,\mathrm{cm\,s^{-2}}$ and redshift of $1 + z \approx 1.259$. Note that the model can still be considered equivalent to any other pair of M_{GR} and R_{GR} that satisfy equation (1).

To set up the initial envelope, the inner portion is composed of an iron substrate between column depths of $10^8 \lesssim y \lesssim 10^{12}\,\mathrm{g\,cm^{-2}}$ ($10^6 \lesssim \rho \lesssim 10^9\,\mathrm{g\,cm^{-3}}$). Nuclear reactions are not calculated in the substrate, and it primarily acts as a thermal sink during bursts, representing the ocean of prior burst ashes. Above the substrate, between $10^4 \lesssim y \lesssim 10^8\,\mathrm{g\,cm^{-2}}$ ($10^4 \lesssim \rho \lesssim 10^6\,\mathrm{g\,cm^{-3}}$), we add

^{4}He to represent leftover fuel from the tail of the previous outburst. The size of this fuel layer could be varied in future studies.

Crustal heating from the inner boundary evolves with accretion rate according to $L_{crust} = Q_b \dot{M}$, where Q_b is the specific energy yield from reactions in the crust. We adopt a value of $Q_b = 0.3$ MeV nucleon^{-1}, following the best-fitting model from G06. To initialize the envelope, we use this value with the long-term average accretion rate of $10^{-11}\,\mathrm{M_\odot\,yr^{-1}}$ (G06), until the layer is in thermal equilibrium.

The observed burst energetics indicate an average hydrogen composition of $\langle X \rangle \approx 0.1$ at the time of ignition (G06). The initial composition of the accreted fuel, however, is less well-constrained because there is a degeneracy between the X_0 and Z_{CNO} that result in the above $\langle X \rangle$ due to stable hot-CNO burning. For the purposes of this study, we choose a metallicity of $Z_{CNO} = 0.02$, following the best-fitting model from G06. Initial tests of X_0 in the range of 0.49–0.62 (from the best-fitting 1σ range of G06) indicated that our models required a lower value to reproduce the observed burst timings. The model presented here has $X_0 = 0.44$, with the remaining composition being $Y_0 = 0.54$.

Accretion discs can cause anisotropies by scattering and blocking X-ray emission from the NS surface, changing the apparent luminosity to an observer (Fujimoto 1988; He & Keek 2016). The strength of the anisotropy is dependent on i, the inclination of the binary system to the observer's line of sight. The inclination is typically poorly constrained in LMXBs, with the value for SAX J1808.4-3658 inferred to be $50° \lesssim i \lesssim 80°$ (Chakrabarty & Morgan 1998; Bildsten & Chakrabarty 2001; Wang et al. 2001; Ibragimov & Poutanen 2009). The effect is represented as a scaling factor, ξ, which we include in the conversion between the model luminosity and the observed flux, given by

$$L = 4\pi d^2 \xi F. \tag{4}$$

Independent factors are used for the persistent emission (ξ_p) and the burst emission (ξ_b), because the emitting region is not necessarily the same for both mechanisms. We used ξ_p when calculating $\dot{M}$ from persistent flux, and ξ_b when calculating the burst flux from the model burst luminosity.

2.5 Accretion history

We inferred the accretion history of the 2002 October outburst from the observed persistent flux, F_p. Accretion on to the surface generates a luminosity of

$$L_{acc} = -\dot{M}\phi \quad \mathrm{erg\,s^{-1}}, \tag{5}$$

where $\phi = -c^2 z/(1 + z) \approx -0.2\,c^2$ is the gravitational potential at the NS surface. Using equations (3) and (4), the accretion rate in the model frame is then given by

$$\dot{M} = \frac{-4\pi d^2 \xi_p (1 + z)}{\phi} F_p. \tag{6}$$

We can rewrite this expression in terms of a conversion constant,

$$\dot{M} = \frac{-c_1(1 + z)}{\phi} F_p, \quad c_1 = 4\pi d^2 \xi_p. \tag{7}$$

For this model, we chose $d = 3.5$ kpc (G06), and $\xi_p = 1.1$, which is the predicted anisotropy factor for an inclination of $55° \lesssim i \lesssim 60°$ (fig. 8 of He & Keek 2016). Thus, we have a conversion constant of $c_1 \approx 1.612 \times 10^{45}\,\mathrm{cm^2}$. Note that this model is still applicable to other combinations of d and ξ_p that preserve c_1.

Figure 1. The accretion rate estimated from *RXTE* observations during the 2002 October outburst, as a fraction of the Eddington-limited rate ($\dot{M}_{Edd} = 1.75 \times 10^{-8}$ M$_\odot$ yr^{-1}), assuming $d = 3.5$ kpc and ξ_p. Time has been zeroed to the start of PCA observations at MJD 52562.07296, and shown are the times of the four observed bursts. Linear interpolation has been used within the high-resolution PCA data, while a toy curve has been inserted as a stand-in for the rise, which was not observed by the PCA. Note that while ASM data have been plotted for reference, the rise curve was not explicitly fit to them.

PCA observations did not commence until the peak of the outburst, and so the precise onset of accretion is ambiguous. For reference, the rise of the subsequent 2005 June outburst was observed to last 5 d (Hartman et al. 2008). For the 2002 October outburst, only 1-d average count rates from the ASM are available, constraining the rise length to 4–5 d. With these considerations, we substituted a toy curve for the accretion rise with a length of approximately 4 d. We expect that differences in the chosen onset should primarily influence the first burst or two, which likely went unobserved (Section 3.2). Nevertheless, we plan to investigate the effect of rise length and shape in a future study.

Combined with the PCA data and equation (6), we thus obtained a continuous $\dot{M}(t)$ curve for input to the model (Fig. 1).

2.6 Burst properties

We ran KEPLER with the above inputs, and obtained a sequence of bursts over the course of the outburst. We then extracted the burst light curves from the modelled NS surface luminosity, and calculated their properties in a process similar to Lampe et al. (2016).

The recurrence time, Δt, is the time from one burst to the next. The burst energy, E_{nuc}, is obtained by integrating over the light curve. This translates to the observed burst fluence, E_b, via

$$E_b = \frac{1}{4\pi d^2 \xi_b} E_{nuc}. \tag{8}$$

Similar to equation (7), this expression can also be written in terms of a scaling factor,

$$E_b = \frac{1}{c_2} E_{nuc} \quad \text{with } c_2 = 4\pi d^2 \xi_b. \tag{9}$$

The profiles of PRE burst light curves in KEPLER noticeably deviate from observations, likely due to the simple atmosphere (see Model *Zm* in Woosley et al. 2004). In our model, the surface luminosity exceeded the Eddington limit by up to a factor of 2, followed by a steep drop (Fig. 2). Luminosity in excess of Eddington should drive further radius expansion, or even be converted into a wind, a mechanism that these models lack. In order to compare our results

Figure 2. The raw simulated burst light curves, in the frame of the model. PRE bursts often exceed the Eddington luminosity in KEPLER, presumably because the physics of photospheric expansion and contraction are not accurately captured by the simple atmosphere, in addition to the lack of an outflow/wind mechanism. Low surface resolution causes the visible steps during the peak, due to convection switching on, as it spreads through these outer zones. To compare with the observations, we manually truncated the light curves at L_{Edd} (⁴He) during analysis (Section 2.6; Fig. 5).

with observations, we manually truncated the light curves at the Eddington luminosity for pure helium, $L_{Edd} = 3.53 \times 10^{38}$ erg s^{-1}, the inferred limit reached for PRE bursts (Kuulkers et al. 2003).

3 RESULTS

We present here a model that closely reproduced the observed burst times. We must emphasize that we present this not as a best-fitting model to the observations, but to demonstrate the feasibility of modelling bursts under varying accretion rates. A more detailed and systematic exploration of model parameters is planned as a future study.

The model produced a sequence of 10 X-ray bursts over the course of the outburst (Table 2; Fig. 3). The bursts have been assigned labels to aid in discussion. Seven are labelled sequentially from B1 to B7,

Table 2. Burst properties for the model with $X_0 = 0.44$ (for a description of burst labels, see Section 3). Observed bursts and their modelled counterparts are highlighted in bold.

Burst	t (h)	Δt (h)	E_b (10^{-6} erg cm^{-2})
P3	-33.1 ± 0.3	—	11.0 ± 0.3
P2	-24.7 ± 0.3	8.4 ± 0.3	2.11 ± 0.06
P1	-8.6 ± 0.5	16.1 ± 0.5	2.57 ± 0.08
B1	9.0 ± 0.8	17.6 ± 0.5	2.73 ± 0.08
O1	8.2^a	–	2.620 ± 0.021
B2	22.7 ± 0.9	13.6 ± 0.4	2.53 ± 0.08
B3	35.6 ± 0.9	12.9 ± 0.4	2.45 ± 0.07
B4	53.3 ± 1.1	44.3 ± 0.8^b	2.67 ± 0.08
O4	53.6	45.4	2.649 ± 0.018
B5	76.5 ± 1.3	23.2 ± 0.7	2.89 ± 0.09
O5	74.7	21.1	2.990 ± 0.017
B6	109.9 ± 1.6	33.4 ± 1.0	3.44 ± 0.10
O6	104.5	29.8	3.460 ± 0.022
B7	165 ± 2	54.6 ± 1.6	4.34 ± 0.13

Notes. ^{a}Observational time uncertainties are <1 s and are excluded for clarity.
^{b}Total interval between B1 and B4 to compare with the observed value.

Figure 3. Upper panel: Fluence, E_b, of the modelled burst sequence against the four observed bursts. Lower panel: the time-varying accretion rate over the event, as a fraction of the Eddington-limited accretion rate. The vertical grey bands indicate when the telescope was collecting data. Note that extra bursts predicted by the model fall outside these observing windows. As is typical for KEPLER models, the first burst was anomalously energetic, and its off-axis E_b is indicated next to the arrow. The fluences have been calculated from the burst energy with the scaling factor $c_2 = 4\pi d^2 \xi_b \approx 1.05 \times 10^{45} \, \mathrm{cm}^2$, chosen such that the RMS error with observations is minimized (Section 3.2).

where B1 is the closest in time to the first observed burst. The observed bursts are labelled O1, O4, O5, and O6, to correspond with their closest model bursts. Three bursts occur prior to B1, and are labelled P1, P2, and P3, in reverse order from B1, in anticipation of future simulations that may produce more or fewer such bursts.

We have applied GR corrections to the modelled burst properties, such that they correspond to a distant observer (Section 2.3). Error bars for the observational data are 1σ uncertainties. We have set the model uncertainties to 3 per cent, which is the typical 1σ variation in modelled burst trains when all input parameters (including $\dot{M}$) are held constant (Section 3.1). The uncertainties in the modelled burst arrival times were obtained by propagating the 3 per cent uncertainty in Δt along the burst train (the uncertainty for the first burst was simply taken to be that of the following recurrence interval). Additional model uncertainties due to the observational uncertainties in distance, inclination, etc. are not considered for the purpose of this paper, because we are not yet attempting parameter estimation.

3.1 Varying versus averaged accretion rates

We performed a comparison test in order to check for any difference in results between a varying and an averaged accretion rate. For each burst interval in the model, Δt_v, we calculated the average accretion rate, $\langle \dot{M} \rangle$. We then independently restarted the simulation at the beginning of each interval, with $\dot{M}$ now fixed at $\langle \dot{M} \rangle$. Once a sequence of 10–15 bursts were produced at each $\langle \dot{M} \rangle$, we calculated the mean recurrence times, $\langle \Delta t \rangle$. We then compared these with the original recurrence times via the ratio $\langle \Delta t \rangle / \Delta t_v$ and calculated the average slope of $\dot{M}(t)$ for each interval, given by $\Delta \dot{M} / \Delta t_v$ (Fig. 4). The standard deviation in burst properties for each $\langle \Delta t \rangle$ was typically ≈ 3 per cent, and we have adopted this as the standard model uncertainty. We have excluded the first burst interval (P3-P2) from this comparison because it is abnormally short, with a recurrence time almost half the length of the next, despite a lower accretion rate. We attribute this to extra heating from the very large preceding burst, P3.

Figure 4. The result of using a constant $\langle \dot{M} \rangle$ in place of a varying $\dot{M}(t)$. The vertical axis is the ratio of the recurrence time from a constant accretion rate ($\langle \Delta t \rangle$) to the recurrence time from a varying accretion rate (Δt_v). On the horizontal axis is the average slope of $\dot{M}(t)$ for the interval. The trendline is a weighted least-squares regression that has been forced through the point $(0,1)$, and has a slope of 0.195.

The variables $\langle \Delta t \rangle / \Delta t_v$ and $\Delta \dot{M} / \Delta t_v$ have a Spearman's rank correlation coefficient of $r_s = 0.60$, with a p-value of $p = 0.12$. A least-squares linear regression, weighted with the uncertainties and forced through the point $(0,1)$, returns a slope of 0.195. The correlation appears to have a low significance, although it is somewhat inconclusive with such a small sample. The scatter may be due in part to the small-scale variations in $\dot{M}(t)$ within each interval, given that $\Delta \dot{M} / \Delta t_v$ is itself an approximation to the slope. Nevertheless, this tentatively suggests that using a constant $\langle \dot{M} \rangle$ may systematically overestimate Δt_v when $\dot{M}(t)$ is increasing, and underestimate Δt_v when $\dot{M}(t)$ is decreasing. This discrepancy could have a significant effect on predictions, because burst properties are strongly dependent on the recurrence time. It's possible that the relationship is dependent on other parameters such as Q_b, X_0, and Z_{CNO}. With further study, a correction factor could account for this systematic discrepancy when an averaged accretion rate cannot be avoided, as with the semi-analytic models of G06.

3.2 Comparison with observed bursts

The four observed burst times are matched to within 0.85, 0.27, 1.8, and 5.4 h, respectively. The rms error for the four bursts is 2.88 h, in comparison to their recurrence times of $18 \lesssim \Delta t \lesssim 33$ h. An additional source of discrepancy here may be the gaps in PCA data. Because we have used linear interpolation in the PCA gaps, any unobserved variation in $\dot{M}$ is not captured in the model.

Six extra bursts were predicted in addition to the four observed. Three of these (P1–P3) were during the accretion rise, preceding O1. Two intervening bursts (B2, B3) fell between O1 and O4, and a final burst (B7) fell during the outburst tail, ≈ 2.5 d after O6. All extra bursts fell within gaps in the PCA data, and so cannot immediately be ruled out by observations. The two intervening bursts, B2 and B3, agree with the predictions of G06, who concluded that two bursts likely occurred between O1 and O4. Our model also suggests that several other bursts may have been missed.

When converted from E_{nuc} via equation (9), the burst fluences, E_b, have a minimized rms error with the observations when $c_2 \approx 1.05 \times 10^{45} \, \mathrm{cm}^2$. Following this scaling, the rms error of the four bursts is $0.11 \times 10^{-6} \, \mathrm{erg \, cm}^{-2}$, in comparison to their fluence

range of $2.7 \lesssim E_b \lesssim 3.4 \, \mathrm{erg \, cm^{-2}}$. Overall, the model reproduces the observed trend of increasing E_b with recurrence time (Fig. 3). This is consistent with larger fuel layers accumulating due to a lowering accretion rate and a cooling envelope.

From equations (7) and (9), we have

$$\frac{\xi_p}{\xi_b} = \frac{c_1}{c_2}, \tag{10}$$

$$d = \sqrt{\frac{c_1}{4\pi\xi_p}} = \sqrt{\frac{c_2}{4\pi\xi_b}}, \tag{11}$$

where for this model $c_1/c_2 \approx 1.536$. According to the thin accretion disc model from He & Keek (2016, Model a), this anisotropy ratio occurs at an inclination of $i = 67.7°$, with $\xi_p = 1.85$ and $\xi_b = 1.21$. Using equation (11), this corresponds to a distance of $d = 2.7$ kpc, which is outside the range proposed by G06 of 3.5 ± 0.1 kpc. This may be due to the inherent differences between the models, that the anisotropy was not considered in their study, or because our model is not yet a global best fit to the system. Additionally, a thin accretion disc may not be realistic for a transiently accreting LMXB, and other disc geometries produce different relationships between ξ_p and ξ_b. For instance, a concave disc (model d of He & Keek 2016) predicts the same anisotropy ratio at $i = 62.6°$, with $\xi_p = 1.34$ and $\xi_b = 1.10$, corresponding to a distance of $d = 3.1$ kpc, closer to the accepted value.

Following Eddington-truncation and scaling, the burst light curves broadly reproduce the observed profiles, featuring sharp rise times and gradual decays (Fig. 5). The plateau peak has been fixed at L_{Edd}, and after scaling c_2 to match the observed E_b, it approximately agrees with the observed peaks.

4 CONCLUSION

We performed simulations of a neutron star envelope during an accretion outburst, using the time-dependent $\dot{M}(t)$ inferred from the 2002 October event of SAX J1808.4-3658. A composition of $Z_{CNO} = 0.02$ and $X_0 = 0.44$ reproduced the four observed burst arrival times with an rms error of 2.88 h, with recurrence times of $18 \lesssim \Delta t \lesssim 33$ h. The modelled sequence contained 10 bursts, six of which did not correspond to observed bursts: three during the accretion rise; two between the first and second observed bursts; and one during the tail, ≈ 2.5 d after the fourth observed burst. These extra bursts fell during times when the source was not being observed by *RXTE*, which had a duty cycle of 38 per cent for the outburst.

Due to the limitations of KEPLER when simulating PRE bursts, the model-predicted luminosities were manually truncated at L_{Edd} in order to compare other light-curve features to observations. The light curves have rapid rise times (<1 s) and fast decays (≈ 20 s), in agreement with the observed characteristics of helium-rich bursts. The burst fluences were reproduced with an rms error of $0.11 \times 10^{-6} \, \mathrm{erg \, cm^{-2}}$ after scaling by $c_2 = 4\pi d^2 \xi_b \approx 1.05 \times 10^{45} \, \mathrm{cm^2}$.

To obtain similar recurrence times to the best-fitting models of G06, we required a lower hydrogen fraction of $X_0 = 0.44$, in comparison to $X_0 = 0.54$. A difference is perhaps unsurprising given the degrees of complexity between the models. For example, the semi-analytic model of G06 solves for thermonuclear stability with a one-zone approximation, does not evolve bursts with time, and uses a simple expression for energy yield to calculate the total burst energy. Furthermore, our results suggest that using averaged accretion rates may overestimate the recurrence times for an increasing

Figure 5. The light curves for each modelled burst (solid green curve) against its observed counterpart (black bins). The model luminosity has been converted to an observed flux using the scaling factor $c_2 = 4\pi d^2 \xi_b = 1.05 \times 10^{45} \, \mathrm{cm^2}$, chosen such that the burst fluences best match the observations (Section 3.2).

$\dot{M}$, and underestimate recurrence times for a decreasing $\dot{M}$. Efforts are currently underway to improve the semi-analytic model of G06 and its application to the 2002 October outburst.

As mentioned, the model presented here is not yet a global fit to the data, and so posterior constraints on the system parameters are not yet possible. In a future study, a more systematic matching could be performed by varying Z_{CNO}, Q_b, M, R, c_1, c_2, and the assumed accretion onset. The strength of the crustal heating, Q_b, may itself evolve with $\dot{M}$. In addition to the burst times, other properties such as the peak luminosities, light-curve profiles, and E_b could be incorporated into the fitting routine to obtain a global likelihood value.

Matching the properties of multiple bursts over a single accretion event provides a new test bed for multizone models. Our simulations are the first to adopt a time-dependent $\dot{M}$ and demonstrate that existing burst models can be extended to transient accretion regimes. This may help to further constrain LMXB properties, including the NS mass and radius, the strength of crustal heating, and the distance and inclination of the system. Furthermore, constraining the fuel composition provides information about the composition of the companion star, and thus the evolutionary history of the binary. This could improve our understanding of the SAX J1808.4-3658 system, and more generally, the origin of AMXPs.

REFERENCES

Belian R. D., Conner J. P., Evans W. D., 1976, ApJ, 206, L135
Bildsten L., Chakrabarty D., 2001, ApJ, 557, 292
Chakrabarty D., Morgan E. H., 1998, Nature, 394, 346
Chakrabarty D., Morgan E. H., Muno M. P., Galloway D. K., Wijnands R., van der Klis M., Markwardt C. B., 2003, Nature, 424, 42
Cumming A., Bildsten L., 2000, ApJ, 544, 453
Cyburt R. H. et al., 2010, ApJS, 189, 240
Cyburt R. H., Amthor A. M., Heger A., Johnson E., Keek L., Meisel Z., Schatz H., Smith K., 2016, ApJ, 830, 55
Dorman B., Arnaud K. A., 2001, ASPC, 238, 415
Fujimoto M. Y., 1988, ApJ, 324, 995
Fujimoto M. Y., Hanawa T., Miyaji S., 1981, ApJ, 247, 267
Galloway D. K., 2006, in AIP Conf. Proc. Vol. 840, The Transient Milky Way: A Perspective for Mirax. AIP Publishing, New York, p. 50
Galloway D. K., Cumming A., 2006, ApJ, 652, 559 (G06)
Galloway D. K., Muno M. P., Hartman J. M., Psaltis D., Chakrabarty D., 2008, ApJS, 179, 360
Galloway D. K., Goodwin A. J., Keek L., 2017, PASA, 34, e019
Grindlay J., Gursky H., Schnopper H., Parsignault D. R., Heise J., Brinkman A. C., Schrijver J., 1976, ApJ, 205, L127
Hartman J. M. et al., 2008, ApJ, 675, 1468
Hartman J. M., Watts A. L., Chakrabarty D., 2009, ApJ, 697, 2102
Heger A., Langer N., Woosley S. E., 2000, ApJ, 528, 368
Heger A., Cumming A., Galloway D. K., Woosley S. E., 2007, ApJ, 671, L141
He C.-C., Keek L., 2016, ApJ, 819, 47
Ibragimov A., Poutanen J., 2009, MNRAS, 400, 492
in 't Zand J. J. M., Heise J., Muller J. M., Bazzano A., Cocchi M., Natalucci L., Ubertini P., 1998, A&A, 331, L25
Jahoda K., Swank J. H., Giles A. B., Stark M. J., Strohmayer T., Zhang W. W., Morgan E. H., 1996. in Siegmund O. H., Gummin M. A., eds, Proc. SPIE Conf. Ser. Vol. 2808, EUV, X-Ray, and Gamma-Ray Instrumentation for Astronomy VII. SPIE, Bellingham, p. 59
Joss P. C., 1977, Nature, 270, 310
Keek L., Heger A., 2011, ApJ, 743, 189
Kuulkers E., den Hartog P. R., in't Zand J. J. M., Verbunt F. W. M., Harris W. E., Cocchi M., 2003, A&A, 399, 663
Lampe N., Heger A., Galloway D. K., 2016, ApJ, 819, 46
Levine A. M., Bradt H., Cui W., Jernigan J. G., Morgan E. H., Remillard R., Shirey R. E., Smith D. A., 1996, ApJ, 469, L33
Lewin W. H. G., Vacca W. D., Basinska E. M., 1984, ApJ, 277, L57
Lewin W. H. G., Paradijs J. V., Taam R. E., 1993, Space Sci. Rev., 62, 223
Patruno A., Watts A. L., 2012, preprint (arXiv:1206.2727)
Patruno A., Bult P., Gopakumar A., Hartman J. M., Wijnands R., van der Klis M., Chakrabarty D., 2012, ApJ, 746, L27
Patruno A. et al., 2017, ApJ, 841, 98
Rauscher T., Heger A., Hoffman R. D., Woosley S. E., 2002, ApJ, 576, 323
Strohmayer T., Bildsten L., 2003, preprint (arXiv:astro-ph/0301544)
Tawara Y. et al., 1984, ApJ, 276, L41
Ubertini P., Bazzano A., Cocchi M., Natalucci L., Heise J., Muller J. M., in 't Zand J. J. M., 1999, ApJ, 514, L27
Wang Z. et al., 2001, ApJ, 563, L61
Weaver T. A., Zimmerman G. B., Woosley S. E., 1978, ApJ, 225, 1021
Wijnands R., 2004, Nucl. Phys. B Proc. Suppl., 132, 496
Wijnands R., van der Klis M., 1998, Nature, 394, 344
Woosley S. E., Taam R. E., 1976, Nature, 263, 101
Woosley S. E. et al., 2004, ApJS, 151, 75

This paper has been typeset from a TEX/LATEX file prepared by the author.

Chapter 5

Multi-epoch X-ray burst modelling: MCMC with large grids of 1D simulations

Multi-epoch X-ray burst modelling: MCMC with large grids of 1D simulations

Zac Johnston,[1,2]* Alexander Heger,[1,2,3,4] Duncan K. Galloway[1,2]

[1]*School of Physics and Astronomy, Monash University, Victoria 3800, Australia*

[2]*Monash Centre for Astrophysics, Monash University, Victoria 3800, Australia*

[3]*Department of Astronomy, Shanghai Jiao Tong University, Shanghai 200240, China*

[4]*School of Physics and Astronomy, University of Minnesota, Minneapolis, Minnesota 55455, USA*

ABSTRACT

Type-I X-ray bursts are recurring thermonuclear explosions on the surface of accreting neutron stars. Matching observed bursts to computational models can help to constrain system properties, such as the neutron star mass and radius, crustal heating rates, and the accreted fuel composition, but systematic parameter studies to date have been limited. We apply Markov chain Monte Carlo methods to 1D burst models for the first time, and obtain system parameter estimations for the 'Clocked Burster', GS 1826−238, by fitting multiple observed epochs simultaneously. We explore multiple parameters which are often held constant, including the neutron star mass, crustal heating rate, and hydrogen composition. To improve the computational efficiency, we precompute a grid of 3840 KEPLER models – the largest set of 1D burst simulations to date – and by interpolating over the model grid, we can rapidly sample burst predictions. We obtain estimates for a CNO metallicity of $Z_{\mathrm{CNO}} = 0.010^{+0.005}_{-0.004}$, a hydrogen fraction of $X_0 = 0.74^{+0.02}_{-0.03}$, a distance of $d\sqrt{\xi_{\mathrm{b}}} = 6.5^{+0.4}_{-0.6}$, and a system inclination of $i = 69^{+2\,\circ}_{-3}$.

Key words: X-rays: bursts – stars: neutron – stars: individual: GS 1826-238 – methods: numerical

1 INTRODUCTION

Type-I thermonuclear X-ray bursts are recurring flashes observed from accreting neutron stars (for reviews, see Lewin et al. 1993; Strohmayer & Bildsten 2006; Galloway & Keek 2017). In the host low-mass X-ray binary (LMXB) systems, a neutron star accretes material from a companion star with a mass of $M \lesssim 1\,\mathrm{M}_\odot$, which accumulates as a $\sim 10\,\mathrm{m}$ envelope on the neutron star surface. Under the weight of accreting material, the base of the envelope is compressed by the extreme surface gravity of $g \sim 10^{14}\,\mathrm{cm\,s^{-2}}$ to the point of thermonuclear runaway. Within seconds, the layer is heated to $\sim 10^9\,\mathrm{K}$, generating a burst of X-rays before cooling to background levels over the following seconds to minutes. New fuel is accreted on top of the 'ashes', and the cycle repeats.

X-ray bursts have been the target of numerical calculations since the 1970s (e.g., Joss 1978; Taam & Picklum 1979), and their diverse behaviour has been studied with a variety of computational models (e.g., Fujimoto et al. 1987; Woosley et al. 2004; Keek et al. 2012). By exploring model parameters

and comparing the predictions with observations, the neutron star system properties can be inferred (e.g., Cumming 2003; Galloway et al. 2004; Keek & Heger 2017; Johnston et al. 2018).

One-dimensional (1D) burst codes are the best tools currently available for this purpose. With adaptive nuclear reaction networks and treatments for convective transport (e.g., Woosley et al. 2004), they can produce detailed simulations of burst energetics not possible in semi-analytic or one-zone models. Multi-dimensional burst simulations are also under active development, but computational costs limit the calculations to $\lesssim 1\,\mathrm{s}$ of simulation time (e.g., Zingale et al. 2015; Cavecchi et al. 2016).

Nevertheless, targeted parameter explorations using 1D models have been relatively limited. Small sets of models are typically used, and many parameters, such as the gravity, fuel composition, and crustal heating, are often held constant. Due to the relatively unexplored parameter space, obtaining robust constraints on system properties is difficult. To encourage more directed modelling efforts, Galloway et al. (2017, hereafter, G17) presented a set of standardised burst observations. Their reference data set included three

*E-mail: zac.johnston@monash.edu

epochs of bursts from GS 1826−238, famously dubbed the 'clocked burster' (e.g., Ubertini et al. 1999). The system's reliability has made it a popular target for modelling (e.g., Galloway et al. 2004; Heger et al. 2007), and particularly for the study of the nuclear *rp*-process (e.g., Schatz et al. 1998; Fisker et al. 2008). The first study to make use of the G17 data set was Meisel (2018, hereafter, M18), who performed the first extended comparison of MESA burst models to GS 1826−238. M18 demonstrated the benefit of fitting multiple epochs by ruling out parameter combinations which otherwise agreed with individual epochs.

GS 1826−238 was discovered as a transient source with the *Ginga* X-ray telescope in 1988 (Makino 1988), and X-ray bursts were later discovered in 1997 (Ubertini et al. 1997, 1999). The system orbital period is not precisely known, but is thought to be roughly 2 h (Homer et al. 1998), implying a hydrogen-rich mass donor, consistent with the long-tailed bursts observed (in't Zand et al. 2009). Despite the popularity of GS 1826−238 for modelling, ambiguity persists regarding the system properties. For example, Galloway et al. (2004, hereafter, G04) modelled bursts observed between 1997 and 2002 using a semi-analytic ignition model (SETTLE, first used in Cumming & Bildsten 2000). They reported that an accreted CNO mass fraction of $Z_{\rm CNO} = 0.001$ best reproduced the trend of recurrence time, Δt, versus accretion rate, $\dot{M}$, but that the observed α ratios were only consistent with a higher metallicity of $Z_{\rm CNO} = 0.02$. Using KEPLER, Heger et al. (2007, hereafter, H07) found good lightcurve agreement[1] for $Z_{\rm CNO} = 0.02$, and M18 found agreement for both $Z_{\rm CNO} = 0.01$ and 0.02 using MESA. The accretion rates are typically inferred to be in the range $\dot{M} = 0.05$–$0.08\ \dot{M}_{\rm Edd}$ (where $\dot{M}_{\rm Edd} = 1.75 \times 10^{-8}\ {\rm M}_\odot\ {\rm yr}^{-1}$; Heger et al. 2007; Galloway et al. 2008, 2017), but M18 reported improved model fits using twice as large accretion rates of $\dot{M} = 0.1$–$0.17\ \dot{M}_{\rm Edd}$.

The inconclusive estimates are, we suggest, partly due to the limited parameter explorations to date, in addition to degeneracies between the model predictions. For example, the metallicity is often fixed at $Z_{\rm CNO} = 0.02$, with an accreted hydrogen fraction of $X_0 = 0.7$. G04 and H07 used fixed crustal heating rates of $Q_{\rm b} = 0.1$ and $0.15\ {\rm MeV\ nucleon}^{-1}$, respectively, whereas M18 considered $Q_{\rm b} = 0.1, 0.5$, and $1.0\ {\rm MeV\ nucleon}^{-1}$. Both H07 and M18 assumed a fixed neutron star mass of $M = 1.4\ {\rm M}_\odot$ and a radius of $R = 11.2\ {\rm km}$, whereas G04 assumed $M = 1.4\ {\rm M}_\odot$ and $R = 10\ {\rm km}$. The earlier estimates for $\dot{M}$ did not account for the possible effect of anisotropic emission (§ 2.6), which is dependent on the system inclination and disc morphology (Fujimoto 1988). Using the disc models of He & Keek (2016), M18 inferred an approximate inclination of 65–80°, suggesting that the X-ray emission is preferentially beamed *away* from the line of sight, allowing for larger $\dot{M}$. To fully account for the complex dependencies between these model parameters and predictions, a more comprehensive analysis is required.

Markov chain Monte Carlo (MCMC) methods are algorithms capable of sampling complex probability distributions (for a comprehensive introduction, see MacKay 2003).

The use of Bayesian statistics in astrophysics has seen a rapid expansion in recent years, but its application to X-ray burst modelling has been minimal. Most recently, Goodwin et al. (2019) applied MCMC methods to the semi-analytic burst code, SETTLE, to model bursts from the transient accretor, SAX J1808.4−3658. This system is also included in the G17 data set, as an example of helium bursts triggered during an accretion event. Pairing a semi-analytic model with MCMC is beneficial due to the computational speed required for drawing thousands of sequential samples. By contrast, 1D burst models can take several days to compute, and are, on their own, unsuitable for MCMC methods.

As we show here for the first time, this computational barrier can be overcome with the use of pre-compiled model grids. For burst properties that vary smoothly over the model parameters, interpolation can be used to sample bursts between existing models with little computational cost. We present the first application of MCMC methods to large grids of 1D burst models. By constructing a grid of 3840 KEPLER simulations, we are able to rapidly sample burst properties across twelve parameters. Using the data set from G17, we fit three epochs of burst data simultaneously, and obtain probability distributions for the system parameters of GS 1826−238.

In Section 2, we describe the KEPLER code and its recent updates, the epoch data used, the construction and interpolation of the model grid, and the setup of the MCMC model. In Section 3, we describe the model results, the posterior distributions, the predicted burst properties, and lightcurve comparisons. In Section 4, we discuss and compare the parameter estimates to previous works, discuss the limitations of the model, and describe potential improvements to the model. In Section 5, we provide concluding remarks and the future outlook.

2 METHODS

2.1 An update on Kepler

KEPLER (Weaver et al. 1978) is a one-dimensional (1D) hydrodynamics code capable of simulating a variety of regimes in stellar evolution and explosive nucleosynthesis (e.g., Woosley et al. 2002; Menon & Heger 2017). It has prominently been used for modelling X-ray bursts, reproducing observed behaviour, including burst energetics, recurrence times, and lightcurves (Woosley et al. 2004; Heger et al. 2007; Keek et al. 2012; Lampe et al. 2016). Because KEPLER has steadily been modified and improved over time, some descriptions in earlier works are now out of date. We here briefly summarise notable changes to the code and model setup.

To aid reproducibility and comparisons to other burst models, we used V2.2 of JINA REACLIB, the public database of nuclear reaction rates[2] (Cyburt et al. 2010).

KEPLER burst models published prior to Johnston et al. (2018) used a setup file which erroneously multiplied the opacities by a factor of ≈1.5. The original intent was to approximate the time-dilation effects of general relativity (GR; Appendix B) by artificially slowing down thermal transport.

---

[1] we note that these models were discovered to use inadvertently large opacities; see § 2.1

[2] https://jinaweb.org/reaclib/db/

Table 1. The observed burst data from three epochs of GS 1826−238, adapted from the reference set of G17. As an additional system constraint, we have included the Eddington flux, F_{Edd}, taken from the peak of a PRE burst observed in 2014 June (Chenevez et al. 2016). To estimate the average recurrence time Δt, G17 collected multiple bursts from each epoch, from which we have obtained the burst rate ν. They then extracted average lightcurves, from which the peak flux, F_{peak}, and fluence, f_{b}, could be determined. The persistent flux F_{p} was averaged for each epoch, and the values listed here have incorporated the bolometric corrections estimated by G17.

Epoch	ν (day^{-1})	F_{peak} ($10^{-9}\,\mathrm{erg\,s^{-1}\,cm^{-2}}$)	f_{b} ($10^{-6}\,\mathrm{erg\,cm^{-2}}$)	F_{p} ($10^{-9}\,\mathrm{erg\,s^{-1}\,cm^{-2}}$)	F_{Edd} ($10^{-9}\,\mathrm{erg\,s^{-1}\,cm^{-2}}$)
1998 Jun	4.67 ± 0.06	30.9 ± 1.0	1.102 ± 0.011	2.108 ± 0.015	–
2000 Sep	5.746 ± 0.014	29.1 ± 0.5	1.126 ± 0.016	2.85 ± 0.03	–
2007 Mar	6.799 ± 0.008	28.4 ± 0.4	1.18 ± 0.04	3.27 ± 0.04	–
2014 Jun	–	–	–	–	40 ± 3

Figure 1. An example burst train simulated using KEPLER (upper panel), and the extracted and stacked burst lightcurves (lower panel) from which average properties can be calculated (§ 2.4).

While the idea was soon abandoned and removal of the factor was intended, it mistakenly remained in the setup file which was adapted for subsequent studies. This error was discovered[3] and amended for the models presented in Johnston et al. (2018). The boosted opacities reduced the thermal conductivity, resulting in an artificially hotter envelope and shorter burst recurrence times. This discrepancy likely explains why comparisons from other burst codes, for example MESA (M18) and SHIVA (José et al. 2010), reported longer recurrence times than the equivalent KEPLER models. This issue should be kept in mind when making comparisons to previous KEPLER models (e.g., Heger et al. 2007; Lampe et al. 2016).

During the setup phase of the model envelope, before accretion and nuclear reactions are switched on, the thermal profile is initialised near equilibrium, in order to minimise simulation 'burn-in'. In previous KEPLER studies (and to our knowledge, all other burst studies in the literature), the only heat source included was the crustal heating, $Q_{\mathrm{b}} \approx 0.15\,\mathrm{MeV\,nucleon^{-1}}$, as a boundary condition at the base of the model, at a column depth of $y \approx 1 \times 10^{12}\,\mathrm{g\,cm^{-2}}$. Additional heat generated by nuclear reactions, $Q_{\mathrm{nuc}} \approx 5\,\mathrm{MeV\,nucleon^{-1}}$, around $y \sim 10^{8}\,\mathrm{g\,cm^{-2}}$, was assumed to largely escape the surface, and was neglected from the setup calculations. Minor heating of the deeper ocean, once the full nuclear simulation began, was expected to stabilise after the initial few bursts. For example, Woosley et al. (2004) discarded the first three bursts from analysis to address this 'thermal inertia', in addition to the related effect of 'chemical inertia' (§ 2.4).

During testing, however, we discovered that nuclear heating can indeed alter the thermal profile enough to influence burst ignition. Models which do not account for Q_{nuc} during setup begin comparatively colder, producing a steadily-changing burst sequence as the envelope is heated by nuclear reactions towards a steady thermal state. This burn-in period can persist for dozens of bursts – much longer than previously assumed. To address this issue, we added during setup a heat source of $Q_{\mathrm{nuc}} = 5\,\mathrm{MeV\,nucleon^{-1}}$ at a depth of $y \approx 8 \times 10^{8}\,\mathrm{g\,cm^{-2}}$ with a Gaussian spread of $y \approx 8 \times 10^{7}\,\mathrm{g\,cm^{-2}}$. This heat source is switched off once the full nuclear calculations begin (further detail is provided in Appendix A). The model burn-in was effectively eliminated, and the burst sequence was stabilised within the first few bursts, as was originally assumed. Further study is still required to explore the optimal configuration of this 'pre-heating', which is likely to depend on other model parameters, such as the composition, accretion rate, and crustal heating.

2.2 Observed data

We used observations of bursts from GS 1826−238 for three epochs: 1998 June, 2000 September, and 2007 March (Table 1). These observations were provided as part of a reference set for burst modelling by G17, using data from the MINBAR catalogue[4]. The epochs were selected by G17 for their burst consistency and the availability of high-precision X-ray lightcurves from the *RXTE* satellite.

---

[3] by Adam Jacobs, Michigan State University, pers. comm.

[4] http://burst.sci.monash.edu/minbar

We also included the peak bolometric flux of $F_{peak} = (40 \pm 3) \times 10^{-9} \, \mathrm{erg\,s^{-1}\,cm^{-2}}$, observed from a photospheric radius expansion (PRE) burst in 2014 (Chenevez et al. 2016). We have assumed that F_{peak} corresponds to the local Eddington luminosity, F_{Edd}, for a mixed hydrogen/helium envelope (§ 2.7), and that F_{Edd} is common to the 1998–2007 epochs.

2.3 Model grid

To model the bursts of GS 1826−238, we computed a regular grid of 3840 KEPLER simulations over five model parameters: the accretion rate, $\dot{m}$, the accreted hydrogen mass fraction X_0, the accreted CNO-metallicity mass fraction, Z_{CNO}, the crustal heating rate, Q_b, and the surface gravitational acceleration, g. Note that for the KEPLER model parameters, we give the local accretion rate per unit area, $\dot{m}$, because the global accretion rate, $\dot{M} = 4\pi R^2 \dot{m}$, depends on the choice of R (§ 2.6). The numerical parameters controlling zone resolution were held constant, following convergence tests to ensure consistent burst sequences. The grid points for each parameter are listed in Table 2.

Following a trial parameter exploration, we chose a grid that approximately covered the observed recurrence times of $3 \lesssim \Delta t \lesssim 6 \, \mathrm{h}$. The model grid represents over 100 000 CPU hours, and is the largest collection of 1D burst models to date, with the previous largest containing 464 KEPLER models (Lampe et al. 2016).

Each model generated a sequence of 30–40 bursts (Fig. 1). Simulating a long sequence ensures a consistent burst train, and reduces the effect of model burn-in (§ 2.1). The entire grid contained a total of approximately 138 000 bursts. Modelling such large collections of bursts has been made possible by improved CPU speeds, which have reduced the computational cost from $\approx 24 \, \mathrm{h}$ per burst in 2003 to $\approx 1 \, \mathrm{h}$ per burst.

2.4 Extracting model bursts

Analysing bursts from KEPLER models has been detailed previously (Woosley et al. 2004; Heger et al. 2007; Keek & Heger 2011; Lampe et al. 2016). We briefly summarise our analysis routine, for which we developed a new open source PYTHON package, PYBURST[5].

The procedure identified bursts from local maxima in the model lightcurve, from which the the individual lightcurves were extracted and analysed (Fig. 1). The peak luminosity, L_{peak}, was taken from the lightcurve maximum. The burst energy, E_b, was the time-integrated luminosity over the lightcurve. The recurrence time, Δt, was the time since the previous burst as measured peak-to-peak, giving a burst rate of $\nu = 1/\Delta t$.

Our inclusion of a nuclear heat source during initialisation of the envelope had substantially reduced thermal burn-in (§ 2.1). Despite this improvement, the first burst still ignites in a chemically pristine envelope, which lacks the complex ashes later accumulated. Due to 'chemical inertia', the models typically also require a 'chemical burn-in' of several bursts to reach a quasi-stable bursting pattern.
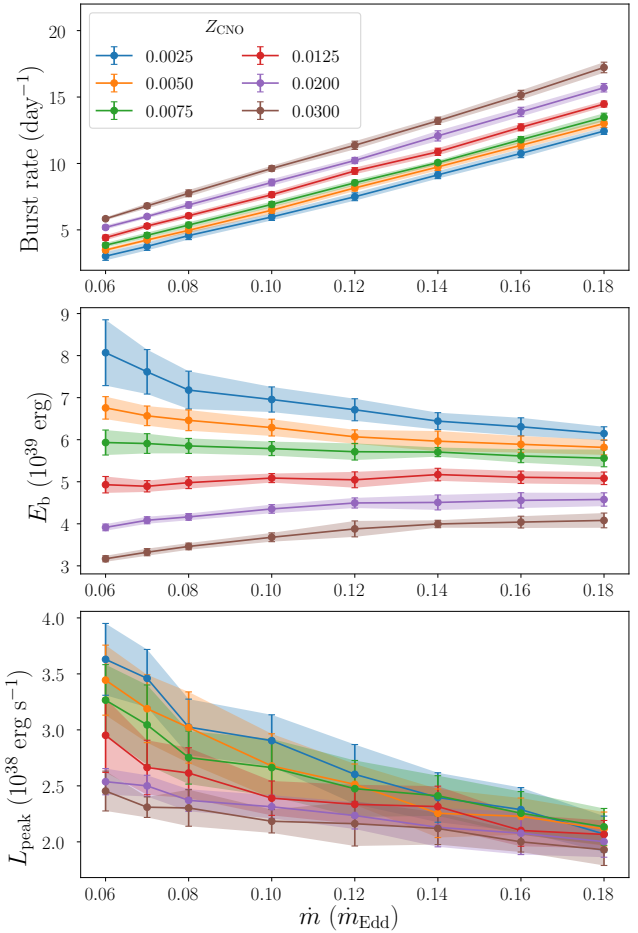
Figure 2. A subset of 48 KEPLER models, from a total grid of 3840, illustrating the dependence of the burst properties on accretion rate, $\dot{m}$, and metallicity, Z_{CNO}. The subset is a grid slice through $X_0 = 0.64$, $Q_b = 0.2 \, \mathrm{MeV\,nucleon^{-1}}$, and $g = 2.256 \times 10^{14} \, \mathrm{cm\,s^{-2}}$. Each point was averaged from a sequence of 20–30 bursts (§ 2.4), and the error bars are 1σ standard deviations. The overall smooth and monotonic behaviour allows for the use of interpolation. The solid lines and shaded uncertainty regions were linearly interpolated between the models. Multivariate linear interpolation was used across all grid parameters for MCMC sampling.

To minimize these combined burn-in effects, we excluded the first 10 bursts of each model from analysis. In previous studies, only ≈ 3 had typically been discarded (Woosley et al. 2004).

Modelled and observed X-ray bursts are occasionally followed by short waiting-time bursts ($\Delta t \lesssim 45 \, \mathrm{min}$), which are thought to be triggered by the ignition of unburned hydrogen (Keek & Heger 2017). These unusually weak bursts were not included in the data set of G17, and we excluded them from our analysis using the threshold of $\Delta t < 45 \, \mathrm{min}$.

The extracted properties of the remaining 20–30 bursts of each model were averaged, and the 1σ standard deviations were adopted as the model uncertainties. This process produced tabulated burst properties across the model grid parameters (Fig. 2), which could be sampled with minimal computational cost.

Table 2. The parameters of the model grid. Every combination was simulated, forming a regular grid in 5 dimensions. All models used a preheating value of $Q_{\rm nuc} = 5\,{\rm MeV\,nucleon}^{-1}$ (§ 2.1). The accretion rate, $\dot{m}$, is the local rate per unit area, because the global accretion rate, $\dot{M}$, is dependent on the choice of R (§ 2.6). The Eddington-limited accretion rate, $\dot{m}_{\rm Edd} = 8.775 \times 10^4\,{\rm g\,cm}^{-2}\,{\rm s}^{-1}$ (assuming $M = 1.4\,{\rm M}_\odot$, $R = 10\,{\rm km}$, and $X = 0.7$), is simply used as a common reference point between models, and was not corrected for the g or X_0 of each model. The values of g correspond to Newtonian surface gravities for masses of $M = 1.4$, 1.7, 2.0, and $2.6\,{\rm M}_\odot$, for a reference radius of $R = 10\,{\rm km}$, but the actual mass and radius are free parameters (§ 2.6).

Parameter	Name	Units	Grid Points	N
$\dot{m}$	Local accretion rate	$(\dot{m}_{\rm Edd})$	0.06, 0.07, 0.08, 0.10, 0.12, 0.14, 0.16, 0.18	8
X_0	H mass fraction	–	0.64, 0.67, 0.70, 0.73, 0.76	5
$Z_{\rm CNO}$	CNO mass fraction	–	0.0025, 0.005, 0.0075, 0.0125, 0.02, 0.03	6
$Q_{\rm b}$	Crustal heating	$({\rm MeV\,nucleon}^{-1})$	0.0, 0.2, 0.4, 0.6	4
g	Surface gravity	$(10^{14}\,{\rm cm\,s}^{-2})$	1.858, 2.256, 2.654, 3.451	4
			Total	3840

2.5 Grid interpolation

Computational speed is critical for the large-scale sampling used in MCMC methods. The smooth and monotonic behaviour of the burst properties (Fig. 2) allowed for the use of interpolation between model predictions. Using multivariate linear interpolation, we constructed a continuous function of the burst properties across the five grid parameters, $\dot{m}$, $Q_{\rm b}$, X_0, $Z_{\rm CNO}$, and g.

Any point within the grid bounds of Table 2 could then be quickly ($\ll 1\,{\rm s}$) sampled to predict the burst properties of interest: ν, $L_{\rm peak}$, and $E_{\rm b}$. In contrast to the roughly 40–90 h required for each KEPLER simulation, this approach granted a considerable efficiency gain.

The interpolated burst properties were in the local neutron star frame of the KEPLER models. In order to compare with the observed data (Table 1) we converted these values to observable quantities. These calculations first accounted for the fact that KEPLER uses Newtonian gravity. A KE-PLER model with a given Newtonian surface gravity, g, can be considered equivalent to a neutron star with an equal g under GR, but a different 'true' mass and radius (e.g., Woosley et al. 2004; Keek & Heger 2011). In the interest of clarity, we will here use quantities that are already corrected for GR[6]. The basic corrections for Newtonian quantities are given in Appendix B.

2.6 Free parameters: mass, distance, and anisotropy

In addition to the grid parameters (Table 2), a set of 'free' parameters are used for calculating observables.

For a given surface gravity, g, we can freely choose the neutron star mass, M, which determines the corresponding radius, R. Both M and R then determine the gravitational redshift, z, which quantifies the neutron star compactness, given by

$$z = \frac{1}{\sqrt{1 - 2GM/(c^2 R)}} - 1, \qquad (1)$$

where G is the gravitational constant and c is the speed of light. The distance to the system, d, is freely chosen.

Finally, we can choose the anisotropy factors, $\xi_{\rm b}$ and $\xi_{\rm p}$. These parameters represent the deviation of the observed

---

[6] These corrections are not applicable to codes which already use GR surface gravity, for example, MESA

flux, F, from an isotropic neutron star luminosity, L, caused by the scattering and blocking of light by the accretion disc (e.g., Fujimoto 1988; He & Keek 2016). The anisotropy parameters are defined with

$$F_{\rm b} = \frac{L_{\rm b}}{4\pi d^2 \xi_{\rm b}}, \qquad F_{\rm p} = \frac{L_{\rm p}}{4\pi d^2 \xi_{\rm p}}, \qquad (2)$$

where the subscripts 'b' and 'p' correspond to the burst and persistent emission, respectively.

Because $\xi_{\rm b}$ and $\xi_{\rm p}$ are degenerate with distance, we combined them into independent parameters: a modified distance, $d\sqrt{\xi_{\rm b}}$, and the anisotropy ratio, $\xi_{\rm p}/\xi_{\rm b}$. We can later retrieve the absolute values for $\xi_{\rm b}$, $\xi_{\rm p}$, and d by choosing an accretion disc model which relates the anisotropy to the system inclination, i (§ 3.3).

2.7 Transforming to observable quantities

The free parameters can then be used to calculate observables from the local burst properties. We here signify observed quantities with the subscript '∞'.

The burst rate as seen by a distant observer is time-dilated with

$$\nu_\infty = \frac{\nu}{1 + z}. \qquad (3)$$

The observed peak flux is given by

$$F_{\rm peak,\infty} = \frac{L_{\rm peak}}{4\pi d^2 \xi_{\rm b}(1 + z)^2}. \qquad (4)$$

The observed fluence is given by

$$f_{\rm b,\infty} = \frac{E_{\rm b}}{4\pi d^2 \xi_{\rm b}(1 + z)}, \qquad (5)$$

where we use the redshift factor of $1 + z$ instead of $(1 + z)^2$, because fluence is time-integrated.

The two remaining observables, $F_{\rm Edd}$ and $F_{\rm p}$, are not predicted from the model grid, but are calculated analytically from the given parameters.

The observed Eddington flux is given by

$$F_{\rm Edd,\infty} = \frac{L_{\rm Edd}}{4\pi d^2 \xi_{\rm b}(1 + z)^2}, \qquad (6)$$

where $L_{\rm Edd}$ is the local Eddington luminosity, given by

$$L_{\rm Edd} = \frac{8\pi G m_{\rm p} c(1 + z)M}{\sigma_{\rm T}(X_0 + 1)}, \qquad (7)$$

where $m_{\rm p}$ is the mass of the proton and $\sigma_{\rm T}$ is the Thompson scattering cross section. This equation assumes that the

radiation pressure is exerted on the electrons of an ionized plasma, and includes a scaling factor of $2/(X_0 + 1)$ to account for the charge per mass for a given composition of hydrogen and helium.

The observed persistent flux is given by

$$F_{p,\infty} = \frac{L_p}{4\pi d^2 \xi_p (1+z)^2},$$

(8)

where $L_p = -4\pi R^2 \dot{m}\phi$ is the local accretion luminosity, and $\phi = -c^2 z/(1+z) \approx -0.2\,c^2$ is the gravitational potential at the neutron star surface.

2.8 Multi-epoch model

To model bursts from multiple epochs, we used our interpolated grid (§ 2.5) to predict the observed burst properties (§ 2.7) of three separate epochs from GS 1826−238 (§ 2.2). Our multi-epoch model contained both epoch-independent and epoch-dependent parameters.

The accreted composition and neutron star properties, X_0, Z_{CNO}, g, and M, are expected to remain unchanged between the observed epochs, and so global parameters are used. Global parameters were also used for the distance and anisotropy, $d\sqrt{\xi_b}$ and ξ_p/ξ_b, although the anisotropy factors ξ_b and ξ_p could feasibly evolve due to changes in the accretion disc. We leave the testing of epoch-dependent parameters of ξ_b and ξ_p for a future study. The six epoch-independent parameters are thus X_0, Z_{CNO}, g, M, $d\sqrt{\xi_b}$, and ξ_p/ξ_b.

Converesely, the accretion rate is expected to evolve between epochs, and so we used three parameters, $\dot{m}_1$, $\dot{m}_2$, and $\dot{m}_3$, where the subscripts 1–3 correspond to the 1998, 2000, and 2007 epochs, respectively. The crustal heating efficiency is predicted to depend on accretion rate (Cumming et al. 2006), and we similarly use three parameters, $Q_{b,1}$, $Q_{b,2}$, and $Q_{b,3}$. These were the first burst models to vary Q_b between accretion epochs.

The 12 parameters could then be separated into the grid parameters, $\dot{m}_i$, $Q_{b,i}$, X_0, Z_{CNO}, and g, for $i = 1, 2, 3$, and the free parameters, M, $d\sqrt{\xi_b}$, and ξ_p/ξ_b. For a given point in parameter space, the former were used to interpolate the local burst properties from the grid, and the latter were used to transform these properties into observables.

We could then apply MCMC methods to draw samples from the multi-epoch model, and compare the predictions to observations to obtain probability distributions over the parameter space.

2.9 MCMC methods

A detailed description of MCMC algorithms is beyond the scope of this paper, but many introductions are available (e.g., MacKay 2003), and we provide here a brief summary.

The target distribution to be sampled – the posterior probability distribution, or simply the 'posterior' – represents the probabilities over model parameters, given a set of data we wish to model. The posterior sampled with MCMC is given by

$$p(\theta|D) = p(D|\theta)p(\theta),$$

(9)

where θ represents the model parameters and D is the data

to be modelled. The likelihood function, or simply the 'likelihood', $p(D|\theta)$, represents the probability of observing the data given the model predictions. The prior distribution, or simply the 'prior', $p(\theta)$, represents any existing beliefs or constraints on the parameters.

An MCMC simulation consists of an ensemble of 'walkers' in parameter space. For each walker, the likelihood is evaluated at its location, and a new step is randomly drawn from a proposal distribution, typically a Gaussian centred on the walker. The likelihood is then evaluated at the proposed step, which is either accepted or rejected based on the relative probability of the two points. Through this repeated process, the walkers 'explore' the parameter space of the posterior distribution. After a sufficiently large number of steps, each point in the chain of steps represents an independent sample drawn from the posterior distribution. The density of the walkers thus corresponds to the probability density.

A major advantage of MCMC methods is the ability to 'marginalise' over uninteresting parameters, by projecting the probability density onto a subset of dimensions. We used marginalisation to produce 1D and two-dimensional (2D) distributions for the parameters of interest.

For our MCMC calculations, we used the open-source PYTHON package EMCEE[7] (V2.2.1), an affine-invariant ensemble sampler (Foreman-Mackey et al. 2013). We constructed the posterior function defined in Eq. (9), and used EMCEE to generate a chain of samples drawn from it. This function takes the 12 multi-epoch model parameters (§ 2.8) as input, and calculates a posterior likelihood using the prior, $p(\theta)$, and the model likelihood, $p(D|\theta)$.

For the grid parameters, the limits were set by the boundaries of the model grid (Table 2). For the free parameters, we imposed limits of $1.0 \leq M \leq 2.2\,\mathrm{M}_\odot$, $1 \leq d\sqrt{\xi_b} \leq 15\,\mathrm{kpc}$, and $0.1 \leq \xi_p/\xi_b \leq 10$. The prior distribution for each parameter was set to $p(\theta) = 0$ outside these boundaries.

We used flat prior distributions for all parameters except Z_{CNO}, setting $p(\theta) = 1$ everywhere inside the boundaries. For Z_{CNO}, we estimated a prior distribution using a process similar to Goodwin et al. (2019). From a simulated catalogue of Milky Way stars, constructed to represent the underlying distributions of the Gaia DR2 catalogue (Rybizki et al. 2018), we took a sample of 100 000 stars located within 15 arcmin of GS 1826−238, and between a distance of 5–9 kpc. We then fit a beta distribution to [Fe/H], obtaining the values of $\alpha = 10.1$ and $\beta = 3.5$ after translating to the interval $[-3.5, 1]$, which contained the vast majority of star samples. We used this as the prior distribution for $\log_{10}(Z_{\mathrm{CNO}}/0.01)$, where we have assumed a solar CNO metallicity of 0.01 (Lodders et al. 2009). This distribution was applied inside the grid bounds of $0.0025 \leq Z_{\mathrm{CNO}} \leq 0.03$, which roughly corresponds to $-0.6 \leq \log_{10}(Z_{\mathrm{CNO}}/0.01) \leq 0.5$.

For a given sample point in parameter space, θ, the local burst properties were interpolated from the model grid (§ 2.5), from which the observables were predicted (§ 2.7). The likelihood function, $p(D|\theta)$, from Equation (9), was then evaluated by comparing these predictions with the observed

---

[7] https://emcee.readthedocs.io/en/v2.2.1/

data, D, using

$$p(D|\theta) = \prod_x \frac{1}{2\pi(\sigma^2 + \sigma_0^2)} \exp\left[-\frac{(x - x_0)^2}{2(\sigma^2 + \sigma_0^2)}\right], \quad (10)$$

where x is iterated over each observable of each epoch, σ is the uncertainty, and the subscript '0' signifies the corresponding observed value from Table 1.

The MCMC model used an ensemble of 1000 walkers, which were initialised in a small 'hyperball' in parameter space. The algorithm was run for 20 000 steps, resulting in a total of 2×10^7 individual samples. The average computation time for each sample was ≈ 0.012 s, for a total of ≈ 560 CPU hours split over 8 cores. For comparison, each KEPLER simulation costs roughly 40–70 CPU hours.

We discarded the first 1000 steps as burn-in, after which the walkers had spread out across the domain of each parameter. To check convergence, we estimated the autocorrelation time (τ) at multiple steps in the chain[8], to ensure the total chain length was longer than $10\,\tau$.

3 RESULTS

We present here the distributions and estimates from the burst matching procedure, and discuss the implications for the system properties. The 2D marginalised posteriors for $\dot{m}_i$, $Q_{b,i}$, X_0, and Z_{CNO}, are shown in Fig. 3, with the 1D posteriors shown along the diagonal. The maximum likelihood estimates for the 1D posteriors are listed in Table 3. Unless otherwise stated, the uncertainties given for 1D parameter estimates are 68 per cent credible intervals, and the 2D contour levels are 38, 68, 87, and 95 per cent credible regions.

There is a strong correlation visible between the accretion rates of each epoch, and the crustal heating of each epoch. These correlations are expected, because for given ratios between the epoch burst properties, similar ratios are required between the epoch parameters. For example, the persistent flux, F_p, is calculated using Eq. (8), and is proportional to $\dot{m}$.

The CNO mass fraction, Z_{CNO}, is correlated with the hydrogen fraction X_0 – a common feature of such model-observation comparisons (e.g., Galloway & Cumming 2006; Goodwin et al. 2019). The correlation arises because multiple pairs of X_0 and Z_{CNO} result in the same reduced hydrogen fraction at ignition.

The distributions of some parameters, for example $Q_{b,1}$, X_0, M, and Z_{CNO}, appear to be truncated by the prior boundaries. These limits could bias the results, potentially underestimating the full extent of the distributions. Some of these boundaries were chosen as expected natural limits, whereas others are simply limited by the model grid. For example, X_0 is truncated at the upper grid limit of $X_0 = 0.76$. Although models with larger X_0 could be added, they would substantially exceed the primordial mass fraction from Big Bang nucleosynthesis (Makki et al. 2019).

On the other hand, $Q_{b,1}$ appears truncated at the upper grid limit of 0.6 MeV nucleon^{-1}. This value is larger than

[8] using PYTHON code adapted from https://dfm.io/posts/autocorr/

Table 3. Maximum likelihood estimates for each 1D marginalised posterior. In addition to the twelve parameters explored directly by the MCMC routine (§ 2.8), we include the derived neutron star properties, R and z (§ 3.4), and the system properties predicted using a disc model of anisotropy, i, ξ_b, ξ_p, and d (§ 3.3). The subscripts of 1–3 correspond to the 1998, 2000, and 2007 epochs, respectively. The accretion rates are given as fractions of fixed reference points, $\dot{m}_{Edd} = 8.775 \times 10^4 \, \mathrm{g\,cm^{-2}\,s^{-1}}$ and $\dot{M}_{Edd} = 1.75 \times 10^{-8} \, \mathrm{M_\odot\,yr^{-1}}$.

| Parameter	Units	Estimate
$\dot{m}_1$	($\dot{m}_{Edd}$)	$0.083^{+0.013}_{-0.011}$
$\dot{m}_2$	($\dot{m}_{Edd}$)	$0.114^{+0.016}_{-0.017}$
$\dot{m}_3$	($\dot{m}_{Edd}$)	$0.132^{+0.018}_{-0.02}$
$Q_{b,1}$	(MeV nucleon^{-1})	$0.36^{+0.10}_{-0.2}$
$Q_{b,2}$	(MeV nucleon^{-1})	$0.17^{+0.10}_{-0.14}$
$Q_{b,3}$	(MeV nucleon^{-1})	$0.15^{+0.1}_{-0.11}$
X_0	(Mass fraction)	$0.74^{+0.02}_{-0.03}$
Z_{CNO}	(Mass fraction)	$0.010^{+0.005}_{-0.004}$
$d\sqrt{\xi_b}$	(kpc)	$6.5^{+0.4}_{-0.6}$
ξ_p/ξ_b	–	$1.57^{+0.15}_{-0.19}$
i	(deg)	69^{+2}_{-3}
ξ_b	–	$1.22^{+0.05}_{-0.06}$
ξ_p	–	$2.0^{+0.2}_{-0.4}$
d	(kpc)	$5.8^{+0.3}_{-0.4}$
g	($10^{14}\,\mathrm{cm\,s^{-2}}$)	$2.8^{+0.4}_{-0.6}$
M	($\mathrm{M_\odot}$)	> 1.7
R	(km)	11.3 ± 1.3
z	–	$0.39^{+0.07}_{-0.07}$
$\dot{M}_1$	($\dot{M}_{Edd}$)	$0.098^{+0.012}_{-0.014}$
$\dot{M}_2$	($\dot{M}_{Edd}$)	$0.132^{+0.016}_{-0.02}$
$\dot{M}_3$	($\dot{M}_{Edd}$)	$0.153^{+0.017}_{-0.02}$

the typical assumed heating of ≈ 0.15 MeV nucleon^{-1} (e.g., Heger et al. 2007), although the amount of crustal heating emerging into the envelope is poorly constrained, and a total of 1–2 MeV nucleon^{-1} is potentially available (Haensel & Zdunik 2008). The CNO metallicity, Z_{CNO}, is also slightly limited by the lower grid boundary of $Z_{CNO} = 0.0025$. A future study could extend the model grid in these parameters, and examine the effect on the posteriors.

3.1 Predicted observables

The distribution of burst properties predicted over the MCMC simulation corresponds to the posterior predictive distribution. This distribution represents the expected observations according to our model, given the posteriors of the model parameters. A random sample of 20 000 points were selected from the chain, and the predicted triplets of observables were extracted for each point using the multi-epoch model (§ 2.8). The distribution peaks and 68 per cent credible intervals are plotted with the original observed data in Fig. 4. The predicted burst properties are consistent with

Figure 3. Posterior distributions for eight of the twelve MCMC parameters (the remaining parameters are shown in Fig. 6 and 7). The 2D contour levels indicate the 38, 68, 87, and 95 per cent credible regions. Along the diagonal are the 1D marginalised posteriors, with the 68 per cent credible interval shaded. $\dot{m}_i$ is given as a fraction of the fixed reference point of $\dot{m}_{\rm Edd} = 8.775 \times 10^4 \, {\rm g\,cm^{-2}\,s^{-1}}$, and the units for $Q_{{\rm b},i}$ are MeV nucleon^{-1}. The maximum likelihood estimates for the 1D posteriors are summarised in Table 3.

the observed data, within uncertainties. Note that this comparison should not be confused with the 'best-fit', which MCMC methods are ill-suited to finding.

3.2 Crustal heating and accretion rate

By using independent crustal heating rates between epochs, we can examine the constraints on $Q_{\rm b}$ as a function of $\dot{m}$. Theoretical models predict that the effective crustal heat-

ing is stronger at low accretion rates, and weaker at higher accretion rates due to neutrino losses (Cumming et al. 2006).

The posteriors of $Q_{\rm b}$ and $\dot{m}$ for each epoch are plotted in Fig. 5. There is significant overlap between the distributions, particularly between the 2000 and 2007 epochs, which have similar estimates for $Q_{\rm b}$. The 1998 epoch, with the lowest inferred $\dot{m}$, covers similar $Q_{\rm b}$ but is overall consistent with larger values, with a 68 per cent 2D credible region extending
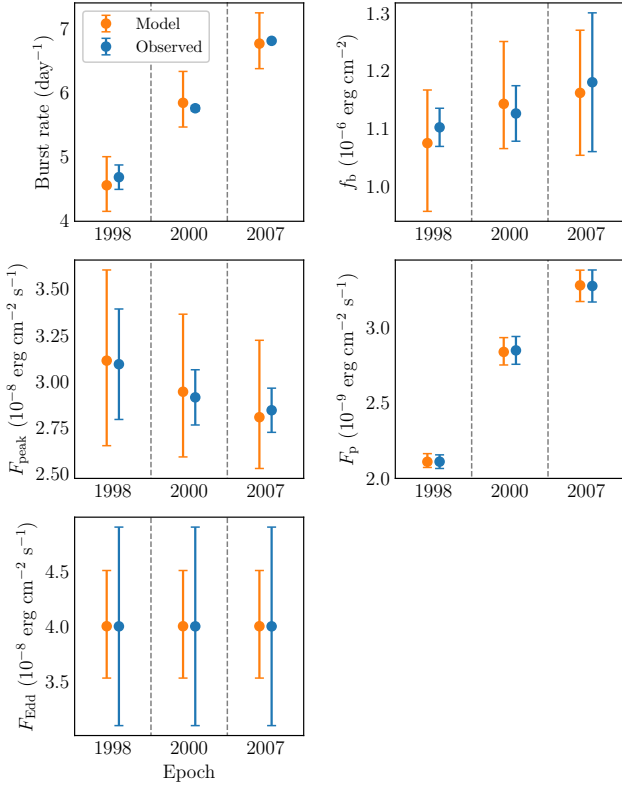
Figure 4. The distributions of the model-predicted observables from the MCMC chain (i.e., the posterior predictive distribution; orange points), plotted against the observed data (blue points). The central points are the distribution peaks, and the error bars are 68 per cent intervals. Each sample of the MCMC chain predicts these observables using the multi-epoch model (§ 2.8). The observed data are consistent with all predicted distributions, within uncertainties.

up to the grid boundary of $0.6\,\mathrm{MeV\,nucleon^{-1}}$, compared to $\approx 0.3\,\mathrm{MeV\,nucleon^{-1}}$ for 2000 and 2007.

This comparison, though inconclusive, suggests an anticorrelation between Q_{b} and $\dot{m}$, as expected, but further investigation is needed. Modelling burst epochs which span a larger range of accretion rates could help to constrain this relationship.

3.3 Distance and inclination

From $\xi_{\mathrm{p}}/\xi_{\mathrm{b}}$ we can derive constraints on the system inclination by adopting a disc model for anisotropy. Disc models have been presented by He & Keek (2016), which predicted the anisotropy according to the system inclination for multiple disc morphologies. We used their model of a thin, flat disc (*Disc a*) to predict the inclination, i, using $\xi_{\mathrm{p}}/\xi_{\mathrm{b}}$. The disc model also predicts ξ_{p} and ξ_{b}, from which we could obtain the absolute distance, d. The posteriors for these quantities are plotted in Fig. 6, and the maximum likelihood estimates are listed in Table 3.

These estimates depend on the assumptions of the thin disc model, and only flat priors were used for $d\sqrt{\xi_{\mathrm{b}}}$ and $\xi_{\mathrm{p}}/\xi_{\mathrm{b}}$. Exploring other priors, and other disc models, could yield different constraints.

Figure 5. 2D posteriors for the crustal heating rate and accretion rate, for each epoch of the model. For clarity, only the 68 and 95 per cent contours are shown.

Figure 6. (a) The distance and anisotropy parameters from the MCMC simulation, and (b) the system inclination and absolute distance, predicted using the model *Disc a* from He & Keek (2016).

3.4 Neutron star properties

We extract distributions for the neutron star properties using the MCMC parameters of M and g. The neutron star radius is calculated by solving

$$g = \frac{GM}{R^2\sqrt{1 - 2GM/(c^2 R)}} \tag{11}$$

for R, given M and g.

The gravitational redshift, z, is then calculated from M and R using Eq. (1). The posteriors for these quantities are plotted in Fig. 7, and the maximum likelihood estimates are listed in Table 3.

The highest probability density for M is against the upper boundary of $M = 2.2\,\mathrm{M_\odot}$, indicating that the distribution is truncated. This upper limit was informed by the

Figure 7. Posteriors for the gravitational parameters. The contour levels and shaded credible intervals are the same as Fig. 3. The mass and surface gravity are parameters of the MCMC model, whereas the radius and redshift are derived from these quantities (§ 3.4). The sharp boundaries which are visible in multiple contours correspond to the upper limit of M.

largest known neutron star mass (Linares et al. 2018), suggesting a possible bias in our model towards large masses. Additionally, the distribution for g is constrained by both the upper and lower model grid boundaries.

We note that only flat prior distributions were used for M and g, and thus did not include any expectations from theoretical EOS predictions, or from mass estimates of similar bursting systems (e.g., Özel et al. 2012). Exploring additional prior distributions, and expanding the model grid in g, is required before drawing stronger conclusions.

Despite the limitations, these results represent a step towards constraining the neutron star mass and radius using 1D burst models.

3.5   Global accretion rate

The model grid uses the local accretion rate per unit area, $\dot{m}$. The global accretion rate, $\dot{M} = 4\pi R^2 \dot{m}$, depends on the neutron star radius, which is determined by g and the free parameter of M. Combining the posterior samples of $\dot{m}_i$ and R, we obtain estimates for $\dot{M}_i$, which are listed in Table 3. We give $\dot{M}_i$ as a fraction of the fixed Eddington rate, $\dot{M}_{\rm Edd} = 1.75 \times 10^{-8}\,{\rm M}_\odot\,{\rm yr}^{-1}$, which is the equivalent of $\dot{m}_{\rm Edd} = 8.775 \times 10^4\,{\rm g\,cm^{-2}\,s^{-1}}$ for $R = 10\,{\rm km}$. We note again that this Eddington value is simply used as common reference points for convenience, and does not represent the 'true' Eddington limit.

3.6   Lightcurve sample

The burst data used by the MCMC model is an incomplete description of the full burst lightcurve. The two quantities extracted from the lightcurves were the fluence, $f_{\rm b}$, and the peak flux, $F_{\rm peak}$. To test whether the full lightcurves of KEPLER models remain consistent with the observations, we performed an additional set of simulations.

We took a random sample of 30 points from the MCMC chain, and for each point generated three new KEPLER models using the sampled parameters $\dot{m}_i$, $Q_{\rm b,i}$, X_0, $Z_{\rm CNO}$, and g. The result was a set of 90 KEPLER simulations, representing a sample of 30 epoch triplets from the posterior distribution.

The modelled bursts were extracted using the same procedure as the original grid (§ 2.4). We calculated average burst lightcurves for each model, and converted them to observable fluxes using the corresponding samples of M and $d\sqrt{\xi_{\rm b}}$ (§ 2.7). These lightcurves are plotted with the observed lightcurves in Fig. 8.

There is good agreement between the modelled and observed lightcurves, particularly considering that the MCMC model was fitting the fluence and peak flux, and not the full lightcurves. This comparison suggests that these scalar quantities may be sufficient proxies for the overall lightcurve – at least for bursts with similar lightcurve morphologies.

Nevertheless, some lightcurve information is still lost with this method. For example, the morphology of the decay tail is not considered, which encodes further information about the rp-process, cooling of the envelope layers in't Zand et al. (2009), and possible interactions between the burst flux and the disc (Worpel et al. 2015). Fitting additional lightcurve data, or even the entire lightcurve itself (§ 4), should remain a goal for future model comparisons.

3.7   Modelling single epochs

To test the benefit of fitting multiple epochs simultaneously, we performed three additional MCMC models, each fitting the data of a single epoch. The posteriors for the single-epoch chains are shown in Fig. 9 (coloured histograms), along with the original multi-epoch posteriors (black histograms).

Compared to fitting the epochs separately, the posterior distributions were generally more constrained when all three epochs were fit simultaneously. An exception appears to be g, although all four distributions for this parameter are heavily truncated at the boundaries, potentially interfering with the results. The parameter constraints also remain overall consistent between the multi-epoch and single-epoch chains.

This comparison supports the approach that fitting multiple epochs simultaneously can help to improve the degeneracies between the system parameters (as tested by M18).

4   DISCUSSION

We constrained system parameters for GS 1826−238 by comparing multi-epoch observations to the most extensive set of 1D model predictions to date. All central values discussed here correspond to the maximum likelihood estimates
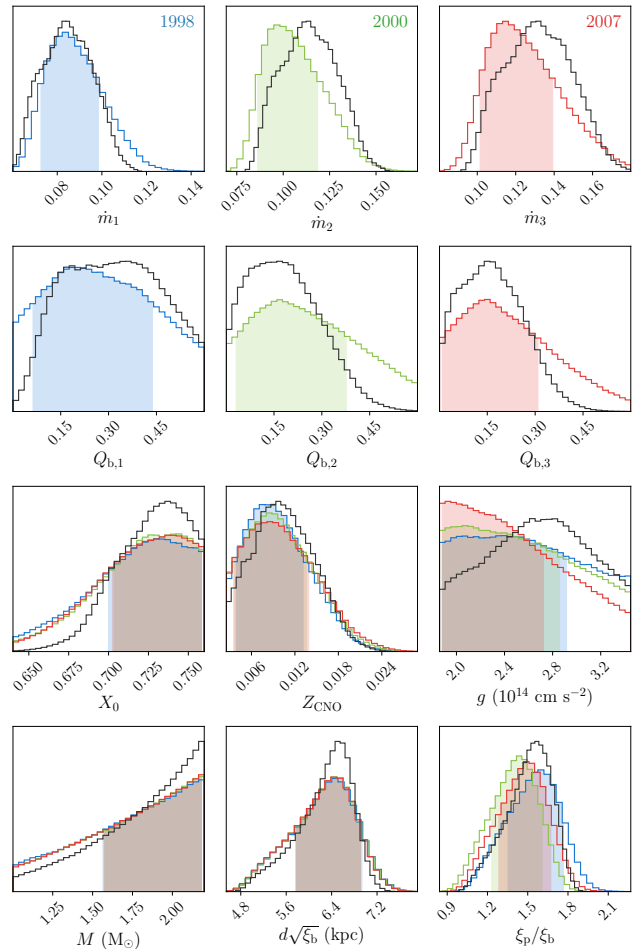
Figure 8. A comparison of the average lightcurves from 30 triplets of KEPLER simulations (blue curves, 90 models total) and the observed epoch lightcurves (black histograms). The parameters for these additional KEPLER models were drawn from 30 random samples of the MCMC chain. There is overall good agreement between the modelled and observed lightcurves, despite the MCMC procedure fitting only the fluence and peak flux from the observed lightcurve. For clarity, only the average curves are plotted, and the 1σ variations cover a wider range, as reflected by the $F_{\rm peak}$ values in Fig. 4.

Figure 9. The posteriors for three additional MCMC models, each fitting only one epoch (coloured histograms), along with the original multi-epoch model (black histogram). The shaded regions are 68 per cent credible intervals.

of the 1D marginalised posteriors, with 68 per cent credible intervals (Table 3).

The global accretion rate estimates of $\dot{M} = 0.098^{+0.012}_{-0.014}$, $0.132^{+0.016}_{-0.02}$, and $0.153^{+0.017}_{-0.02} \, \dot{M}_{\rm Edd}$, are roughly double those initially suggested by G17 of 0.0513, 0.0692, and $0.0796 \, \dot{M}_{\rm Edd}$, respectively, although those estimates did not account for anisotropy or different values of gravity. Conversely, our central values are only ≈ 10 per cent smaller than those reported by M18 of $\dot{m} = 0.11, 0.15$, and $0.17 \, \dot{m}_{\rm Edd}$, and are consistent within 2σ. Planned comparisons of MESA and KEPLER models could test whether their predictions are now more consistent, given the improvements to KEPLER described in § 2.1.

The crustal heating estimates for the 2000 and 2007 epochs of $0.17^{+0.10}_{-0.14}$ and $0.15^{+0.1}_{-0.11}$ MeV nucleon^{-1} are centred near the canonical value of $Q_{\rm b} = 0.15$ MeV nucleon^{-1}, though with broad credible intervals. On the other hand, the estimate for the 1998 epoch, with a lower accretion rate (§ 3.2), is roughly double, at $0.36^{+0.10}_{-0.2}$ MeV nucleon^{-1}, although 0.15 MeV nucleon^{-1} still lies within 1σ. In agreement with M18, crustal heating above $Q_{\rm b} \approx 0.5$ MeV nucleon^{-1} is disfavoured in all epochs, although the upper limit of our model grid is $Q_{\rm b} = 0.6$ MeV nucleon^{-1}, compared to the 1.0 MeV nucleon^{-1} considered by M18.

The CNO metallicity of $Z_{\rm CNO} = 0.01^{+0.005}_{-0.004}$ is centred

on the assumed solar value of 0.01. The chosen prior distribution was also centred near 0.01 (§ 2.9). The result broadly supports a solar metallicity, whereas previous studies typically used higher values of $Z_{\rm CNO} = 0.02$ (e.g., M18; H07). Values below ≈ 0.005 are disfavoured, such as the low-metallicity of $Z_{\rm CNO} = 0.001$ suggested by G04. Our model grid, however, only extends down to $Z_{\rm CNO} = 0.0025$, and could be expanded in future studies.

The accreted hydrogen fraction of $X_0 = 0.74^{+0.02}_{-0.03}$ is larger than the commonly-assumed value of $X_0 = 0.7$, which lies slightly outside 1σ, but still within 2σ of our estimate. The 1σ credible interval extends up to $X_0 \approx 0.76$, at odds with M18, who reported poor model fits for $X_0 = 0.75$.

Studies which do not account for burst anisotropy in their distance estimates are implicitly reporting $d\sqrt{\xi_{\rm b}}$. Our value of $d\sqrt{\xi_{\rm b}} = 6.5^{+0.4}_{-0.6}$ kpc is consistent with the previous estimates of 6 kpc (M18), 6.1 kpc (G17), and (6.07 ± 0.18) kpc (H07). Our distance is larger than the estimate of (5.7 ± 0.2) kpc from Chenevez et al. (2016), which was obtained using the same 2010 $F_{\rm Edd}$, but assumed a fixed mass of $M = 1.4 \, M_\odot$ and a radius of $R = 10$ km. Our distance is also consistent with earlier upper limits of 8 kpc

(in 't Zand et al. 1999) and 7.5 ± 0.5 kpc (Kong et al. 2000), but is larger than the more recent upper limit of 4.0–5.5 kpc (Zamfir et al. 2012).

The anisotropy ratio of $\xi_p/\xi_b = 1.57^{+0.15}_{-0.19}$ agrees with the original estimate of $\xi_p/\xi_b = 1.55$ from H07, but not with the value of 3.5 from M18, although both of these studies explored fewer model parameters. Using a flat disc model from He & Keek (2016), we obtained from ξ_p/ξ_b a system inclination of $i = 69^{+2°}_{-3°}$. This inclination is consistent with the upper limit of $i \lesssim 70°$ from low-amplitude optical modulations (Homer et al. 1998), and the range of 40–70° suggested by Mescheryakov et al. (2004) from models of the disc size.

The gravitational redshift of $z = 0.39 \pm 0.07$ is larger than the commonly-assumed value of $z = 0.26$ from $M = 1.4\,M_\odot$ and $R = 11.2$ km, and agrees with the value of 0.42 from M18, but is outside the inferred range of $z = 0.19$–0.28 for GS 1826−238 reported by Zamfir et al. (2012). The radius of $R = 11.3 \pm 1.3$ km is consistent with the rough upper limit of 9.0–13.2 km suggested by Zamfir et al. (2012).

This study is the first application of MCMC methods to 1D burst models featuring adaptive nuclear networks. As such, simplifying assumptions have been made and some care should be taken when interpreting the results.

It should be emphasized that our posterior statistics are fully dependent on the assumptions contained in the KEPLER models and our interpolation between their predictions. Although it is currently among the most advanced codes for simulating X-ray bursts, KEPLER is still the subject of ongoing refinements (§ 2.1), and is inherently limited to spherical symmetry. Additionally, using linear interpolation to sample between the models (§ 2.5) may introduce artificial 'kinks' at the grid points, potentially affecting the resulting distributions. Our comparisons of the posterior predictive distribution to the observed data (§ 3.1), and a sample of full lightcurves (§ 3.6), however, suggest that the interpolated models are not behaving unexpectedly.

The posterior distributions were truncated by some of the model grid boundaries, notably X_0, Z_{CNO}, Q_b, and g (Fig. 3). Some of these boundaries are physically-motivated, for example $X_0 < 0.76$ and $Q_b > 0$ MeV nucleon^{-1}, whereas others could realistically be extended, for example below $Z_{CNO} = 0.0025$ and above $Q_b = 0.6$ MeV nucleon^{-1}. Large extensions of the five-dimensional model grid, however, are limited by computational costs.

The information we included in our prior distributions (§ 2.9) was relatively limited. All parameters except for Z_{CNO} used flat priors. Aside from providing no additional constraints, flat priors could give undue weight to physically unrealistic regions of parameter space. For example, all combinations of M and g – and by extension, the corresponding R and z – were considered equally likely under the prior assumptions. This may have contributed to the possible bias towards large M (§ 3.4).

The limitations discussed above can be investigated and improved upon in future work. Linear interpolation, while computationally fast, has limited accuracy. Other interpolation methods, such as cubic splines, could be explored, but care should be taken to avoid introducing artefacts. A parameter sensitivity study could also identify which grid parameters can afford fewer model points, reducing the total number of simulations required.

The observed values of F_{peak} and f_b were taken from

observed lightcurves. The full burst lightcurves, however, encode additional information about the rates of heating and cooling, and the extent of rp-process burning in the tail. Fitting whole lightcurves could improve, or even significantly reshape, the posteriors. Implementing this approach, however, poses certain challenges. Whereas interpolating scalar quantities is straightforward, it is unclear how best to do so for lightcurves. If the lightcurves significantly change in morphology, interpolation could introduce nonphysical features. A possible alternative is to use machine learning to efficiently predict lightcurves between models, such as the methods recently applied to gravitational waveforms of neutron star mergers (Easter et al. 2019). Extra parametrizations of the lightcurve could also be used, by fitting curves to the burst tail (in't Zand et al. 2009). Nevertheless, our test of a limited sample of full lightcurves (§ 3.6) suggests that f_b and F_{peak} may still serve as reasonable representations of the lightcurve.

A key benefit of MCMC methods is their ability to efficiently handle large numbers of parameters. Additional parameters not used in this work could also be explored. For example, using epoch-dependent anisotropy ratios, ξ_p/ξ_b, could test for possible changes in the accretion disc properties between epochs. When calculating the Eddington flux, F_{Edd}, we assumed that the hydrogen fraction was equal to the accreted fraction, X_0, but expansion of the outer layers during PRE may expose deeper hydrogen-poor layers, increasing F_{Edd}. This hypothesis could be tested by including the hydrogen composition for F_{Edd} as an additional parameter.

5 CONCLUSION

We carried out Markov chain Monte Carlo (MCMC) simulations to model multi-epoch X-ray bursts from GS 1826−238. By precomputing a grid of 3840 KEPLER models, we interpolated the predicted burst properties and efficiently sampled the model parameter space. Applying the Bayesian framework of MCMC allowed us to systematically examine the relationships between the model parameters and the predicted burst properties. We obtained probability distributions for the properties of GS 1826−238, including the accretion rates, crustal heating rates, accreted composition, and surface gravity.

This work represents the most comprehensive use of 1D models on a burst source to date. We have explored model parameters which are often held constant in burst models, including the crustal heating, accreted hydrogen composition, surface gravity, the neutron star mass and radius, and the gravitational redshift. By using epoch-dependent parameters of Q_b, we have also tested the dependence of crustal heating on accretion rate (§ 3.2), suggesting a preference for stronger crustal heating at lower accretion rates.

Although we have focused on GS 1826−238, the methods presented here are applicable to other X-ray burst observations. Once the model grids are established, they can also be reused for similar systems. By incorporating new epoch data and expanding the grid parameters, we can analyse additional sources suggested by G17, such as the helium-burster, 4U 1820−30. Preliminary work is already underway

to model PRE bursts from 4U 1820−30with a new model grid, which we plan to present in a future publication.

This work demonstrates the largely uncharted potential of using 1D burst models for the parameter estimation of neutron star systems.

SOFTWARE PACKAGES USED

pyburst, emcee, matplotlib, scipy, astropy, chainconsumer

REFERENCES

Cavecchi Y., Levin Y., Watts A. L., Braithwaite J., 2016, Monthly Notices of the Royal Astronomical Society, 459, 1259
Chenevez J., et al., 2016, The Astrophysical Journal, 818, 135
Cumming A., 2003, The Astrophysical Journal, 595, 1077
Cumming A., Bildsten L., 2000, The Astrophysical Journal, 544, 453
Cumming A., Macbeth J., in 't Zand J. J. M., Page D., 2006, The Astrophysical Journal, 646, 429
Cyburt R. H., et al., 2010, The Astrophysical Journal Supplement Series, 189, 240
Easter P. J., Lasky P. D., Casey A. R., Rezzolla L., Takami K., 2019, Physical Review D, 100, 043005
Fisker J. L., Schatz H., Thielemann F.-K., 2008, The Astrophysical Journal Supplement Series, 174, 261
Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, Publications of the Astronomical Society of the Pacific, 125, 306
Fujimoto M. Y., 1988, The Astrophysical Journal, 324, 995
Fujimoto M. Y., Sztajno M., Lewin W. H. G., van Paradijs J., 1987, The Astrophysical Journal, 319, 902
Galloway D. K., Cumming A., 2006, The Astrophysical Journal, 652, 559
Galloway D. K., Keek L., 2017, arXiv:1712.06227 [astro-ph]

Galloway D. K., Cumming A., Kuulkers E., Bildsten L., Chakrabarty D., Rothschild R. E., 2004, The Astrophysical Journal, 601, 466
Galloway D. K., Muno M. P., Hartman J. M., Psaltis D., Chakrabarty D., 2008, The Astrophysical Journal Supplement Series, 179, 360
Galloway D. K., Goodwin A. J., Keek L., 2017, Publications of the Astronomical Society of Australia, 34, e019
Goodwin A. J., Galloway D. K., Heger A., Cumming A., Johnston Z., 2019, arXiv:1907.00996 [astro-ph]
Haensel P., Zdunik J. L., 2008, Astronomy and Astrophysics, 480, 459
He C.-C., Keek L., 2016, The Astrophysical Journal, 819, 47
Heger A., Cumming A., Galloway D. K., Woosley S. E., 2007, The Astrophysical Journal Letters, 671, L141
Homer L., Charles P. A., O'Donoghue D., 1998, Monthly Notices of the Royal Astronomical Society, 298, 497
Johnston Z., Heger A., Galloway D. K., 2018, Monthly Notices of the Royal Astronomical Society, 477, 2112
José J., Moreno F., Parikh A., Iliadis C., 2010, The Astrophysical Journal Supplement Series, 189, 204
Joss P. C., 1978, The Astrophysical Journal Letters, 225, L123
Keek L., Heger A., 2011, The Astrophysical Journal, 743, 189
Keek L., Heger A., 2017, The Astrophysical Journal, 842, 113
Keek L., Heger A., in't Zand J. J. M., 2012, The Astrophysical Journal, 752, 150
Kong A. K. H., Homer L., Kuulkers E., Charles P. A., Smale A. P., 2000, Monthly Notices of the Royal Astronomical Society, 311, 405
Lampe N., Heger A., Galloway D. K., 2016, The Astrophysical Journal, 819, 46
Lewin W. H. G., Paradijs J. V., Taam R. E., 1993, Space Science Reviews, 62, 223
Linares M., Shahbaz T., Casares J., 2018, The Astrophysical Journal, 859, 54
Lodders K., Palme H., Gail H.-P., 2009, in Martienssen W., Trümper J., eds, , Vol. 4B, Solar System. Springer Berlin Heidelberg, Berlin, Heidelberg, pp 712–770, doi:10.1007/978-3-540-88055-4_34, http://materials.springer.com/lb/docs/sm_lbs_978-3-540-88055-4_34
MacKay D. J. C., 2003, Information Theory, Inference, and Learning Algorithms. Cambridge University Press
Makino F., 1988, International Astronomical Union Circular, 4653, 2
Makki T. R., El Eid M. F., Mathews G. J., 2019, Modern Physics Letters A, 34, 1950194
Meisel Z., 2018, The Astrophysical Journal, 860, 147
Menon A., Heger A., 2017, Monthly Notices of the Royal Astronomical Society, 469, 4649
Mescheryakov A. V., et al., 2004, Astronomy Letters, 30, 751
Özel F., Psaltis D., Narayan R., Santos Villarreal A., 2012, The Astrophysical Journal, 757, 55
Rybizki J., Demleitner M., Fouesneau M., Bailer-Jones C., Rix H.-W., Andrae R., 2018, Publications of the Astronomical Society of the Pacific, 130, 074101
Schatz H., et al., 1998, Physics Reports, 294
Strohmayer T., Bildsten L., 2006, in Lewin W., van der Klis M., eds, , Vol. 39, Compact Stellar X-ray Sources. Cambridge University Press, pp 113–156
Taam R. E., Picklum R. E., 1979, The Astrophysical Journal, 233, 327
Ubertini P., et al., 1997, International Astronomical Union Circular, 6611, 1
Ubertini P., Bazzano A., Cocchi M., Natalucci L., Heise J., Muller J. M., in 't Zand J. J. M., 1999, The Astrophysical Journal Letters, 514, L27
Weaver T. A., Zimmerman G. B., Woosley S. E., 1978, The Astrophysical Journal, 225, 1021

Woosley S. E., Heger A., Weaver T. A., 2002, Reviews of Modern Physics, 74, 1015

Woosley S. E., et al., 2004, The Astrophysical Journal Supplement Series, 151, 75

Worpel H., Galloway D. K., Price D. J., 2015, The Astrophysical Journal, 801, 60

Zamfir M., Cumming A., Galloway D. K., 2012, The Astrophysical Journal, 749, 69

Zingale M., Malone C. M., Nonaka A., Almgren A. S., Bell J. B., 2015, The Astrophysical Journal, 807, 60

in 't Zand J. J. M., Heise J., Kuulkers E., Bazzano A., Cocchi M., Ubertini P., 1999, Astronomy and Astrophysics, 347, 891

in't Zand J. J. M., Keek L., Cumming A., Heger A., Homan J., Méndez M., 2009, Astronomy and Astrophysics, 497, 469

APPENDIX A: MODEL PRE-HEATING

Excessive burn-in can occur during simulations if nuclear heating, Q_{nuc}, is neglected during the model setup phase (§ 2.1). Addressing the issue is not straightforward, however, because nuclear heating occurs throughout the envelope at difference rates, depending on the local conditions. By contrast, the flux from crustal heating, Q_{b}, is simply implemented as a lower boundary condition. Predicting Q_{nuc} in advance is difficult prior to running the full simulation with a nuclear network.

We added a heat source during the setup of the envelope before the full burst simulation begins. For a chosen Q_{nuc}, the total heat flux is given by $F_{\mathrm{nuc}} = Q_{\mathrm{nuc}}\dot{m}$, which we distributed throughout the envelope as a Gaussian function, centred at a column depth of $y = 8 \times 10^8 \,\mathrm{g\,cm}^{-2}$ with a spread of $\sigma = 8 \times 10^7 \,\mathrm{g\,cm}^{-2}$. For comparison, Q_{b} is implemented at the lower model boundary of $y \sim 10^{12} \,\mathrm{g\,cm}^{-2}$.

We tested this model setup with a heating of $Q_{\mathrm{nuc}} = 5\,\mathrm{MeV\,nucleon}^{-1}$, approximately the energy yield for hydrogen burning. We used model parameters of $X_0 = 0.73$, $Z_{\mathrm{CNO}} = 0.005$, $\dot{m} = 0.2$, and $Q_{\mathrm{b}} = 0.05\,\mathrm{MeV\,nucleon}^{-1}$. The burn-in was largely eliminated from the resulting burst simulation (Fig. A1), in contrast to an identical model without preheating (effectively, $Q_{\mathrm{nuc}} = 0\,\mathrm{MeV\,nucleon}^{-1}$). Previous studies typically discarded only the first ≈ 3 bursts to account for model burn-in (e.g., Woosley et al. 2004). We demonstrate, however, that a 10–20 per cent discrepancy persists between the recurrence times even after 50 bursts. Nevertheless, further investigation is required into the sensitivity of models to preheating. Other bursting regimes, for example helium bursts, may require additional testing of the heating rates and depths.

APPENDIX B: GR CORRECTIONS

KEPLER uses Newtonian gravity to calculate the gravitational acceleration, given by

$$g = \frac{GM}{R^2}. \tag{B1}$$

In the GR regime of a neutron star, however, the surface gravity is given by

$$g = \frac{GM(1+z)}{R^2}, \tag{B2}$$

where z is the gravitational redshift given in Eq. (1).

Figure A1. The evolution of recurrence time for models with and without nuclear preheating. Previous burst models did not include nuclear heating when initialising the envelope ($Q_{\mathrm{nuc}} = 0\,\mathrm{MeV\,nucleon}^{-1}$, blue points). When nuclear heating is included, the systematic burn-in is effectively eliminated ($Q_{\mathrm{nuc}} = 5\,\mathrm{MeV\,nucleon}^{-1}$, orange points).

Because the envelope is a thin shell ($\Delta R \ll R$), the surface gravity is approximately constant throughout the envelope. The Newtonian KEPLER model is equivalent to neutron stars under GR with different M and R, but with the same g. There is a contour of M and R pairs which satisfy this constraint. For a chosen M and R, the Newtonian KEPLER quantities can be corrected to the equivalent GR values. A more detailed description of these corrections can be found in appendix B of Keek & Heger (2011).

Keek & Heger (2011) defined the mass and radius ratios between the two regimes,

$$\varphi = \frac{M_{\mathrm{g}}}{M_{\mathrm{k}}}, \quad \xi = \frac{R_{\mathrm{g}}}{R_{\mathrm{k}}}, \tag{B3}$$

where we here signify the Newtonian and GR quantities with the subscripts 'k', and 'g', respectively. Setting the requirement that g must be equal under the two regimes, the above ratios are related by

$$\xi^2 = \varphi(1+z), \tag{B4}$$

where z is evaluated for M_{g} and R_{g}.

The ratio of the neutron star surface areas is given by ξ^2, and so the GR-corrected luminosity is given by

$$L_{\mathrm{g}} = \xi^2 L_{\mathrm{k}} = \varphi(1+z)L_{\mathrm{k}}. \tag{B5}$$

For a given accretion rate per unit area, $\dot{m}$, the global accretion rate, $\dot{M} = 4\pi R^2 \dot{m}$, is also scaled by the area ratio,

$$\dot{M}_{\mathrm{g}} = \xi^2 \dot{M}_{\mathrm{k}} = \varphi(1+z)\dot{M}_{\mathrm{k}}. \tag{B6}$$

Both regimes are in the same local reference frame, and so Δt and ν are not time-dilated.

These GR-corrected quantities were used to calculate the predicted observables in § 2.6 and 2.7.

APPENDIX C: MODEL DATA

The data used in this work are publicly available at Mendeley Data[9]. Included is a data table of the analysed model grid, listing the input parameters and summary burst properties of each model as Newtonian KEPLER quantities (i.e., *not* corrected for GR). Additionally, the full MCMC chains are included as 3D arrays, containing 1000 walkers × 12 parameters × 20 000 steps (including the initial 1000 burn-in steps which were discarded from our analysis). Further information on how to load and use this data is provided in the data repository.

The software tools used to extract the model burst properties, analyse the model grid, and manage the MCMC models, have been collected under a PYTHON package called PYBURST, which can be downloaded from https://github.com/zacjohnston/pyburst/.

This paper has been typeset from a TEX/LATEX file prepared by the author.

Chapter 6

Multi-epoch MCMC Models of a Helium Burster

In Chapter 5 we applied Markov Chain Monte Carlo (MCMC) methods to one-dimensional (1D) burst models, to match multi-epoch observations of the "Clocked Burster", GS 1826−238. We demonstrated the potential for using precomputed model grids to efficiently obtain probability distributions over system parameters. We present here a provisional extension of these methods to 4U 1820−30, a helium-accreting system which exhibits photospheric radius-expansion (PRE) bursts. Modelling PRE bursts poses additional challenges to mixed hydrogen/helium (H/He) bursts, and for the scope of this project, the grid parameters were kept limited in comparison to our study of GS 1826−238. We therefore note that the resulting parameter estimates and model predictions should be considered provisional, and further work is needed to address the limitations discussed below.

6.1 4U 1820−30

4U 1820−30 is a low-mass X-ray binary (LMXB) in the globular cluster NGC 6624, and was the source in which X-ray bursts were discovered (Grindlay et al., 1976). The companion star is a white dwarf and has one of the shortest known orbital periods of $11\,\mathrm{min}$ (King and Watson, 1986; Stella, Priedhorsky, and White, 1987). This places the system in the class of ultracompact binaries with periods of $\lesssim 1\,\mathrm{h}$. The compact orbit indicates that the material accreted onto the neutron star is hydrogen-poor, with previous estimates for a hydrogen mass fraction of $X_0 \lesssim 0.1$ (Cumming, 2003), and $X_0 = 0.0$ is often assumed (e.g., Suleimanov et al., 2017). The accreted fuel triggers helium bursts, in contrast to the mixed H/He bursts of GS 1826−238 we modelled in Chapter 5. Helium bursts are frequently characterised by photospheric radius-expansion (PRE), which is thought to occur when the burst luminosity reaches the local Eddington limit, L_{Edd} (e.g., Kuulkers et al., 2003).

Because 4U 1820−30 resides in a globular cluster, we enjoy the benefit of independent distance measurements. With optical observations, Kuulkers et al. (2003) obtained a distance to NGC 6624 of $d = (7.6 \pm 0.4)\,\mathrm{kpc}$, and Valenti, Ferraro, and Origlia (2007) obtained $(8.4 \pm 0.6)\,\mathrm{kpc}$ using near-infrared measurements.

A superburst was observed from 4U 1820−30 in September 1999 (Strohmayer and Brown, 2002). These rare, energetic ($\sim 10^{42}\,\mathrm{erg}$) bursts are thought to result from the ignition of a deep carbon ocean (Woosley and Taam, 1976). The occurrence of a superburst in 4U 1820−30 suggests

Table 6.1: The multi-epoch burst data from 4U 1820$-$30, used by our MCMC routine. The values are adapted from Table 2 of G17. We have assumed F_{Edd} corresponds to the observed bolometric peak flux, F_{peak}. The F_{p} values include the bolometric corrections from G17. Because only two bursts were observed for 2009, we have assumed an instrument timing uncertainty of $1\,\mathrm{s}$, although the burst-to-burst variation is likely larger ($\approx 25\,\mathrm{s}$ for 1997).

| Epoch | ν | F_{Edd} | F_{p} |
	(day^{-1})	$(10^{-9}\,\mathrm{erg\,s^{-1}\,cm^{-2}})$	$(10^{-9}\,\mathrm{erg\,s^{-1}\,cm^{-2}})$
1997 May	8.95 ± 0.02	61 ± 2	5.4 ± 0.7
2009 June	12.6850 ± 0.0019	56.6 ± 1.4	8.54 ± 0.09

that carbon is steadily accumulated during the nuclear processing of the accreted fuel (Cumming and Bildsten, 2001).

Alongside GS 1826$-$238, a multi-epoch dataset for 4U 1820$-$30 was included in Galloway, Goodwin, and Keek (2017, hereafter G17) as a target for PRE burst modelling. The system is thus a natural choice for extending the methods from Chapter 5.

6.2 Methods

The methods for this study consist of the multi-epoch observed data (Section 6.2.1), the construction of the model grid (Section 6.2.2), the interpolated multi-epoch model (Section 6.2.3), and the MCMC methods (Section 6.2.4).

For ease of comparison between different models and studies, all accretion rates are given as a fraction of the "canonical" Eddington-limited rate for solar composition, $\dot{m}_{\mathrm{Edd}} = 8.775 \times 10^4\,\mathrm{g\,cm^{-2}\,s^{-1}}$. This value assumes $X_0 = 0.7$ and Newtonian gravity for $M = 1.4\,\mathrm{M_\odot}$ and $R = 10\,\mathrm{km}$. The "true" Eddington rate is a factor of 1.7 larger for pure helium, and depends on the neutron star mass and radius.

6.2.1 Multi-epoch Data

To model bursts from 4U 1820$-$30, we again used multi-epoch data from the reference set provided by G17. The burst data was from two accretion epochs, observed on 4 May 1997 and 12 June 2009 (Table 6.1). In this study we fit the observed burst rate, ν, Eddington-limited flux, F_{Edd}, and persistent flux, F_{p}. We assumed that F_{Edd} corresponds to the observed peak burst flux, F_{peak}, and that F_{p} corresponds to the accretion luminosity, L_{acc}.

For this initial study, we did not fit the observed fluence, f_{b}, because the behaviour of the neutron star atmosphere near the Eddington luminosity remains poorly understood. It is unclear how the total burst energy, E_{b}, translates to an observed f_{b}, particularly because KEPLER lacks a detailed treatment of the photosphere, and exhibits anomalous super-Eddington luminosities during PRE (Woosley et al., 2004; Johnston, Heger, and Galloway, 2018). Because of this model limitation, we only used KEPLER models to predict ν, and calculated F_{Edd} and F_{p} analytically (Section 6.2.3).

Table 6.2: The parameters of the KEPLER model grid. Every combination was iterated, totalling 168 simulations. The Q_b step size of 0.025 corresponds to the span between 0.025–0.15, and the step size of 0.05 corresponds to the span between 0.20–0.40. Parameters which were held constant are $X_0 = 0.0$, $Z_{CNO} = 0.015$, and $g = 1.858 \times 10^{14}$ cm s^{-2}. The accretion rates are given as a fraction of $\dot{m}_{Edd} = 8.775 \times 10^4$ g cm^{-2} s^{-1} (assuming $M = 1.4\,M_\odot$, $R = 10$ km, and $X = 0.7$), which is simply used as a common reference point to other models, and does not represent the "true" $\dot{m}_{Edd}$.

Parameter	Units	Range	Step size	N
$\dot{m}$	$(\dot{m}_{Edd})$	0.175–0.500	0.025	14
Q_b	(MeV nuc^{-1})	0.01, 0.025–0.15, 0.2–0.40	0.025, 0.05	12
			Total	168

6.2.2 Model Grid

To obtain burst predictions that could be sampled quickly with MCMC methods, we precomputed a grid of KEPLER models. Because this study represents the first extension of our MCMC methods to a new bursting regime, we reduced the total number of models by using constant values for the accreted hydrogen fraction, X_0, the CNO metallicity, Z_{CNO}, and the surface gravity g. We set the hydrogen composition to $X_0 = 0.0$, which is commonly assumed for this source due to its ultra-compact orbit (Stella, Priedhorsky, and White, 1987). The CNO metallicity was set to $Z_{CNO} = 0.015$, between commonly-used values of 0.01 and 0.02 (e.g., Cumming, 2003; Heger et al., 2007; Meisel, 2018), although the influence of Z_{CNO} will be reduced due to the absence of hydrogen. The remaining mass fraction of 0.985 was assigned to helium. The surface gravity was set to $g = 1.858 \times 10^{14}$ cm s^{-2}, corresponding to a gravitational neutron star mass of $M = 1.4\,M_\odot$ and a radius of $R = 11.2$ km, although the final choice of M and R remains a free parameter (Section 6.2.3). Exploring variations of these parameters remains a goal for a future study (Section 6.4.1).

We varied two model parameters: the local accretion rate, $\dot{m}$, and the crustal heating rate, Q_b. Following a limited parameter exploration, using the observed recurrence times of $1 \lesssim \Delta t \lesssim 3$ h as a guide, we chose a regular grid of values between $0.01 \leq Q_b \leq 0.4$ MeV nucleon^{-1} and $0.175 \leq \dot{m} \leq 0.5\,\dot{m}_{Edd}$, resulting in 168 KEPLER models (Table 6.2). Despite the reduced size of the 4U 1820$-$30 model grid compared to GS 1826$-$238 (Chapter 5), it is nevertheless the largest set of 1D models of hydrogen-poor bursts to date.

In order to minimise the effect of model burn-in (see Section 3.2), a sequence of 40–50 bursts were produced for each model, and the first 30 were excluded from analysis. The remaining 10–20 bursts were extracted using the same methods described in Section 2.2 and Chapter 5, using our PYTHON package PYBURST. The average burst properties were calculated, and the standard deviation was adopted as the uncertainty. We thus obtained a tabulated set of model predictions over the grid of $\dot{m}$ and Q_b (Figure 6.1). Linear interpolation could then be used to rapidly ($\ll 1$ s) sample burst properties anywhere across the grid.

Figure 6.1: Burst properties for a subset of the full grid of 168 models of pure-helium bursts. For clarity, only every second grid point in Q_b is plotted. Each point corresponds to a KEPLER simulation of 40–50 bursts, from which the average burst properties were calculated for the final 10–20 bursts. In contrast to the model grid from Chapter 5, we fit only the burst rate to the observed data. The burst energy, E_b, is shown here simply for illustration, and was calculated after truncating the burst lightcurve at an assumed Eddington limit of $L_{Edd} = 3.5 \times 10^{38} \, \mathrm{erg \, s^{-1}}$. The error bars and interpolated shaded regions are 1σ standard deviations, but are generally too small to be visible here.

6.2.3 Multi-epoch Modelling

To predict the multi-epoch properties of the observed data (Table 6.1), we used an approach similar to the multi-epoch model for GS 1826−238 in Chapter 5. For the model grid parameters of accretion rate and crustal heating we used epoch-dependent parameters, $\dot{m}_1$, $\dot{m}_2$, $Q_{b,1}$, and $Q_{b,2}$, where the subscripts 1 and 2 correspond to the 1997 and 2009 epochs, respectively. The remaining "free" parameters were epoch-independent: the neutron star mass, M, the anisotropy-modified distance, $d\sqrt{\xi_b}$, and the anisotropy ratio, ξ_p/ξ_b. We note again that unlike our model for GS 1826−238, we used fixed values of $X_0 = 0.0$, $Z_{CNO} = 0.015$, and $g = 1.858 \times 10^{14}\,\mathrm{cm\,s^{-2}}$ (Section 6.2.2). Our multi-epoch model for 4U 1820−30 thus consisted of seven parameters: $\dot{m}_1$, $\dot{m}_2$, $Q_{b,1}$, $Q_{b,2}$, M, $d\sqrt{\xi_b}$, and ξ_p/ξ_b.

For a given choice of these parameters, the three observed quantities were predicted for both epochs: the burst rate, ν, the Eddington flux, F_{Edd}, and the persistent flux, F_p. These observables were predicted using the same procedure described in Chapter 5. The burst rate was interpolated from the model grid for the given $\dot{m}_i$ and $Q_{b,i}$, for $i = 1, 2$, and F_{Edd} and F_p were again calculated directly. These calculations included GR-corrections to account for the Newtonian gravity used in KEPLER (Section 2.3).

6.2.4 MCMC Method

We used MCMC methods to sample the parameter space of our multi-epoch model (Section 6.2.3) and compare the predictions to the observed data (Section 6.2.1). The MCMC routine was adapted directly from the model applied to GS 1826−238 in Chapter 5, but with modified parameters, priors, and observed data. We again used the open-source PYTHON ensemble sampler, EMCEE[1] (Foreman-Mackey et al., 2013).

The prior distribution was set to $p(\theta) = 0$ outside the parameter boundaries. For the parameters $\dot{m}_i$ and $Q_{b,i}$, the boundaries were those of the model grid (Table 6.2). For the free parameters, we set limits of $1.0 \leq M \leq 3.5\,\mathrm{M_\odot}$, $1 \leq d\sqrt{\xi_b} \leq 15\,\mathrm{kpc}$, and $0.1 \leq \xi_p/\xi_b \leq 10$. We used a larger upper limit than $M = 2.3\,\mathrm{M_\odot}$ used in Chapter 5, because the posterior distributions were found to be strongly truncated. Values above this limit exceed the most massive neutron stars observed to date (e.g., Linares, Shahbaz, and Casares, 2018; Cromartie et al., 2019), but for this initial study of 4U 1820−30, we wished to explore the possible extent of the model bias towards large masses.

We applied flat (i.e., uniform) prior distributions within these limits for all parameters except $d\sqrt{\xi_b}$, for which we utilised two distance estimates for the globular cluster NGC 6624. These estimates were $d = (7.6 \pm 0.4)\,\mathrm{kpc}$ from optical measurements (Kuulkers et al., 2003) and $d = (8.4 \pm 0.6)\,\mathrm{kpc}$ from near-infrared measurements (Valenti, Ferraro, and Origlia, 2007). We used the joint distribution of these constraints with a flat prior for ξ_b to obtain a Gaussian prior of $d\sqrt{\xi_b} = (7.85 \pm 0.33)\,\mathrm{kpc}$.

Given a sample point in parameter space, the observables for ν, F_{Edd}, and F_p were predicted with the multi-epoch model. The predictions were

---

[1] https://emcee.readthedocs.io/en/v2.2.1

Figure 6.2: Estimates for the autocorrelation time, τ, for each parameter at multiple points along the chain. Above $N \sim 10^4$, the estimates for τ are consistently growing slower than $\tau = N/10$, indicating that the chain has progressed beyond $10\,\tau$.

compared with the observed data using the likelihood function given by

$$\ln\left[p(D|\theta)\right] = -\frac{1}{2}\sum_x \left\{ \frac{(x-x_0)^2}{\sigma^2 + \sigma_0^2} + \ln\left[2\pi(\sigma^2 + \sigma_0^2)\right] \right\}, \qquad (6.1)$$

where x and its uncertainty, σ, were iterated over the predicted values for each epoch, and the subscript '0' signifies the corresponding observed values.

The MCMC chain consisted of 1000 walkers initialised in a small "hyper ball" in parameter space. The sampler was run for 50 thousand steps, producing 50 million samples in total. The first 1000 steps were discarded as burn-in.

To test the sampler convergence, the autocorrelation time τ was estimated for each parameter at multiple steps along the chain[2] (Figure 6.2). For large numbers of samples, $N \gtrsim 10^4$, the estimates for τ begin to converge toward a final value. Although the chain is not long enough to obtain converged values for τ itself, the estimates are consistently growing slower than the $\tau = N/10$ line, indicating that the total chain length is larger than $10\,\tau$.

6.3   Results

The MCMC chain was analysed using the open source PYTHON package CHAINCONSUMER[3] (Hinton, 2016). The marginalised one-dimensional (1D) and two-dimensional (2D) posterior distributions for all seven parameters are plotted in Figure 6.3. The maximum likelihood estimates for the 1D posteriors are listed in Table 6.3. We discuss here the general results

---

[2]using a PYTHON routine adapted from https://dfm.io/posts/autocorr
[3]https://samreay.github.io/ChainConsumer

Table 6.3: The maximum likelihood estimates for the 1D marginalised posteriors, with 68% credible intervals. The quantities of d, i, ξ_b, and ξ_p were calculated using a disc anisotropy model. The neutron star properties of R and z were calculated from M and the fixed value of $g = 1.858 \times 10^{14}\,\mathrm{cm\,s^{-2}}$. The global accretion rates, $\dot{M}_i$, were calculated from $\dot{m}_i$ and R. The accretion rates are given as fractions of the canonical Eddington rate, $\dot{m}_{\mathrm{Edd}} = 8.775 \times 10^4\,\mathrm{g\,cm^{-2}\,s^{-1}}$ and $\dot{M}_{\mathrm{Edd}} = 1.75 \times 10^{-8}\,\mathrm{M_\odot\,yr^{-1}}$.

Parameter	Units	Estimate
$\dot{m}_1$	($\dot{m}_{\mathrm{Edd}}$)	$0.208^{+0.05}_{-0.005}$
$\dot{m}_2$	($\dot{m}_{\mathrm{Edd}}$)	$0.37^{+0.06}_{-0.07}$
$Q_{\mathrm{b},1}$	($\mathrm{MeV\,nucleon^{-1}}$)	$0.36^{+0.04}_{-0.17}$
$Q_{\mathrm{b},2}$	($\mathrm{MeV\,nucleon^{-1}}$)	$0.04^{+0.11}_{-0.02}$
M	($\mathrm{M_\odot}$)	$2.3^{+0.3}_{-0.2}$
$d\sqrt{\xi_b}$	(kpc)	7.9 ± 0.3
ξ_p/ξ_b	–	$1.8^{+0.4}_{-0.3}$
d	(kpc)	6.9 ± 0.3
i	(deg)	74^{+2}_{-4}
ξ_b	–	1.33 ± 0.09
ξ_p	–	$2.4^{+0.7}_{-0.5}$
R	(km)	$15.1^{+0.9}_{-1}$
z	–	0.36 ± 0.03
$\dot{M}_1$	($\dot{M}_{\mathrm{Edd}}$)	$0.34^{+0.06}_{-0.04}$
$\dot{M}_2$	($\dot{M}_{\mathrm{Edd}}$)	$0.54^{+0.13}_{-0.08}$

Figure 6.3: Marginalised posterior distributions for all seven MCMC parameters. The 1D posteriors are along the diagonal, with 68% credible intervals shaded. The 2D contour levels are 38, 68, 87, and 95% credible regions. The units for $\dot{m}$ are $\dot{m}_{\mathrm{Edd}} = 8.775 \times 10^4 \, \mathrm{g\,cm^{-2}\,s^{-1}}$, Q_{b} are MeV nucleon^{-1}, M are $\mathrm{M}_{\odot}$, and $d\sqrt{\xi_{\mathrm{b}}}$ are kpc. The maximum likelihood estimates for the 1D posteriors are listed in Table 6.3.

of the MCMC posteriors, and the additional system properties we can derive from them. We further discuss the specific parameter estimates and comparisons to previous works in Section 6.4.

In the H/He models of GS 1826−238 in Chapter 5, we found that some of the posteriors were limited by the parameter boundaries, and the same issue is evident for the 4U 1820−30 distributions here. Most of the 38% contours of the 2D posteriors for $Q_{b,1}$ reach the model grid upper limit of $0.4 \, \mathrm{MeV \, nucleon^{-1}}$, and all of the 87% contours for $\dot{m}_2$ are truncated at the upper limit of $\dot{m} = 0.5 \dot{m}_{\mathrm{Edd}}$. On the other hand, the 95% 2D contours between the free parameters of M, $d\sqrt{\xi_b}$, and ξ_p/ξ_b lie completely within the boundaries, indicating that the distributions for these parameters are not artificially constrained.

In contrast to the GS 1826−238 models in Chapter 5, each epoch pair of $\dot{m}_i$ and $Q_{b,i}$ are tightly correlated. This difference may be partly due to the lack of hydrogen burning, reducing the burst-to-burst variability of the pure helium models. The typical standard deviation of Δt for the helium burst models was $\approx 0.9 \, \%$, in comparison to $\approx 4 \, \%$ for the H/He models. The larger variability of the H/He bursts, and the additional influence of X_0 and Z_{CNO} on burst ignition, possibly "washed out" the correlation.

The anti-correlation itself between the posteriors of $Q_{b,i}$ and $\dot{m}_i$ can be understood from their combined effect on the burst ignition depth, y_{ig}, and recurrence time, Δt. Increasing Q_b decreases the depth at which nuclear burning is unstable, y_{ig}. The recurrence time is then determined by the time taken to accrete the column of material, $\Delta t = y_{\mathrm{ig}}/\dot{m}$. For a given observed Δt used by the MCMC model, an increase in $\dot{m}$ can be compensated by a decrease in Q_b.

6.3.1 Predicted Observables

The posterior predictive distribution is given by the distribution of observables predicted by the MCMC model. As a consistency check, we can compare these predicted distributions to the original observed data to ensure the model is behaving as expected. We took a random sample of 20 000 points from the MCMC chain, and extracted the multi-epoch burst properties predicted by the model. The peaks and 68% intervals for the distributions are plotted against the observed data in Figure 6.4. The observed values are consistent with the prediction distributions within the uncertainties. This consistency indicates that the model predictions are behaving normally in the MCMC simulation.

6.3.2 Crustal heating and accretion rate

The epoch-dependant crustal heating parameters, $Q_{b,i}$, allow us to examine the relationship between Q_b and $\dot{m}$. The 2D posteriors for each epoch are plotted in Figure 6.5. As discussed above, these narrow contours arise due to the degeneracy between Q_b and $\dot{m}$ for a given Δt, and should not be confused with the underlying relationship between the two quantities.

There is some overlap between the tails of the 1D posteriors for Q_b, although the 1997 epoch favours larger values of $0.36^{+0.04}_{-0.17} \, \mathrm{MeV \, nucleon^{-1}}$, whereas 2009 is consistent with smaller values of $0.04^{+0.11}_{-0.02} \, \mathrm{MeV \, nucleon^{-1}}$. The 1D estimates are similarly separated in $\dot{m}$, with lower values for 1997
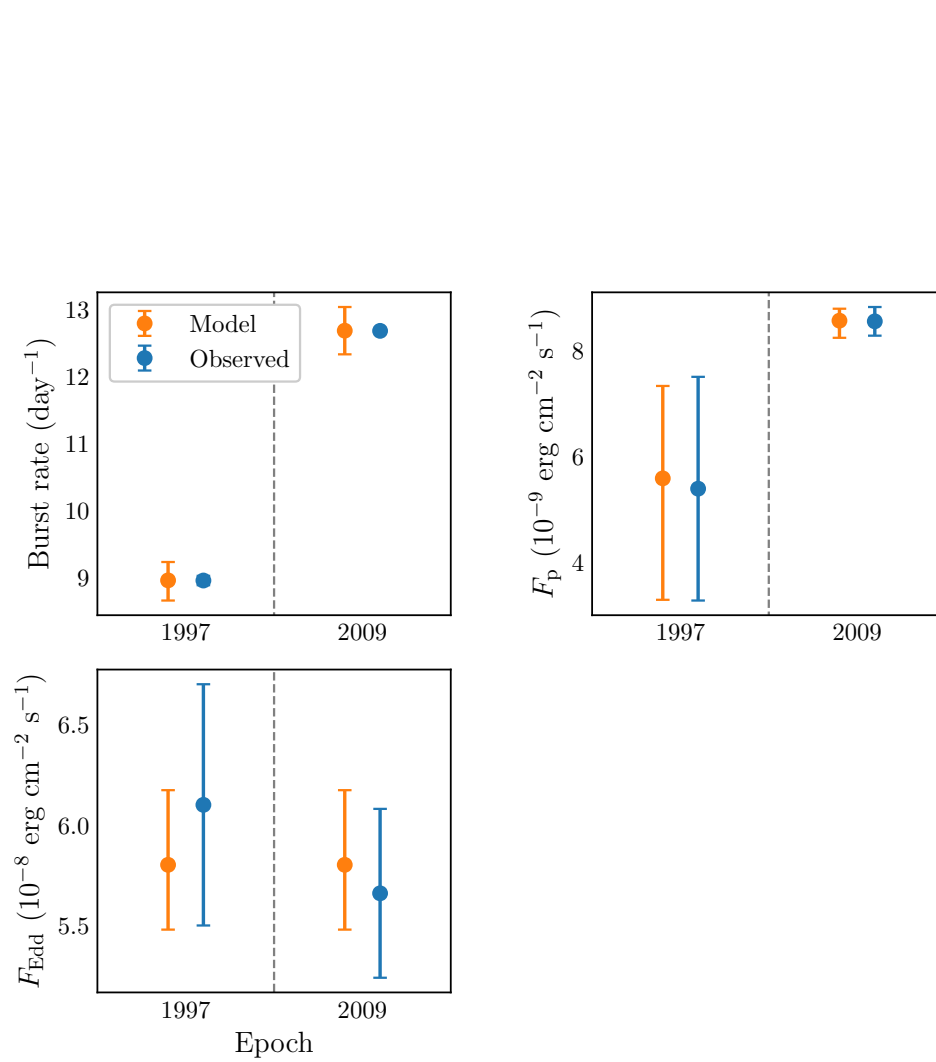
Figure 6.4: Distributions of the predicted burst properties (i.e., the posterior predictive distribution; orange points), and the observed epoch data (blue points). The error bars are 68% credible intervals. The observed data are consistent with the predicted distributions to within uncertainties.

Figure 6.5: The posteriors for Q_b and $\dot{m}$ for each epoch. For clarity, only the 68 and 95% contour levels are shown. The distribution for each epoch is strongly degenerate, but the 1997 epoch is overall consistent with lower accretion rates and lower crustal heating than 2009.

of $.208^{+0.05}_{-0.005}\,\dot{m}_{\mathrm{Edd}}$, and higher values for 2009 of $0.37^{+0.06}_{-0.07}\,\dot{m}_{\mathrm{Edd}}$. As with the GS $1826-238$ models in Chapter 5, this comparison offers a tentative indication towards an underlying relationship between Q_b and $\dot{m}$.

6.3.3 Distance and Inclination

Using the same approach as the GS $1826-238$ models in Chapter 5, we obtained posteriors for the system inclination, i, and the absolute distance, d, by choosing a model for the disc anisotropy. We used the model *Disc a* for a thin flat disc from He and Keek (2016, see Figure 2.3 from Chapter 2) to calculate these quantities from the anisotropy ratio, ξ_p/ξ_b, and the modified distance, $d\sqrt{\xi_b}$. The posteriors for these quantities are plotted in Figure 6.6, and the 1D estimates are listed in Table 6.3, in addition to individual estimates for ξ_b and ξ_p.

The burst anisotropy factor of $\xi_b = 1.33 \pm 0.09$ indicates that the burst emission is preferentially beamed *away* from the observer. This effect causes the source to appear dimmer, and at a larger inferred distance of $d\sqrt{\xi_b} = 7.9 \pm 0.3\,\mathrm{kpc}$ when isotropic emission is assumed. Given this estimate of ξ_b, the actual distance is closer, at $d = 6.9 \pm 0.3\,\mathrm{kpc}$.

We note that these estimates are dependent on the disc model from He and Keek (2016), and using other models for anisotropy could produce different estimates.

Figure 6.6: Posteriors for the distance and anisotropy parameters of the MCMC model (left panel), and the inclination and absolute distance calculated from them using the *disc a* anisotropy model of He and Keek (2016, right panel). The 2D contour levels are 38, 68, 87, and 95 % credible regions. The 1D shaded intervals are 68% credible intervals. The maximum likelihood estimates for the 1D posteriors are listed in Table 6.3.

6.3.4   Neutron Star Properties

Using the same procedure as Chapter 5, we calculate the neutron star radius, R, and the gravitational redshift, z, using the MCMC parameter of M and the fixed value of $g = 1.858 \times 10^{14} \, \mathrm{cm \, s^{-2}}$. The 1D posteriors for these quantities are plotted in Figure 6.7, and the estimates are listed in Table 6.3.

Due to the larger upper boundary of $M = 3.5 \, \mathrm{M_\odot}$, compared to $M = 2.2 \, \mathrm{M_\odot}$ in Chapter 5, the model was free to explore larger masses, resulting in a posterior of $2.3^{+0.3}_{-0.2}$. This value is larger than the typical observed range of 1–2 $\mathrm{M_\odot}$ (Özel, Gould, and Güver, 2012; Miller, 2013), and also exceeds both the largest observed mass of $2.14^{+0.10}_{-0.09} \, \mathrm{M_\odot}$ (Cromartie et al., 2019) and the maximum mass of $2.17 \, \mathrm{M_\odot}$ inferred from the neutron star merger GW170817 (Margalit and Metzger, 2017). We further discuss these results in Section 6.4.

6.3.5   Global Accretion Rate

We calculate the global accretion rates given by $\dot{M} = 4\pi R^2 \dot{m}$ for each sample of R and $\dot{m}$. The 1D marginalised estimates for $\dot{M}_\mathrm{i}$ are listed in Table 6.3, given as a fraction of the canonical Eddington-limited rate, $\dot{M}_\mathrm{Edd} = 1.75 \times 10^{-8} \, \mathrm{M_\odot \, yr^{-1}}$, which is equivalent to $\dot{m}_\mathrm{Edd} = 1.492 \times 10^5 \, \mathrm{g \, cm^{-2} \, s^{-1}}$ assuming $R = 10 \, \mathrm{km}$. Once again, this Eddington rate is simply used as a common reference point, and is not adjusted for each sample of M and R.

6.3.6   Lightcurve Sample

The MCMC routine only compared a single quantity extracted from the observed lightcurve: the peak flux, F_peak, which was assumed to be the Eddington flux, F_Edd. As with our models for GS 1826−238 in Chapter 5, we

Figure 6.7: Posterior distributions for the neutron star properties. The neutron star mass, M, is an original parameter of the MCMC model, whereas the radius, R, and redshift, z, are calculated from M and g. In contrast to the GS 1826$-$238 model from Chapter 5, we used a fixed gravity of $g = 1.858 \times 10^{14}\,\mathrm{cm\,s^{-2}}$, and thus only a single contour line of M and R was explored. The shaded region is the 68% credible interval.

produced a limited sample of full burst lightcurves to check for consistency with the observations.

We took a random sample of 30 points from the MCMC chain, and computed a new epoch pair of KEPLER models for each point of $\dot{m}_\mathrm{i}$ and $Q_\mathrm{b,i}$, resulting in a total of 60 models. The model luminosities were truncated at the corresponding Eddington limit, L_Edd, and the train of burst lightcurves was extracted. These lightcurves were then transformed to observable fluxes using the sampled parameters of M and $d\sqrt{\xi_\mathrm{b}}$. Instead of calculating an average burst lightcurve for each model sequence, we simply compare the last two bursts of each model. The total sample of 120 model lightcurves are shown in Figure 6.8 with the epoch observations.

The collective model lightcurves are broadly consistent with the observations, and the length of the PRE phase and the cooling of the tail are reproduced. KEPLER has limited atmosphere physics for modelling the PRE process, and with our simple truncation at L_Edd, the PRE phase itself is poorly reproduced. Nevertheless, this comparison suggests that the modelled bursts remain broadly consistent with observations, even given that F_Edd is the only lightcurve quantity being matched by the MCMC routine.

6.4 Discussion

We compare here the parameter estimates obtained for 4U 1820$-$30 with previous studies and predictions. Our reported values are the maximum likelihoods for the 1D marginalised posteriors, listed in Table 6.3, where the uncertainties are the 68% credible intervals.

We obtain accretion rates of $\dot{m}_1 = 0.208^{+0.05}_{-0.005}$ and $\dot{m}_2 = 0.37^{+0.06}_{-0.07}\,\dot{m}_\mathrm{Edd}$ for the 1997 and 2009 epochs, respectively. These values are larger than the initial 0.144 and 0.226 $\dot{m}_\mathrm{Edd}$ suggested by Galloway, Goodwin, and Keek (2017), although their values did not include anisotropy and thus implicitly assumed $\xi_\mathrm{p} = 1$. Our estimate for a persistent anisotropy of $\xi_\mathrm{p} = 2.4^{+0.7}_{-0.5}$ implies larger local accretion rates, because the resulting emission is preferentially beamed away from the observer. On the other hand, $\dot{m}_1$ is consistent

Figure 6.8: The full burst lightcurves from an additional 30 epoch pairs of KE-PLER models (blue curves), and the observed lightcurves (black histograms). The model parameters were taken from a random sample of the MCMC chain, and the lightcurves were truncated at L_{Edd}, because KEPLER has limited ability to model the PRE burst phase. Instead of the average lightcurves used in Chapter 5, only the last two bursts from each model are shown. There is good overall agreement, particularly for the burst tail. The only lightcurve quantity matched by the MCMC routine was F_{Edd}, which we assumed to correspond to F_{peak}.

with the range of $\dot{m} \approx 0.2$–$0.26\,\dot{m}_{\mathrm{Edd}}$ reported by Cumming (2003), for models of an earlier epoch with a slightly longer recurrence time of $\Delta t = 3.2\,\mathrm{hr}$, compared to $\Delta t \approx 2.7\,\mathrm{h}$ for 1997.

Our estimates for crustal heating are $Q_{\mathrm{b},1} = 0.36^{+0.04}_{-0.17}$ and $Q_{\mathrm{b},2} = 0.04^{+0.11}_{-0.02}\,\mathrm{MeV\,nucleon^{-1}}$. Although broad, these values are roughly consistent with the range of 0.1–$0.2\,\mathrm{MeV\,nucleon^{-1}}$ used by Cumming (2003), and the typical expected rate of $Q_{\mathrm{b}} \approx 0.15\,\mathrm{MeV\,nucleon^{-1}}$ (e.g., Cumming et al., 2006). The total heating in the crust is predicted to be approximately 1–2 $\mathrm{MeV\,nucleon^{-1}}$ (Haensel and Zdunik, 2008), but the outflowing flux into the envelope depends on the thermal transport properties of the crust (e.g., Brown and Cumming, 2009), shallow heating (Deibel et al., 2015), and Urca neutrino cooling (Schatz et al., 2014).

Similar to the results for GS 1826−238 in Chapter 5, we again obtain an unusually large neutron star mass of $M = 2.3^{+0.3}_{-0.2}\,\mathrm{M}_\odot$ and a very large radius of $R = 15.1^{+0.9}_{-1}\,\mathrm{km}$. The mass is larger than the typical expected range of $1 \lesssim M \lesssim 2\,\mathrm{M}_\odot$ (Miller, 2013), and exceeds all previous estimates for 4U 1820−30, including $M = 1.29^{+0.19}_{-0.07}\,\mathrm{M}_\odot$ (Shaposhnikov and Titarchuk, 2004) and $M = 1.58 \pm 0.06\,\mathrm{M}_\odot$ (Güver et al., 2010). It even exceeds the two largest observed neutron star masses of $2.27^{+0.17}_{-0.15}\,\mathrm{M}_\odot$ (Linares, Shahbaz, and Casares, 2018), and $2.14^{+0.10}_{-0.09}\,\mathrm{M}_\odot$ (Cromartie et al., 2019). The radius is unrealistically large when compared to typical values predicted from equation of state models (e.g., Özel and Freire, 2016), and in light of recent constraints from the first gravitational wave observations of a neutron star merger (Abbott et al., 2018; Most et al., 2018). We again emphasise the limitations of this initial study, and further investigation will be required before more robust estimates for the neutron star properties can be obtained.

Our gravitational redshift of $z = 0.36 \pm 0.03$ is similar to the redshift inferred for GS 1826−238 of $z = 0.39 \pm 0.07$ (Chapter 5), and is consistent with the fit of $z = 0.36$ assuming $d = 7.0\,\mathrm{kpc}$ from Shaposhnikov and Titarchuk (2004). It is slightly smaller than the value of $z \approx 0.43$ from the mass and radius estimates of Güver et al. (2010).

Our distance estimate of $d\sqrt{\xi_{\mathrm{b}}} = 7.9 \pm 0.3\,\mathrm{kpc}$ is almost unchanged from the prior distribution of $d\sqrt{\xi_{\mathrm{b}}} = 7.85 \pm 0.33\,\mathrm{kpc}$, and is consistent with the posterior of $d\sqrt{\xi_{\mathrm{b}}} = 7.47 \pm 0.38\,\mathrm{kpc}$ reported by Özel et al. (2016), who also used a similar prior. It is larger, however, than the value of $d\sqrt{\xi_{\mathrm{b}}} = 6.5 \pm 0.5\,\mathrm{kpc}$ from Suleimanov et al. (2017).

Using a model for disc anisotropy from He and Keek (2016), we obtained estimates for an inclination of $i = 74^{+2°}_{-4}$, and anisotropy factors of $\xi_{\mathrm{b}} = 1.33 \pm 0.09$ and $\xi_{\mathrm{p}} = 2.4^{+0.7}_{-0.5}$. The inclination is larger than the earlier estimate of 35–$50°$ inferred from ultraviolet modulations in the companion (Anderson et al., 1997), but is consistent with the value of 73–$80°$ inferred from radius expansion models with varied mass and radius (Shaposhnikov and Titarchuk, 2004). The persistent anisotropy is slightly larger than the range of 1.5–2 from Cumming (2003).

6.4.1 Future Work

This study is the first extension of our multi-epoch MCMC models to a helium burster, 4U 1820−30. Because of the simplifications made, these results should be considered a proof-of-concept for future efforts.

To reduce the number of precomputed KEPLER simulations in the grid, we used fixed values for X_0, Z_{CNO}, and g (Section 6.2.2). Although the short orbital period of $11\,\mathrm{min}$ indicates a hydrogen-poor accreted composition, semi-analytic models suggest a small hydrogen fraction of $X_0 \lesssim 0.10$ may still be possible (Cumming, 2003). Expanding our model grid beyond $X_0 = 0.0$ and $Z_{\mathrm{CNO}} = 0.015$ could improve the constraints on the accreted composition.

The particularly large neutron star mass of $M = 2.3^{+0.3}_{-0.2}\,\mathrm{M_\odot}$ and radius of $R = 15.1^{+0.9}_{-1}\,\mathrm{km}$, along with our results for GS 1826$-$238 (Chapter 5), suggests a possible bias in the model towards large masses and radii. A possible contributing factor is that we assumed a fixed surface gravity of $g = 1.858 \times 10^{14}\,\mathrm{cm\,s^{-2}}$, and so only a single contour of M and R was explored, potentially excluding more realistic values. Using a variable parameter for g when modelling GS 1826$-$238 did not prevent large mass estimates, although a more typical radius of $R = 11.3 \pm 1.3\,\mathrm{km}$ was obtained. Another possibility is that the flat priors used for M and g do not appropriately "penalise" unrealistic values. Future studies could incorporate existing mass constraints into the prior. For example, the mass distribution of observed neutron star populations of $M = 1.46 \pm 0.21\,\mathrm{M_\odot}$ (Özel, Gould, and Güver, 2012), or theoretical predictions from equation of state models, as used in Goodwin et al. (2019).

Only the burst recurrence time, Δt, was predicted with the model grid, while F_{Edd} and F_{p} were calculated analytically. The dataset from Galloway, Goodwin, and Keek (2017) does include burst fluences, and modelling these in future studies may help to break the degeneracy between Q_{b} and $\dot{m}$ (Figure 6.3). However, the appropriate method to predict fluence remains unclear, because KEPLER can only crudely reproduce PRE lightcurves (Section 6.3.6). A simple cutoff at the Eddington luminosity could be adopted (Figure 6.8), but the value of L_{Edd} relies on the choice of M and R during MCMC sampling. Integrating over the model lightcurves during the MCMC routine, instead of storing precomputed fluences in the grid, could severely impact computational efficiency.

Similar to our posteriors for GS 1826$-$238, some of the parameter distributions for 4U 1820$-$30 are truncated at the boundaries of the grid (Figure 6.3). For example, $Q_{\mathrm{b},1}$ is limited at $0.4\,\mathrm{MeV\,nucleon^{-1}}$, and $\dot{m}_2$ at $0.5\,\dot{m}_{\mathrm{Edd}}$. The model grid could simply be extended, but preliminary models with larger Q_{b} and $\dot{m}$ transitioned to stable burning. Because of the rectangular structure of the model grid, regions of parameter space with low probability density must be simulated, for example large pairs of Q_{b} and $\dot{m}$ (Figure 6.5). An irregular grid structure could instead be used for the parameter space of interest, but this approach may slow down interpolation calculations and reduce the efficiency of the MCMC routine. Alternatively, separate grids could be used to model each epoch, each covering a smaller suitable region of parameter space.

6.5 Conclusion

We have presented the application of MCMC methods to 1D models of hydrogen-poor PRE bursts. By comparing samples from a precomputed model grid with multi-epoch data, we obtained system parameters for the

ultra-compact helium accretor, 4U 1820−30. Our parameter estimates were generally consistent with previous works, although the anomalous neutron star mass of $M = 2.3^{+0.3}_{-0.2}\,\mathrm{M}_\odot$ indicates that more robust priors should be explored. Future studies should also aim to include X_0, Z_{CNO}, and g in the model grid, and extend the range of models in Q_{b} and $\dot{m}$. Despite these limitations, the posterior distributions of the model predictions were consistent with the data (Figure 6.4), and our comparison of model lightcurves, truncated at L_{Edd}, also revealed broad agreement (Figure 6.8).

This study demonstrates that the MCMC models developed in Chapter 5 for GS 1826−238 can be extended to other systems and bursting regimes. With further refinement, these methods represent a promising avenue for constraining the properties of accreting neutron star systems.

Chapter 7

Conclusion

The results presented in this thesis represent valuable steps forward in the modelling of X-ray bursts. In this concluding chapter, we summarise the main results from Chapters 3, 4, 5, and 6, and discuss the possible directions of future work.

7.1   Summary

In Chapter 3, we presented improvements to the setup and analysis of KE-PLER burst models. It was discovered that previous KEPLER burst models had inadvertently been using incorrect opacities (Section 3.1). An opacity multiplication factor of ≈ 1.5 had mistakenly remained in the setup files, which led to artificially hotter thermal profiles and increased burst rates. This error likely contributed to discrepancies that had been noticed between the recurrence times predicted by KEPLER and other burst codes.

Another issue we uncovered was extended model burn-in (Section 3.2). Previously, models were assumed to reach a steady limit cycle within the first few bursts. We found, however, that systematic trends in the burst properties could persist for tens of bursts, potentially affecting the entire model sequence. To address this issue, we tested the addition of a nuclear heat source during the thermal setup of the envelope, which was previously neglected. We found that including this nuclear "preheating" helped the envelopes begin closer to thermal equilibrium, and largely removed the burn-in.

Following our improvements to the KEPLER burst models, we performed the first direct comparisons between 1D burst codes for the same input parameters (Section 3.3). Using an existing set of five MESA burst models, we compared the predicted burst rates, energetics, and lightcurves to KEPLER. Although we found generally consistent predictions for low accretion rates, at higher accretion rates KEPLER produced energetics that were ≈ 10–$40\,\%$ larger. Additionally, there was a systematic offset of 1 burst per day between the predicted burst rates.

In Chapter 4, we presented the first burst simulations to use time-dependent accretion rates. By allowing the accretion rate to vary continuously with time, we modelled four observed bursts from a transient accretion episode of SAX J1808.4$-$3658. We successfully reproduced the observed burst timings and and fluences, and predicted additional bursts during windows when the source was not being observed, in line with previous models. We also computed models using average accretion rates instead of continuously-varying rates, to compare with the method previously used for models of this system. The comparison suggested a possible

systematic bias when using averaged accretion rates, resulting in larger recurrence times when $\dot{M}$ is increasing, and smaller recurrence times when $\dot{M}$ is decreasing.

In Chapter 5, we presented the first application of Markov Chain Monte Carlo (MCMC) methods to large grids of burst models. We precomputed a grid of 3840 KEPLER simulations across five model parameters. This was the largest set of 1D burst models to date, and varied parameters that are often left fixed, including the crustal heating, Q_b, and the surface gravity, g. By interpolating the burst predictions over the grid, we could efficiently sample the parameter space using MCMC methods. Using multi-epoch observations of the Clocked Burster, GS 1826−238, we obtained posterior probability distributions for the system parameters. With epoch-dependent crustal heating, we could test for a dependence of Q_b on $\dot{m}$, and found that lower accretion rates were consistent with larger Q_b, although there was significant overlap between the posteriors. This study demonstrated the possibility of constraining system properties using multi-epoch burst data.

In Chapter 6, we extended our MCMC methods from Chapter 5 to a helium bursting source, 4U 1820−30. We precomputed a grid of 168 hydrogen-poor simulations, and fit observed photospheric radius expansion (PRE) bursts from two epochs of 4U 1820−30. We thus obtained posterior distributions for the system parameters. The predicted distributions of the observables were consistent with the data, and the posterior constraints on the system parameters generally agreed with previous estimates. A sample comparison of the full burst lightcurves suggested that the overall morphology is reproduced, despite the known limitations of KEPLER for simulating PRE lightcurves. This initial extension of our multi-epoch MCMC routine to another system demonstrated its feasibility as a generalised approach to burst modelling.

7.2   Future Work and Outlook

Our improvements to the KEPLER model setup in Chapter 3 can still be developed further. The implementation of nuclear preheating largely reduced the model burn-in, but small systematic trends in Δt remained for pure He models. The preheating setup was only tested for a depth of $y = 8 \times 10^7 \, \mathrm{g \, cm^{-2}}$ and a strength of $Q_\mathrm{nuc} = 5 \, \mathrm{MeV \, nucleon^{-1}}$, and these values should ideally be adjusted depending on the burst conditions for each model. Predicting the nuclear energy production prior to actually computing the KEPLER simulation may prove difficult, and so an iterative approach may be needed.

Additional comparisons between KEPLER, MESA, and other burst codes should also be pursued. Our comparison focused on five MESA models of varying $\dot{m}$, with fixed values for X_0, Z_CNO, Q_b, and g. The discrepancies we found in Section 3.3 may behave differently for other input parameters. In-depth comparisons of the burst trains, thermal profiles, convective regions, and ashes composition will also help determine the fundamental differences between the burst codes.

Our simulation in Chapter 4 demonstrated the feasibility of modelling bursts during unstable accretion episodes. Only a single model matched to the observations was presented as a test case, but systematic parameter

studies (such as that carried out in Chapter 5) are needed to constrain the system properties. In particular, determining the accreted composition of X_0 and Z_{CNO} could help to constrain the evolutionary history of the binary system. Extending these methods to other transient accretors could also further test the model capabilities for reproducing observed burst properties.

Our multi-epoch MCMC models in Chapter 5 pose promising avenues for future work. As this study was the first implementation of these methods, multiple simplifying assumptions were made. Flat prior distributions were used for all parameters except Z_{CNO}. Priors could be explored which are informed by theoretical expectations, such as distributions over the neutron star mass and radius based on equation of state models. This may help to address the unusually large posteriors for mass, which were truncated at our chosen upper limit of $M = 2.2\,M_\odot$. Of the observables matched by the MCMC routine, the fluence and peak flux were derived from the burst lightcurve, but the full lightcurve itself was not compared with the models. A challenge with implementing such a comparison, however, is the question of how to interpolate lightcurves between the model grid points. Alternatively, additional parametrisations of the lightcurve could be used as further constraints, such as exponential or power law fits to the decay tail. Care should be taken that these quantities behave smoothly over the model grid. Finally, several posteriors were truncated by the boundaries of the model grid, for example g, Z_{CNO}, and Q_b. The existing model grid can be extended in these parameters, to better span the parameter space.

Most of the improvements proposed above also apply to our extension to 4U 1820−30 in Chapter 6. In comparison to the models of GS 1826−238, this study was kept to a limited scope, and lays the groundwork for an expanded project in future. Only $\dot{m}$ and Q_b were varied for the model grid, to limit the total number of simulations. Expanding the grid to include the parameters explored for GS 1826−238 – X_0, Z_{CNO}, and g – is a natural next step. KEPLER lightcurves exhibit super-Eddington luminosities during PRE. To avoid the ambiguity regarding how to correctly extract burst fluences, we interpolated only the burst rate from the model grid. Fitting the observed fluences may help to break the strong degeneracies between $\dot{m}$ and Q_b seen in the posteriors. A possible first test is to simply calculate fluences after applying a flat truncation at the Eddington luminosity, as used for the sample lightcurve comparison.

In closing, we have presented multiple contributions to the modelling of thermonuclear X-ray bursts. The methods developed here serve as a step towards obtaining robust constraints for the properties of accreting neutron stars.

Bibliography

Abbott, B. P. et al. (2018). GW170817: Measurements of Neutron Star Radii and Equation of State. *Physical Review Letters* 121, p. 161101. DOI: `10.1103/PhysRevLett.121.161101`.

Anderson, Scott F. et al. (1997). Time-Resolved Ultraviolet Observations of the Globular Cluster X-Ray Source in NGC 6624: The Shortest Known Period Binary System. *The Astrophysical Journal Letters* 482, pp. L69–L72. DOI: `10.1086/310672`.

Babushkina, O. P. et al. (1975). Hard X-ray bursts in June 1971. *Soviet Astronomy Letters* 1, pp. 32–34.

Belian, R. D., J. P. Conner, and W. D. Evans (1976). The discovery of X-ray bursts from a region in the constellation Norma. *The Astrophysical Journal Letters* 206, pp. L135–L138. DOI: `10.1086/182151`.

Bildsten, Lars (1997). Thermonuclear Burning on Rapidly Accreting Neutron Stars. *arXiv:astro-ph/9709094*. arXiv: astro-ph/9709094.

Brown, Edward F. and Lars Bildsten (1998). The Ocean and Crust of a Rapidly Accreting Neutron Star: Implications for Magnetic Field Evolution and Thermonuclear Flashes. *The Astrophysical Journal* 496.2, p. 915. DOI: `10.1086/305419`.

Brown, Edward F. and Andrew Cumming (2009). Mapping crustal heating with the cooling light curves of quasi-persistent transients. *The Astrophysical Journal* 698.2, pp. 1020–1032. DOI: `10.1088/0004-637X/698/2/1020`.

Cavecchi, Yuri et al. (2013). Flame propagation on the surfaces of rapidly rotating neutron stars during Type I X-ray bursts. *Monthly Notices of the Royal Astronomical Society* 434, pp. 3526–3541. DOI: `10.1093/mnras/stt1273`.

Cavecchi, Yuri et al. (2015). Rotational effects in thermonuclear type I bursts: equatorial crossing and directionality of flame spreading. *Monthly Notices of the Royal Astronomical Society* 448, pp. 445–455. DOI: `10.1093/mnras/stu2764`.

Cavecchi, Yuri et al. (2016). Fast and slow magnetic deflagration fronts in type I X-ray bursts. *Monthly Notices of the Royal Astronomical Society* 459.2, pp. 1259–1275. DOI: `10.1093/mnras/stw728`.

Clark, G. W. et al. (1976). Recurrent brief X-ray bursts from the globular cluster NGC 6624. *The Astrophysical Journal* 207, pp. L105–L108. DOI: `10.1086/182190`.

Cornelisse, R. et al. (2003). Six years of BeppoSAX Wide Field Cameras observations of nine galactic type I X-ray bursters. *Astronomy and Astrophysics* 405, pp. 1033–1042. DOI: `10.1051/0004-6361:20030629`.

Cromartie, H. T. et al. (2019). Relativistic Shapiro delay measurements of an extremely massive millisecond pulsar. *Nature Astronomy*, pp. 1–5. DOI: `10.1038/s41550-019-0880-2`.

Cumming, Andrew (2003). Models of Type I X-Ray Bursts from 4U 1820-30. *The Astrophysical Journal* 595, pp. 1077–1085. DOI: `10.1086/377446`.

Cumming, Andrew and Lars Bildsten (2001). Carbon Flashes in the Heavy-Element Ocean on Accreting Neutron Stars. *The Astrophysical Journal Letters* 559, pp. L127–L130. DOI: 10.1086/323937.

Cumming, Andrew et al. (2006). Long Type I X-Ray Bursts and Neutron Star Interior Physics. *The Astrophysical Journal* 646, pp. 429–451. DOI: 10.1086/504698.

Cyburt, R. H. et al. (2016). Dependence of X-Ray Burst Models on Nuclear Reaction Rates. *The Astrophysical Journal* 830, p. 55. DOI: 10.3847/0004-637X/830/2/55.

Cyburt, Richard H. et al. (2010). The JINA REACLIB Database: Its Recent Updates and Impact on Type-I X-ray Bursts. *The Astrophysical Journal Supplement Series* 189, pp. 240–252. DOI: 10.1088/0067-0049/189/1/240.

Deibel, Alex et al. (2015). A Strong Shallow Heat Source in the Accreting Neutron Star MAXI J0556-332. *The Astrophysical Journal Letters* 809, p. L31. DOI: 10.1088/2041-8205/809/2/L31.

Fisker, Jacob Lund, Hendrik Schatz, and Friedrich-Karl Thielemann (2008). Explosive Hydrogen Burning during Type I X-Ray Bursts. *The Astrophysical Journal Supplement Series* 174.1, p. 261. DOI: 10.1086/521104.

Fisker, Jacob Lund et al. (2006). The Importance of 15O(α,γ)19Ne to X-Ray Bursts and Superbursts. *The Astrophysical Journal* 650.1, p. 332. DOI: 10.1086/507083.

Foreman-Mackey, Daniel et al. (2013). emcee: The MCMC Hammer. *Publications of the Astronomical Society of the Pacific* 125.925, pp. 306–312. DOI: 10.1086/670067.

Fujimoto, M. Y., T. Hanawa, and S. Miyaji (1981). Shell flashes on accreting neutron stars and X-ray bursts. *The Astrophysical Journal* 247, pp. 267–278. DOI: 10.1086/159034.

Fujimoto, M. Y. and D. Sugimoto (1979). Asymptotic Strength of Thermal Pulses in the Helium Shell Burning. *Publications of the Astronomical Society of Japan* 31, pp. 1–10.

Fujimoto, Masayuki Y. (1988). Angular distribution of radiation from low-mass X-ray binaries. *The Astrophysical Journal* 324, pp. 995–1000. DOI: 10.1086/165955.

Galloway, Duncan K., Adelle J. Goodwin, and Laurens Keek (2017). Thermonuclear Burst Observations for Model Comparisons: A Reference Sample. *Publications of the Astronomical Society of Australia* 34, e019. DOI: 10.1017/pasa.2017.12.

Galloway, Duncan K. and Laurens Keek (2017). Thermonuclear X-ray bursts. *arXiv:1712.06227 [astro-ph]*. arXiv: 1712.06227.

Galloway, Duncan K. et al. (2008). Thermonuclear (Type I) X-Ray Bursts Observed by the Rossi X-Ray Timing Explorer. *The Astrophysical Journal Supplement Series* 179.2, p. 360. DOI: 10.1086/592044.

Goodwin, A. J. et al. (2019). A Bayesian Approach to Matching Thermonuclear X-ray Burst Observations with Models. *arXiv:1907.00996 [astro-ph]*. arXiv: 1907.00996.

Grindlay, J. et al. (1976). Discovery of intense X-ray bursts from the globular cluster NGC 6624. *The Astrophysical Journal Letters* 205, pp. L127–L130. DOI: 10.1086/182105.

Gupta, Sanjib et al. (2007). Heating in the Accreted Neutron Star Ocean: Implications for Superburst Ignition. *The Astrophysical Journal* 662.2, p. 1188. DOI: 10.1086/517869.

Güver, Tolga et al. (2010). The Mass and Radius of the Neutron Star in 4U 1820-30. *The Astrophysical Journal* 719, pp. 1807–1812. DOI: 10.1088/0004-637X/719/2/1807.

Haensel, P. and J. L. Zdunik (2008). Models of crustal heating in accreting neutron stars. *Astronomy and Astrophysics* 480, pp. 459–464. DOI: 10.1051/0004-6361:20078578.

Hanawa, T. and M. Y. Fujimoto (1984). Thermal response of neutron stars to shell flashes. *Publications of the Astronomical Society of Japan* 36.2, p. 199.

Hansen, C. J. and H. M. van Horn (1975). Steady-state nuclear fusion in accreting neutron-star envelopes. *The Astrophysical Journal* 195, p. 735. DOI: 10.1086/153375.

He, C.-C. and L. Keek (2016). Anisotropy of X-Ray Bursts from Neutron Stars with Concave Accretion Disks. *The Astrophysical Journal* 819, p. 47. DOI: 10.3847/0004-637X/819/1/47.

Heger, A. et al. (2003). How Massive Single Stars End Their Life. *The Astrophysical Journal* 591, pp. 288–300. DOI: 10.1086/375341.

Heger, Alexander, Andrew Cumming, and S. E. Woosley (2007). Millihertz Quasi-periodic Oscillations from Marginally Stable Nuclear Burning on an Accreting Neutron Star. *The Astrophysical Journal* 665, pp. 1311–1320. DOI: 10.1086/517491.

Heger, Alexander et al. (2007). Models of Type I X-Ray Bursts from GS 1826-24: A Probe of rp-Process Hydrogen Burning. *The Astrophysical Journal Letters* 671, pp. L141–L144. DOI: 10.1086/525522.

Hinton, Samuel R. (2016). ChainConsumer. *The Journal of Open Source Software* 1, p. 00045. DOI: 10.21105/joss.00045.

Hoffman, Jeffrey A., Herman L. Marshall, and Walter H. G. Lewin (1978). Dual character of the rapid burster and a classification of X-ray bursts. *Nature* 271.5646, pp. 630–633. DOI: 10.1038/271630a0.

Johnston, Zac, Alexander Heger, and Duncan K Galloway (2018). Simulating X-ray bursts during a transient accretion event. *Monthly Notices of the Royal Astronomical Society* 477.2, pp. 2112–2118. DOI: 10.1093/mnras/sty757.

Johnston, Zac, Alexander Heger, and Duncan K. Galloway (2019). Multi-epoch X-ray burst modelling: MCMC with large grids of 1D simulations. *arXiv:1909.07977 [astro-ph]*. arXiv: 1909.07977.

José, Jordi and Margarita Hernanz (1998). Nucleosynthesis in Classical Novae: CO versus ONe White Dwarfs. *The Astrophysical Journal* 494, pp. 680–690. DOI: 10.1086/305244.

José, Jordi et al. (2010). Hydrodynamic models of Type I X-ray bursts: metallicity effects. *The Astrophysical Journal Supplement Series* 189.1, pp. 204–239. DOI: 10.1088/0067-0049/189/1/204.

Joss, P. C. (1977). X-ray bursts and neutron-star thermonuclear flashes. *Nature* 270, pp. 310–314. DOI: 10.1038/270310a0.

— (1978). Helium-burning flashes on an accreting neutron star - A model for X-ray burst sources. *The Astrophysical Journal Letters* 225, pp. L123–L127. DOI: 10.1086/182808.

Keek, L. and A. Heger (2011). Multi-zone Models of Superbursts from Accreting Neutron Stars. *The Astrophysical Journal* 743.2, p. 189. DOI: 10.1088/0004-637X/743/2/189.

— (2016). Carbon production on accreting neutron stars in a new regime of stable nuclear burning. *Monthly Notices of the Royal Astronomical Society: Letters* 456.1, pp. L11–L15. DOI: 10.1093/mnrasl/slv167.

Keek, L., A. Heger, and J. J. M. in't Zand (2012). Superburst Models for Neutron Stars with Hydrogen- and Helium-rich Atmospheres. *The Astrophysical Journal* 752, p. 150. DOI: 10.1088/0004-637X/752/2/150.

King, A. R. and M. G. Watson (1986). The shortest period binary star? *Nature* 323.6084, p. 105. DOI: 10.1038/323105a0.

Koike, O. et al. (1999). Rapid proton capture on accreting neutron stars - effects of uncertainty in the nuclear process. *Astronomy and Astrophysics* 342, pp. 464–473.

Koike, Osamu et al. (2004). Final Products of the rp-Process on Accreting Neutron Stars. *The Astrophysical Journal* 603, pp. 242–251. DOI: 10.1086/381354.

Kuulkers, E. et al. (2003). Photospheric radius expansion X-ray bursts as standard candles. *Astronomy and Astrophysics* 399, pp. 663–680. DOI: 10.1051/0004-6361:20021781.

Lamb, D. Q. and F. K. Lamb (1978). Nuclear burning in accreting neutron stars and X-ray bursts. *The Astrophysical Journal* 220, pp. 291–302. DOI: 10.1086/155905.

Lampe, Nathanael, Alexander Heger, and Duncan K. Galloway (2016). The Influence of Accretion Rate and Metallicity on Thermonuclear Bursts: Predictions from KEPLER Models. *The Astrophysical Journal* 819, p. 46. DOI: 10.3847/0004-637X/819/1/46.

Lapidus, I. I. and R. A. Sunyaev (1985). Angular distribution and polarization of X-ray-burster radiation (during stationary and flash phases). *Monthly Notices of the Royal Astronomical Society* 217, pp. 291–303. DOI: 10.1093/mnras/217.2.291.

Lewin, W. H. G., W. D. Vacca, and E. M. Basinska (1984). Precursors to X-ray bursts - The result of expansion and subsequent contraction of the neutron star's photosphere. *The Astrophysical Journal* 277, pp. L57–L60. DOI: 10.1086/184202.

Lewin, W. H. G. et al. (1976). Discovery of X-ray bursts from several sources near the galactic centre. *Monthly Notices of the Royal Astronomical Society* 177, 83P–92P. DOI: 10.1093/mnras/177.1.83P.

Lewin, Walter H. G., Jan Van Paradijs, and Ronald E. Taam (1993). X-ray bursts. *Space Science Reviews* 62.3-4, pp. 223–389. DOI: 10.1007/BF00196124.

Liebendörfer, Matthias, Stephan Rosswog, and Friedrich-Karl Thielemann (2002). An Adaptive Grid, Implicit Code for Spherically Symmetric, General Relativistic Hydrodynamics in Comoving Coordinates. *The Astrophysical Journal Supplement Series* 141, pp. 229–246. DOI: 10.1086/339872.

Linares, M., T. Shahbaz, and J. Casares (2018). Peering into the Dark Side: Magnesium Lines Establish a Massive Neutron Star in PSR J2215+5135. *The Astrophysical Journal* 859.1, p. 54. DOI: 10.3847/1538-4357/aabde6.

Maraschi, L. and A. Cavaliere (1977). X-ray bursts of nuclear origin? *Highlights in Astronomy*. Vol. 4, p. 127.

Margalit, Ben and Brian D. Metzger (2017). Constraining the Maximum Mass of Neutron Stars from Multi-messenger Observations of GW170817. *The Astrophysical Journal* 850.2, p. L19. DOI: 10.3847/2041-8213/aa991c.

Meisel, Zach (2018). Consistent Modeling of GS 1826-24 X-Ray Bursts for Multiple Accretion Rates Demonstrates the Possibility of Constraining *rp*-process Reaction Rates. *The Astrophysical Journal* 860.2, p. 147. DOI: 10.3847/1538-4357/aac3d3.

Meisel, Zach and Alex Deibel (2017). Constraints on Bygone Nucleosynthesis of Accreting Neutron Stars. *The Astrophysical Journal* 837, p. 73. DOI: 10.3847/1538-4357/aa618d.

Meisel, Zach, Grant Merz, and Sophia Medvid (2019). Influence of Nuclear Reaction Rate Uncertainties on Neutron Star Properties Extracted from X-Ray Burst Model–Observation Comparisons. *The Astrophysical Journal* 872.1, p. 84. DOI: 10.3847/1538-4357/aafede.

Miller, M. Coleman (2013). Astrophysical Constraints on Dense Matter in Neutron Stars. *arXiv:1312.0029 [astro-ph, physics:nucl-th]*. arXiv: 1312.0029.

Most, Elias R. et al. (2018). New Constraints on Radii and Tidal Deformabilities of Neutron Stars from GW170817. *Physical Review Letters* 120, p. 261103. DOI: 10.1103/PhysRevLett.120.261103.

Özel, Feryal and Paulo Freire (2016). Masses, Radii, and the Equation of State of Neutron Stars. *Annual Review of Astronomy and Astrophysics* 54.1, pp. 401–440. DOI: 10.1146/annurev-astro-081915-023322.

Özel, Feryal, Andrew Gould, and Tolga Güver (2012). The mass and radius of the neutron star in the bulge low-mass X-ray binary KS 1731–260. *The Astrophysical Journal* 748.1, p. 5. DOI: 10.1088/0004-637X/748/1/5.

Özel, Feryal et al. (2016). The dense matter equation of state from neutron star radius and mass measurements. *The Astrophysical Journal* 820.1, p. 28. DOI: 10.3847/0004-637X/820/1/28.

Paradijs, J. van, W. Penninx, and W. H. G. Lewin (1988). On the relation between X-ray burst properties and the persistent X-ray luminosity. *Monthly Notices of the Royal Astronomical Society* 233, pp. 437–450. DOI: 10.1093/mnras/233.2.437.

Parikh, A. et al. (2009). Impact of uncertainties in reaction Q values on nucleosynthesis in Type I X-ray bursts. *Physical Review C* 79, p. 045802. DOI: 10.1103/PhysRevC.79.045802.

Parikh, Anuj et al. (2008). The Effects of Variations in Nuclear Processes on Type I X-Ray Burst Nucleosynthesis. *The Astrophysical Journal Supplement Series* 178, pp. 110–136. DOI: 10.1086/589879.

Paxton, Bill et al. (2015). Modules for experiments in stellar astrophysics (MESA): binaries, pulsations, and explosions. *The Astrophysical Journal Supplement Series* 220.1, p. 15. DOI: 10.1088/0067-0049/220/1/15.

Peng, Fang, Edward F. Brown, and James W. Truran (2007). Sedimentation and Type I X-Ray Bursts at Low Accretion Rates. *The Astrophysical Journal* 654.2, p. 1022. DOI: 10.1086/509628.

Rakavy, G., G. Shaviv, and Z. Zinamon (1967). Carbon and Oxygen Burning Stars and Pre-Supernova Models. *The Astrophysical Journal* 150, p. 131. DOI: 10.1086/149318.

Rauscher, T. et al. (2002). Nucleosynthesis in Massive Stars with Improved Nuclear and Stellar Physics. *The Astrophysical Journal* 576, pp. 323–348. DOI: 10.1086/341728.

Schatz, H. et al. (1998). rp-Process Nucleosynthesis at Extreme Temperature and Density Conditions. *Physics Reports* 294. DOI: 10.1016/S0370-1573(97)00048-3.

Schatz, H. et al. (2001). End Point of the rp Process on Accreting Neutron Stars. *Physical Review Letters* 86.16, p. 3471. DOI: 10.1103/PhysRevLett.86.3471.

Schatz, H. et al. (2014). Strong neutrino cooling by cycles of electron capture and β- decay in neutron star crusts. *Nature* 505.7481, pp. 62–65. DOI: 10.1038/nature12757.

Schwarzschild, M. and R. Härm (1965). Thermal Instability in Non-Degenerate Stars. *The Astrophysical Journal* 142, p. 855. DOI: 10.1086/148358.

Shaposhnikov, Nickolai and Lev Titarchuk (2004). On the Nature of the Flux Variability during an Expansion Stage of a Type I X-Ray Burst: Constraints on Neutron Star Parameters for 4U 1820–30. *The Astrophysical Journal Letters* 606.1, p. L57. DOI: 10.1086/421015.

Shara, M. M. (1982). Localized thermonuclear runaways and volcanoes on degenerate dwarf stars. *The Astrophysical Journal* 261, pp. 649–660. DOI: 10.1086/160376.

Stella, L., W. Priedhorsky, and N. E. White (1987). The discovery of a 685 second orbital period from the X-ray source 4U 1820 - 30 in the globular cluster NGC 6624. *The Astrophysical Journal Letters* 312, pp. L17–L21. DOI: 10.1086/184811.

Strohmayer, Tod and Lars Bildsten (2006). New Views of Thermonuclear Bursts. *Compact Stellar X-ray Sources*. Ed. by Walter Lewin and Michiel van der Klis. Vol. 39. Cambridge University Press, pp. 113–156.

Strohmayer, Tod E. and Edward F. Brown (2002). A Remarkable 3 Hour Thermonuclear Burst from 4U 1820-30. *The Astrophysical Journal* 566, pp. 1045–1059. DOI: 10.1086/338337.

Sugimoto, D. and M. Y. Fujimoto (1978). A General Theory for Thermal Pulses of Finite Amplitude in Nuclear Shell- Burnings. *Publications of the Astronomical Society of Japan* 30, pp. 467–482.

Suleimanov, Valery F. et al. (2017). Basic parameters of the helium-accreting X-ray bursting neutron star in 4U 1820-30. *Monthly Notices of the Royal Astronomical Society* 472.4, pp. 3905–3913. DOI: 10.1093/mnras/stx2234.

Sztajno, M. et al. (1987). Constraints on the mass-radius relation of the neutron star in 4U 1746-37/NGC 6441. *Monthly Notices of the Royal Astronomical Society* 226, pp. 39–55. DOI: 10.1093/mnras/226.1.39.

Taam, R. E. (1980). X-ray bursts from thermonuclear runaways on accreting neutron stars. *The Astrophysical Journal* 241, pp. 358–366. DOI: 10.1086/158348.

Taam, R. E. and R. E. Picklum (1978). Nuclear fusion and carbon flashes on neutron stars. *The Astrophysical Journal* 224, p. 210. DOI: 10.1086/156367.

Tawara, Y. et al. (1984). A very long X-ray burst with a precursor from XB 1715-321. *The Astrophysical Journal* 276, pp. L41–L44. DOI: 10.1086/184184.

Valenti, E., F. R. Ferraro, and L. Origlia (2007). Near-Infrared Properties of 24 Globular Clusters in the Galactic Bulge*. *The Astronomical Journal* 133.4, p. 1287. DOI: 10.1086/511271.

Wallace, R. K. and S. E. Woosley (1981). Explosive hydrogen burning. *The Astrophysical Journal Supplement Series* 45, pp. 389–420. DOI: 10.1086/190717.

Wallace, R. K., S. E. Woosley, and T. A. Weaver (1982). The thermonuclear model for X-ray transients. *The Astrophysical Journal* 258, pp. 696–715. DOI: 10.1086/160119.

Watts, Anna L. (2012). Thermonuclear Burst Oscillations. *Annual Review of Astronomy and Astrophysics* 50.1, pp. 609–640. DOI: 10.1146/annurev-astro-040312-132617.

Weaver, T. A., G. B. Zimmerman, and S. E. Woosley (1978). Presupernova evolution of massive stars. *The Astrophysical Journal* 225, pp. 1021–1029. DOI: 10.1086/156569.

Wiescher, M. et al. (2010). The Cold and Hot CNO Cycles. *Annual Review of Nuclear and Particle Science* 60.1, pp. 381–404. DOI: 10.1146/annurev.nucl.012809.104505.

Woosley, S. E., A. Heger, and T. A. Weaver (2002). The evolution and explosion of massive stars. *Reviews of Modern Physics* 74, pp. 1015–1071. DOI: 10.1103/RevModPhys.74.1015.

Woosley, S. E. and Ronald E. Taam (1976). γ-ray bursts from thermonuclear explosions on neutron stars. *Nature* 263.5573, pp. 101–103. DOI: 10.1038/263101a0.

Woosley, S. E. et al. (2004). Models for Type I X-Ray Bursts with Improved Nuclear Physics. *The Astrophysical Journal Supplement Series* 151, pp. 75–102. DOI: 10.1086/381533.

Zingale, M. et al. (2015). Comparisons of Two- and Three-Dimensional Convection in Type I X-Ray Bursts. *The Astrophysical Journal* 807, p. 60. DOI: 10.1088/0004-637X/807/1/60.