UNIVERSITY OF SOUTHAMPTON

School of Mathematical Sciences
Group of Applied Mathematics and Theoretical Physics

From data to dense matter: Inferring neutron star physics using computational analyses

by

Marcella Wijngaarden
MSc
ORCiD: 0000-0003-2804-4075

A thesis for the degree of
Doctor of Philosophy

February 2022

University of Southampton

Abstract

Faculty of Social Sciences
School of Mathematical Sciences

Doctor of Philosophy

From data to dense matter: Inferring neutron star physics using computational analyses

by Marcella Wijngaarden

Neutron stars host a range of interesting matter states from matter at supranuclear densities in their cores to possible neutron superfluidity and exotic nucleonic matter in their crusts. In this thesis, I use two distinct approaches probing different density and temperature regimes of the neutron star, probing the surface and crustal layers using X-ray data and the deep interior at high densities using gravitational waves. In the first part, we work on improvements to the theoretical numerical models of cooling neutron stars, which can be confronted with observations to infer interior properties. In particular, we investigate the effect of diffusive nuclear burning in the low-density envelope on cooling simulations and provide new analytic temperature relations. We find that a time-varying envelope composition should be taken into account and that the envelope can produce a non-negligible heat flux that would otherwise be interpreted as originating from the interior. In the second part, we develop a new computational gravitational wave data analysis method in order to simultaneously find and extract information about both the inspiral and post-merger remnant of a binary neutron star (BNS) coalescence. The code models the inspiral using solutions to general relativity whereas any unmodelled signal component is captured by sine-Gaussian wavelets. For the first time, we perform hybrid analyses of the full simulated BNS signal including both the inspiral and the post-merger component and find that some features are only extracted using the full analysis.

Contents

List of Figures

List of Tables

Declaration of Authorship

I declare that this thesis and the work presented in it is my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Parts of this work have been published as:
   M. J. P. Wijngaarden, Wynn C. G. Ho, Philip Chang, Craig O. Heinke, Dany Page, Mikhail Beznogov, and Daniel J. Patnaude. Diffusive nuclear burning in cooling simulations and application to new temperature data of the Cassiopeia A neutron star. *MNRAS*, 484:974–988, March 2019.
   M. J. P. Wijngaarden, Wynn C. G. Ho, Philip Chang, Dany Page, Rudy Wijnands, Laura S. Ootes, Andrew Cumming, Nathalie Degenaar, and Mikhail Beznogov. The effect of diffusive nuclear burning in neutron star envelopes on cooling in accreting systems. *MNRAS*, 493(4):4936–4944, April 2020.
   Marcella Wijngaarden, Katerina Chatziioannou, Andreas Bauswein, and Neil Cornish et. al. Analyzing of the full gravitational wave signal from neutron star binary coalescences. *to be submitted to Phys. Rev. D.*, 2021

Signed:..... Date:.....

Acknowledgements

There were many times I thought I would never reach this moment and I am grateful to many that I did.

First of all, I would like to thank my supervisor Wynn Ho. Thank you for always being a great support to me and teaching me an incredible amount about science and working in science. It is exceptional how you believed in me from the start and pushed me to get out of my comfort zone and present at and travel to many places. Thanks to your support and encouragement, my world is a lot larger now!

I would also like to thank the Gravity Group in Southampton, with many incredible scientists that are also great at explaining complicated things. A particular big thanks to Nils Andersson, for also always being a great support navigating many situations, keeping things light even when they seem heavy to me and being a great second supervisor.

I started a pre-doctoral project at the Flatiron Institute during an unusual time and thank the CCA and Will Farr for their support. I am especially very grateful to my project supervisor Katerina Chatziioannou, who continues to work with me long after. Thank you for welcoming me into the exciting world between neutron stars and gravitational waves data analysis. I learned so much from working with you and you truly make the world of academia a better place.

During my time as a PhD student, I have been lucky to meet many amazing people. In particular, I am happy to have met such a wonderful friend in Southampton and have made such great memories with Emma Osborne! I would also like to thank the people I have met in Southampton as well as those I enjoyed conferences and (summer) schools with and my friends at home in Amsterdam.

I thank everyone in science who worked with me and everybody who believed in me, supported me and reached out during the years, both in difficult and good times. I see you and I hope to pay it forward.

I would like to thank everybody at Project Cece, which we started during my masters and grew out of hand (in a good way!) during my PhD.

But mostly, I am so grateful to my parents, Karl & Nienke Wijngaarden, and my sister, Melissa. My being able to do this work is due to all of us. You believed in me when I didn't and never allowed me to give up. Unlike many in the field, I was not intrigued by the night sky or interested in science from a young age. But when I did discover how magical real physics actually was, my family took me seriously as a 15 year old girl. I will forever remember and cherish the many conversations about physics with my dad and how we excited eachother with new facts we discovered. Thank you for

taking me seriously. While I write these words, my father is in the hospital due to an aorta dissection and everything feels surreal. I am so lucky with you all, to many more exciting adventures in the future.

A lot has happened in life during the time I was also pursuing this PhD. Thank you to everyone who has been part of these difficult, beautiful, confusing, challenging and also joyeus times with me, especially you Adam.

Chapter 1

Introduction

1.1 Neutron stars

Since their discovery neutron stars (NSs) have provided scientists with a unique laboratory to study physics in an extreme conditions. The extreme densities in their cores allows for the study of matter at densities close to and exceeding the nuclear saturation density ($\rho_{\rm nuc} \sim 2.7 \times 10^{14}$ g cm^{-3}). The low temperatures and high densities (up to the order of $\sim 10^{15}$ g cm^{-3}) present in the neutron star cores cannot be accessed in Earth-based experiments (such as particle colliders). Therefore, neutron stars are currently our only window into this regime of dense matter physics.

Besides ultra dense matter in their cores, neutron stars host a range of unique matter states, including neutron superfluidity, relativistic degenerate lattice structures and possible exotic shapes of nuclei, as the density increases rapidly in the short ($\sim$1km) outer layers spanning from the star's surface to its core. Due to its compactness, the neutron star is located in an exceptionally strong gravitational field which can be used to test the theory of general relativity in the strong field limit. Neutron stars also support the strongest, stable magnetic fields (typically $\sim 10^8 - 10^{13}$ G) known to exist in the universe.

What makes neutron stars so interesting, is also what makes them so complex. None of the neutron star features can be studied in isolation and many areas - fluid dynamics, general relativity, nuclear physics, etc. - of sometimes uncertain physics have to be taken into account. Therefore, the study of neutron stars is truly a interdisciplinary effort captivating e.g., astrophysicists, nuclear and particle physicists, relativists and computational experts.

The interdisciplinary effort to gain a better understanding of the interior properties of neutron stars through observations of macroscopic properties includes studies of heavy pulsar observations (Antoniadis et al., 2013; Cromartie et al., 2019; Fonseca

et al., 2021), surface cooling after accretion or in isolation (Potekhin et al., 2015; Wijnands et al., 2017), NICER X-ray pulse-profiling (Miller et al., 2019; Riley et al., 2019; Miller et al., 2021; Riley et al., 2021), nuclear theory and experiment and studies combining the information from multiple channels (Raaijmakers et al., 2020; Dietrich et al., 2020; Landry et al., 2020; Al-Mamun et al., 2021; Raaijmakers et al., 2021; Biswas, 2021; Pang et al., 2021; Legred et al., 2021).

In this thesis I focus on two separate approaches for the study of neutron stars concentrating on different temperature regimes (cooling cold sources vs. a hot post-merger remnant), data sources (electromagnetic radiation vs. gravitational waves) and state in the neutron star's life (stars that have been around for many years vs. a short-lived neutron star remnant after the merger of two neutron stars). I will discuss these two approaches in the next sections.

1.2 Neutron star thermal evolution from electromagnetic data

Neutron stars are one of the possible remnants of a supernova explosion at the end of the live of a massive star $\gtrsim 8\ M_\odot$. When left undisturbed the neutron star will continue to cool over its lifetime, as its core does not contain fuel like typical solar-like stars. Immediately following its creation, the star will cool down very quickly through neutrino emission and the central region can reach temperatures of 10^{8-9} K within days. When the interior temperature has cooled below $\sim 10^8$, after $\sim 10^{3-5}$ years, the thermal evolution becomes dominated by photon cooling (Page et al., 2004; Potekhin et al., 2015).

The rate at which the neutron star cools down is set by its internal structure and composition. The observed cooling rate can thus be used to infer properties of the dense interior by comparison to theoretical predictions. Here I describe the key physics components to construct neutron star cooling models which can be compared with observations. In this description we will assume spherical symmetry and low magnetic field strengths, which is appropriate for the neutron star sources modeled in this thesis.

1.2.1 Composition

Due to the large density gradient in the outer layers of the neutron star, its composition is not homogeneous and it is necessary to consider how the composition changes with increasing density. In this section I describe the general neutron star composition structure as broadly illustrated in Figure 1.1.

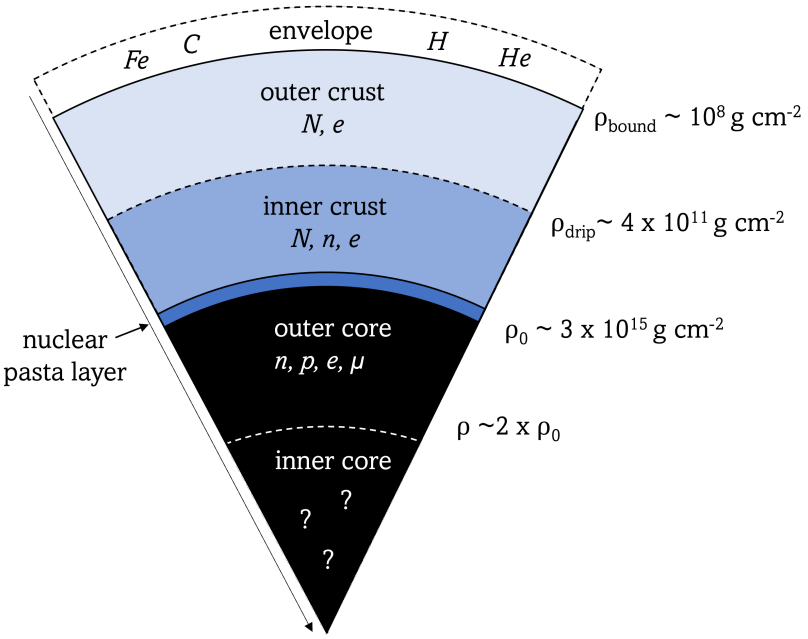

FIGURE 1.1: Schematic representation of the general neutron star structure (not to scale). In this figure, ρ_{drip} and ρ_0 denote the neutron drip density and the nuclear density, respectively. The primary constituents of each layer are indicated where N, n, p, e^- and μ^- represent nuclei, neutrons, protons, electrons and muons, respectively. The nuclear pasta layer is part of the inner crust and highlighted in the left figure in dark blue.

Starting at low densities, the envelope makes up a thin layer of only ~ 10 meters consisting of heavy and/or light elements such as iron, carbon, helium or hydrogen elements. As nucleonic matter isn't fully degenerate yet in this region, it contains a strong temperature gradient and acts as a thermal insulator between the degenerate interior and the surface. The heat conducting properties of the envelope are very sensitive to the elements making up its composition. As the envelope sets the relation between the interior temperature and the observed surface temperature, it is important for the interpretation of the observed surface temperature evolution to accurately model its composition (see Section 1.2.4 and Chapter 2 and 3).

The outer crusts consists of degenerate, relativistic electrons and fully ionized atoms. With increasing density, the region becomes more neutron rich as electron captures take place converting protons into neutrons. At the *neutron drip* density of $\rho_{drip} \approx 4 \times 10^{11}$ g cm^{-3}, the pressure is large enough for the neutrons to "drip" out of the nuclei and this marks the transition into the inner crust region (Haensel and Zdunik, 2003; Chamel and Haensel, 2008). Here, matter consists of electrons, free neutrons, which may become superfluid, and neutron-rich nuclei which form heavy isotopes (Haensel and Zdunik, 1990; Douchin and Haensel, 2001). At the bottom of the inner crust, the high pressure forces the clusters of nuclei to take on exotic shapes which are reminiscent of pastas such as spaghetti, gnocchi and lasagna (Alcain and

Dorso, 2017; Caplan et al., 2018; Newton et al., 2021). This region is commonly referred to as the *pasta layer* (see Figure 1.1).

The core makes up the largest part of the neutron star containing ∼99% of its mass. When the nuclear saturation density ($\rho_{\rm nuc} \sim 2.7 \times 10^{14}$ g cm^{-3}) is reached, which approximately marks the transition to the outer core, the nuclei break down into individual nucleons predominantly as degenerate neutrons but also a fraction of other particles such electrons, protons and muons. Neutron stars can also have an inner core where the density exceeds several times the nuclear saturation density and the composition is unknown. Several models for the composition and behaviour of matter at these densities have been proposed, including the existence of exotic particles such as hyperons and deconfined quarks.

1.2.2 The neutron star structure

Due to the compactness of the neutron star it induces a strong gravitational field and its structure must be described in the framework general relativity. For a spherical non-rotating star, its structure is described by the Tolman-Oppenheimer-Volkoff (TOV) equation for hydrostatic equilibrium,

$$\frac{dP}{dr} = -\frac{Gm\rho}{r^2}\left[1+\frac{P}{\rho c^2}\right]\left[1+\frac{4\pi r^3 P}{mc^2}\right]\left[1-\frac{2Gm}{rc^2}\right]^{-1} = -(\rho c + P)\frac{d\Phi}{dr} \tag{1.1}$$

where r is the radial coordinate, P is the pressure, G is the gravitational constant, ρ is the gravitational mass density, c is the speed of light, and m is the enclosed mass function

$$\frac{dm}{dr} = 4\pi r^2 \rho \tag{1.2}$$

and the metric function $\Phi(r)$ is a measure of the gravitational potential described as

$$\frac{d\Phi}{dr} = \frac{1}{c^2}\frac{Gm}{r^2}\left[1+\frac{4\pi r^3 P}{mc^2}\right]\left[1-\frac{2Gm}{rc^2}\right] = -\frac{dP}{dr}\left[\rho+\frac{P}{c^2}\right]^{-1}. \tag{1.3}$$

A critical ingredient for any theoretical model of the neutron star structure is the *equation of state*, which relates the pressure p to density ρ at a given temperature T and is set by the star's microscopic properties (such as particle content). When supplemented with an equation of state, integrating Equations 1.1-1.3 yields the global structure of the star (describing macroscopic properties such as total mass, radius and

surface temperature). Cooling simulations are often applied to old, isolated neutron stars or neutron stars in accreting binary systems (see Chapter 2 and 3, respectively), which are cold $\sim 10^{6-8}$K from the particle physics perspective compared to typical excitation energies in nuclei $\sim 10^{10}$ K. Therefore, in these simulations the equation of state is simplified to a relation between pressure and density, P = P(ρ).

1.2.3 Thermal evolution

After composing the neutron star structure, we can consider how its radial temperature profile evolves over time. Ultimately, the predicted temperature at the surface for the simulations can be compared with observed neutron star temperature data. As the temperature evolution is set by the neutron stars microscopic properties, we can constrain different models for the neutron star microphysics.

Assuming spherical symmetry and including relativistic effects, the one-dimensional temperature evolution of the star is governed by the equations for energy balance (Equation 1.4) and heat transport (Equation 1.5). The energy balance equation effectively balances any heat sources with photon and neutrino cooling rates in each radial zone:

$$\frac{d(Le^{2\Phi})}{dr} = -\frac{4\pi r^2 e^{\Phi}}{\sqrt{1-2Gm/rc^2}}\left(C_v\frac{dT}{dt} + e^{\Phi}(Q_\nu - Q_h)\right), \tag{1.4}$$

where L is the internal luminosity, Φ is the gravitational potential, C_v denotes the specific heat per unit volume, Q_ν and Q_h are, respectively, the neutrino emissivity and heating rate per unit volume. The heat transport equation can be written as:

$$\frac{dTe^{\phi}}{dr} = -\frac{1}{\kappa 4\pi r^2}\frac{Le^{\phi}}{\sqrt{1-2Gm/rc^2}}, \tag{1.5}$$

where κ is the thermal conductivity. Note that in both equations all physical properties are a function of the radial (r) and time (t) coordinate and together the equations describe how the temperature at each radial coordinate changes over time. Noting that the flux $F = L/(4\pi r^2)$, Equation 1.4 and Equation 1.5 can be combined into a single heat diffusion equation.

While the equations governing the star's thermal evolution may seem clear-cut to solve numerically, they couple to the star's complex microphysics through the specific heat per unit volume C_v and thermal conductivity κ as well as through the heating (Q_h) and neutrino cooling (Q_ν) rate. Both quantities are strongly affected by the

composition, density and temperature and change rapidly in different regions of the star (see Figure 1.1)).

The total specific heat per unit volume is a sum of the contributions by different particles in the region, i.e. $C_v = \sum_i C_{v,i}$, where i denotes the particle species (e.g., neutrons, protons, electrons etc.). The total specific heat of the star is dominated by the contribution from the core, which makes up about 90% of the neutron star mass, and its properties depend on its equation of state. Therefore, constraints on the core specific heat could potentially help constrain the equation of state and vice versa (Cumming et al., 2017). In the crust, which contains a large density gradient, the specific heat from high to low densities is composed of degenerate free neutrons, the nuclear lattice and a degenerate electron gas. At low densities, for low temperatures, the specific heat is dominated by the contribution from electrons as a strongly degenerate, relativistic gas, while for higher temperatures ($T \sim 10^{8-9}$ K) contributions from the ion lattice become relevant as it decrystallizes into a plasma state. Vibrations of the crystal lattice can also contribute through lattice vibrations. When moving to higher densities in the crust, the contribution from free unpaired neutrons become significant and the dominant feature ot the specific heat. If the temperature is below a critical value, the neutrons can form pairs and become superfluid. In that case, their contribution to the specific heat completely drops and becomes negligible. Therefore, studying cooling rates in this regime of the crust can lead to constraints on superfluid properties, which is an area of research in itself (Chamel, 2017; Sauls et al., 2020; Andersson, 2021).

The thermal conductivity also depends on the microscopic properties of the neutron star as its set by the scattering rate of the heat carrying particles. A region with high scattering rate corresponds to a region where the thermal conductivity is low and thus delays the overall cooling rate of the star. Multiple collision types contribute to the total scattering rate, depending on the conditions in a region of the neutron star these include photon, electron-electron, electron-ion, electron-impurity, electron-phonon, diffusive neutron-neutron and neutron-nuclear cluster scattering (Potekhin et al., 2015).

1.2.4 Heat blanketing envelopes in cooling simulations

While the envelope contains a negligible amount of the neutron star mass and radius, it plays a critical role in constraining neutron star interior physics from observations of the surface temperature. In cooling simulations the neutron star interior structure and temperature evolution (Equations 1.4 and 1.5) is calculated up to the bottom of the envelope at $\rho_b = 10^8$ g cm^{-3}, at which they are coupled to pre-calculated analytic temperature relations for the envelope. This upper boundary is chosen such that the luminosity at the boundary is equal to the photon luminosity at the surface

$L(\rho_b) = L_s = 4\pi R^2 \sigma T_s^4$. Temperature relations from fits to envelope numerical simulations can then be used to map the surface temperature (T_s) to the corresponding temperature at the bottom of the envelope (T_b), commonly referred as the T_b-T_s relationship, for a given envelope composition parameterized by the light element column density y_{light} (g cm^{-2}) (Potekhin et al., 1997).

Using T_b-T_s relations for the envelope is common in stellar evolution codes and has significant numerical advantages. As the conductive properties of the envelope depend on its composition, the accuracy of the technique depends strongly on the chemical composition of the envelope, which is often unknown (see Beznogov et al. 2021 for a recent review on envelope physics). If more lighter elements (such as H and He) are present in the envelope, the thermal conductivity of the envelope can increase strongly, leading to a higher surface temperature for the same T_b, as is illustrated in Figure 1.2. When lighter elements are present in the envelope, the electron-ion collisions become less frequent and the thermal conductivity of the envelope increases. In this case, the observed X-ray temperature will be larger while the interior temperature remains the same Therefore, it is important to use the appropriate envelope composition in order to infer the interior temperature of the neutron star from its surface temperature.

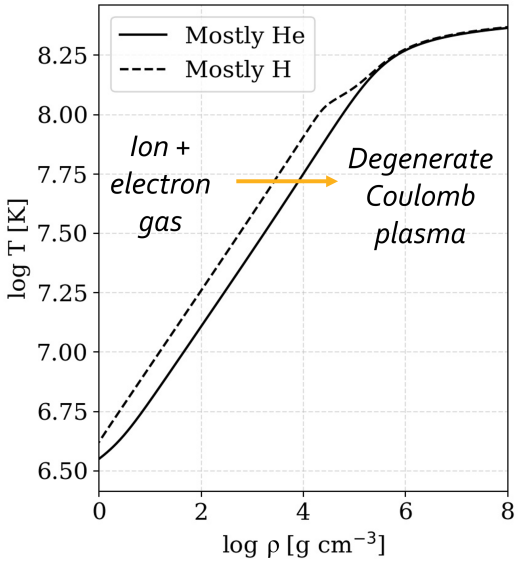

FIGURE 1.2: Illustrative envelope temperature profiles for an envelope composed predominantly of helium (solid) and hydrogen (dashed). A strong temperature gradient is observed as the envelope transitions from an ion - electron gas at low densities to a plasma where the electrons are fully degenerate. Both temperature profiles correspond to the same temperature at the bottom of the envelope (T_b), but have different surface temperatures illustrating the importance of the envelope composition.

Another crucial assumption for this technique is that the envelope composition does not change over the time, as that would change the required T_b-T_s relation. However, it has been shown that the diffusive tail of the light element layer in the envelope can

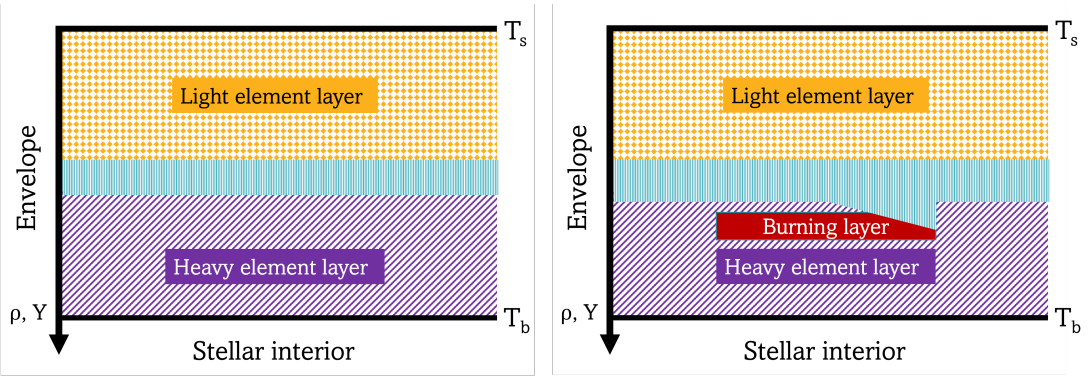

FIGURE 1.3: Illustrative representation of neutron star envelopes which are highly stratified due to strong gravity such that heavier elements sink down to deeper layers. The typical envelope model includes a transition density layer at which it transitions from a light to heavy element composition. The right image represents a model used in this thesis, where the possibility of nuclear burning of the diffusive tails of light elements is included in the envelope models.

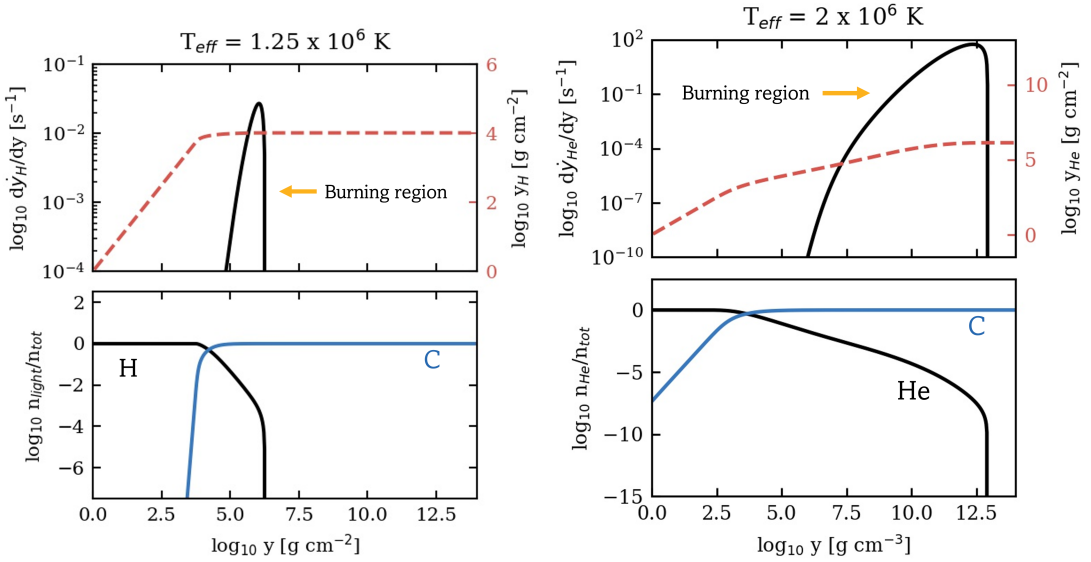

FIGURE 1.4: Example envelope composition profiles (bottom) for a H-C (left) and He-C (right) composition. The top plots show the nuclear burning rates at a given depth (black) and illustrates how the burning takes place in diffusive tail of the light element component. Additionally, the integrated total light element column density is shown (red, dashed).

reach depths where the temperature and density are high enough to start nuclear burning (see illustrative model of the envelope in Figure 1.3). The burning rates depend on the composition and temperature of the envelope and could drive diffusive downwards currents of lighter elements, potentially significantly changing the envelope composition over time. In Figure 1.4 I show example composition profiles of the envelope with the burning region indicated In Chapter 2 and 3 we investigate the effect that including diffusive nuclear burning has on the composition of the envelope and for the study of neutron stars using cooling simulations.

1.3 A post-merger remnant from gravitational wave data

1.3.1 Neutron star tidal deformability

The detection of the first merging binary neutron stars (GW170817) consisted of a premerger GW signal from the binary neutron star inspiral (Abbott et al., 2017e, 2018a), where the two orbiting neutron stars emit gravitational waves as they loose orbital energy and angular momentum and eventually collide. The premerger signal carries information on the neutron star tidal deformation as the tidal interactions accelerate the inspiral, which can be used to infer physical properties of the star and equation of state (see Chatziioannou 2020 for a recent review).

The tidal deformability, which describes how easily the neutron star is deformed when subjected to an external field, can be defined as

$$\lambda \equiv -\frac{Q_{ij}}{\epsilon_{ij}} \tag{1.6}$$

where Q_{ij} is the induced quadrupole and ϵ_{ij} is the perturbing tidal field. In relation to the dimensionless tidal Love number k_2, the strong dependence on the neutron star radius R in the tidal deformability becomes apparent as

$$\lambda = \frac{2}{3} k_2 R^5 \tag{1.7}$$

where both k_2 and R depend on the internal structure of the star (Hinderer et al., 2010).

1.3.2 Measuring the tidal deformability

For a binary coalescence, the dominant contribution to the time-varying quadruple moment that sets the gravitational wave emission is the binary orbital dynamics (Blanchet, 2014). The most prominent observable effect from the tidal deformation in neutron stars, is due to the additional induced quadruple moment affecting the binding energy of the system and increasing the rate of energy lost as gravitational wave emission.

The change in binary orbital frequency over time (df/dt) is directly related to the the rate of energy emission as gravitational waves (dE/dt) and the systems binding energy. These terms are generally expressed as post-Newtonian expansions (i.e., a series expansion in small binary velocity compared to the speed of light) (Blanchet,

2014). To leading order, the leading-order quadruple is proportional to $\sim r^2$, where r is the distance between the binary components. The tidal correction relative to the leading order term is proportional to $\sim r^{-5}$ (Chatziioannou, 2020). It can be seen that the effect of tidal interactions is small at early stages of the inspiral when the component separation is large. Conversely, as the separation decreases as the binary components spiral in, the effect from tidal interactions becomes stronger which explains why their imprint is primarily seen in the late-stages of the emitted gravitational wave signal.

As the system will loose energy faster due to tidal interactions, the inspiral phase speeds up and the overall system merges faster than a similar system without tidal interactions. This effect is more prominent for larger tidal deformabilities leading to a shorter inspiral phase. Therefore, larger tidal deformabilities (corresponding to lower mass neutron stars) are easier to measure from the inspiral gravitational wave signal.

1.3.3 Post-merger phase

Predictions describing how the post-merger remnant evolves after the collision of the two neutron stars, depend on their masses and the neutron star equation of state. For the majority of realistic equation of states and neutron stars with usual binary masses, a quasi-stable hypermassive neutron star supported by differential rotation and thermal effects is expected to be formed which may survive for ∼10-1000 ms (Baumgarte et al., 2000). During this time, the remnant is expected to emit gravitational waves in frequencies between 1500 and 4000 Hz (e.g., Ruffert and Janka 2001; Shibata 2005) which is an interesting range for current and future gravitational wave detectors. For neutron star mergers with even lower binary masses and for some equation of states, is it also possible for a supra-massive or stable neutron star remnant to form (e.g., Margalit and Metzger 2017).

Detection of a post-merger signal would probe a different regime of the NS equation of state, as the more massive remnant will have higher maximum densities than its pre-merger components (Ruffert and Janka, 2001). This would allow for gaining insights in possible high-density phenomena such as phase-transitions (e.g., Shibata et al. 2005). Also, detection of such a post-merger remnant gravitational wave signal would allow us to finally distinguish between the different predicted scenarios for the remnant.

The exact physics governing the post-merger phase is not well understood and numerical simulations are complicated by effects such as thermal effects, turbulences, magnetohydrodynamical instabilites and possible phase transitions. Despite the uncertainties in the simulations, they make a robust empirical prediction for the frequency of the peak of the post-merger spectrum which correlates with properties of

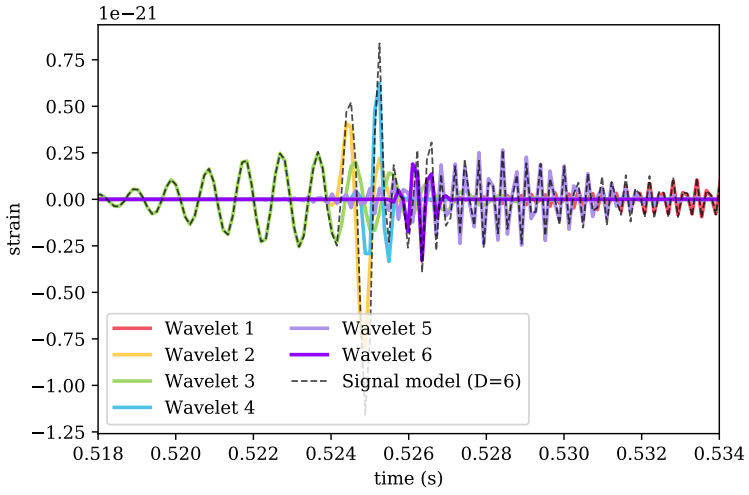

FIGURE 1.5: Illustration of the decomposition of a `BayesWave` gravitational wave model in the time-domain (plotted at arbitrary times). The signal model is the sum of a variable number of sine-Gaussian (wavelet) functions (D), in this case it is composed of 6 wavelets.

the neutron star, such as its radius and tidal deformability (e.g., Bauswein and Janka 2012; Bauswein et al. 2012). These empirical relations can be applied to constrain the neutron star equation of state using a measurement of the dominant post-merger spectrum feature.

1.3.4 Full BNS signal analysis

A particular challenge in gravitational wave astronomy is distinguishing weak signals of which the morphology is not well understood, such as a possible post-merger signal after a binary neutron star merger, from noise. In those cases, it is advantageous to use a search method that makes minimal assumptions about the signal morphology. `BayesWave`[1] uses a Bayesian multi-component approach that simultaneously characterizes the noise (including non-stationary and non-Gaussian features), as well as instrumental glitches and short-duration sources of gravitational waves. The key idea behind `BayesWave` is that all these components with complex morphologies can be reconstructed as a combination of a variable dimension of 'simple' sine-Gaussian (wavelet) functions (see illustration in Figure 1.5). The parameters defining the shape of these wavelet functions as well as the number of wavelet functions are found using a Reversible Jump Markov Chain Monte Carlo algorithm.

As the morphology of the post-merger signal is not well understood, it is very suitable to model with minimal assumptions about the signal. This method has been shown to be successful in recovering simulated post-merger signals in gravitational waves data (Chatziioannou et al., 2017) and has been applied to the observed binary neutron star

[1] `BayesWave` is open source and maintained here: `https://git.ligo.org/lscsoft/bayeswave/`

merger GW170817 for which no post-merger signal was recovered (Abbott et al., 2017b, 2019b). Using this method, only the high frequency ($\sim 1024 - 4096$ Hz) data where the post-merger peak signal is expected is used while the lower frequency data containing the pre-merger phase of the signal is discarded.

So far only separate analyses of the GW signal have been performed focusing either on the pre-merger or on the post-merger part of the data Abbott et al. (2019b, 2017b, 2019a, 2020). This has the advantage that the computational cost is significantly reduced, as the post-merger signal requires a large sampling rate (typically 8194 Hz) but does not need a long data duration (e.g. ~ 1 s is enough) while the premerger analysis uses long duration data segments (typically ~ 64 s) but does not need a large sampling rate ($\sim$2048 Hz). Since no post-merger signal has been detected yet, treating the pre-merger separately did not bias the inferred binary parameters. However, with future detector upgrades a detection of a post-merger is anticipated, motivating the need for a full, coherent analysis method. In Chapter 4 we develop an analysis method that simultaneously analyzes the inspiral and post-merger part of the data using traditional matched-filtering techniques for the inspiral and wavelets for the post-merger features of the signal.

1.4 Thesis outline

The goal of my PhD research is to develop and use improved computational models and analyses methods to infer neutron star (NS) interior properties from current and future data sources. Specifically, I am using complimentary techniques using observations of electromagnetic and gravitational wave radiation that probe different layers of the neutron star and allow for inferring neutron star physics in different states. This thesis is organized into two parts looking at two extremes of the neutron star state as schematically shown in Figure 1.6. The first part focuses on improving and applying computational models of the thermal evolution of relatively undisturbed, cooling neutron stars. The second part of the thesis focuses on computing and applying new computationally techniques to analyze the hot, potentially short-lived remnant of the merger of two neutron stars using gravitational wave data.

Chapter 2 consists of detailed modeling of the neutron star thermal evolution which can be compared to multiwavelength observations of thermal radiation. One of the challenges affecting this technique, is to infer the interior temperature from observations of surface emission. The relation between surface temperature and interior temperature is set by heat conducting properties of a thin outer envelope. Most state of the art models consider the envelope to consist of chemically pure layers separated by narrow transition bands driven by diffusive mixing. However, it is possible for elements to diffuse to depths where the density and temperature are

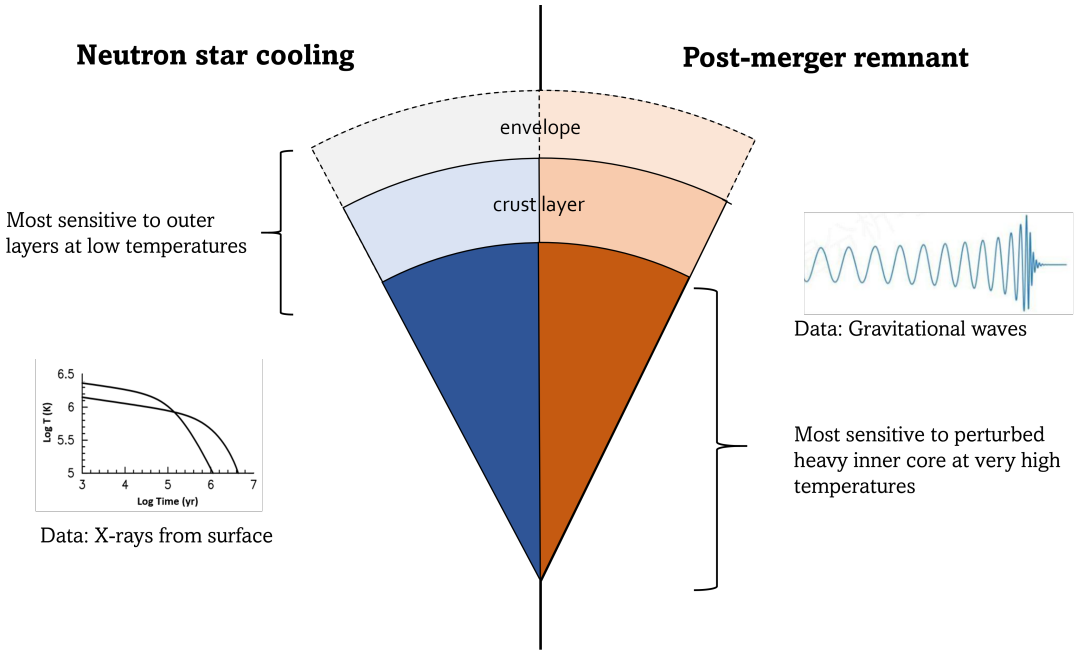

FIGURE 1.6: Schematic overview of the two complimentary parts of this thesis that focus on neutron star cooling studies (left) and binary neutron star post-merger studies (right). The studies probe different regimes of the neutron star state and use different data sources. In neutron star cooling studies we are most sensitive to the surface and outer layers of the stars which host a variety of exotic matter states. In studies of the post-merger remnant after a binary neutron star merger, we gain more information about the ultradense neutron star core in a hot, disturbed state.

sufficient to ignite nuclear burning, thus changing the chemical composition and heat conduction properties. I used numerical models of the envelope that include the effects of diffusive nuclear burning to create grids in the envelope parameter space. This allowed me to calculate new analytic relations between the neutron star surface and interior temperature for varying chemical compositions, and these relations can directly be used in neutron star cooling simulations. I performed simulations to model cooling of isolated neutron stars using these new temperature relations and allowing for a time-variable envelope composition instead of static envelope compositions. This work was published in Wijngaarden et al. (2019).

In Chapter 3 I consider another scenario where envelope models are important: The cooling of crusts of neutron stars in X-ray binaries after the crust has been lifted out of thermal equilibrium with the core during an accretion outburst. After the accretion halts, the crust cools back to thermal equilibrium over timescales of years to decades. Thus it is possible to cover the full relaxation time with observations and probe the deeper layers of the crust. All previous cooling codes assume that no heat is generated in the envelope and use the envelope temperature relations as a boundary condition for the interior cooling. Due to the typically large post-outburst temperatures, I investigated whether this assumption is valid. I coupled a cooling code to our envelope models and applied this to accreting systems. We found that the luminosity

generated in the envelope due to diffusive nuclear burning is not always negligible and that the burning rates are sufficiently large for the envelope composition to change significantly during the cooling phase. This work was published in Wijngaarden et al. (2020).

In Chapter 4 I turn my attention to another facet to study these same objects but in completely different conditions. The Chapters described above consist of computational models that can help to constrain neutron star properties by comparison with electromagnetic observations, while now I consider the study of a newly created neutron star using gravitational waves. A post-merger gravitational wave signal from a rapidly-rotating, strongly deformed neutron star could provide new probes of neutron star properties and nuclear matter, but its physics and signal morphology is much less understood than the pre-merger phase. Therefore, it is unreasonably computationally expensive to search for a post-merger signal by matching the data with (poorly understood) physical models. In Chapter 4 I develop and apply methods for constraining NS interior properties from possible detections of gravitational wave signals from NS post-merger remnants. I implement data analysis methods that use a model-agnostic approach while also utilizing more confident knowledge about the pre-merger phase.

Part I

Neutron star interior physics from electromagnetic observations

Chapter 2

Effect of diffusive nuclear burning on isolated neutron star cooling behaviour [1]

Abstract

A critical relation in the study of neutron star cooling is the one between surface temperature and interior temperature. This relation is determined by the composition of the neutron star envelope and can be affected by the process of diffusive nuclear burning (DNB), which occurs when elements diffuse to depths where the density and temperature are sufficiently high to ignite nuclear burning. We calculate models of H-He and He-C envelopes that include DNB and obtain analytic temperature relations that can be used in neutron star cooling simulations. We find that DNB can lead to a rapidly changing envelope composition and prevents the build-up of thermally stable hydrogen columns $y_H \gtrsim 10^7$ g cm^{-2}, while DNB can make helium envelopes more transparent to heat flux for surface temperatures $T_s \gtrsim 2\times 10^6$ K. We perform neutron star cooling simulations in which we evolve temperature and envelope composition, with the latter due to DNB and accretion from the interstellar medium. We find that a time-dependent envelope composition can be relevant for understanding the long-term cooling behaviour of isolated neutron stars. We also report on the latest *Chandra* observations of the young neutron star in the Cassiopeia A supernova remnant; the resulting 13 temperature measurements over more than 18 years yield a ten-year cooling rate of $\approx 2\%$. Finally, we fit the observed cooling trend of the Cassiopeia A neutron star with a model that includes DNB in the envelope.

[1] This chapter is based on publication Wijngaarden et al. (2019), see Publications and author contributions.

2.1 Introduction

The physical properties, such as composition and structure, of neutron stars (NSs) are still uncertain. Valuable information about the neutron star interior can be obtained by comparing observations of thermal radiation of cooling NSs with theoretical models of how a NS cools over time. Isolated NSs cool via neutrino and photon emission after they are formed hot in a supernova explosion (see review by Potekhin et al. 2015). Additionally, cooling models can be used to study the cooling of accretion-heated NSs in binary systems (for a recent review see Wijnands et al. 2017).

Observations of several isolated neutron stars suggest that the envelopes of younger pulsars ($<10^4-10^5$ yr) may consist of light elements such as hydrogen or helium, whereas the envelopes of older pulsars ($>10^5$ yr) may be composed of heavier elements (see, e.g., Yakovlev et al. 2001; Potekhin 2014). Evidence for the presence of carbon in a NS envelope was found for the first time using spectral observations of the young NS in Cassiopeia A (Ho and Heinke, 2009). Since then, several other NS spectra are found to be fit by carbon atmosphere spectra (see De Luca 2017, and references therein). These observations indicate a possible evolution of the composition of the envelope through nuclear burning, after the initial composition is set by fallback material onto the NS after a supernova explosion. The NS in Cassiopeia A (we hereafter refer to the NS as Cas A) is an intriguing object not only because of the chemical composition of its envelope, but also because surface temperatures derived from X-ray observations suggest that the source might be cooling rapidly (Heinke and Ho 2010; c.f. Posselt et al. 2013). Constraining cooling models with multiple temperature observations can provide fundamental knowledge on the properties of the NS interior, such as critical temperatures for superconductivity and superfluidity (e.g., Shternin et al. 2011; Page et al. 2011; Ho et al. 2015). Other possibilities include fast or rotationally induced neutrino cooling (Negreiros et al., 2013; Taranto et al., 2016), magnetic field decay (Bonanno et al., 2014), slow thermal relaxation (Blaschke et al., 2012, 2013), stellar fluid oscillations (Yang et al., 2011), and transition to axions or quark matter (Noda et al., 2013; Sedrakian, 2013; Leinson, 2014; Hamaguchi et al., 2018).

A major source of uncertainty in cooling models is the chemical composition of the thin outer envelope which extends from just below the surface down to a boundary mass density of $\rho_b \sim 10^8 - 10^{10}$ g cm^{-3}. The envelope acts as a thermal insulator due to its relatively poor thermal conductivity and thus sets a strong temperature gradient between the surface temperature T_s and the interior temperature at the bottom of the envelope T_b [$\equiv T(\rho_b)$]. The composition affects the thermal conductivity and thus how transparent the envelope is to the heat flux. This means that a measured surface temperature can correspond to different interior temperatures depending on the composition of the envelope.

Several models of the heat blanketing envelope have been developed after the work by Gudmundsson et al. (1983), who calculated the thermal structure and conductivity of an iron envelope. Later models went beyond the single element envelope and considered the envelope as multiple shells of light elements (H, He, C) separated by an abrupt boundary (Potekhin et al., 1997, 2003). These models have subsequently been used in many interior studies. Most recently, Beznogov et al. (2016) developed new models for binary ion mixtures (H-He, He-C and C-Fe) where ion diffusion is taken into account but the possible effect of nuclear burning due to diffusion is not included.

Numerical calculations of the neutron star envelope including the effects of diffusive nuclear burning (DNB) have been performed by Chang and Bildsten (2003, 2004); Chang et al. (2010). The key idea of DNB is that the tail of the light element layer may diffuse down to depths where the density and temperature is large enough to ignite burning. The depletion of the light elements at this depth can subsequently drive a light element current from the surface to the burning layer. In Chang and Bildsten (2003) and Chang and Bildsten (2004), a H-C envelope was modelled, and it was found that diffusive nuclear burning can readily consume all of the initial hydrogen in the NS envelope within $\sim 10^5$ yr. Motivated by the detection of carbon on the surface of Cas A, Chang et al. (2010) calculated models for DNB in He-C envelopes and found that any primordial helium is likely depleted during the hot ($T_b > 10^8$ K), early evolution of the NS.

Typically, NS cooling models calculate the NS structure and thermal evolution from the centre out to the bottom of the envelope, where they are coupled to theoretical envelope temperature relations (between T_b and T_s) based on the envelope models above. Diffusive nuclear burning can change the composition of the envelope and hence its thermal conductivity. Therefore, it can affect the temperature relations between T_s and T_b. We use the work of Chang and Bildsten (2003, 2004); Chang et al. (2010) and calculate models including diffusive nuclear burning for H-He and He-C envelopes and provide analytic envelope relations (based on those presented in Beznogov et al. 2016) that can be used in cooling models. We use these temperature relations to investigate the effect of diffusive nuclear burning and time-variable envelopes on neutron star cooling. We also report on four new *Chandra* observations of the NS in Cassiopeia A resulting in the longest spanning (>18 years) dataset currently available. We then use the new temperature relations including DNB to model the observed cooling of Cas A.

2.2 Theory of diffusive nuclear burning

2.2.1 The physics of DNB

We summarize the key physics input for diffusive nuclear burning as calculated in Chang et al. (2010). For a full description of the physics we refer the reader to this work. The neutron star envelope, at $\rho < 10^8 - 10^{10}$ g cm^{-3}, has a thickness ($\sim$ a few hundred metres) that is much smaller than the stellar radius and can be well approximated as a plane-parallel layer with a constant downward gravitational acceleration g. The thermal profile is obtained by solving the heat diffusion equation for a constant flux in radiative equilibrium

$$\frac{\partial T}{\partial z} = -\frac{3\kappa\rho}{16T^3}T_s^4, \tag{2.1}$$

where z is the depth, T_s is the effective (non-redshifted) surface temperature and κ is the opacity. In the outermost layer of the envelope the radiative opacity is dominated by free-free absorption and Thomson scattering. We use the radiative opacities of Potekhin and Yakovlev (2001) and the analytic formulae for the conductive opacity given by Potekhin et al. (1999).

We consider the case where a trace ion species, hereafter referred to as 'light elements', is immersed in a fixed background 'heavy' ion species. The equation for hydrostatic equilibrium, which is dominated by electric and gravitational forces (see Chang et al. 2010) is used to obtain the number density profiles of ion species in the absence of nuclear burning. We require overall charge neutrality such that, $n_e = Z_{\rm light}n_{\rm light} + Z_{\rm heavy}n_{\rm heavy}$, where n_e is the electron number density, and $Z_{\rm light}$, $n_{\rm light}$, $Z_{\rm heavy}$ and $n_{\rm heavy}$ are the charge number and number density of the light (trace) and heavy (background) element species, respectively. The diffusive tail of the light element layer may extend deep down into the underlying heavier element layer. When lighter elements diffuse to depths where they are consumed by nuclear burning, the depletion of light elements at this depth induces a diffusive light element current from the surface to the burning layer. Because the timescale at which the local column density changes is much longer than the local nuclear timescale, a steady state approximation is used for the light element current induced by DNB ($\partial n_{\rm light}/\partial t = 0$). The light element current is defined as

$$J_{\rm light} = -\mathcal{D}\frac{dn_{\rm light}}{dz} + n_{\rm light}w_{\rm light} = n_{\rm light}v_{\rm light}, \tag{2.2}$$

where $v_{\rm light}$ and $w_{\rm light}$ are the light element relative velocity and drift velocity, respectively. The diffusion coefficient $\mathcal{D}$ is calculated as in Chang and Bildsten (2004) (similar to Brown et al. 2002)

$$\mathcal{D} \approx 10^{-3} \frac{A_{\rm light}^{0.1} T_6^{1.3}}{Z_{\rm light}^{1.3} Z_{\rm heavy}^{0.3} \rho_5^{0.6}} {\rm cm^2 s^{-1}}, \tag{2.3}$$

where $T_6 = T/10^6$ K and $\rho_5 = \rho/10^5$ g cm^{-3}. The atomic mass and charge number are denoted by A and Z, respectively. From the light element continuity equation and the definition of the current, one can derive the general case of the steady-state diffusion equation (which is derived in detail in Chang and Bildsten 2004)

$$\frac{\partial^2 f_{\rm light}}{\partial z^2} + \frac{\partial f_{\rm light}}{\partial z} \frac{\partial \ln n_{\rm light,0}}{\partial z} = \frac{f_{\rm light}}{\mathcal{D}\tau_{\rm light}}, \tag{2.4}$$

where $n_{\rm light,0}$ is the number density of the trace ion in the absence of nuclear burning obtained by solving the hydrostatic balance equations numerically, $f_{\rm light}$ is the correction factor to the number density due to nuclear burning (i.e., $f_{\rm light} \equiv n_{\rm light}/n_{\rm light,0}$) and $\tau_{\rm light}$ is the local lifetime against nuclear capture set by the local nuclear burning rate. With increasing depth and temperature, the ion scale height ($=[d\ln n_{\rm light,0}/dz]^{-1}$) rapidly becomes larger than the nuclear scale height ($=\sqrt{\mathcal{D}\tau_{\rm light}}$). Therefore, at the burning layer where the second term can effectively be dropped, Equation 2.4 represents a diffusion equation for $f_{\rm light}$ with a nuclear-driven source.

In our models, hydrogen can be consumed steadily by the proton capture reactions of the CNO-cycle and pp chain reactions. Helium is consumed by proton- and alpha capture reactions and the triple alpha process, although the reaction rates of the triple alpha process were found to be negligible in the parameter space of interest (see Chang et al. 2010) and thus they are ignored. Nuclear reaction rates and experimental reaction cross section values are obtained from the NACRE[2] (Angulo et al., 1999) and REACLIB[3] database compilations.

2.2.2 Model description

We model the neutron star envelope for ^{1}H-^{4}He and ^{4}He-^{12}C mixtures following the work of Beznogov et al. (2016), but including the physics of DNB. These envelope compositions were chosen to be able to compare our solutions with those found in Beznogov et al. (2016), but in reality other ions may be present in neutron star envelopes. In particular, one might consider a H-He-C envelope, with an intervening helium layer between the hydrogen and carbon components. The effects of such an envelope are strongly dependent on the assumed size of each layer. For simplicity in the present work, we limit our study of DNB to two component envelopes. We use a surface gravity of $g_s = 2.43 \times 10^{14}$ cm s^{-2} in our calculations, but the resulting

[2] http://pntpm.ulb.ac.be/Nacre/nacre.htm
[3] http://nucastro.org/reaclib.html

temperature relations at other surface gravities can be approximated using the temperature scaling relation by Gudmundsson et al. (1983).

We start our numerical calculations with a given surface temperature T_s and fraction $x_{\rm light}$ of the light element species present at the photosphere. The partial pressure $p_{\rm light}$ and column density $y_{\rm light}$ of the ion components at the photosphere is set by $x_{\rm light}$ as $y_{col} = y_{\rm light} + y_{\rm bkg} = x_{\rm light} \times y_{col} + x_{\rm bkg} \times y_{col}$. We then integrate the system of equations in Section 2.2.1 to obtain profiles for the internal temperature, ion densities and nuclear burning rates. A given photospheric light element fraction can correspond to different values of the total light element column density depending on the surface temperature. Therefore, we perform a large number of numerical calculations for varying T_s and $x_{\rm light}$ to obtain sufficient coverage in the resulting T_s-T_b-$y_{\rm light}$ parameter space that is used for fitting analytic functions.

We note that in the analytic temperature relations calculated by Beznogov et al. (2016), the amount of light elements is characterized by their effective transition density $\rho*$. This is an artificially chosen density at which a total mass of light elements is contained in the outer shell $\rho < \rho*$. Note that the light element column density $y_{\rm light}$ can be converted to $\rho*$ making use of the relevant equation of state.

2.2.3 Parameter ranges

We fit the numerical results of the models described above to obtain analytic $T_s - T_b$ relations that can be coupled to thermal evolution codes for the neutron star interior. Note that the $T_s - T_b$ relations are actually relations for T_s, T_b and $y_{\rm light}$ (or equivalently, $\rho*$) for a given surface gravity, i.e. $T_b(T_s, y_{\rm light})$. We are limited in our parameter ranges by both physical and numerical constraints. Physical constraints arise for example from the fact that the light element column density cannot be larger than the envelope-crust boundary column density y_b or smaller than the photospheric density. For the H-He envelopes, we let $\rho_b = 10^8$ g cm^{-3} and increase ρ^* up to $\sim 10^7$ g cm^{-3}. For He-C envelopes we use $\rho_b = 10^{10}$ g cm^{-3} and let ρ^* vary up to $\sim 10^8$ g cm^{-3}. The surface temperature range for all compositions is 0.3×10^6 K$< T_s < 3 \times 10^6$ K, such that it encompasses the observed temperatures for both isolated neutron stars and transient neutron stars in quiescence.

2.3 Burning rates

The reaction rates due to DNB depend on the composition of the envelope, the amount of burning material available and the temperature in the burning region. Using the nuclear reactions that consume hydrogen and helium discussed in Section 2.2, the total burning rate is

$$\zeta_{\rm DNB} = \frac{y_{\rm light}}{\tau_{\rm light,col}} = \int dz \frac{A_{\rm light} n_{\rm light} m_p}{\tau_{\rm light}(n_{\rm light}, n_{\rm sub}, T)}, \tag{2.5}$$

where $A_{\rm light}$ is the atomic number of the light element species, $\tau_{\rm light,col}$ is the characteristic time for the total integrated light element column to be consumed and $\tau_{\rm light}$ is the local lifetime of the light element species due to DNB. The local number densities of the α- or p-capturing substrate are denoted by $n_{\rm sub}$.

In Figure 2.1 we plot the resulting burning rates for several values of $y_{\rm light}$ for both the H-He and He-C envelope models. For hydrogen burning (dashed lines in Figure 2.1), we see that the burning rates are more sensitive to the amount of material available for burning (as characterised by the hydrogen column density) than to the temperature (characterised by the surface temperature). The strong dependence on column density is not surprising, as the burning region is located in the exponentially suppressed diffusive hydrogen tail. Interestingly, Figure 2.1 shows that DNB can alter the hydrogen column density significantly for the entire range of plotted T_s. Especially for intermediate and large hydrogen column densities ($y_{\rm H} \gtrsim 10^5$ g cm^{-2}) the hydrogen column density lifetime is $<< 10^4$ yr. It is clear that the helium burning rates drop off rapidly with decreasing temperatures, which is expected as the dominant burning processes are the highly temperature sensitive α-capturing reactions. We find that the helium burning rate effectively switches off for $T_s \lesssim 10^6$ K. At temperatures $T_s \gtrsim 10^6$ K the burning rate is sufficiently large to change the helium column density within years.

Using the burning rates calculated above, we illustrate the evolution of the light element column over time for a fixed surface temperature in Figure 2.2 and 2.3 for hydrogen and helium, respectively. We use an initial hydrogen column density of $y_{\rm light} = 10^6$ g cm^{-2} and vary the surface temperature between $0.4 \times 10^6 < T_{\rm s} < 3.2 \times 10^6$ K. For all temperatures, the sharpest drop in hydrogen column density occurs within the first ~ 10 years due to the strong density dependence of the burning rate. After that time, the column decreases more gradually. For helium burning we use an initial helium column density of $y_{\rm light} = 10^{10}$ g cm^{-2} and vary the surface temperature between $2 \times 10^6 < T_{\rm s} < 3.2 \times 10^6$ K. It is clear that the helium column change is more strongly coupled to the temperature and that for $T_{\rm s} \lesssim 10^6$ K the column does not change significantly within 100 years. At higher temperatures, Figure 2.3 illustrates that DNB can even substantially alter the helium column on relatively short timescales of $< 10 - 100$ years.

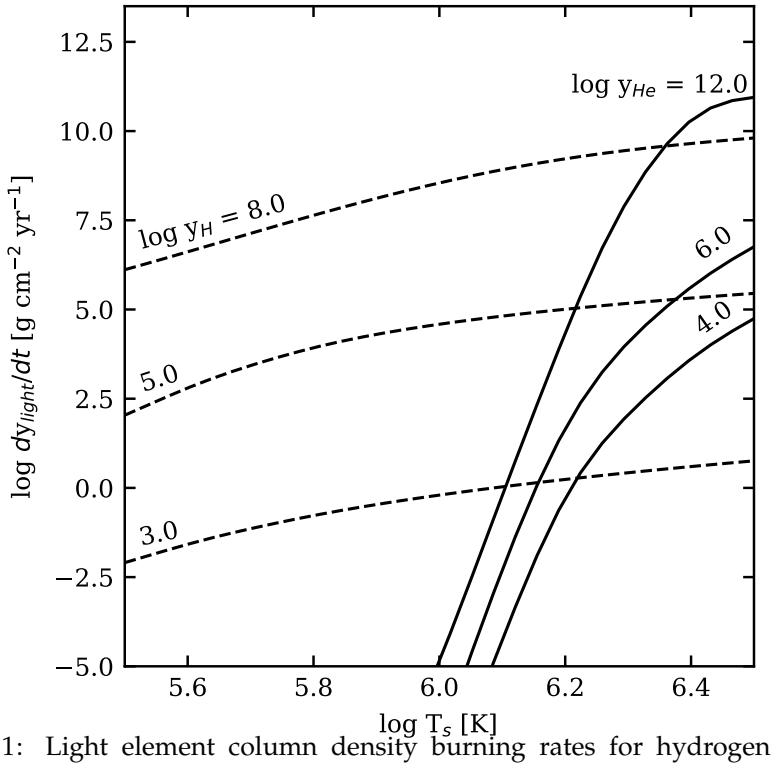

FIGURE 2.1: Light element column density burning rates for hydrogen burning (dashed) and helium burning (solid). For each line the value of log y_{light} is indicated. Note that log dy_{light}/dt is the same as log ζ_{DNB}; see Equation 2.5.

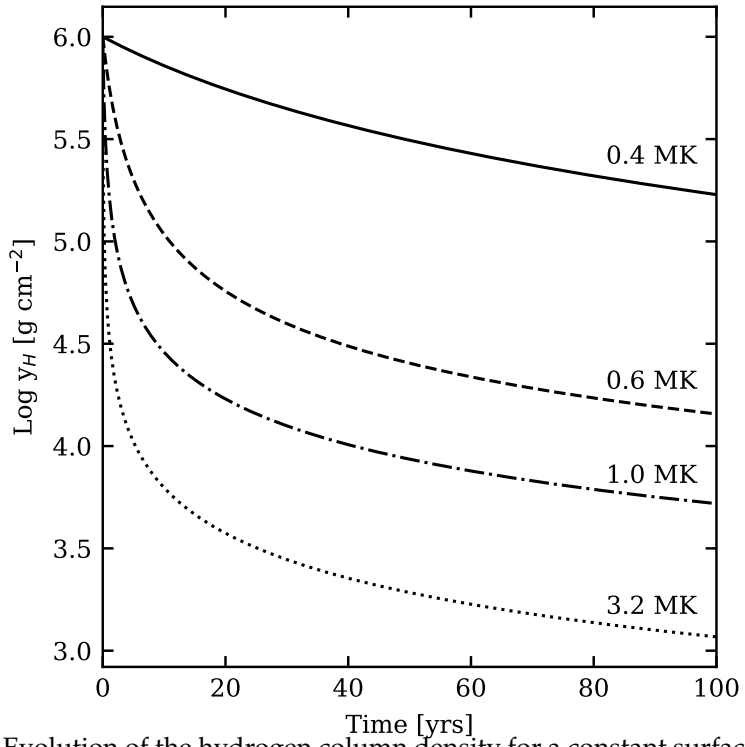

FIGURE 2.2: Evolution of the hydrogen column density for a constant surface temperature [(0.4, 0.6, 1.0, 3.2)$\times$ 10^6 K] and initial y_H = 10^6 g cm^{-2}.

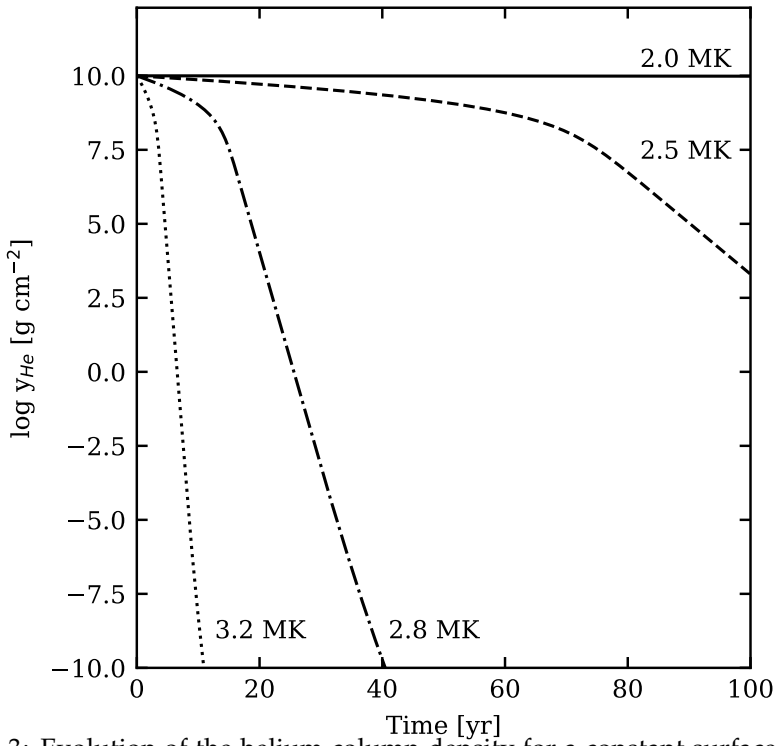

FIGURE 2.3: Evolution of the helium column density for a constant surface temperature [(2, 2.5, 2.8, 3.2)$\times$ 10^6 K] and initial $y_{\rm He}$ = 10^{10} g cm^{-2}.

2.4 Static envelope temperature relations

We present analytic temperature relations that include diffusive nuclear burning in the envelope. The fit functions and parameters for each of the mixtures are given in Appendix B.1. In this section, we review their main characteristics and implications.

The temperature relations are most sensitive to the composition in a small region in the envelope which is commonly referred to as the sensitivity strip. The location of the sensitivity strip in the envelope is determined by where the radiative κ_r and conductive κ_c opacity are comparable, i.e., $\kappa_r \approx \kappa_c$. Even a small amount of light elements in this region can have a large effect on the resulting temperature profile in the envelope. The composition outside of the sensitivity strip barely affects the temperature relation. Therefore, in order to understand the effect that DNB has on the temperature relations, $T_{\rm b}(T_{\rm s}, y_{\rm light})$, we will consider for which values of $y_{\rm light}$ the composition in the sensitivity strip is altered by DNB.

In Figure 2.4, we show the $T_{\rm s}$-$y_{\rm H}$ parameter space for a H-He envelope; it is important to note that what is illustrated in Figures 2.4 and 2.5 is light element column density, not total column density. For small hydrogen columns and high surface temperatures (lower right blue region), the bulk of the hydrogen remains close to the surface and does not reach the sensitivity strip; thus the temperature relations resemble those of a pure helium envelope and are only a function of $T_{\rm s}$. The region enclosed by solid lines

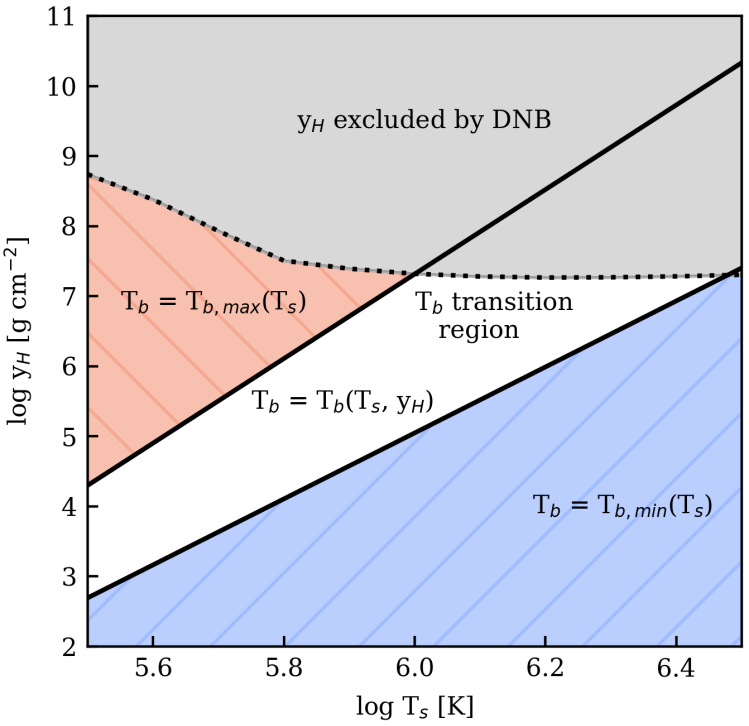

FIGURE 2.4: T_s-$y_{\rm H}$ parameter space for the H-He envelope model with $\rho_b = 10^8$ g cm^{-3}. The solid lines enclose the region where the boundary temperature is sensitive to the hydrogen column density, i.e. $T_{\rm b} = T_{\rm b}(T_s, y_{\rm H})$. In this region, the boundary temperature transitions from that corresponding to a pure hydrogen envelope (red), where $T_{\rm b} = T_{\rm b,max}(T_{\rm s})$, to that of a pure helium envelope (blue), where $T_{\rm b} = T_{\rm b,min}(T_{\rm s})$. Hydrogen column densities outside of the transition region have negligible effect on the $T_{\rm b}$-$T_{\rm s}$ relations. The grey shaded region shows the parameter space that is excluded due to DNB (see text). Note that at large $T_{\rm s}$, some excluded hydrogen columns overlap with the transition region. This means that the spread in $T_{\rm b}$ is smaller when DNB is taken into account.

indicates temperatures and hydrogen columns which affect the composition of the sensitivity strip and thus the resulting $T_{\rm b}$. For large hydrogen columns (red region), hydrogen makes up most of the composition in the sensitivity strip, and increasing the column further does not change the composition inside the sensitivity strip. In this case, the temperature relations resemble a pure hydrogen envelope and are unaffected by the size of the hydrogen column. Hydrogen DNB occurs in the diffusive tail and the reaction rates are highly sensitive to the amount of material available at the burning depth (see Section 2.3). Therefore, DNB does not impact the composition distribution in the sensitivity strip for small hydrogen columns. For large H columns ($y_{\rm H} \gtrsim 10^7$ g cm^{-2}), the time for hydrogen to diffuse down to the burning depth is shorter than the rate at which it is consumed. In other words, the burning reactions are not limited by the diffusion time but by the nuclear reaction rate. Thus, hydrogen at the burning depth is consumed rapidly, which prevents the build up of large thermally stable hydrogen columns ($y_{\rm H} \gtrsim 10^7$ g cm^{-2}).

In Figure 2.5, we show $T_{\rm s}$-$y_{\rm He}$ for the He-C envelope model. Diffusive helium burning

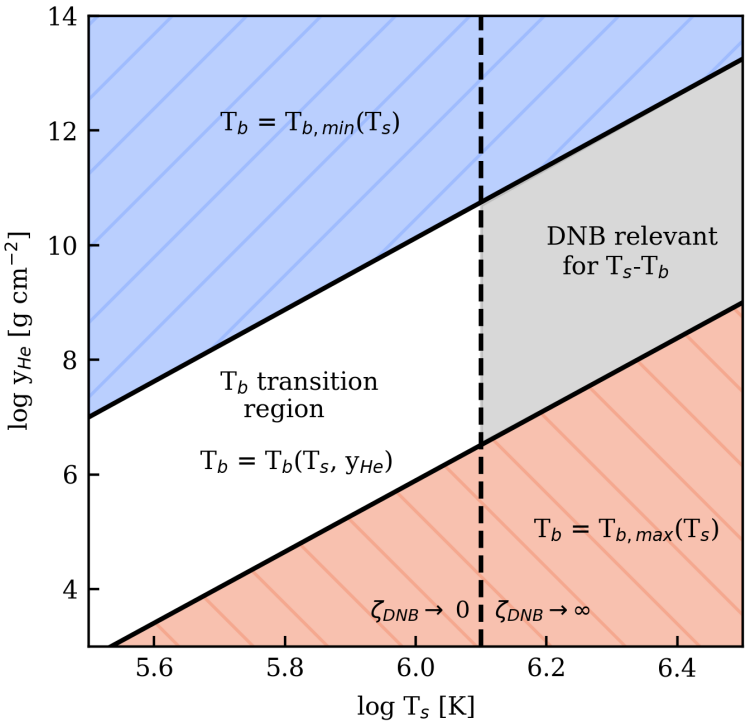

FIGURE 2.5: T_s-$y_{\rm He}$ parameter space for a He-C envelope model with $\rho_b = 10^{10}$ g cm^{-3}. The solid lines enclose the region where the boundary temperature is sensitive to the helium column density, i.e. $T_{\rm b} = T_{\rm b}(T_s, y_{\rm He})$. In this region, the boundary temperature transitions from that corresponding to a pure helium envelope (blue), where $T_{\rm b} = T_{\rm b,min}(T_{\rm s})$, to that of a pure carbon envelope (red), where $T_{\rm b} = T_{\rm b,max}(T_{\rm s})$. Helium column densities outside of the transition region have a negligible effect on the $T_{\rm b}$-$T_{\rm s}$ relations. The dashed line indicates the surface temperature above which the DNB rate $\zeta_{\rm DNB}$ is significant. For helium column densities in the transition region where $\zeta_{\rm DNB}$ is significant (highlighted in grey), DNB alters the $T_{\rm b}$-$T_{\rm s}$ relations considerably.

takes place deep in the envelope (10^8 g cm$^{-2} \lesssim$ y $\lesssim 3 \times 10^{12}$ g cm^{-2}) and reaction rates are highly sensitive to the temperature (see Section 2.3). At $T_{\rm s} < 1.25 \times 10^6$ MK (denoted by the dashed line in Figure 2.5), burning rates are too low to affect the envelope composition and temperature relations. Above this surface temperature, DNB prevents the helium column from extending deeper into the envelope beyond the burning depth. As a result, more helium resides at lower densities in the sensitivity strip compared to the case when DNB is not taken into account. Thus a He-C envelope is more transparent to heat flux with DNB for sufficiently high temperatures.

Figures 2.4 and 2.5 emphasise that the effect of hydrogen on the conductivity of the envelope is very different than that of helium, as first noted by Beznogov et al. (2016). Typically, increasing the amount of light elements leads to a higher thermal conductivity of the envelope, such that for a given $T_{\rm s}$ the corresponding $T_{\rm b}$ is lower. An exception to this is the case of hydrogen, which has the opposite behaviour, as its larger radiative opacities results in a lower thermal conductivity. Therefore, increasing the amount of hydrogen, leads to a larger $T_{\rm b}$.

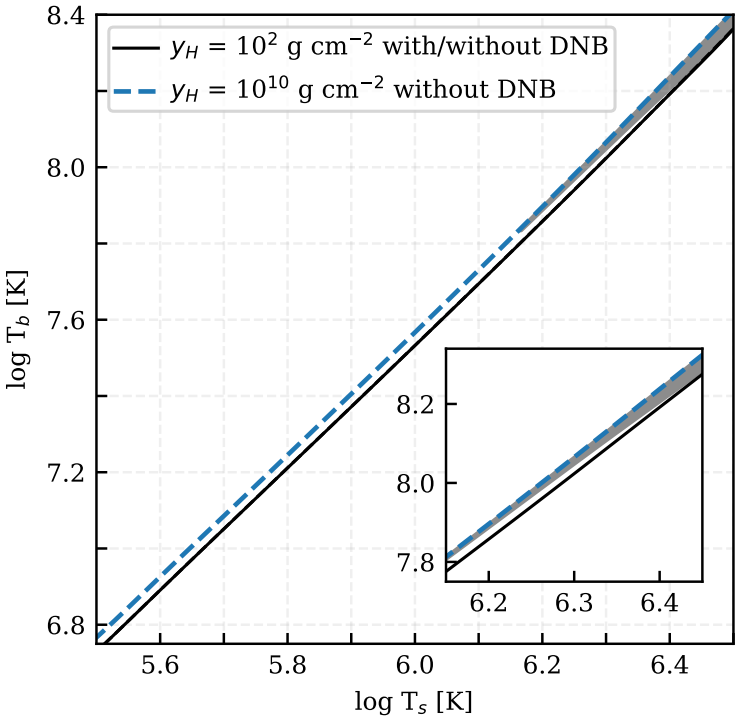

FIGURE 2.6: The T_s-T_b relation for the H-He envelope with $\rho_b = 10^8$ g cm^{-3}. The solid line corresponds to both the model with and without DNB. The blue dashed line corresponds to an hydrogen column of 10^{10} g cm^{-2} when DNB is not included. The grey shaded region (see also inset) corresponds to the grey region in Figure 2.4 and shows the region in the $T_s - T_b$ relations where DNB is reaction rate limited (see text). Note that the spread in T_b for a H-He envelope is small compared to other compositions and becomes even smaller at large temperatures due to DNB.

The resulting $T_s - T_b$ relations for a H-He envelope including DNB are shown in Figure 2.6. For comparison, the temperature relations where DNB is neglected are also plotted. It is clear that including DNB has a negligible effect on the thermal conductivity of the H-He envelope as the resulting temperature relations are indistinguishable. However, including DNB in the H-He envelope leads to an upper limit on the size of the hydrogen column, making the range in T_b for a given T_s smaller than when DNB is ignored (the difference is indicated by the shaded region in Figure 2.6).

For a He-C envelope, the temperature relations are shown in Figure 2.7. As expected, these relations are unaffected by DNB at low temperatures. At higher temperatures ($T_s > 2 \times 10^6$ K; see Section 2.3) where helium burning is significant, DNB changes the temperature relations. In particular, for a given surface temperature, the boundary temperature is lower with DNB; equivalently, the same boundary temperature corresponds to a higher surface temperature. The He-C envelope is a better conductor and more transparent to the heat flux at large temperatures when DNB is taken into account since helium burning at large depths implies a greater helium abundance at shallower depths (in the sensitivity strip) for a given helium column (see above).

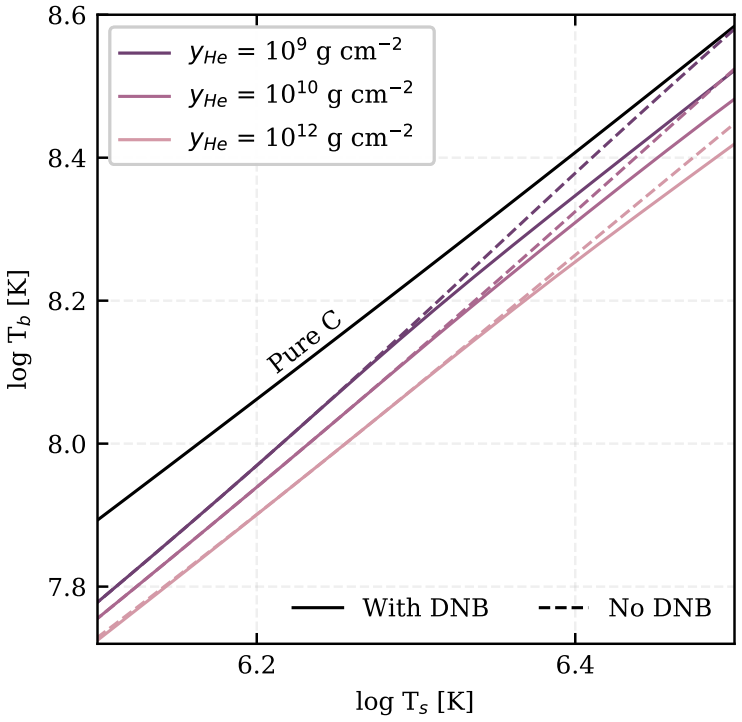

FIGURE 2.7: Comparison of the T_s-T_b relation including DNB (solid) and without DNB (dashed) for the He-C envelope with $\rho_b = 10^{10}$ g cm^{-3}. A small range in T_s is shown since DNB is negligible at lower T_s (see Figure 2.5).

2.5 Time variable envelopes

Envelope temperature relations, such as the ones presented in Section 2.4, are used to calculate the NS temperature evolution while assuming that the envelope composition does not change (see, e.g., Potekhin et al. 2015 for a review of NS cooling). The burning rates in Section 2.3 indicate that modelling the envelope with a fixed composition over the course of the NS lifetime can be inaccurate. In this section, we explore time-variable envelope compositions due to DNB and their effect on NS temperature evolution. Instead of calculating a cooling curve with a fixed envelope composition, we use a fixed initial composition and evolve both heat diffusion in the star and envelope composition.

2.5.1 Cooling model

We model the thermal evolution of the NS using the general relativistic cooling code *NSCool*[4] (Page, 2016) for a NS with mass $M = 1.4\ M_{\rm Sun}$ and radius $R = 11.5$ km, which is consistent with a surface gravity $g_s = 1.8 \times 10^{14}$ cm s^{-2}. We use the nuclear A18+δv+UIX* equation of state (Akmal et al., 1998). Neutron superfluidity is taken into account using the 'SFB' $^1S_0$ gap model Schwenk et al. (2003), and proton

[4] An older version of this code is available at `http://www.astroscu.unam.mx/neutrones/NSCool/`

superconductivity is modelled using the 'CCDK' $^1S_0$ gap model (Chen et al., 1993). Possible pairing of neutrons in the spin triplet-state is not considered in this section but will be explored in the next one. While uncertainties in the superfluid properties of neutrons and superconducting properties of protons affect the cooling curves, we do not expect this to significantly change our results, as we use the same assumptions for all cooling curves and are only interested in the differences resulting from varying envelope compositions. The energy transport and conservation equations are solved up to the outer boundary defined as the bottom of the envelope: $\rho_b = 10^8$ g cm^{-3} and $\rho_b = 10^{10}$ g cm^{-3} for a H-He and He-C envelope, respectively. We use the $T_s - T_b$ models including DNB as described in Section 2.4 to relate the surface temperature and the boundary temperature.

2.5.2 Variable envelope composition

We use fits of the numerical burning rates to calculate the evolution of the light element column density $y_{\rm light}(t)$. In addition to solving equations describing the interior thermal profile, we evaluate the change in $y_{\rm light}$ at each time step. We model two competing processes that affect the evolution of column density. On the one hand, the amount of light elements decreases due to DNB when there is enough material available to burn and the temperature is sufficiently large. On the other hand, the amount of light elements increases due to accretion from the interstellar medium (ISM) or from a companion star. A mass accretion rate $\dot{M}_{\rm acc}$ is equivalent to an increasing light element column density $\dot{y}_{\rm acc}$ via

$$\dot{y}_{\rm acc} \approx \dot{M}_{\rm acc} \, \frac{g_s}{4\pi G M}. \tag{2.6}$$

The total change in $y_{\rm light}$ is obtained by combining the decreasing effect from DNB and the increasing effect from accretion (from the ISM). Thus the change in $y_{\rm light}$ in a time interval dt is

$$\Delta y_{\rm light} = \frac{dy_{\rm light}}{dt} dt = (\, \dot{y}_{\rm acc} - \zeta_{\rm DNB} \,)\, dt, \tag{2.7}$$

where $\zeta_{\rm DNB}$ is calculated as in Equation 2.5 for a grid of $T_s - T_b - y_{\rm light}$ and interpolated between gridpoints. Note that our steady state approximation is not strictly valid at early, hot stages immediately after NS birth, when the NS cools faster than diffusion through the envelope. A calculation of neutron star interior temperature evolution and diffusion in the envelope is needed for a self-consistent solution, which is beyond the scope of this work.

2.5.3 ISM accretion rates

Whether the NS is accreting from the ISM and at what rate depends on the local density of the ISM and properties of the NS (e.g., its mass, radius and magnetosphere; see, e.g. Blaes and Madau 1993). Furthermore, when considering central compact objects (CCOs), the NS is within a supernova remnant, for which the distribution and composition differ from the ISM (see, e.g. Vink 2012; Lopez and Fesen 2018).

Following classical accretion physics (see, e.g. Shapiro and Teukolsky (1983)), the accretion rate onto the NS surface can be described using the Bondi formula

$$\dot{M} = \frac{4\pi\lambda_a G^2 M^2 \rho_{\rm ISM}}{\left(v^2 + c_s^2\right)^{3/2}}, \tag{2.8}$$

where $\rho_{\rm ISM}$ is the local interstellar mass density, v is relative speed between the NS and ISM, and c_s is the speed of sound. The parameter λ_a can vary between 0.25 and 1 (Hunt, 1971). Here we take c_s = 10 km s^{-1} and λ_a = 1.

For simplicity, we consider three constant values of mass accretion rate for our models, i.e., we do not account for potential short time scale variability in the accretion rate. In the present work, we consider only the effect of a light element column density on the resulting cooling curve and ignore the effect of heat generated by nuclear reactions (see Section 2.7). We can obtain the highest accretion rate considered here ($\dot{M}_{\rm acc} = 10^{-15}\ M_{\rm Sun}$ yr^{-1}) by using Equation 2.8 for an isolated 1.4 $M_{\rm Sun}$ NS with $v = 20$ km s^{-1} and ISM number density $n = 1$ cm^{-3}. Lower accretion rate values (10^{-18} and $10^{-21}\ M_{\rm Sun}$ yr^{-1}) are possible for higher velocities (e.g., 200 or 1000 km s^{-1}) and/or lower ISM densities. These low accretion rates are further motivated by possible suppression of the Bondi accretion rate such that NSs accrete at lower rates (see, e.g., Illarionov and Sunyaev 1975; Perna et al. 2003; Krumholz et al. 2005).

2.5.4 Cooling curves

For H-He envelopes, we show the cooling curves and hydrogen column evolution in Figure 2.8 for an initial hydrogen column density of $y_{\rm H} = 10^6$ g cm^{-2}; note that the column evolution curves here and in Figures 2.9 and 2.10 are smoothed (averaged over short time intervals) in order to remove short timescale fluctuations due to the finite time steps of the numerical simulations. The first conclusion that is clear from this plot is that the cooling curve is almost unaffected by changes in hydrogen column density, as we already found in Section 2.4. Nevertheless, it is interesting to see how the hydrogen column changes over time for varying ISM accretion rates. Note that for both the H-He and He-C models, we have artificially set the lower limit for the light element column at 100 g cm^{-2}. Changes in the light element column below this limit

will not affect the cooling curve, as they occur at low densities outside of the sensitivity strip. If DNB is the only mechanism for depleting surface hydrogen, the column density initially drops sharply due to high temperatures, until hydrogen reaction rates become limited by the amount of material available in the burning region. In other words, when the temperature has dropped significantly, the burning is limited by the time it takes for hydrogen to diffuse to the burning depth. Thus for most of such a NS's life (when $T_s \lesssim 10^6$ K), DNB in the NS envelope is not limited by the nuclear reaction rate, but by the diffusion time. For the two scenarios with higher accretion rates we find that, after a period of accreting and burning at similar rates ($\gtrsim 10^5$ years), the accretion rate dominates, and the hydrogen column grows.

In Figure 2.9, we show cooling curves and helium column density evolution for an He-C envelope with an initial helium column density of $y_{He} = 10^{10}$ g cm^{-2} for different accretion rates. In all cases where $\dot{M}_{acc}$ is less than the maximum accretion rate, more than 99% of the initial helium is consumed at the hot early stages and the helium column density drops to $y_{He} < 2 \times 10^6$ g cm^{-2} within 1 year. During this time, helium is more rapidly burned than it is accreted. Depending on the accretion rate, the temperature has dropped enough at $t \sim 1 - 100$ years for the burning rate to become comparable to the accretion rate. After this time (when $\dot{y}_{acc} > \zeta_{DNB}$), the helium column gradually increases as matter accumulates from the ISM, and the temperature is too low for DNB to sufficiently counteract. The scenario with no accretion results in an effectively constant low column density after the initial helium column is depleted. The resulting cooling curve at times > 1 year is the same as a cooling curve with no helium column.The cooling curves for the scenarios with larger accretion rates show short timescale variations, corresponding to variations in helium column. Although the short timescale changes in column density can be large, the effect on the surface temperature is small ($\sim$ a few percent), as the $T_s - T_b$ relations are only sensitive to changes in the column density that affect the composition in the sensitivity strip.

2.5.5 Implications for photospheric composition

In the discussion above, we focus on the effect of DNB on the composition of the envelope and the resulting temperature relations between the surface and the interior. We now consider the implications of the envelope evolution for the composition at the photosphere, which is uncertain for isolated NSs and is important for interpreting observations of NS surface radiation.

In Figures 2.8 and 2.9, the (artificially set) minimum light element column of $y_{light} = 100$ g cm^{-2} is large enough such that the photosphere, and corresponding spectrum, will always be dominated by hydrogen and helium, respectively. Due to this large lower limit, these figures do not directly address whether or not heavier elements can be present in the photosphere. Nevertheless Figure 2.8 shows that, even at the low

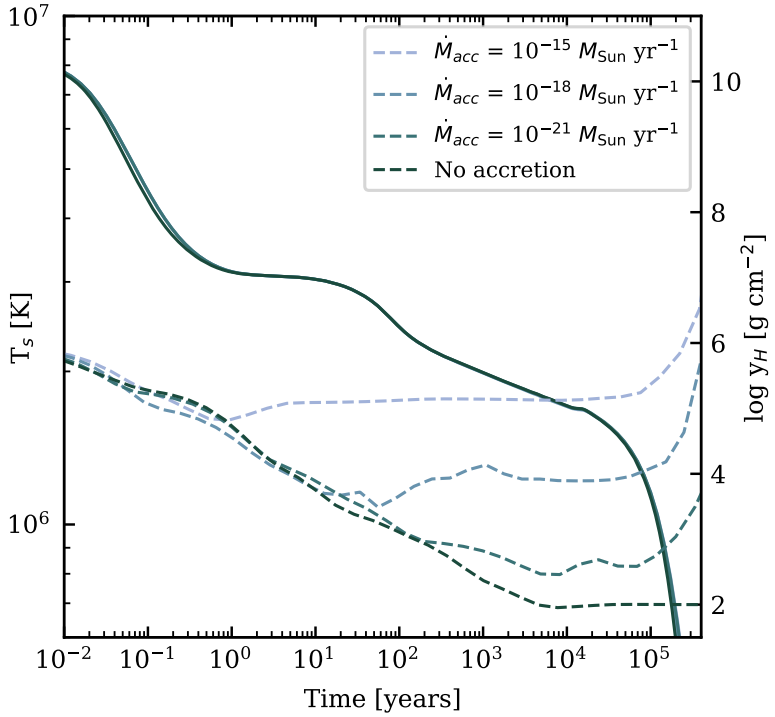

FIGURE 2.8: Cooling curves with initial hydrogen column density $y_H = 10^6$ g cm^{-2} (solid curves) for different mass accretion rates. Dashed lines show the corresponding evolution of envelope composition (right axis).

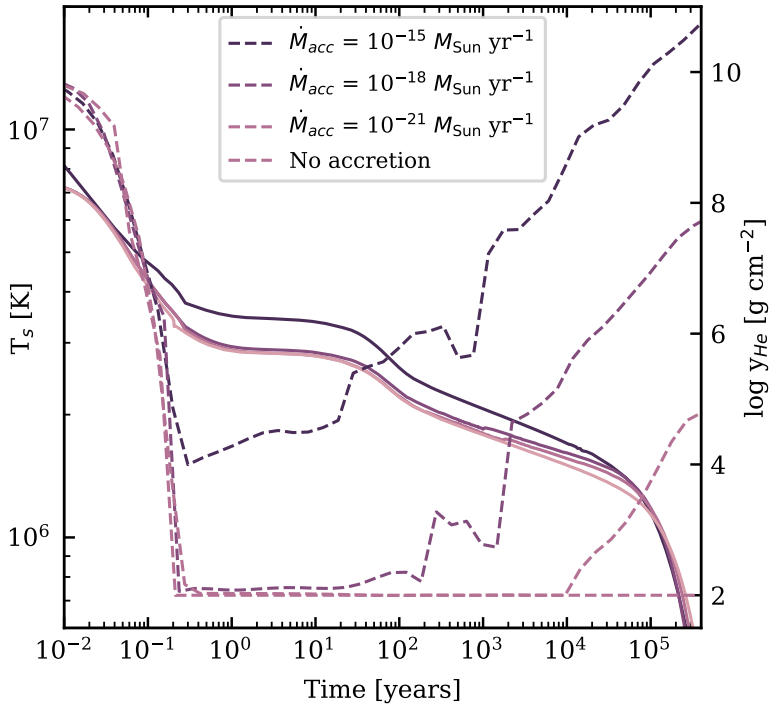

FIGURE 2.9: Cooling curves with initial helium column density $y_{He} = 10^{10}$ g cm^{-2} (solid curves) for different mass accretion rates. Dashed lines show the corresponding evolution of envelope composition (right axis).

accretion rate of $10^{-21} M_{\rm Sun}$ yr^{-1}, the chosen minimum column is never reached, as the hydrogen column plateaus and increases after a few thousand years. Thus for a H-He envelope, the photosphere will be dominated by hydrogen even for very low accretion rates. On the other hand, Figure 2.9 shows that, for accretion rates $< 10^{-15} M_{\rm Sun}$ yr^{-1}, there will be a period after $\sim$1 month, where the spectrum of the He-C envelope can be dominated by carbon.

As accreted material from the ISM consists mainly of hydrogen, and thus accretion will mostly make the photosphere (and envelope) more hydrogen rich, it is interesting to consider the evolution of a H-C envelope. Therefore, we compute a H-C envelope model with $\rho_b = 10^{10}$ g cm^{-3} and $g_s = 2.43 \times 10^{14}$ cm s^{-2}, for which the possible nuclear reactions are proton-proton captures and proton captures onto carbon (see Chang and Bildsten 2004). In Figure 2.10 we show the column evolution for an initial hydrogen column of $y_{\rm H} = 10^6$ g cm^{-2} using a lower limit of 0.1 g cm^{-2}. Note that this lower limit corresponds to a hydrogen fraction at the photosphere of $\sim$5% and thus a carbon dominated spectrum. Only slightly larger hydrogen columns are enough to result in a hydrogen spectrum, as a hydrogen column of $\sim$ 1.5 g cm^{-2} corresponds to a hydrogen fraction of $\sim$90% at the photosphere. Figure 2.10 shows that due to the highly efficient capture onto carbon, the initial hydrogen column is rapidly depleted within an hour for accretion rates $< 10^{-15} M_{\rm Sun}$ yr^{-1}. For these accretion rates, a period can exist when the spectrum is that of a carbon atmosphere. An accretion rate of $10^{-20} M_{\rm Sun}$ yr^{-1} leads to an optically thick hydrogen atmosphere after $\sim$1000 year. For higher accretion rates, the hydrogen column will effectively always be large enough to maintain a hydrogen atmosphere. On the other hand, accretion rates $< 10^{-21} M_{\rm Sun}$ yr^{-1} are too low to build up an optically thick hydrogen column and will result in a carbon atmosphere after the initial hydrogen column is consumed. These results are relevant for the understanding of the neutron star in Cas A, of which the spectra can be well described by a carbon atmosphere (Ho and Heinke, 2009), as well as other NSs for which a carbon atmosphere is suggested (see Section 2.7). We will consider temperature relations for H-C envelopes and their application to these neutron stars in future work.

2.6 Application to Cassiopeia A

2.6.1 *Chandra* data analysis

We process all *Chandra* ACIS-S GRADED data using Chandra Interactive Analysis of Observations (CIAO) 4.10 and Chandra Calibration Database (CALDB) 4.8.1, following the procedure described in Heinke and Ho (2010). Spectra are extracted using SPECEXTRACT from a 4 pixel circular region centered on the source, while background is taken from a 5−8 pixel annulus centered on the source. Observations

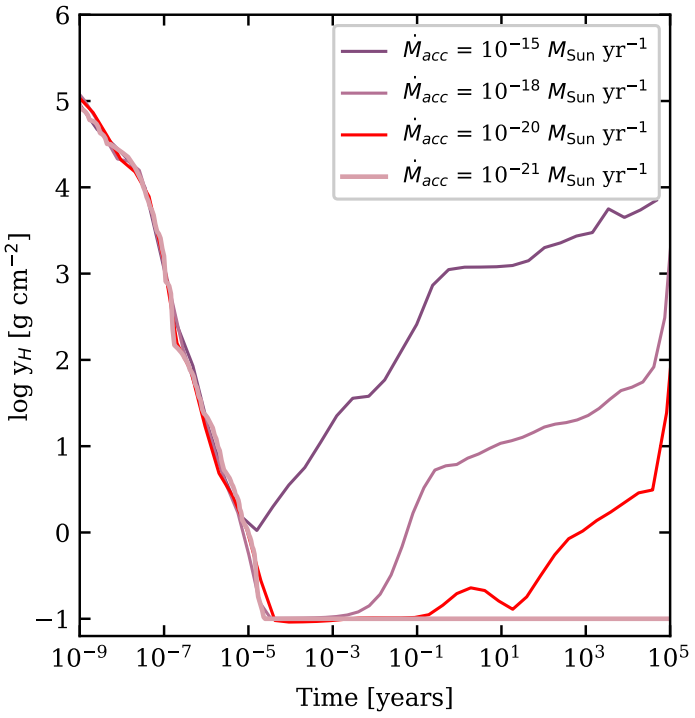

FIGURE 2.10: Evolution of the hydrogen column with an initial column density $y_H = 10^6$ g cm^{-2} for different mass accretion rates.

taken close in time (see Table 2.1) are merged using COMBINE_SPECTRA and DMGROUP with a minimum of 200 counts per energy bin.

Spectra are fit with XSPEC 12.10.0c (Arnaud, 1996) and a four component model consisting of PILEUP, TBABS, SPEXPCUT, and NSX. For PILEUP (Davis, 2001), we set the grade migration parameter to 0.27 for observations through 2004, which have a frame time of 3.24 s, and to 0.24 for observations starting from 2007, which have a frame time of 3.04 s; the maximum number of photons is set to 5, and the point-spread-function fraction is set to 0.95 (see Heinke and Ho 2010, for details). To model photoelectric absorption, we use TBABS with abundances from Wilms et al. (2000). To model interstellar dust scattering, we use SPEXPCUT with exponent index $\alpha = -2$ and characteristic energy $E_{cut}(\mathrm{keV}) = [0.49 N_H (10^{22}\ \mathrm{cm}^{-2})]^{1/2}$ (Predehl et al., 2003), where N_H is the absorption column determined from TBABS. For NSX (Ho and Heinke, 2009), we assume a partially ionized carbon atmosphere and fix the distance to 3.4 kpc (Reed et al., 1995) and normalization to 1.

2.6.2 Spectral fits

To determine the neutron star mass M and radius R for use in spectral fitting, we first allow fit parameters M, R, N_H, and T_s to vary but tie them between observations, so that each has the same value across all observations. The resulting fit produces best-fit values of $M = 1.44 \pm 0.15\ M_{Sun}$, $R = 13.78 \pm 0.36$ km,

$N_{\rm H} = (1.623 \pm 0.030) \times 10^{22}$ cm^{-2}, and $T_{\rm s} = (1.728 \pm 0.030) \times 10^{6}$ K, with $\chi^2 = 1646$ for 1458 degrees of freedom (dof). If we instead allow $T_{\rm s}$ to have different values between observations, then the best-fit values are $M = 1.65 \pm 0.16\,M_{\rm Sun}$, $R = 12.94 \pm 0.34$ km, and $N_{\rm H} = (1.674 \pm 0.031) \times 10^{22}$ cm^{-2}, with $\chi^2/{\rm dof} = 1522/1446$. If we allow both $N_{\rm H}$ and $T_{\rm s}$ to have different values between observations, then the best-fit values are $M = 1.46 \pm 0.16\,M_{\rm Sun}$ and $R = 13.69 \pm 0.36$ km, with $\chi^2/{\rm dof} = 1497/1434$. A f-test between the first and second fits ($\Delta\chi^2 = 124$) and between the second and third fits ($\Delta\chi^2 = 25$) yield probabilities of 9.8×10^{-19} and 0.024, respectively, which suggest that allowing $T_{\rm s}$, and possibly $N_{\rm H}$, to change between observations is warranted. Therefore in the spectral fits performed below, we untie $T_{\rm s}$ and fix the neutron star mass and radius to $M = 1.65\,M_{\rm Sun}$ and $R = 12.9$ km, respectively.

We show the M-R confidence contour in Fig. 2.11, assuming a distance of 3.4 kpc. The best-fit values from our previous work (Heinke and Ho, 2010) and used in subsequent works (Shternin et al., 2011; Elshamouty et al., 2013; Ho et al., 2015) are $M = 1.65\,M_{\rm Sun}$ and $R = 10.3$ km. It is important to note that, while the M-R constraints we obtain are somewhat constraining, there are caveats. First, there are calibration issues with the data, which are discussed in Section 2.7 (see also Posselt et al. 2013). Second, the magnetic field strength of Cas A is assumed here to be low enough as to not affect its atmospheric emission. This may not necessarily be the case, as three other CCOs have surface magnetic fields $\sim 10^{10} - 10^{11}$ G (Gotthelf et al., 2013; Ho, 2013; De Luca, 2017), which is somewhat larger than accommodated for in our atmosphere model (c.f. Potekhin et al. 2014). Finally, it is possible that the NS surface is not uniform in temperature, and hot spots could dominate the observed emission (see, e.g., Gotthelf et al. 2010; Bogdanov 2014), although pulsations from this emission are not detected in Cas A (Murray et al., 2002; Halpern and Gotthelf, 2010).

Next, we perform two fits of all spectra with the same set of models but fixing the parameters $M = 1.65\,M_{\rm Sun}$ and $R = 12.9$ km. The first fit allows $T_{\rm s}$ to vary between observations but holds the absorption parameter fixed at $N_{\rm H} = 1.67 \times 10^{22}$ cm^{-2}. This fit yields $\chi^2/{\rm dof} = 1522/1449$, and the temperature evolution is shown in Fig. 2.12 and Table 2.1. The second fit allows both $T_{\rm s}$ and $N_{\rm H}$ to change between observations. This fit yields $\chi^2/{\rm dof} = 1498/1436$, and temperature and absorption column evolutions are shown in Figs. 2.13 and 2.14, respectively (see also Table 2.1). A f-test between these two fits ($\Delta\chi^2 = 24$) yields a probability of 0.043, so allowing $N_{\rm H}$ to vary is possibly justified (see Section 2.7). For comparison to our previously published results, we also perform a fit using fixed $(M, R, N_{\rm H}) = (1.65\,M_{\rm Sun}, 10.3\ {\rm km}, 1.73 \times 10^{22}\ {\rm cm}^{-2})$. Results are shown in Fig. 2.12 and Table 2.1 and have $\chi^2/{\rm dof} = 1562/1449$. The approximately constant offset in temperatures (e.g., $\sim 2.13/1.87 = 1.14$) between the two fixed-$N_{\rm H}$ fit results plotted in Fig. 2.12 is primarily due to the difference in assumed R and its anti-correlation with

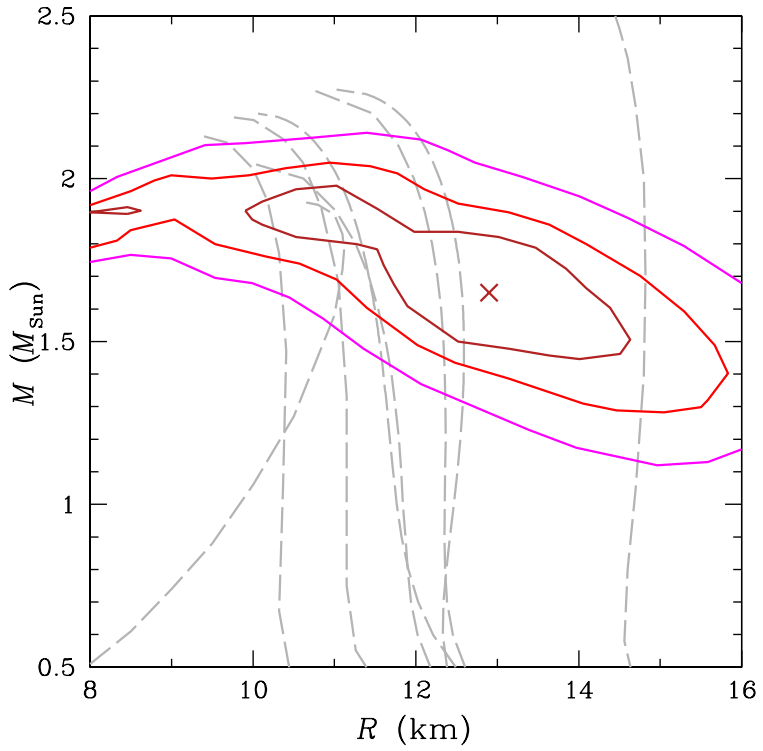

FIGURE 2.11: Neutron star mass M and radius R. Solid lines are the 1σ, 90%, and 99% confidence contours on M and R obtained from fitting *Chandra* spectra of the neutron star in the Cassiopeia A supernova remnant, assuming $N_{\rm H}$ does not change between observations. Cross marks the best-fit $(M, R) = (1.65\,M_{\rm Sun}, 12.9\,{\rm km})$. Dashed lines indicate mass-radius relation obtained from a sample of theoretical nuclear equations of state (see Douchin and Haensel 2001; Lattimer and Prakash 2001; Potekhin et al. 2013; Baldo et al. 2014, and references therein).

$T_{\rm s}$, i.e., for constant luminosity, $T_{\rm s} \propto R^{-1/2}$, which yields $(12.9\,{\rm km}/10.3\,{\rm km})^{1/2} = 1.12$, as well as a difference in gravitational redshift.

Finally, we perform a linear fit to each set of declining temperatures $T_{\rm s}$. For $(M, R, N_{\rm H}) = (1.65\,M_{\rm Sun}, 10.3\,{\rm km}, 1.73 \times 10^{22}\,{\rm cm}^{-2})$, the slope is $-13 \pm 1\,{\rm K\,d}^{-1}$ ($\approx -2.2 \pm 0.2\%$ per 10 years) with $\chi^2/{\rm dof} = 14.8/11$. For $(M, R, N_{\rm H}) = (1.65\,M_{\rm Sun}, 12.9\,{\rm km}, 1.67 \times 10^{22}\,{\rm cm}^{-2})$, the slope is $-11 \pm 1\,{\rm K\,d}^{-1}$ ($\approx -2.1 \pm 0.2\%$ per 10 years) with $\chi^2/{\rm dof} = 15.7/11$. For $(M, R) = (1.65\,M_{\rm Sun}, 12.9\,{\rm km})$ and changing $N_{\rm H}$, the slope is $-14 \pm 1\,{\rm K\,d}^{-1}$ ($\approx -2.7 \pm 0.3\%$ per 10 years) with $\chi^2/{\rm dof} = 6.2/11$. We discuss uncertainties beyond statistical ones in our measured temperature decline rates and compare to rates measured in previous works in Section 2.7.

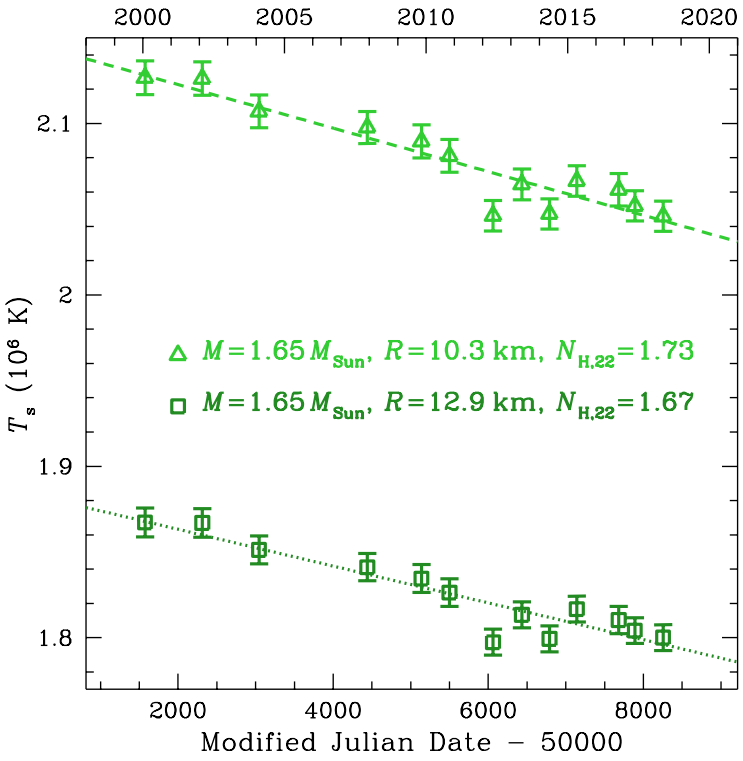

FIGURE 2.12: Surface temperature of the neutron star in the Cassiopeia A supernova remnant as measured from *Chandra* ACIS-S GRADED spectra over the past 18 years. Triangles indicate T_s obtained using the best-fit neutron star mass $M = 1.65\,M_{Sun}$ and radius $R = 10.3$ km from Heinke and Ho (2010) and absorption column $N_{H,22} \equiv N_H/10^{22}\,cm^{-2} = 1.73$ from Elshamouty et al. (2013), while squares indicate T_s obtained using updated best-fit $M = 1.65\,M_{Sun}$, $R = 12.9$ km, and $N_{H,22} = 1.67$ (see text). Error bars are 1σ. Dashed and dotted lines show linear fits to each set of T_s.

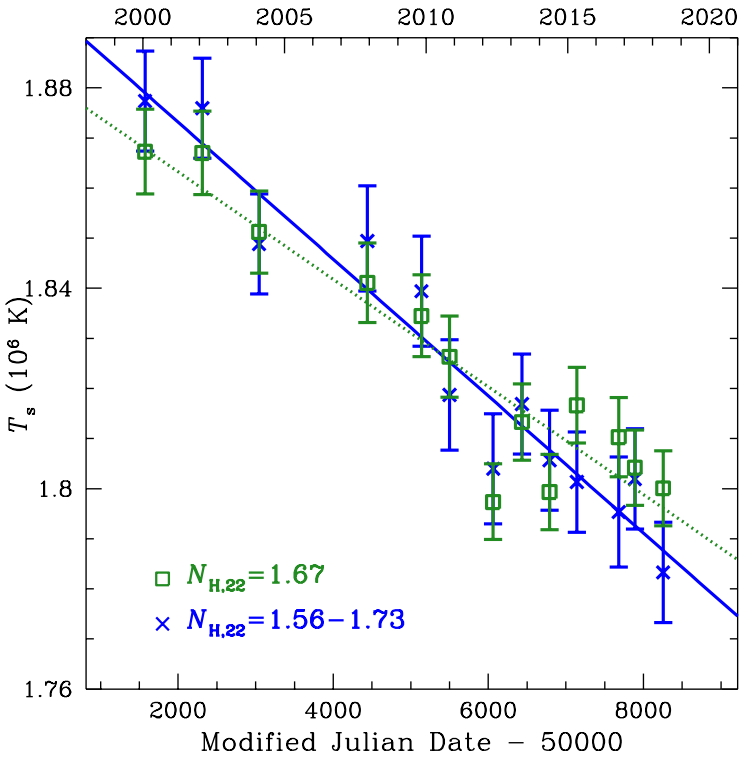

FIGURE 2.13: Surface temperature of the neutron star in the Cassiopeia A supernova remnant as measured from *Chandra* ACIS-S GRADED spectra over the past 18 years. Data points indicate T_s obtained using best-fit neutron star mass $M = 1.65\,M_{Sun}$ and radius $R = 12.9$ km and a constant absorption column $N_{H,22} \equiv N_H/10^{22}$ cm$^{-2} = 1.67$ (squares) and changing N_H (crosses; see Fig. 2.14). Error bars are 1σ. Dotted and solid lines show linear fits to each set of T_s.

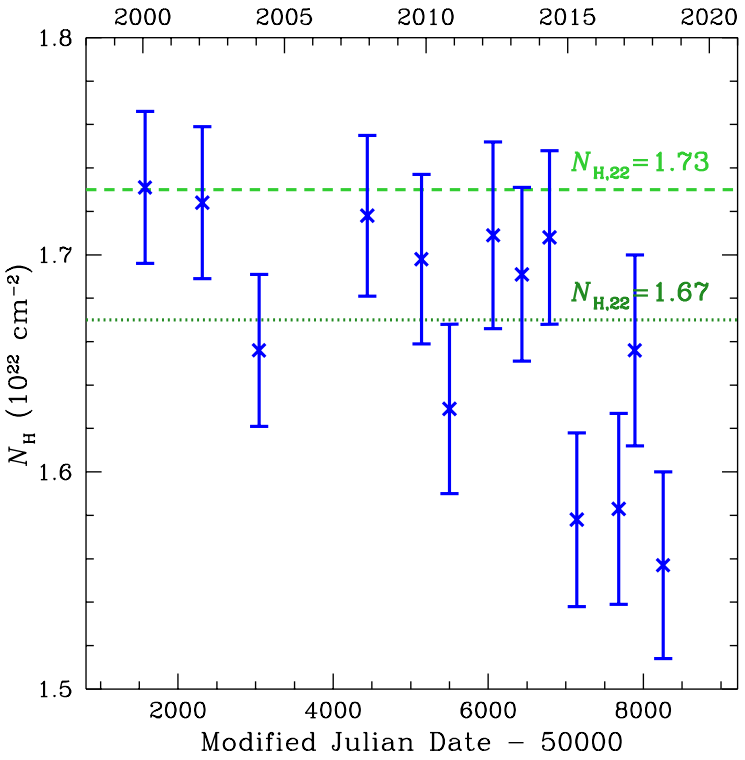

FIGURE 2.14: Absorption column $N_{\rm H}$ as measured from a fit to *Chandra* spectra of the Cassiopeia A neutron star. Data points indicate $N_{\rm H}$ using best-fit neutron star mass $M = 1.65\,M_{\rm Sun}$ and radius $R = 12.9$ km and changing $T_{\rm s}$, while dotted and dashed lines indicate constant $N_{\rm H,22} \equiv N_{\rm H}/10^{22}\ {\rm cm}^{-2} = 1.67$ and 1.73, respectively, used in other spectral fits (see Figs. 2.12 and 2.13). Error bars are 1σ.

TABLE 2.1: Surface temperature $T_{s,6}$ (10^6 K), absorbed 0.5–7 keV flux F_{-13}^{abs} (10^{-13} erg cm^{-2} s^{-1}), and absorption column $N_{H,22}$ (10^{22} cm^{-2}) determined from model fits to *Chandra* ACIS-S GRADED spectra of the neutron star in the Cassiopeia A supernova remnant. Three spectral fit results are shown, two with constant N_H and one with changing N_H. Each set of 13 temperatures and fluxes are fit to a linear decline, with decline rate and fit statistic as shown. For merged ObsIDs, the MJD listed is that of the first ObsID. Number in parentheses is 1σ uncertainty in last digit.

		$N_{H,22}$	1.73		1.67		see below		
		M (M_{Sun})	1.65		1.65		1.65		
		R (km)	10.3		12.9		12.9		
		χ^2/dof	1562/1449		1522/1449		1498/1436		
ObsID	Date	MJD	$T_{s,6}$	F_{-13}^{abs}	$T_{s,6}$	F_{-13}^{abs}	$T_{s,6}$	F_{-13}^{abs}	$N_{H,22}$
114	2000 Jan 30	51573.4	2.127(10)	7.4(2)	1.867(8)	7.3(1)	1.877(10)	7.3(2)	1.73(3)
1952	2002 Feb 6	52311.3	2.126(10)	7.4(2)	1.867(8)	7.3(1)	1.876(10)	7.4(2)	1.72(3)
5196	2004 Feb 8	53043.7	2.107(10)	7.1(1)	1.851(8)	7.0(2)	1.849(10)	7.0(2)	1.66(3)
9117/9773	2007 Dec 5/8	54439.9	2.098(9)	7.0(2)	1.841(8)	6.9(2)	1.849(11)	6.9(2)	1.72(4)
10935/12020	2009 Nov 2/3	55137.9	2.090(10)	6.9(2)	1.834(8)	6.8(2)	1.839(11)	6.8(1)	1.70(4)
10936/13177	2010 Oct 31/Nov 2	55500.2	2.081(10)	6.8(2)	1.826(8)	6.7(2)	1.819(11)	6.6(1)	1.63(4)
14229	2012 May 15	56062.4	2.046(9)	6.4(2)	1.797(8)	6.3(1)	1.804(11)	6.4(1)	1.71(4)
14480	2013 May 20	56432.6	2.064(9)	6.6(2)	1.813(8)	6.5(1)	1.817(10)	6.5(1)	1.69(4)
14481	2014 May 12	56789.1	2.047(9)	6.3(2)	1.799(7)	6.2(1)	1.806(10)	6.3(1)	1.71(4)
14482	2015 Apr 30	57142.5	2.066(9)	6.7(1)	1.817(8)	6.6(1)	1.801(10)	6.6(1)	1.58(4)
19903/18344	2016 Oct 20/21	57681.2	2.061(9)	6.6(2)	1.810(8)	6.5(2)	1.795(11)	6.5(1)	1.58(4)
19604	2017 May 16	57889.7	2.052(9)	6.5(2)	1.804(7)	6.4(2)	1.802(10)	6.4(2)	1.66(4)
19605	2018 May 15	58253.7	2.046(9)	6.4(2)	1.800(7)	6.4(1)	1.783(10)	6.4(1)	1.56(4)
		10-year decline rate	$2.2 \pm 0.2\%$	$7.5 \pm 1.0\%$	$2.1 \pm 0.2\%$	$7.3 \pm 0.9\%$	$2.7 \pm 0.3\%$	$7.7 \pm 0.9\%$	
		χ^2/dof	14.8/11	11.0/11	15.7/11	15.8/11	6.2/11	10.9/11	

2.6.3 Cas A cooling curve

The observed temperature drop has been interpreted as a result of neutron superfluidity and proton superconductivity in Cas A (Page et al. 2011; Shternin et al. 2011; Yakovlev et al. 2011; Ho et al. 2015, see Section 4.1 for references on alternative interpretations). With the addition of 4 new data points presented above, we now have a set of 13 temperatures spanning over 18 years. We apply the cooling model described in Section 2.5.1 to the full set of data. We use the temperature data corresponding to the spectral fits with $(M, R, N_{\rm H}) = (1.65\,M_{\rm Sun}, 12.9\ {\rm km}, 1.67 \times 10^{22}\ {\rm cm}^{-2})$. The cooling model makes use of a time-variable He-C envelope model with an initial He column of 10^6 g cm^{-2} and assumes no accretion from the ISM. As we found in Section 2.5, the cooling curve is not sensitive to the initial helium column density, as practically all helium will be depleted within 10 years. Therefore, the cooling curve at the time of the Cas A observations (at age $\approx 320 - 340$ yr) is the same as a cooling curve using a pure carbon envelope.

In order to describe the rapid temperature decline, we follow the previous works (e.g., Page et al. 2011; Shternin et al. 2011) by invoking superfluidity of neutrons in the inner crust through pairing in the $^1S_0$ state and in the core through pairing in the $^3P_2$ state. Within the minimal cooling paradigm (Gusakov et al., 2004; Page et al., 2004), this means that stronger neutrino emission is only allowed as a result of neutron or proton pairing. The exact properties, such as the critical temperature $T_{\rm c}$, of neutron $^3P_2$ superfluidity in the core are uncertain and can be constrained by comparing the cooling curve to the observed temperatures (Page et al., 2011; Shternin et al., 2011; Ho et al., 2015). In our model here, we use superfluidity and superconductivity properties similar to those in Page et al. (2011); specifically, we use a neutron $^3P_2$ superfluid model with a maximum critical temperature $T_{\rm c} \sim 5.9 \times 10^8$ K for $M = 1.7\,M_{\rm Sun}$ (with the APR EOS, resulting in R = 11.4 km) and the 'CCDK' proton $^1S_0$ superconducting model described in Page et al. (2004). Note that our considered theoretical mass and radius using the APR EOS (1.7 Msun and 11.4 km) is within the 90 percent confidence contour of mass and radius obtained from the spectral analysis (see Figure 2.11). The aim of our calculations here is not to find a best fitting model but to show that models fitting the new data are in line with previous results for superfluid properties and can also be explained using the envelope $T_{\rm b} - T_{\rm s}$ relations that include DNB.

To compare the model to the observations, we use the date for the birth of the Cas A NS of 1681 ± 19 (Fesen et al., 2006). The resulting cooling curve is shown in Figure 2.15 and can adequately describe the observed temperatures. It is important to note that several combinations of superfluid and neutron star properties can describe the observed cooling, so the model presented here is not a unique solution and is merely for illustrative purposes.

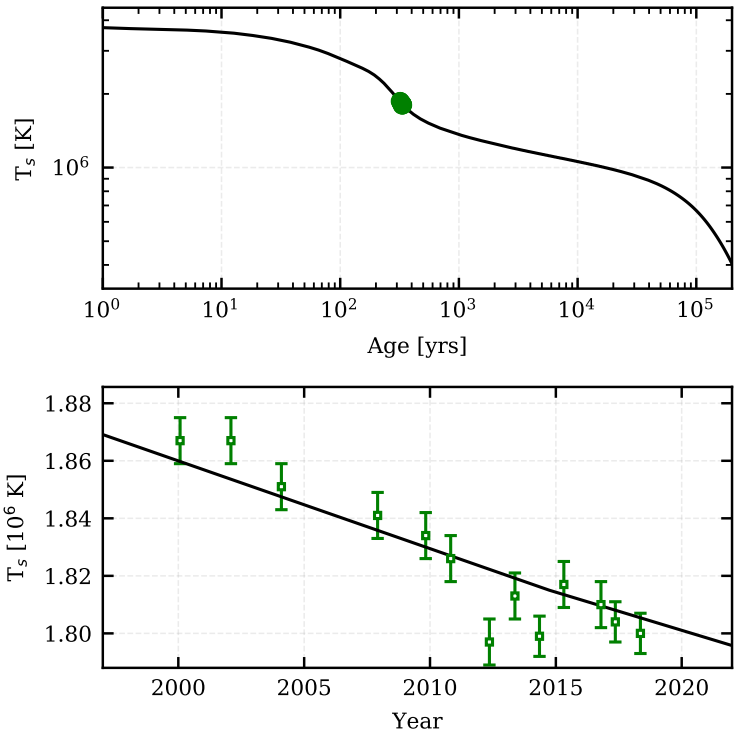

FIGURE 2.15: Example cooling curve for Cas A with $M = 1.7\ M_{\rm Sun}$ and $R = 11.4$ km. Model parameters are similar to those reported in Page et al. (2011). The rapid temperature decline at ~ 340 yr is explained by a neutron $^3P_2$ superfluid in the core with a maximum $T_{\rm c} \approx 5.9 \times 10^8$ K.

2.6.4 Implications of DNB on Cas A photosphere

In Section 5.5, we show that DNB rapidly consumes any initial hydrogen from H-C envelopes for low to moderate accretion rates, thus allowing for an atmosphere composed of carbon to be visible. Effectively, hydrogen atmosphere formation by accretion in the dense post-supernova environment is delayed by DNB to a later, lower density epoch. For the relatively young Cas A NS, this is likely the case. In order for Cas A to have a carbon atmosphere, it must accrete at an average rate lower than $10^{-20}\ M_{\rm Sun}\ {\rm yr}^{-1}$; for larger accretion rates, DNB cannot consume hydrogen fast enough to prevent formation of a hydrogen atmosphere. While DNB can explain a carbon atmosphere, it is unlikely to be responsible for the observed temperature variations in Cas A, as the light element burning rate is negligible when the atmosphere is dominated by carbon.

2.7 Discussion

In this present work, we consider neutron star envelopes consisting of either H-He or He-C and including the physics of diffusive nuclear burning, in order to investigate

how DNB impacts the thermal evolution of NSs. DNB is relevant for the thermal evolution in both a static and time-dependent manner: static because DNB affects the interior composition profile of the envelope and thus can alter the $T_s - T_b$ relations, and time-dependent because it can alter the composition of the envelope over time. In Section 2.3, we calculate the burning rates and find that DNB can alter the composition on short time-scales ($\lesssim 10 - 100$ yr). For H-He, the hydrogen burning rate is strongly coupled to the amount of hydrogen available (characterised by its column density) whereas for the He-C envelope the burning rate is highly sensitive to the temperature and is only significant for $T_s > 10^6$ K.

In Section 2.4, we obtained static analytic relations for the surface temperature T_s and the temperature at the bottom of the envelope T_b. These temperature relations (in Appendix B.1) can be used to model the thermal evolution of isolated and accreting neutron stars. We find the following main differences in the $T_s - T_b$ relations due to DNB:

- For a H-He envelope, we find that when DNB is taken into account, there is a clear upper limit on the size of the hydrogen column. Larger hydrogen columns ($y_H \gtrsim 10^7$ g cm^{-2}) are not thermally stable, as hydrogen is consumed rapidly when DNB is nuclear-limited. For temperatures above $T_s \sim 10^6$ K, this means the spread in T_b is smaller than without DNB (see Figure 2.4).
- For He-C envelopes, DNB is highly sensitive to the temperature, and DNB makes the envelope more transparent to the heat flux at $T_s > 2 \times 10^6$ K for the helium column sizes that affect the sensitivity strip (see Figure 2.5).

We then combine these results in Section 2.5 and use the DNB burning rates (Section 2.3) and the temperature relations (Section 2.4) to calculate NS cooling curves with envelope compositions that evolve over time through nuclear burning and accretion. As the H-He temperature relations only span a small range of T_b, the changing envelope composition does not impact the cooling curves significantly. We find that if DNB is the only burning mechanism and no accretion takes place, an initial hydrogen column of 10^6 g cm^{-2} is depleted after $\sim 5 \times 10^3$ years because the burning rate drops as the hydrogen column decreases. For the He-C envelope all helium will be depleted during the hot early cooling phase (within 1 year) if no accretion takes place. Note that we have only considered illustrative cases and have not included heat generated by the nuclear reactions nor the possibility of sudden accretion events.

For several CCOs, the spectra could be well described by a carbon atmosphere (see, e.g., Klochkov et al. 2016; Suleimanov et al. 2017; Doroshenko et al. 2018). We consider the photosphere of a H-C envelope to investigate the possibility of a carbon atmosphere on isolated neutron stars. Low levels of accretion lead to an optically thick

hydrogen atmosphere, with a hydrogen column of 1.5 g cm^{-2} corresponding to a hydrogen abundance in the atmosphere of 90%. The composition evolution of the H-C envelope in Figure 2.10 for accretion rates $\leq 10^{-20} M_{\rm Sun}$ yr^{-1} indicate that it is possible for Cas A, at an age of 340 years, to not yet accrete enough hydrogen to form an optically thick hydrogen atmosphere. For older NSs, the possibility of a carbon atmosphere becomes less likely, as even a little bit of accretion leads to a hydrogen dominated atmosphere.

Note that we have only considered constant accretion rates here, whereas scenarios with changing accretion rates are more realistic for CCOs, as their environment is highly variable (see, e.g., Vink 2012; Lopez and Fesen 2018). For example, it is possible that the NS initially accretes at very low rates after the nearby environment has been cleared out by the SN shock. Later on, a reverse shock could introduce a higher density of the surrounding medium, and thus a higher accretion rate onto the NS. Early fallback of material could also lead to varying accretion rates. As the variability in the accretion rate and the composition of the accreted material in SNRs are highly model dependent, we have limited our models to illustrative cases with constant rates and compositions.

In all the envelope models currently used in cooling simulations, it is assumed that the flux generated by nuclear reactions in the envelope is negligible compared to the radiative flux from the neutron star interior. However, we find that the assumption of constant flux through the envelope can be invalid in relevant regions of parameter space. The assumption of negligible heat production in the envelope is certainly incorrect for the hot, initial post-supernova stage, when surface temperatures are $\gg 10^6$ K. However, as the energy is produced at low densities ($\rho < 10^8 - 10^{10}$ g cm^{-3}) and the interior is still hot, this heat is radiated from the surface and will not affect cooling at later times. Including nuclear heating at early times will likely lead to an even more rapid consumption of light elements than calculated here. At later times, heat generated by DNB is negligible for H-He envelopes, while heat created by DNB can be important for high temperatures (T $\gtrsim 2 \times 10^6$ K) and large helium columns ($y_{\rm He} \gtrsim 10^{11}$ g cm^{-2}) in the case of He-C envelopes. Interestingly, for H-C envelopes the luminosity generated by DNB can be larger than the interior luminosity for hydrogen columns greater than $y_{\rm H} > 10^6$ g cm^{-2} and surface temperatures higher than $T_{\rm s} \gtrsim 10^6$ K. We will examine the effect of heat generated by nuclear reactions in future work.

In Section 2.6, we present new observations of the NS in the Cassiopeia A supernova remnant and apply our model to this data. We performed spectral fits of 13 *Chandra* ACIS-S GRADED observations of Cas A taken over more than 18 years (see Table 2.1). For constant interstellar absorption ($N_{\rm H} \approx 1.7 \times 10^{22}$ cm^{-2}), these data indicate that the NS temperature is decreasing at a ten-year rate of $2.1 \pm 0.2\%$ (1σ) and a decreasing ten-year flux rate of $\approx 7\%$. However, it is known that there is a contaminant building up on the *Chandra* ACIS-S detector that can affect GRADED mode data if uncorrected

(Posselt et al., 2013; Plucinsky et al., 2016, 2018). Elshamouty et al. (2013) analyse data from different detectors and operating modes and infer a ten-year temperature decline rate of $3.5 \pm 0.4\%$ for ACIS-S GRADED data and a ten-year decline rate of $1.0 \pm 0.7\%$ for HRC-S data, suggesting that at least one dataset is affected by systematic errors. Posselt et al. (2013) analyse an alternative set of 2 *Chandra* ACIS-S subarray observations taken over 6 years in FAINT mode, to minimize the effects of pile-up and contaminant; they find a lower ten-year rate of flux decline and no apparent temperature decline ($10 \pm 5\%$ and $1.3 \pm 1.0\%$, respectively, where errors here are at 90% confidence level). An important result of our work here is that the temperature decline rate ($2.1 \pm 0.2\%$ for constant $N_{\rm H}$ and $2.7 \pm 0.3\%$ for varying $N_{\rm H}$) is significantly lower than values previously reported based on a smaller set of ACIS-S GRADED observations ($3.9 \pm 0.7\%$ per 10 years from Heinke and Ho 2010 and $3.5 \pm 0.4\%$ per 10 years from Elshamouty et al. 2013). Therefore, the rates of temperature decline measured using HRC-S data and ACIS-S subarray and GRADED data all agree at $\sim 90\%$ confidence level.

Very recently, Posselt and Pavlov (2018) follow up their previous analysis by including a new ACIS-S subarray observation, which expands this dataset to 3 measurements in 9 years, and use CALDB 4.7.3[5] (also testing versions up to 4.7.7); their analysis yields 3σ upper limits on the ten-year rate of temperature decline of 2.4% assuming a constant $N_{\rm H}$ and 3.3% assuming a varying $N_{\rm H}$. In our new analysis using 13 measurements in 18 years, as well as the latest *Chandra* calibration and contamination models (CALDB 4.8.1; see footnote 5 and Appendix A.2), the temperature decline we measured using ACIS-S GRADED mode is in agreement with the upper limits found by Posselt and Pavlov (2018) using the subarray observations.

As noted above, our spectral fits suggest $N_{\rm H}$ may be variable (see Table 2.1 and Fig. 2.14). Alp et al. (2018) study the effect on $N_{\rm H}$ of material expelled by a supernova using three-dimensional supernova simulations. They find that different supernova models do not manifest measurable differences in $N_{\rm H}$ at the current age of the Cassiopeia A neutron star. They also analyse the same three Chandra observations (with $N_{\rm H}$ tied across observations) as those used by Posselt and Pavlov (2018) and measure a ten-year rate of temperature decline of $\sim 1 \pm 1\%$ and a supernova ejecta contribution to $N_{\rm H}$ that is indistinguishable from that due to the intervening interstellar medium. While motion of intervening supernova remnant material (or the NS) as the cause of varying $N_{\rm H}$ cannot be ruled out, a varying $N_{\rm H}$ likely reflects incomplete modelling of the detector contaminant (see, e.g., Plucinsky et al. 2018). We note however that in the case of our spectral fits allowing for a varying $N_{\rm H}$, the ten-year rates of temperature and flux decline are actually larger, i.e., $2.7 \pm 0.3\%$ and

[5] CALDB 4.7.3 includes contaminant model N0010 for ACIS-S. Similarly, CALDB 4.7.8 (used in Appendix A.2) includes N0010. On the other hand, the most recent calibration CALDB 4.8.1 (used in Section 2.6) includes the latest contaminant model N0012 for ACIS-S (see `http://cxc.cfa.harvard.edu/caldb4/downloads/Release_notes/CALDB_v4.8.1.html`).

$\approx 8\%$, respectively (see Fig. 2.13), than the case when $N_{\rm H}$ is constant. Recent analysis by Plucinsky et al. (2018) of the ACIS contaminant suggests the build up on the detector is not increasing as much as predicted by model N0010, and this effect is taken into account in the revised model N0012 (see footnote 5). Future refinements of the ACIS-S calibration may aid attempts to more accurately measure the cooling rate of the NS in Cassiopeia A.

Chapter 3

Effect of diffusive nuclear burning on cooling in accreting binary systems [1]

Abstract

Valuable information about the neutron star interior can be obtained by comparing observations of thermal radiation from a cooling neutron star crust with theoretical models. Nuclear burning of lighter elements that diffuse to deeper layers of the envelope can alter the relation between surface and interior temperatures and can change the chemical composition over time. We calculate new temperature relations and consider two effects of diffusive nuclear burning (DNB) for H-C envelopes. First, we consider the effect of a changing envelope composition and find that hydrogen is consumed on short timescales and our temperature evolution simulations correspond to those of a hydrogen-poor envelope within $\sim$100 days. The transition from a hydrogen-rich to a hydrogen-poor envelope is potentially observable in accreting NS systems as an additional initial decline in surface temperature at early times after the outburst. Second, we find that DNB can produce a non-negligible heat flux, such that the total luminosity can be dominated by DNB in the envelope rather than heat from the deep interior. However, without continual accretion, heating by DNB in H-C envelopes is only relevant for $<$1-80 days after the end of an accretion outburst, as the amount of light elements is rapidly depleted. Comparison to crust cooling data shows that DNB does not remove the need for an additional shallow heating source. We conclude that solving the time-dependent equations of the burning region in the envelope self-consistently in thermal evolution models instead of using static temperature relations would be valuable in future cooling studies.

[1] This chapter is based on publication Wijngaarden et al. (2020), see Publications and author contributions.

3.1 Introduction

Neutron stars (NSs) are an excellent laboratory for investigating extreme physics over a large range of densities, including supranuclear densities in their cores. One method to probe NS interiors is comparing theoretical models with observed temperatures of cooling NSs. If left undisturbed, isolated NSs cool over time after they are born hot in supernova explosions (see, e.g., Potekhin et al. 2015). NSs in binary systems can occasionally undergo accretion outbursts, when new material is accreted from the companion star onto the NS surface. During an accretion outburst, the NS crust can be heated out of thermal equilibrium with the core as a sequence of non-equilibrium nuclear reactions in the crust, collectively known as *deep crustal heating*, are triggered when the underlying material is compressed by newly accreted matter (Haensel and Zdunik, 1990; Brown et al., 1998). After the accretion outburst, the NS crust cools back to thermal equilibrium with the core on time scales of years (see, e.g., Wijnands et al. 2017 for a review). This means that the cooling phase can be covered with multiple observations during which the surface temperature can be measured to constrain theoretical cooling curves, which is a useful opportunity to gain insight in the properties and physics of the NS crust and core (e.g., Page and Reddy 2013; Cumming et al. 2017; Brown et al. 2018).

For several crust cooling sources, the high observed surface temperatures $\lesssim 100$ days after the accretion outburst can not be explained using the standard deep crustal heating model but require the presence of an additional shallow heating source (at densities $\rho < 10^{11}$ g cm^{-3}) during outburst. Typically, during the outbursts, an amount of $\sim$1-2 MeV per accreted nucleon of shallow heating is needed to explain the observed cooling curves after the outbursts are over (e.g., Brown and Cumming 2009; Degenaar et al. 2014; Merritt et al. 2016), but for one source (MAXI J0556-332) as much as $\sim$17 MeV nucleon^{-1} was needed (Homan et al., 2014; Deibel et al., 2015; Parikh et al., 2017). This means that the amount of shallow heating can be larger than that due to deep crustal heating, which releases $\sim$2 MeV nucleon^{-1}. One of the outstanding problems in cooling studies is the unknown physical mechanism of this shallow heat (see the discussion in Deibel et al. 2015). Thus, it is useful to explore additional heating mechanisms that are currently not accounted for in the present crustal heating models.

One of the challenges in understanding crust cooling and shallow heating is obtaining the interior temperature (at the bottom of the envelope) from observations of the surface emission. The relation between surface temperature and interior temperature is set by the heat conducting properties of the thin outer envelope ($\rho < \rho_b = 10^8 - 10^{10}$ g cm^{-3}), which is highly sensitive to chemical composition (Gudmundsson et al., 1983; Potekhin et al., 1997; Brown et al., 2002). NS cooling codes typically calculate the interior structure and thermal evolution from the centre of the star out to the bottom of the envelope (at ρ_b). Then, analytic fits of the temperature

relations between the NS surface temperature (T_s) and the temperature at the bottom of the envelope, T_b [$\equiv T(\rho_b)$], are used as a boundary condition. Analytic fits to numerical envelope calculations for different static envelope compositions and envelope sizes have been calculated for one- (Gudmundsson et al., 1983) and two or more chemical component models (see, e.g., Potekhin et al. 1997, 2003; Beznogov et al. 2016; Wijngaarden et al. 2019).

Diffusive nuclear burning (DNB) of hydrogen and helium can affect the envelope T_s-T_b relations and can alter the envelope composition over time. DNB occurs when lighter elements diffuse to depths where the temperature and density are large enough to ignite nuclear burning of these elements. This process has been studied for H-C (Chang and Bildsten, 2003, 2004) and He-C envelopes (Chang et al., 2010). In addition, Wijngaarden et al. (2019) calculated H-He and He-C temperature relations including the effect of DNB and investigated how DNB can alter the envelope composition over time for cooling isolated neutron stars.

Here we obtain analytic temperature relations for H-C envelopes and explore the effect of DNB on the observed cooling of NSs after an accretion outburst. The temperatures in the envelope after accretion outbursts are sufficiently large such that the heat deposited in the envelope by DNB may not be negligible, as was assumed in previous works. We investigate whether the common assumption in cooling studies that no energy is generated in the envelope during quiescence and that its composition does not change significantly during the cooling time is valid when DNB is taken into account. We compare the luminosity generated by DNB in the envelope when the NS is cooling to the interior luminosity, to explore the role of DNB as an additional heating mechanism. We consider how the envelope composition changes over time due to the consumption of lighter elements, which by itself can alter the observed cooling curves as the heat conduction properties of the envelope are sensitive to its chemical composition.

3.2 H-C envelopes

3.2.1 T_s- T_b relations

We calculate new temperature relations that include diffusive nuclear burning for a H-C envelope with $\rho_b = 10^{10}$ g cm^{-3}. These temperature relations were calculated by computing a grid of envelope models for different combinations of surface temperatures and hydrogen columns and fitting an analytic relation to these results. More details on this approach can be found in Appendix B.1 and Wijngaarden et al. (2019). In Figure 3.1, we show the resulting relation between surface temperature and bottom boundary temperature for varying hydrogen column sizes. For a H-C mixture,

increasing the hydrogen column leads to a better heat-conducting envelope (i.e., the same bottom boundary temperature corresponds to a larger surface temperature). Note that increasing the hydrogen column in a H-He envelope has the opposite effect (see, e.g., Beznogov et al. 2016). Analytic fits for the T_s-T_b relations as a function of hydrogen column size (y_H) and scalable by surface gravity are given in Appendix B.1.

The T_s-T_b relation is highly sensitive to the composition in the *sensitivity strip* of the envelope (where the dominant opacity changes between radiative and conductive). Thus when the change in hydrogen column size does not affect the composition in the sensitivity strip, the effect on the T_s-T_b relation is negligible. This is shown in Figure 3.2, where we explore the parameter space of T_s and y_H and show for which values the temperature relations are sensitive to T_s and y_H. The red and blue regions in Figure 3.2 show for which hydrogen column sizes the temperature relations are insensitive to further changes in column size and correspond to the maximum and minimum T_b for a given T_s, respectively.

Including DNB leads to a clear upper limit on the hydrogen column size, as the build-up of larger columns is prevented by rapid hydrogen burning. Similar to what was found for H-He envelopes, we find that hydrogen columns $> 10^7$ g cm^{-2} cannot be sustained when DNB is taken into account. Note that for most temperatures, this means that a large portion of hydrogen columns located in the transition region are unphysical, leading to a smaller range in allowed boundary temperatures for a given surface temperature. The boundary temperatures corresponding to the excluded hydrogen columns (see Figure 3.2) are highlighted in grey in Figure 3.1. As shown by Figure 3.1 and 3.2, the main effect of DNB for the temperature relations is not the alteration of the composition profile in the sensitivity strip, but the introduction of naturally excluded hydrogen column sizes.

3.3 DNB luminosity

We compare the total DNB energy deposited in the envelope to the luminosity flowing from the interior (i.e., from any layer deeper than the envelope). The energy generated by DNB is converted to a total DNB luminosity by integrating the local energy generation rate over the mass in the envelope:

$$L_{\rm DNB} = \int \epsilon \, {\rm d}m = \int \epsilon \, 4\pi R^2 {\rm d}y \ {\rm erg \ s^{-1}}, \tag{3.1}$$

where M is the total mass of the star, and the mass interval dm is converted to column density interval dy following Gudmundsson et al. (1983). The local energy generation rate is calculated as

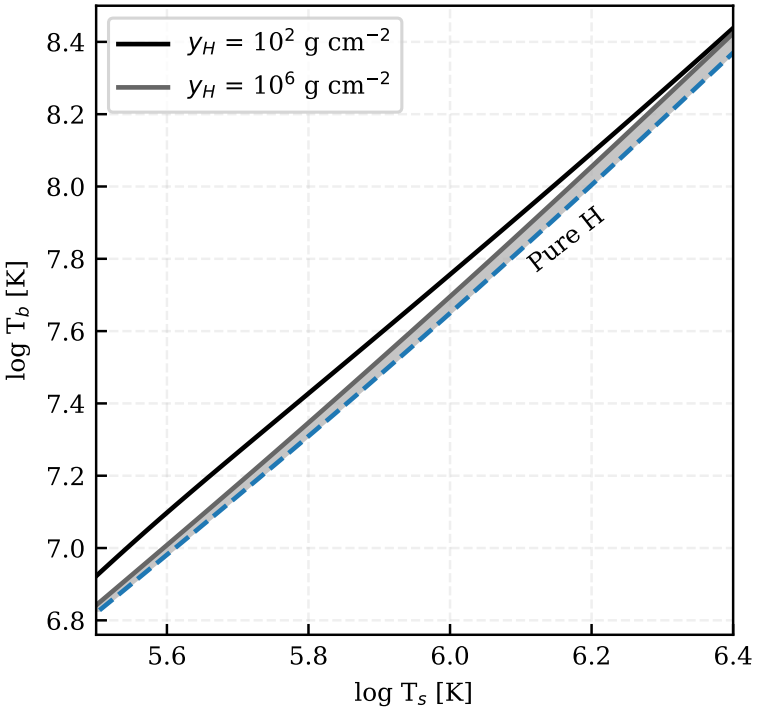

FIGURE 3.1: The T_s-T_b relation for the H-C envelope with $\rho_b = 10^{10}$ g cm^{-3}. For hydrogen column densities of 10^2 g cm^{-2} and 10^6 g cm^{-2}, the solid lines corresponds to both the model with and without DNB, as they are indistinguishable. The blue dashed line corresponds to an hydrogen column of 10^{10} g cm^{-2} when DNB is not included. The grey shaded region corresponds to the grey region in Figure 3.2 and shows the region in the $T_s - T_b$ relations that is excluded by taking DNB into account as the build-up of those hydrogen columns are prevented by nuclear burning.

$$\epsilon = \sum \frac{Q_i r_i}{\rho} \text{ erg g}^{-1} \text{ s}^{-1}, \tag{3.2}$$

where Q_i is the energy released by a reaction of type i, r_i is the reaction rate of reaction type i, and ρ is the local density. The luminosity flowing from the interior through envelope, is typically assumed to be constant, such that the luminosity at the surface is $L_{\text{int}} = L_s = 4\pi R^2 \sigma T_s^4$ which can be used as a boundary condition in thermal evolution codes (see, e.g., Chamel and Haensel 2008; Brown and Cumming 2009; Page and Reddy 2013; Potekhin et al. 2015).

In Figure 3.3 we show conditions (T_s, y_{light}) when the nuclear burning luminosity becomes relevant by computing the ratio $L_{\text{DNB}}/L_{\text{int}}$ for $y_{\text{light}} = y_{\text{H}}$ (left) and $y_{\text{light}} = y_{\text{He}}$ (right). For H-C envelopes, the dominant nuclear reactions are proton captures, which efficiently produce heat for a relatively large part of the y_{H} -T_s parameter space. For surface temperatures above $\sim 10^6$ K and hydrogen columns larger than $\sim 10^5$ g cm^{-2}, the energy deposited by nuclear burning becomes comparable to and larger

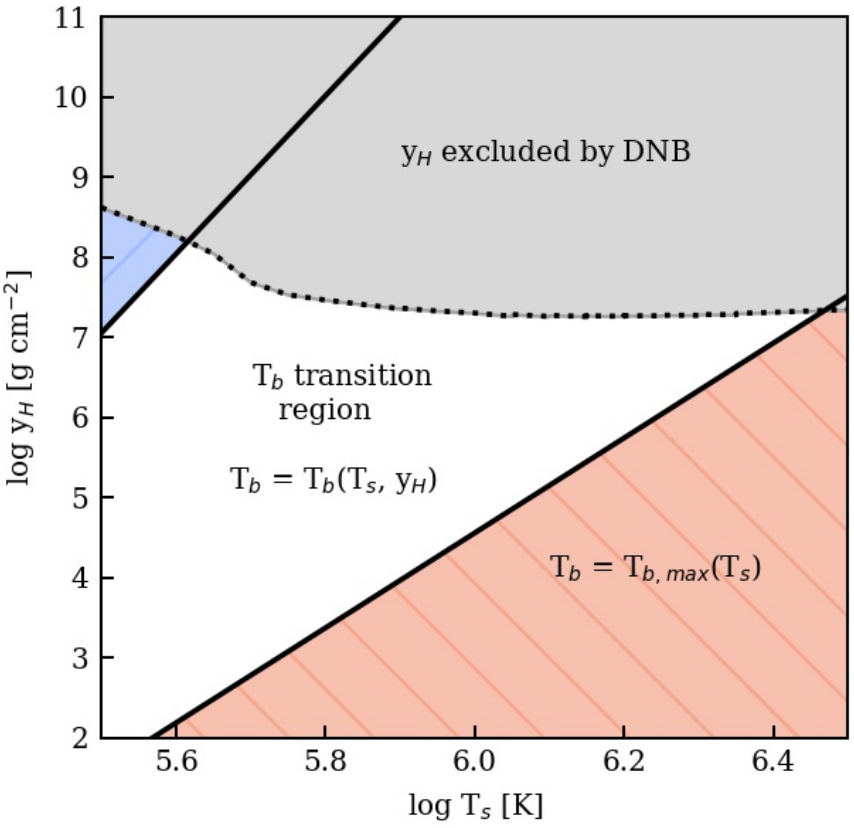

FIGURE 3.2: T_s-y_H parameter space for the H-C envelope model with $\rho_b = 10^{10}$ g cm^{-3}. The solid lines enclose the region where the boundary temperature is sensitive to the hydrogen column density, i.e. $T_b = T_b(T_s, y_H)$. In this region, for a given T_s and increasing y_H, the boundary temperature changes from that corresponding to a pure carbon envelope (red), where $T_b = T_{b,max}(T_s)$, to that of a pure hydrogen envelope (blue), where $T_b = T_{b,min}(T_s)$. Further changes in the hydrogen column outside of the transition region have negligible effect on the T_b-T_s relations. The grey shaded region shows the parameter space that is excluded due to DNB (see text). Note that at large y_H, some excluded hydrogen columns overlap with the transition region. This means that the range in T_b is smaller when DNB is taken into account.

than the interior luminosity (see left panel in Figure 3.3). The hydrogen column densities and surface temperatures where the DNB luminosity becomes non-negligible are in the range inferred from observations of neutron stars in LMXBs after accretion outbursts (see Wijnands et al. 2017 for an observational overview). In Figure 3.4 we show the magnitude of the nuclear burning luminosity for varying hydrogen column (for $T_s = 1.26 \times 10^6$ K and $T_s = 3.16 \times 10^6$ K) in the top panel, and varying surface temperature in the bottom panel (for $y_H = 3.16 \times 10^4$ g cm^{-2} and $y_H = 3.16 \times 10^6$ g cm^{-2}). As an initially large hydrogen column after an accretion episode is not unlikely, the energy deposited in the envelope by diffusive nuclear burning could affect the surface temperature evolution after the outburst.

For He-C envelopes, where the dominant nuclear reaction are highly temperature

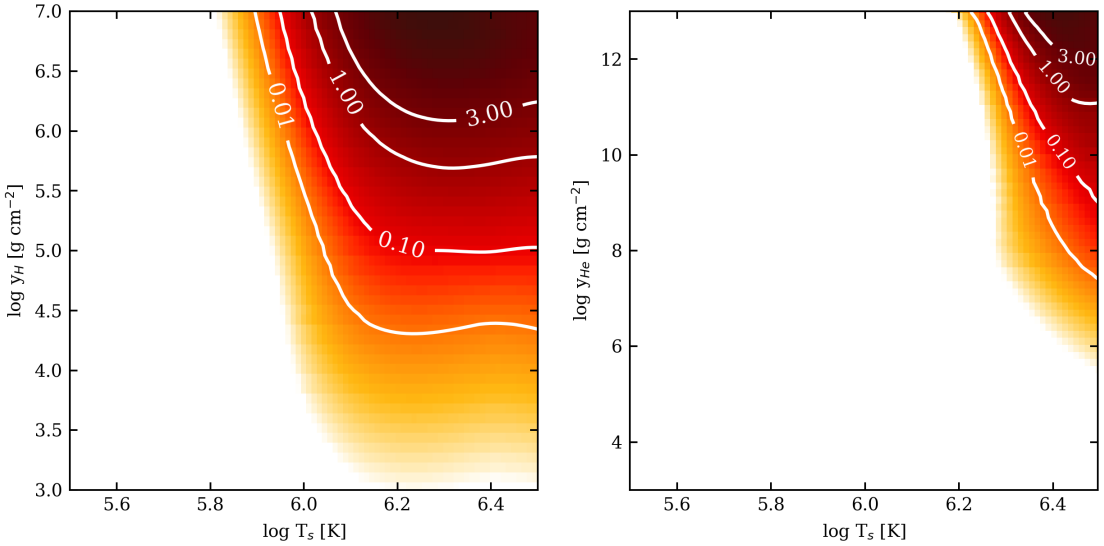

FIGURE 3.3: The explored T_s - y_{light} parameter space for H-C envelopes (left) and He-C envelopes (right). The color is set by the ratio between the DNB luminosity and blackbody luminosity (L_{DNB}/L_{int}, where $L_{int} = L_s = 4\pi R^2 \sigma T_s^4$). The colored region shows where the luminosity due to diffusive nuclear burning in the envelope becomes relevant compared to the blackbody luminosity for varying surface temperatures and light element columns. The contour line corresponding to unity indicates where L_{DNB} becomes larger than L_{int}. The figure parameter ranges are chosen to match relevant values of surface temperatures and light element column densities in crust cooling studies.

sensitive alpha-captures, the heat generated by nuclear reactions is only relevant for very high temperatures, $T_s > 2 \times 10^6$ K, and helium columns, $y_{He} \gtrsim 10^9$ g cm^{-2}, as is shown in the right panel of Figure 3.3. In this regime, the luminosity from nuclear burning becomes larger than 10^{34} erg s^{-1}, as is shown in Figure 3.5. For smaller columns and temperatures, the nuclear burning luminosity rapidly drops and is negligible compared to the interior luminosity (indicated by the dashed lines). The nuclear burning luminosity is relevant compared to the interior luminosity (indicated by the dashed lines) for a smaller range of (and at significantly higher) surface temperatures and column densities compared to H-C envelopes.

3.4 DNB Effect on post-outburst cooling curves

We investigate the heating effect of DNB on the post-outburst temperature evolution of a cooling NS crust. We simulate the accretion outburst to obtain the post-outburst interior luminosity evolution using the relativistic cooling code *NSCool* (Page, 2016) for a NS with mass $M = 1.4\ M_{Sun}$ and radius $R = 11.5$ km. The outer boundary condition (at the bottom of the envelope, ρ_b) of the cooling code is that the interior luminosity at ρ_b is the same as the luminosity at the surface (i.e., no energy is generated or lost in the envelope): $L_s = L_{int}$. Thus, we obtain the surface temperature without DNB ($T_{s,old}$)

directly from $L_{\rm int}$. In the following, we will assume the heat generated by DNB in the low density envelope flows towards the surface and DNB does not affect the interior luminosity. Therefore, we can calculate the heating effect of DNB by post-processing the interior luminosity and surface temperature evolution from *NSCool*.

We convert the surface temperature without DNB ($T_{\rm s,old}$) to a surface temperature with DNB ($T_{\rm s}$) using $L_{\rm s} = L_{\rm int} + x_{\rm out} L_{\rm DNB}$, where $L_{\rm DNB}$ is the DNB luminosity and we introduce $x_{\rm out}$ as the fraction of DNB luminosity that is radiated outwards so that it is easy to investigate the effect of relaxing our assumption of $x_{\rm out} = 1$. The resulting surface temperature including the heat generated by DNB can then be calculated as

$$T_{\rm s} = T_{\rm s,old} \left(1 + \frac{x_{\rm out} L_{\rm DNB}}{L_{\rm int}}\right)^{1/4}. \tag{3.3}$$

For a given M and R, the nuclear burning luminosity for a H-C envelope can be calculated as a function of $T_{\rm b}$ and $y_{\rm H}$ (analogue to Figure 3.3). Before we can calculate the nuclear burning luminosity for a given $T_{\rm b}$, we convert the surface temperature, $T_{\rm s,old}$, to the corresponding $T_{\rm b}$ using the analytic envelope relations for a H-C envelope described in Section 3.2.1.

3.4.1 Effect on cooling curves due to heat generation

In Figure 3.6 and 3.7, we show the resulting surface temperature evolution for an accretion outburst duration of 1 and 10 years, respectively, with constant accretion rates. As the amount of heat deposited on the star during the outburst is sensitive to the outburst duration and accretion rate, we vary either one of these parameters to obtain varying interior luminosity strengths after the outburst. We compare the surface temperature with and without DNB, for two initial hydrogen columns and use the burning rate at each timestep to evolve the hydrogen column size.

Figures 3.6 and 3.7 clearly show that the importance of DNB depends sensitively on the initial hydrogen column after the accretion outburst, as this sets the amount of hydrogen available for burning. For initial hydrogen columns $< 3.2 \times 10^5$ g cm^{-2}, the change in surface temperature due to DNB is $< 5\%$. For larger initial hydrogen columns ($y_{\rm H} \gtrsim 3.2 \times 10^6$ g cm^{-2}), which are typically assumed in cooling studies for a solar mixture of light elements, the initial change in surface temperature is >5 - 50 % depending on the interior luminosity strength set by the accretion history. In the following, we discuss the low and high post-outburst luminosity scenarios for the case when the cooling starts with a large initial hydrogen column ($\log y_{\rm H}$[g cm^{-2}] = 6.5).

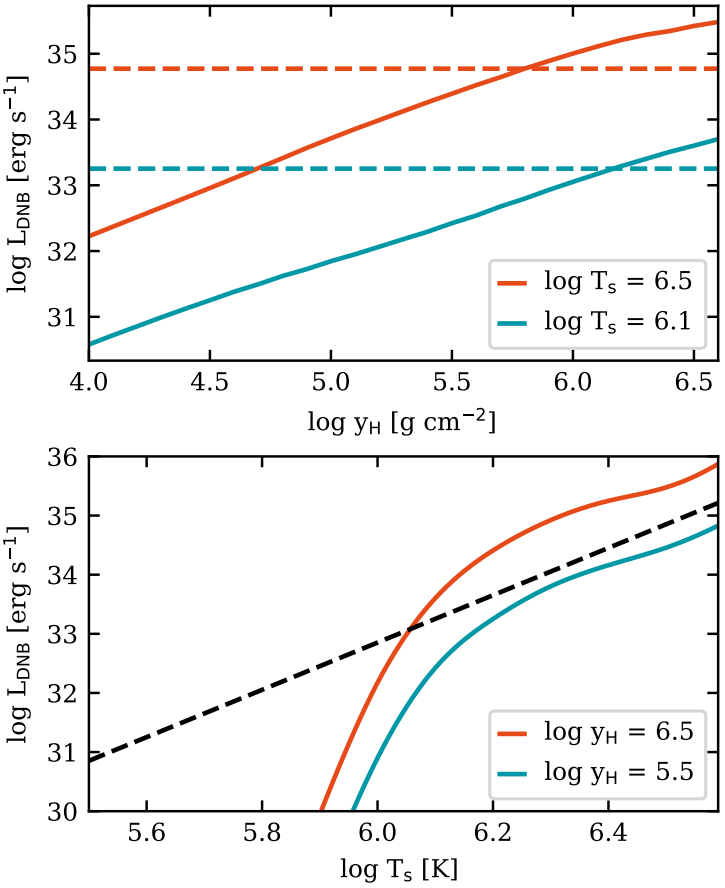

FIGURE 3.4: The DNB luminosities in a H-C envelope, for varying column densities (top) and temperatures (bottom). The luminosity corresponding to the surface temperatures are shown as dashed lines for comparison.

Figure 3.6 corresponds to the scenario of a relatively short accretion period, resulting in low post-outburst cooling luminosities and surface temperatures. Despite the relatively low temperatures, $T^0_{s,old} = 1.2 \times 10^6$ K (corresponding to $kT_s = 103$ eV), when there is a substantial hydrogen column, the heating due to DNB is large and increases the initial post-outburst surface temperature by ∼19% to $T^0_s \sim 1.4 \times 10^6$ K (corresponding to $kT_s = 120$ eV). The absolute increase in initial surface temperature is $\Delta T_{DNB} = 2.2 \times 10^5$ K. For these temperatures, the duration of higher surface temperatures due to heat released by DNB is ∼ 10 days.

For longer accretion outbursts, which result in higher post-outburst luminosities (see Figure 3.7), the initial increase in surface temperature due to DNB is larger (when the initial hydrogen column is large enough, $y_H \gtrsim 1.2 \times 10^6$ g cm^{-2}). In this case, the initial surface temperature increases by 47% corresponding to an absolute increase in temperature of $\Delta T_{DNB} = 7.7 \times 10^5$ K. This means that the initial surface temperature has increased from $T^0_{s,old} = 1.6 \times 10^6$ K (corresponding to $kT_s = 141$ eV), to $Ts^0 \sim 2.4 \times 10^6$ K (corresponding to $kT_s = 204$ eV). However, higher post-outburst luminosities

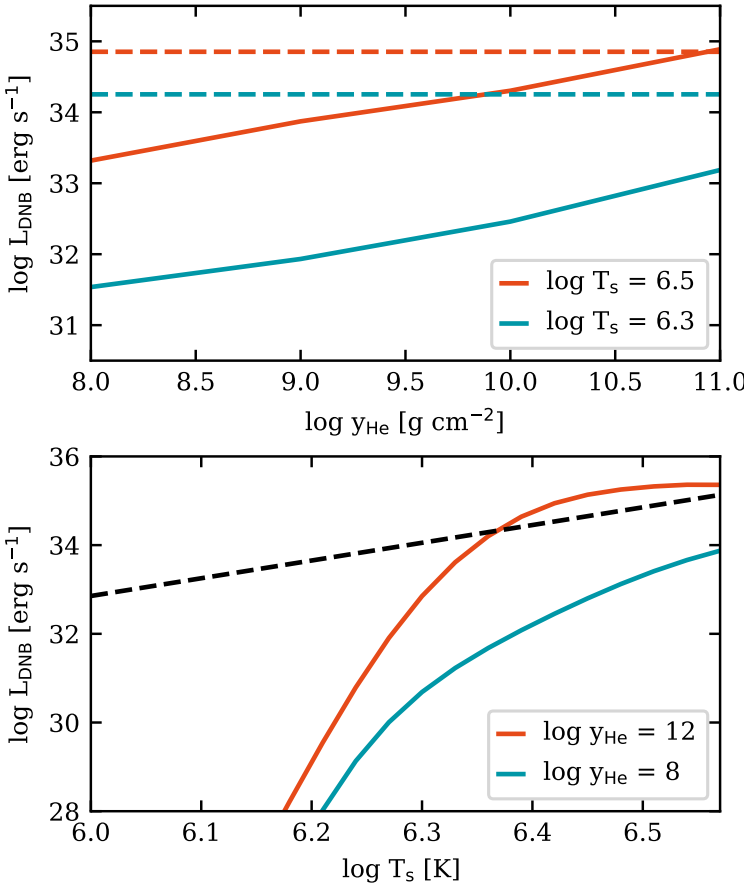

FIGURE 3.5: The DNB luminosities in a He-C envelope, for varying column densities (top) and temperatures (bottom). The luminosity corresponding to the surface temperatures are shown as dashed lines for comparison.

lead to shorter durations (∼hours to days) when DNB contributes significantly to the total luminosity and surface temperature. The total amount of energy available from burning hydrogen is set by the initial hydrogen column, and corresponds to the integral of the elevated luminosity cooling curves, which is constant in all cases with the same initial hydrogen column. For the initial hydrogen column of $y_H \sim 1.2 \times 10^6$ g cm^{-2}, the total amount of energy that is released from hydrogen burning is 10^{38} erg (i.e., ∼7 MeV per nucleon).

Without residual accretion or a large helium buffer that may slow down the depletion of hydrogen, the times when DNB significantly increases the observed surface temperature for a H-C envelope are ∼ 0.1 - 80 days from the end of the accretion outburst, depending on the initial hydrogen column and cooling luminosity. For larger initial cooling luminosities (and thus higher surface temperatures), the initial change in surface temperatures increases, while the relevant duration decreases, as the available hydrogen is consumed more rapidly.

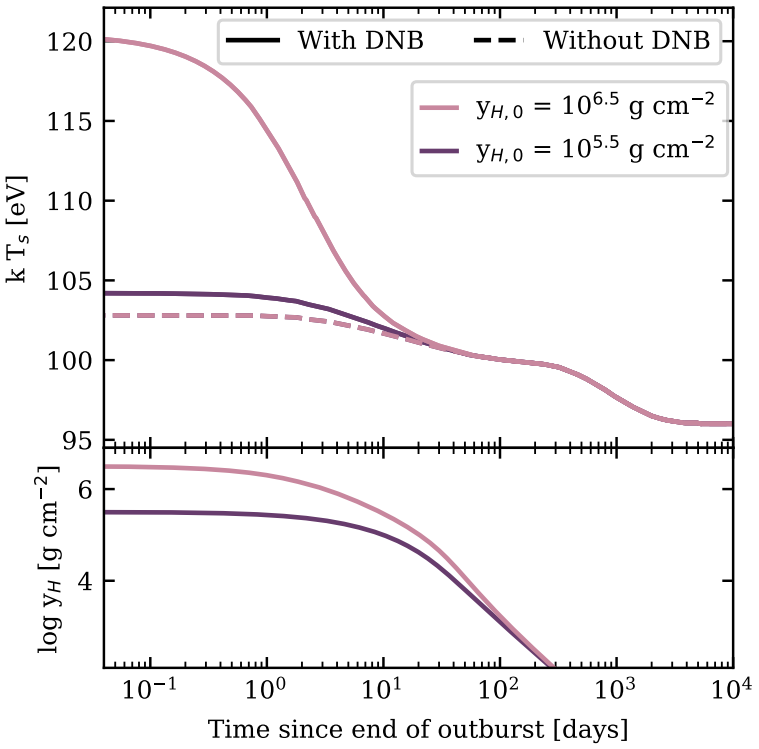

FIGURE 3.6: Comparison of the post-outburst surface temperature evolution with (solid) and without (dashed) additional DNB heating for different initial hydrogen columns (see line colours). The interior luminosity was calculated using an outburst duration of 1 year and $\dot{M} = 10^{17}$ g s$^{-1} \approx 2 \times 10^{-9} M_{\rm Sun}$ yr^{-1}.

3.4.2 Effect on cooling curves due to changing composition

In addition to releasing energy through nuclear reactions, DNB can alter the observed cooling after accretion outbursts by altering the light element column in the envelope over time. In Section 3.2.1, we show the effect on the $T_{\rm s}$-$T_{\rm b}$ temperature relations for varying hydrogen column sizes in H-C envelopes (see Figure 3.1). The decreasing hydrogen column due to DNB, will only affect the cooling curve when the hydrogen column is varied in the transition region (as shown in Figure 3.2). At the relevant temperatures for NSs in quiescence (with initial post-outburst surface temperatures $T_{\rm s} > 8 \times 10^5$ K), this means that the same boundary temperature (at the bottom of the envelope) corresponds to different surface temperatures when the hydrogen column varies between 10^4 g cm$^{-2} \lesssim y_{\rm H} \lesssim 10^7$ g cm^{-2}. When the initial hydrogen column after the accretion outburst is smaller than $y_{\rm H} \sim 10^4$ g cm^{-2}, a further decreasing hydrogen column will not affect the cooling curve. When the initial hydrogen column is larger than $y_{\rm H} \sim 10^4$ g cm^{-2}, the cooling curve transitions from relatively hydrogen-rich envelope temperature relations towards those for a hydrogen-poor envelope. We show this effect in Figure 3.8, where we ignore the effect of heat released by DNB and only show the effect of a changing envelope composition. The dashed cooling curves in Figure 3.8 correspond to fixed hydrogen

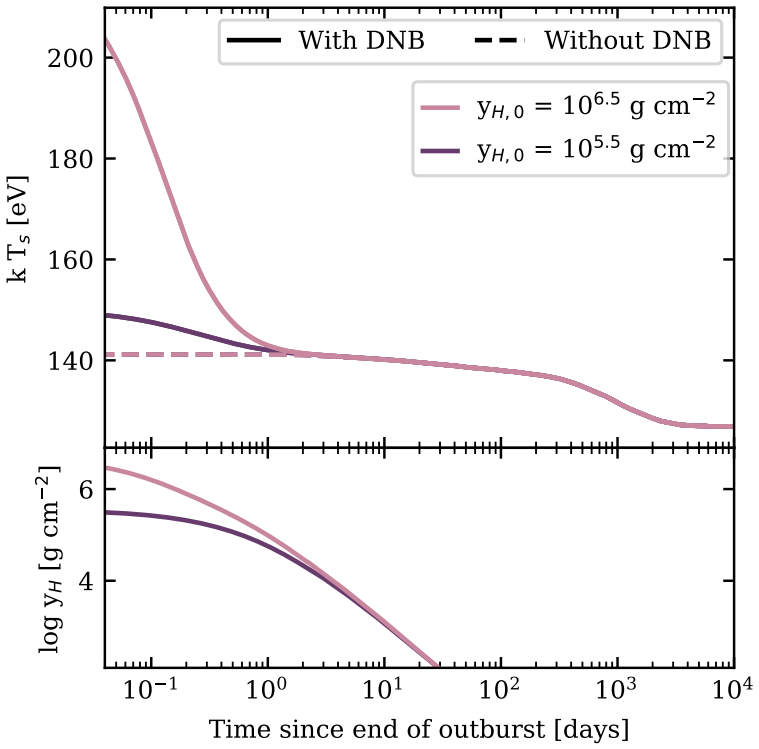

FIGURE 3.7: Comparison of the post-outburst surface temperature evolution with (solid) and without (dashed) additional DNB heating for different initial hydrogen columns (see line colours). The interior luminosity was calculated using an outburst duration of 10 years and $\dot{M} = 10^{17}$ g s^{-1} $\approx 2 \times 10^{-9} M_{\rm Sun}$ yr^{-1}.

columns, while the solid cooling curves correspond to the scenario where the hydrogen column decreases over time due to DNB. When DNB is taken into account, all the plotted cooling curves (irrespective of initial column density) correspond to that of a hydrogen-poor envelope after $\sim$ 50 days. Note that for larger post-outburst luminosities, it can take less time ($\sim$ 1 day) to decrease the hydrogen column below $y_{\rm H} \sim 10^4$ g cm^{-2} (see, e.g., the bottom panel in Figure 3.7).

Currently, the unknown post-outburst envelope composition introduces an uncertainty in cooling models through the $T_{\rm s}$-$T_{\rm b}$ relations and is often left as a free fit variable. By taking into account DNB, it is possible to further constrain the post-outburst envelope composition, although the initial composition that is left after the accretion outburst is still uncertain. Regardless of the initial hydrogen column, we find that our cooling curves correspond to the hydrogen-poor limit within $\sim$100 days.

3.5 Application to crust cooling sources

In this section, we calculate cooling curves including DNB for observed crust-cooling sources in LMXBs. We select sources for which shallow heating was invoked in order

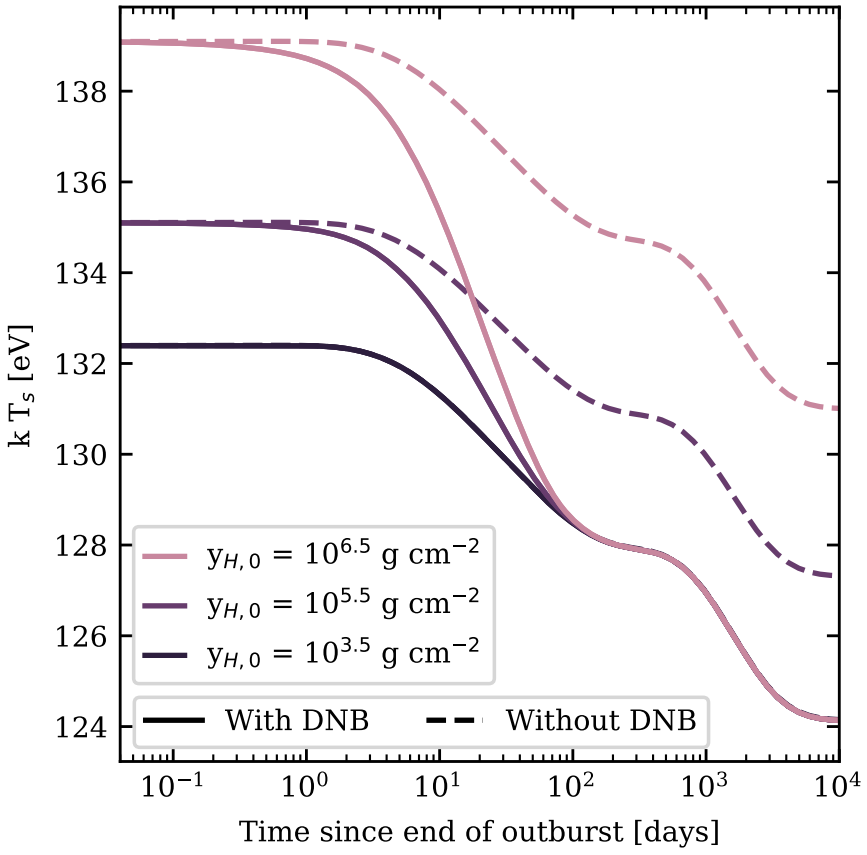

FIGURE 3.8: The effect of a changing hydrogen column due to DNB (while ignoring heat release, see text) on the the post-outburst surface temperature evolution for different (initial) hydrogen columns (see linecolors). The interior cooling was calculated using an outburst duration of 1 years and $\dot{M} = 10^{17}$ g s^{-1} $\approx 2 \times 10^{-9} M_{\rm Sun}$ yr^{-1}.

to explain the observed temperatures (see Section 3.1). We model the sources without shallow heating to show how the quiescent heating due to DNB compares to observations. We use the method described in Section 3.4, but include the time-variable accretion rate described in Ootes et al. (2016) to estimate the effect of DNB luminosity on cooling curves.

3.5.1 Aql X-1

The NS in the Aql X-1 LMXB displays frequent accretion outbursts with typical durations of a few months (see, e.g., Waterhouse et al. 2016). Due to its frequent outbursts, Aql X-1 is a promising source for the study of shallow heating during outbursts. Ootes et al. (2018) modelled its thermal evolution for the period between 1996-2015, which includes 23 accretion outbursts. They found that the observed cooling after multiple outbursts could be explained using different magnitudes and depths of the shallow heating source. New temperature data for the cooling after its 2016 outburst was presented by Degenaar et al. (2019), who further investigated the

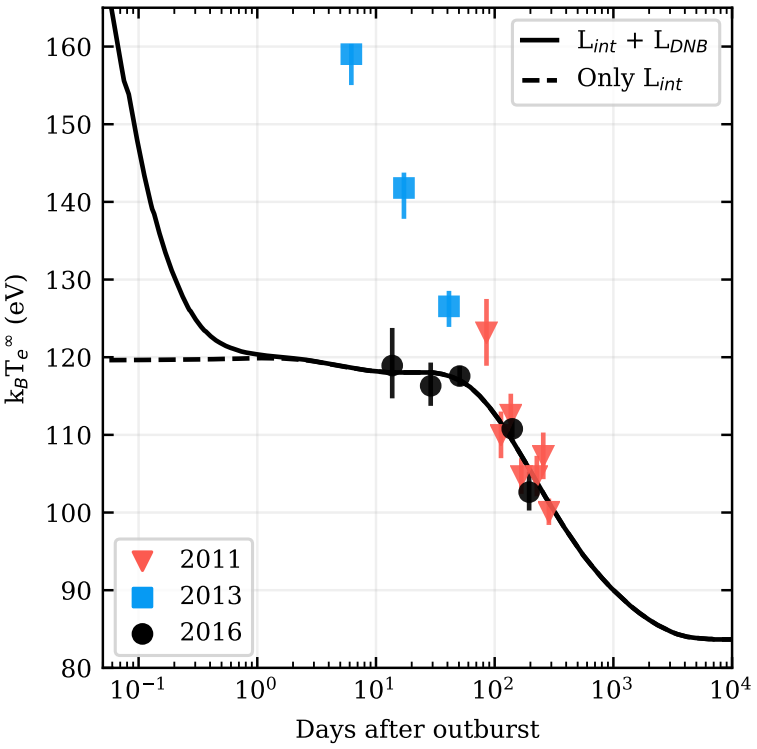

FIGURE 3.9: Illustrative cooling curves for the neutron star in Aql X-1 compared to the inferred temperatures after its 2011 (red triangles), 2013 (blue squares) and 2016 (black circles) outbursts. The cooling curves in this graph correspond to the same cooling model with the same NS properties. The differences between the curves are due the absence/presence of DNB. The dashed line shows a cooling model that includes a shallow heat source with $Q_{\rm shallow}$ = 1.3 MeV nuc^{-1}. The solid line shows the same cooling model when heating due to DNB is included. As the hydrogen column is quickly depleted at these temperatures, the DNB luminosity is only relevant $<$ 1 day.

shallow heating properties by comparing the cooling after its 2011, 2013 and 2016 outbursts. The three outbursts showed a striking similarity in outburst properties such as duration, peak flux and overall shape, but were followed by cooling tracks with very distinct differences in the early time thermal evolution. Degenaar et al. (2019) found that the depth and magnitude of the shallow heating source during the 2016 outburst must have been larger than during the 2013 outburst[2]. This indicates that the shallow heating properties can be different after multiple similar outbursts in the same source. Here, we briefly explore an alternative scenario for the differences in the early time cooling curves by calculating the heating effect from DNB of accreted hydrogen.

We use *NSCool* to calculate an illustrative cooling model that loosely follows the inferred temperatures after the 2016 outburst as well as the late-time cooling of the 2011 and 2013 outbursts. We do not attempt to find a best-fit model, as we only intend to investigate whether a cooling model with the same NS, accretion- and shallow

[2]No strong constraints could be placed on the shallow heating properties during the 2011 outburst due to the lack of observations during the early cooling phase $<$80 days (Ootes et al., 2018; Degenaar et al., 2019).

heating properties could explain the observed differences at early times. We use a mass and radius of M = 1.6 $M_{\rm Sun}$ and R = 11 km, which are consistent with those used in the spectral fits to obtain the temperature data (Ootes et al., 2018; Degenaar et al., 2019). As the three outbursts had similar durations and accretion rates, we use an average accretion rate of $\dot{M} = 10^{18}$ g s^{-1} for an outburst duration of 2.5 months. Our model assumes a core temperature of $T_{\rm core} = 1.1 \times 10^8$ K and includes a shallow heating source of $Q_{\rm shallow} = 1.3$ MeV nuc^{-1} at a maximum depth of $\rho = 2.5 \times 10^{10}$ g cm^{-3}.

In Figure 3.9, we compare the inferred temperatures after the three outbursts to the cooling track given by our model (dashed line) as well as the track given by the same model when luminosity from DNB is added (solid line). When including the DNB luminosity, we assume an initial hydrogen column after the outburst of $y_{\rm H}= 1.2 \times 10^6$ g cm^{-2}. We find that, due to the high temperatures, the hydrogen column drops rapidly to $y_{\rm H}= 3.2 \times 10^4$ g cm^{-2} within 1 day and no significant luminosity is produced by DNB after that time. Therefore, the cooling curves with and without a DNB luminosity are the same after 1 day and DNB of hydrogen alone does not explain the observed differences up to 80 days into the cooling. It is also unlikely that the differences in early time cooling between the outbursts can be explained with DNB and residual low-level accretion, as the accretion rates required produce a luminosity that would dominate the total observed luminosity.

Note that in all our cooling curves including DNB, we made the assumption that the generated heat immediately leaves the star at the surface and has no effect on the interior cooling. Therefore, the cooling curves are only affected during active DNB. An inwards heatflow could potentially prolong the timescale over which DNB affects the cooling curve when accretion is not included, as the heat will reach the surface at later times. However, as the heat is produced at shallow depths, this effect is expected to be small.

3.5.2 MAXI J0556-332

The cooling after the 2012 outburst of MAXI J0556-332 is the most extreme case for which shallow heating has been invoked. The temperatures after the end of the outburst, are by far the hottest ever observed. A very large amount of shallow heating, $\sim$15 MeV nuc^{-1}, was necessary during the outburst to explain the observed cooling after the outburst was over (Homan et al., 2014; Deibel et al., 2015; Parikh et al., 2017), which is about an order of magnitude more than what was needed for other sources. Interestingly, for the cooling after two subsequent reheating outbursts in this source, much less shallow heating was needed (Parikh et al., 2017).

We discuss DNB in relation to this source, as it poses the most extreme scenario in terms of both shallow heating and large post-outburst temperatures. In fact, the source is so hot that besides proton captures onto carbon, another nuclear reaction needs to be included. At temperatures below 10^8 K, a result of the proton capture onto C is the creation of ^{13}N, which decays into ^{13}C. At temperatures $T > 10^8$ K, the time in which ^{13}N captures hydrogen is shorter than the time in which it decays, thus adding another reaction in which additional heat is released. Therefore, we also include the second step of the hot CNO cycle, such that the hydrogen capturing reactions are:

$$\mathrm{H}_1^1 + \mathrm{C}_6^{12} \rightarrow \mathrm{N}_7^{13} + \gamma, \tag{3.4}$$

$$\mathrm{H}_1^1 + \mathrm{N}_7^{13} \rightarrow \mathrm{O}_8^{14} + \gamma, \tag{3.5}$$

where the second reaction deposits an additional amount of energy in the envelope (i.e, 4.63 MeV per reaction).
We find that, even with the inclusion of hydrogen captures onto nitrogen, DNB alone does not produce enough heat by far to explain the large observed temperatures observed in MAXI J0556-332. At these temperatures, the burning rate is so high that large accretion rates are required to maintain the hydrogen column at levels at which enough heat is produced. Even with the inclusion of decaying low-level accretion rates after the outburst, we find that very large initial accretion rates ($\dot{M}_{\rm init} \sim 10^{15}$ g s^{-1}) are required to reach the observed post-outburst temperatures. These accretion rates can be ruled out, as the corresponding accretion luminosity ($\sim 10^{35}$ erg s^{-1}) would be significantly larger than the observed luminosity. We also consider a possible He-C envelope, as the conditions are suitable for helium burning. However, we find that while a large helium column ($y_{\rm He} \sim 10^{12}$ g cm^{-2}) will increase the temperature during the full cooling time, the shift is not large enough to explain the observed temperatures. Smaller helium columns will not make a significant contribution to heat generation (see luminosity in Figure 3.3).

In Figure 3.10, we show the relative change in column size for varying initial light element columns. For H-C envelopes, we find that all hydrogen is rapidly depleted within ∼20 days, regardless of the initial column size after the accretion outburst. For He-C envelopes, where only helium that penetrates deep in the envelope is consumed by captures onto heavier elements, the helium column can decrease but is not completely consumed in the crust-cooling time. It is interesting that observations of a hydrogen atmosphere in this source, may be evidence for the presence of low-level accretion in quiescence as all initial hydrogen would be consumed within 20 days. Note that only a small accretion rate is needed to regain a hydrogen atmosphere (see Chang and Bildsten 2004; Wijngaarden et al. 2019 for a discussion on the effect of DNB on the composition of the atmosphere).

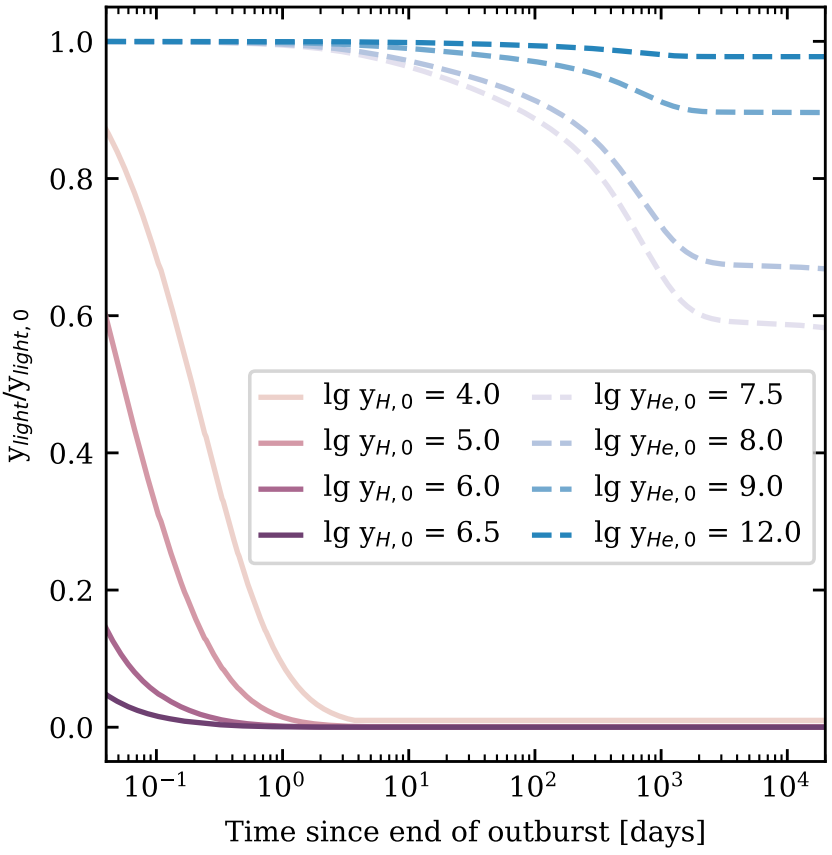

FIGURE 3.10: Evolution of the light element column through DNB of MAXI J0556-332 for hydrogen (in a H-C envelope) and helium (in a He-C envelope) for varying initial column sizes. All initial hydrogen is consumed in the H-C envelope, while initially large helium column sizes decrease but remain present despite the relatively large temperatures.

We note that other nuclear reactions may be taking place which are currently unaccounted for. Liu et al. (2017) calculate cooling curves for MAXI J0556-332 without shallow heating but including the full beta-decay limited hot CNO cycle and while assuming a decaying accretion rate during quiescence. In their calculation, nuclear burning does not occur diffusively but follows a fixed burning rate for the hot CNO cycle (independent of the amount of hydrogen available and the density) in the region of the envelope where $T > 10^9$ K and changes in abundances are not taken into account. They conclude that the hot CNO cycle can significantly increase the surface temperature and find that the hot CNO cycle operates up to $\sim$500 days into quiescence. However, as we show here, when changes in abundance are taken into account (i.e., hydrogen depletion due to nuclear burning) the increase in luminosity due to hydrogen burning effectively shuts off much sooner (within 20 days if no further accretion takes place).

3.6 Discussion

3.6.1 General effects of DNB

We study the effect of diffusive nuclear burning on observed neutron star crust cooling in three ways: 1) by obtaining static temperature relations for H-C envelopes that include DNB (Section 3.2.1), 2) by considering the additional heating due to DNB (Section 3.4.1), and 3) by considering time variable light element columns (Section 3.4.2). It is typically assumed that no heat is generated in the NS envelope during quiescence, but here we find that the DNB luminosity is not always negligible. Specifically, the DNB luminosity is relevant in H-C envelopes for hydrogen columns $\gtrsim 10^5$ g cm^{-2} at surface temperatures $\gtrsim 1$ MK, and for He-C envelopes for helium columns $\gtrsim 10^{11}$ g cm^{-2} at surface temperatures $\gtrsim 2$ MK. Thus, the surface luminosity can be dominated by DNB for relevant parts of the parameter space (see Figure 3.3) which can affect the interpretation of early time surface temperature observations. However, as the light element column decreases rapidly, the DNB luminosity is only relevant at times within $\sim$80 days after the outburst for low post-outburst temperatures and within $\sim$1 day for larger post-outburst temperatures (see Section 3.4).

We find that the conditions in quiescent NSs in LMXBs are such that active DNB can rapidly decrease the light element column size on relatively short timescales (within $\sim$1-100 days). Good data sampling in the early phases ($<$50-100 days) are necessary to see the effects. This result can be relevant for crust cooling studies, as the T_s-T_b relations are not fixed for one column size during the time the crust cools back to thermal equilibrium with the core. Instead, the T_s-T_b relations change accordingly when the light element column size decreases in the transition region (see Figure 3.2). This could manifest in two ways:

- For cooler post-outburst temperatures, the hydrogen column is consumed at a lower rate and the corresponding change in T_s-T_b relations is potentially observable as an initial decrease in surface temperature at $\lesssim 80$ days after the end of the outburst (see Figure 3.8). This initial cooling is dominated by the change in composition and thus transparency of the envelope, rather than cooling of the interior.
- For hotter post-outburst temperatures, the hydrogen column is consumed at a sufficiently short timescale (within hours to a few days) such that the full observable cooling curve corresponds to a static, hydrogen-poor envelope. Limits on the envelope composition due to rapid DNB at early times can be used to constrain the static envelope composition at later times if the post-outburst

accretion rate is assumed to be negligible for replenishing the light element column in the sensitivity strip.

Our main conclusion is that the current static treatment of the neutron star envelope during thermal evolution studies, even on the short timescales considered in cooling studies, may not be appropriate. Processes such as DNB can significantly alter the envelope composition over time and can release additional heat, making it worthwhile to include the burning region of the envelope in future thermal evolution calculations. This would require setting the upper boundary of thermal transport equations to lower densities than $\rho_b = 10^8$ or $\rho_b = 10^{10}$ instead of using static temperature relations to connect the temperature at those densities to the temperature at the surface (see Beznogov et al. 2020 for a more consistent treatment of the envelope applied to young neutron stars).

3.6.2 DNB as shallow heating mechanism?

We find that DNB can increase the initial post outburst surface temperatures to even larger values than the observed initial temperature, without invoking an unknown shallow heat source. However, the temperature increase due to DNB drops too rapidly (<10 days) to explain the large observed surface temperatures at 50-100 days for Aql X-1. We find that without an additional shallow heating source, large accretion rates are necessary to replenish the hydrogen column which would produce an X-ray luminosity that would dominate the total luminosity. Therefore, DNB of hydrogen alone is not able to lift the need for an additional shallow heating source during the outbursts of Aql X-1. DNB in both H-C and He-C envelopes cannot replace shallow heating for MAXI J0556-332, the source with the most extreme post-outburst temperatures (see Section 3.5.2). The large post-outburst accretion rates needed to match these temperatures with DNB, would overwhelm diffusive burning and produce accretion luminosities that are larger than the observed luminosity.

The amount of shallow heating has been found to differ between sources and even after multiple outbursts in the same source. Nuclear burning could generate different amounts of additional heating for different post-outburst envelope compositions (which depend on the accretion history). We examine whether including both DNB and shallow heating for multiple Aql X-1 outbursts (i.e., the 2011, 2013 and 2016 outbursts from Ootes et al. 2018 and Degenaar et al. 2019) could lift the need for a varying amount of shallow heating between outbursts. We find that at the inferred post-outburst temperatures, the hydrogen in the envelope would be consumed within 1 day and DNB no longer produces a significant luminosity at the time of the cooling observations. Furthermore, the required accretion rates to replenish the hydrogen

column sufficiently to explain the early time cooling observations are expected to produce accretion luminosities much larger than the observed luminosity.

It would be interesting to include nuclear burning in the envelope both during the accretion outburst (non-diffusive burning) as well as in quiescence (diffusive burning), as potential heat flow into the neutron star interior is currently not accounted for and could extend the time when DNB is relevant. The work presented here involves calculating independently the evolutions of interior luminosity, DNB luminosity and size of light element column. It would be worthwhile to perform a self-consistent calculation, which includes a cooling code outer boundary condition that can vary with the burning region. Such a study could also include a full nuclear reaction network to evolve the envelope composition beyond two component models considered here and in previous works.

Part II

Neutron star interior physics from gravitational waves

"…and suddenly we can hear the orchestra of the universe instead of only seeing it"

- Prof. Jo van den Brand

Chapter 4

Analyzing of the full gravitational wave signal from neutron star binary coalescences [1]

Abstract

The gravitational wave signal emitted during the coalescence of two neutron stars carries information about the stars' internal structure. During the long inspiral phase the main matter observable is the tidal interaction between the binary components, an effect that can be parametrically modeled with compact-binary solutions to General Relativity. After the binary merger the main observable is frequency modes of the remnant, most commonly giving rise to a short-duration signal accessible only through numerical simulations. The complicated morphology and the decreasing detector sensitivity in the relevant frequencies currently hinder detection of the post-merger signal and motivate separate analyses for the pre-merger and post-merger data. However, planned and ongoing detector improvements could soon put the post-merger signal within reach. In this study we target the whole pre-merger and post-merger signal without an artificial separation at the binary merger. We construct a hybrid analysis that models the inspiral with templates based on analytical calculations and calibrated to numerical relativity and the post-merger signal with a flexible morphology-independent analysis. Applying this analysis to GW170817 we find, as expected, that the post-merger signal remains undetected. We further study simulated signals and find that we can reconstruct the full signal and simultaneously estimate both the pre-merger tidal deformation and the post-merger signal frequency content. Our analysis allows us to study neutron star physics using all the data available and directly test the pre-merger and post-merger signal for consistency thus probing effects such as the onset of the hadron-quark phase transition.

[1] This chapter is based on publication Wijngaarden et al. (2021), see Publications and author contributions.

4.1 Introduction

Understanding of properties of supranuclear matter in the core of neutron stars (NSs), commonly encoded through their Equation of State (EoS) is an ongoing challenge (Lattimer and Prakash, 2016; Özel and Freire, 2016; Oertel et al., 2017; Baym et al., 2018). Electromagnetic and gravitational wave (GW) observations of the coalescence of two NSs provide an additional tool to study the ultra-dense interiors of NSs, as they affect the emitted signals. Two such GW events, GW170817 (Abbott et al., 2017e) and GW190425 (Abbott et al., 2020), have already been detected by LIGO (Aasi et al., 2015) and Virgo (Acernese et al., 2015) and planned, improved GW detectors promise more and louder detections in the future (Abbott et al., 2013). These data will add to the interdisciplinary effort to study the EoS, including observations of macroscopic NS properties such as heavy pulsar observations (Antoniadis et al., 2013; Cromartie et al., 2019; Fonseca et al., 2021), X-ray pulse-profiling with NICER (Miller et al., 2019; Riley et al., 2019; Miller et al., 2021; Riley et al., 2021), and nuclear theory and experiment (e.g. Raaijmakers et al. 2020; Dietrich et al. 2020; Landry et al. 2020; Al-Mamun et al. 2021; Reed et al. 2021; Essick et al. 2020; Raaijmakers et al. 2021; Biswas 2021; Pang et al. 2021; Legred et al. 2021; Essick et al. 2021b,a).

The GW signal from a binary NS (BNS) coalescence consists of two parts: a pre-merger and a post-merger. In the pre-merger phase the binary components inspiral toward each other as they lose orbital energy to GWs and eventually collide (Blanchet, 2014). In the case of GW170817, the detectable pre-merger signal increased in frequency from $\sim$ 23Hz to many hundreds of Hz as the orbital separation between the two NSs rapidly decayed over $\sim$ 2 minutes (Abbott et al., 2019b) before finally merging at a frequency of $\sim$ 1500Hz (Torres-Rivas et al., 2019). During the late stages of the inspiral, tidal interactions accelerate the binary evolution and leave an imprint on the signal that depends on the NS matter properties and can thus be used to infer the EoS (Chatziioannou, 2020). The pre-merger GW signal is typically modeled with compact binary coalescence (CBC) templates that are based on approximate solutions to the General Relativity field equations and numerical relativity (NR) simulations, see Dietrich et al. (2021) for a recent review.

After the inevitable merger, the remnant star evolves in a way that depends on the system mass and the EoS, see Bernuzzi (2020) for a recent review. For most EoSs and NS masses, a hypermassive NS supported by differential rotation and thermal effects is expected to be formed and sustained for $\sim 10 - 100$ms (Baumgarte et al., 2000). During this time, the NS remnant emits a post-merger GW signal with a characteristic peak at $1500 - 4000$Hz (Xing et al., 1994; Ruffert et al., 2001; Shibata et al., 2005; Shibata, 2005; Shibata and Taniguchi, 2006; Oechslin and Janka, 2007; Hotokezaka et al., 2011; Bauswein and Janka, 2012; Bauswein et al., 2012; Bauswein et al., 2012; Hotokezaka et al., 2013; Takami et al., 2014, 2015; Kastaun and Galeazzi, 2015;

Bernuzzi et al., 2015; Bauswein and Stergioulas, 2015; Foucart et al., 2016; Lehner et al., 2016; East et al., 2016; Dietrich et al., 2017) depending on the mass and EoS, a promising frequency range for upcoming GW detectors[2]. While the exact physics governing the post-merger phase is not fully understood and NR simulations are complicated by factors such as thermal effects, turbulence, magnetohydrodynamical instabilites, neutrinos, and possible phase transitions, simulations make a robust empirical prediction for the frequency of the peak of the post-merger spectrum given the NS mass and EoS.

Information extracted from a post-merger signal about NS properties, is complementary to the pre-merger. Analysis of a post-merger signal would probe different density (Bauswein et al., 2020) and temperature (Raithel et al., 2021; Hammond et al., 2021) regimes of the EoS, as the more massive, hot remnant will have a higher maximum density than the pre-merger NSs (Ruffert et al., 2001). This complementarity can lead to insights about high-density phenomena such as phase-transitions between the densities probed by the pre-merger and post-merger signals (Radice et al., 2017; Most et al., 2019; Bauswein et al., 2019a,b; Weih et al., 2020; Bauswein and Blacker, 2020; Liebling et al., 2021; Prakash et al., 2021). Finally, detection of a post-merger signal would allow us to determine the nature of the merger remnant, namely NS or BH, which has implications for a potential electromagnetic counterpart and its interpretation (Margalit and Metzger, 2019).

In the case of GW1708017 and GW190425 only piecewise analyses of the GW signal have been performed focusing either on the pre-merger or on the post-merger part of the signal (Abbott et al., 2019b, 2017b, 2019a, 2020). This reduces the computational cost significantly, as the post-merger signal requires a large sampling rate (typically 8194Hz) while the premerger analysis involves long duration data segments (typically 128s). As no post-merger signal is detectable yet (Abbott et al., 2019b, 2017b), excluding it from the pre-merger analysis did not bias the results for the binary parameters (Dudi et al., 2018). However, looking forward and in the case of a post-merger signal detection, separate analyses will not be able to exploit phase coherence through merger or parameter relations, such as those between the tidal deformability from the pre-merger signal and the frequency peak of the post-merger spectrum (Bauswein et al., 2020; Bernuzzi et al., 2015).

Another reason for separate pre-merger and post-merger analyses is that the morphology and details of the post-merger signal are not well understood, making it difficult to construct a first-principles physical model such as the pre-merger CBC templates. While approximate analytic models for the post-merger signal have been proposed with the aid of NR simulations (Hotokezaka et al., 2013; Bauswein et al.,

[2]Depending on the NS mass and EoS, other possibilities for the post-merger remnant include direct collapse to a black hole (BH) whose ringdown signal is too high-frequency to be detected, and a long-lived NS that emits a signal that can last for minutes or more. In what follows we only consider the most common case, namely a hypermassive NS and a short-duration GW signal.

2016; Bose et al., 2018; Tsang et al., 2019; Breschi et al., 2019; Easter et al., 2020), these are still limited in accuracy as well as by uncertainties of the simulations they are based on (Breschi et al., 2019). An alternative are morphology-independent analyses that are not limited to a specific signal type (Clark et al., 2014, 2016; Chatziioannou et al., 2017). Specifically, a model-agnostic approach with the `BayesWave` algorithm (Cornish and Littenberg, 2015; Cornish et al., 2021; LIGO Scientific Collaboration and Virgo Collaboration, 2018) has been shown to accurately determine the main features of the post-merger signal (Chatziioannou et al., 2017; Torres-Rivas et al., 2019), and in some cases do so more accurately than tailored models (Easter et al., 2020).

`BayesWave` models a GW signal with a sum of sine-Gaussian wavelets. Both the number of wavelets and their parameters are marginalized over, resulting in a flexible analysis that has been applied to a variety of signals (Kanner et al., 2016; Littenberg et al., 2016; Bécsy et al., 2017; Tsang et al., 2018, 2020; Pankow et al., 2018; Millhouse et al., 2018; Ghonge et al., 2020; Dálya et al., 2021; Chatziioannou et al., 2021b). `BayesWave`'s sine-Gaussian wavelets are particularly adept for post-merger signals that are dominated by distinct frequency components, as each component can be approximately modeled by a wavelet. In this context, `BayesWave` has been used to search for a short-duration, high-frequency signal after both GW170817 (Abbott et al., 2019b) and GW190425 (Abbott et al., 2020), returning null results and upper limits on the energy content (Abbott et al., 2019b).

In this paper we construct a hybrid analysis of the full BNS GW signal that targets both the pre-merger and the post-merger data. We simultaneously analyze the full signal using

1. a CBC template to describe the well-modeled pre-merger phase
2. sine-Gaussian wavelets to capture the less well understood post-merger

The parameters of the CBC template and the wavelets are simultaneously sampled over to obtain the combined multidimensional posterior for all components of the model. Those parameters include the pre-merger tidal deformability that quantifies the inspiral tidal deformation and the peak frequency of the post-merger spectrum. Our analysis extracts both simultaneously and allows for direct consistency comparisons.

After validating the hybrid modeled/unmodeled analysis on a toy model based on GW150914 data, we apply it to GW170817 and find that the post-merger signal remains undetected, as expected from the upper limit estimates of Abbott et al. (2019b). We further analyze simulated signals of high signal-to-noise ratio (SNR) for which the post-merger signal is detectable. We show that our analysis can reconstruct

the pre-merger signal as well as the main components of the post-merger signal such as the dominant frequency mode. We simulate signals for which the pre-merger and post-merger parts are consistent (corresponding to a hadronic EoS) and inconsistent (corresponding to an EoS with a strong phase transition in the relevant density scales) and demonstrate that our analysis can extract either behavior.

The rest of the paper is organized as follows. In Sec. 4.2 we describe in detail our analysis and algorithm. In Sec. 4.3 we present a proof-of-principle analysis on GW150914. In Sec. 4.4 we reanalyze the GW170817 data. In Sec. 4.5 we apply our analysis to simulated data. In Sec. 4.6 we conclude.

4.2 Methodology and Models

In order to analyze a GW signal that contains both modeled and unmodeled features we employ a hybrid analysis that makes use of both CBC templates and flexible models for the signal. We base our analysis on the morphology-agnostic data analysis algorithm `BayesWave` (Cornish and Littenberg, 2015; Cornish et al., 2021) by extending it to account for CBC templates similar to Chatziioannou et al. (2021a). In its core functionality `BayesWave` simultaneously models a GW *signal* (referred to as the "signal model" in `BayesWave` literature), instrumental *glitches* (the "glitch model"), and the gaussian detector *noise* (the "PSD model"). `BayesWave` uses minimal assumptions, i.e., no physical model, to describe these components. Instead, the *signal* and the *glitch* are modeled as the sum of a variable number of sine-Gaussian wavelets, whereas the *noise* model describes the power spectral density (PSD) of the Gaussian noise with a variable number of spline points and Lorentzians. For this work and Chatziioannou et al. (2021a) a fourth component has been added, which models GW signals through CBC templates obtained by solving the two-body problem in general relativity; the *CBC* model.

4.2.1 Overview of Models

`BayesWave` uses sampling methods to characterize the multi- and variable-dimension posterior distribution for independent models that target different features of the GW data. The goal is to draw samples from the posterior distribution function $p(M|d)$, the probability that a model M describes the data d and a prior $p(M)$, defined as

$$p(M|d) = \frac{p(M)p(d|M)}{p(d)}, \tag{4.1}$$

where $p(d)$ is the evidence, and $p(d|M)$ expresses the likelihood of the data for a given M. We model the data in each interferometric detector d_I as a linear combination of

multiple components

$$d_I = g_I + h_I^{\rm cbc} + h_I^{\rm w} + n_I \equiv M_I + n_I, \tag{4.2}$$

where g_I denotes any detector glitches[3], the full GW signal ($h_I^{\rm cbc} + h_I^{\rm w}$) consists of the CBC waveform model $h_I^{\rm cbc}$ and any signal (such as the post-merger) that is not included in the CBC template and is captured by wavelets $h_I^{\rm w}$, and n_I is the Gaussian component of the noise. The combination of a signal, glitch, and noise defines the full model M. Given such a model, the likelihood function $p(d_I|M_I)$, in short L_I, can be computed by noting that the residual $r_I = d_I - M_I$ is consistent with Gaussian noise (see also Eq. (3) of Cornish et al. (2021)):

$$\log L_I = -\frac{1}{2}(r_I|r_I) + C_I, \tag{4.3}$$

where C_I is a constant that depends on the PSD of the detector noise $S_{n,I}(f)$, and the noise weighted inner product is defined as

$$(a|b) \equiv 2\int \frac{a^*(f)b(f) + b^*(f)a(f)}{S_{n,I}(f)} df, \tag{4.4}$$

where an asterisk denotes a complex conjugate. The remaining ingredients of the analysis are a specific model for each data component, the corresponding prior, and the sampling procedure.

We express the GW signal at geocenter and then project it onto each detector in the network to compute the detector response in the frequency domain for each detector

$$\begin{aligned} h^{\rm cbc/w}(f) = e^{2\pi i \Delta t(\alpha,\delta)} \\ \left[F^+(\alpha,\delta,\psi)h_+^{\rm cbc/w} + F^\times(\alpha,\delta,\psi)h_\times^{\rm cbc/w}\right]. \end{aligned} \tag{4.5}$$

The intrinsic part of the signal for each GW polarization mode is given by $h_+^{\rm cbc/w}$ and $h_\times^{\rm cbc/w}$ and it is different under the CBC or the wavelet model. The projection then involves a time delay in the arrival of the signal in the different detectors $\Delta t(\alpha,\delta)$, the detector antenna patterns $F^+(\alpha,\delta,\psi)$, $F^\times(\alpha,\delta,\psi)$ for each GW polarization mode, parametrized through the sky location α, δ corresponding to the right ascension and declination respectively, and the polarization angle ψ. The projection is independent of whether the signal is expressed through CBC templates or wavelets so the functional form of Eq. 4.5 and the parameters (α,δ,ψ) are common. The ratio $h_+^{\rm cbc/w}/h_\times^{\rm cbc/w}$ is related to the binary inclination ι and is also shared between the CBC and wavelet models.

[3]For the remainder of the paper we will work with data without instrumental glitches and thus ignore the glitch model g_I. As we will not be making use of the glitch model, we will also refer to the wavelet signal model simply as the *wavelet* model to avoid confusion with the CBC signal model.

By construction `BayesWave` breaks down the full parameter space defined by all models in different blocks and uses a blocked Gibbs sampler to explore the posterior. Each block of parameters is sampled over with different (Reversible jump (Green, 1995)) Markov Chain Monte Carlo (RJMCMC) samplers. Parallel tempering is employed both for efficient posterior exploration and evidence calculation (Lartillot and Philippe, 2006), so all samplers share the same number and temperature of chains and exchanges between the chains proceed for all models simultaneously. In order for this procedure to lead to efficient exploration of the posterior distribution, each block of parameters needs to be selected to be as independent from and uncorrelated to other blocks as possible.

The sampler ensemble consists of 1. a CBC MCMC, Sec. 4.2.2, 2. an RJMCMC sampler for the wavelet intrinsic parameters, Sec. 4.2.3, 3. an MCMC sampler for the common extrinsic parameters for the CBC and wavelet models, Sec. 4.2.4, and 4. an RJMCMC noise model sampler, Sec. 4.2.5. The structure is given in Fig. 4.1 and it is similar to the code used in Chatziioannou et al. (2021a) but instead of nonoverlapping parameter blocks, we now allow for a set of common parameters to be updated in multiple blocks as detailed in the subsequent subsections. Sampling then proceeds iteratively both within each block ($\mathcal{O}(10^2)$ iterations at a time) and between different blocks ($\mathcal{O}(10^4)$ iterations). In the rest of this section, we discuss each (RJ)MCMC, which parameters it updates, and how it relates to the other samplers. The discussion is structured to be as self-contained and detailed as possible; readers primarily interested on the results can find them starting in Sec. 4.3.

4.2.2 CBC-specific parameters: CBC sampler

A non-precessing quasicircular inspiral of two compact objects is characterized by a set of 6 intrinsic source parameters and a set of 7 extrinsic parameters that describe the binary's location and orientation relative to the detectors. The CBC sampler updates the parameters that are specific to $h_+^{\rm cbc}$ and $h_\times^{\rm cbc}$, namely masses, spins, tidal parameters (for BNSs), distance, time, and phase of coalescence. Note that some of these parameters, such as distance and time, are degenerate with parameters updated in the extrinsic sampler as they also affect the wavelet model (see Sec. **??**). In Fig. 4.1 the CBC sampler is depicted by the very last green box "CBC MCMC" and the relevant parameters are collectively denoted by $\theta^{\rm cbc}$.

4.2.2.1 Waveform model and Parameters

The CBC sampler is integrated with the `LALSimulation` (LIGO Scientific Collaboration, Virgo Collaboration, 2018) waveforms suite and can make use of any non-precessing waveform model; we here focus on PhenomD (Husa

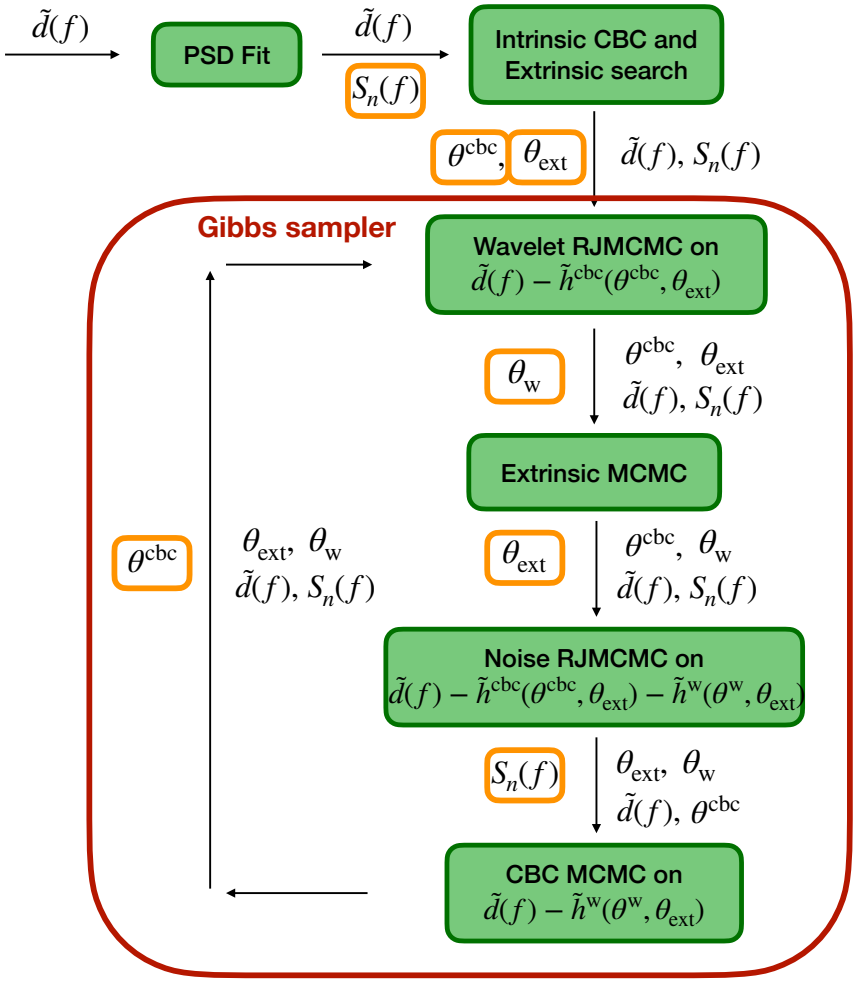

FIGURE 4.1: General code workflow. After some preconditioning, the data are used to obtain a quick fit to the noise PSD $S_n(f)$ and CBC $\theta^{\rm cbc}$ and extrinsic $\theta_{\rm ext}$ parameters. These are then used as a starting point for the blocked Gibbs sampler (red box) which consists of 4 independent samplers. Each iteration of the Gibbs sampler consists of a wavelet RJMCMC that updates the wavelet parameters $\theta^{\rm w}$, an extrinsic MCMC that updates the extrinsic parameters $\theta_{\rm ext}$, a noise RJMCMC that updates $S_n(f)$, and a CBC MCMC that updates $\theta^{\rm cbc}$. Orange rectangles enclose the parameter that is being updated from the preceding green box. Each green block consists of $\mathcal{O}(10^2)$ iterations, while the entire red Gibbs box consists of $\mathcal{O}(10^4)$ iterations.

et al., 2016; Khan et al., 2016) and PhenomDNRT (Dietrich et al., 2017, 2019b,a). The former includes the full inspiral-merger-ringdown (IMR) signal expected from a BBH coalescence. The latter targets the inspiral of BNS systems as it includes the effect of tidal interactions between the NSs, manifesting as a change in the waveform phase. PhenomDNRT models the long inspiral phase and instead of terminating abruptly at merger, it decays rapidly through a window function after the merger frequency (Dietrich et al., 2019b,a). PhenomDNRT is well suited for our purposes regarding BNS signals as the waveform remains smooth and continuous, naturally allowing the resulting CBC (inspiral) + wavelet (post-merger) model to be smooth without any forced stitching.

A total of 9 CBC parameters are updated by the CBC sampler $\theta^{\rm cbc} \equiv (m_1, m_2, \chi_1, \chi_2, \Lambda_1, \Lambda_2, D_L, t_c, \phi_c)$.

- The masses of the binary components are m_1 and m_2 with the convention $m_1 \geq m_2$. We also parameterize the masses through the total mass $M = m_1 + m_2$, chirp mass $\mathcal{M} = (m_1 m_2)^{3/5} / (m_1 + m_2)^{1/5}$, mass ratio $q = m_2/m_1 < 1$, and symmetric mass ratio $\eta = (\mathcal{M}/M)^{5/3} < 1/4$.

- The dimensionless spin components aligned with the orbital angular momentum are χ_1 and χ_2. Current detector sensitivities make individual spin components hard to measure (Pürrer et al., 2016; Chatziioannou et al., 2018b) so a commonly used spin combination is the mass-weighted effective spin parameter,

$$\chi_{\rm eff} = \frac{m_1\chi_1 + m_2\chi_2}{m_1 + m_2}, \tag{4.6}$$

  which is conserved under spin-precession to at least the second post-Newtonian order (Racine, 2008).

- Tidal interactions in BNS signals are encoded by the NS dimensionless tidal deformabilities

$$\Lambda_{1,2} = \frac{2}{3} k_{1,2} \left(\frac{R_{1,2}}{m_{1,2}}\right)^5, \tag{4.7}$$

  where $k_{1,2}$ is the tidal Love number, $R_{1,2}$ is the radius, and $m_{1,2}$ is the mass of each binary component. To leading order, the tidal effects are imprinted in the waveform through the binary tidal deformability (Favata, 2014; Wade et al., 2014)

$$\tilde{\Lambda} = \frac{16}{13} \frac{(12q+1)\Lambda_1 + (12+q)q^4\Lambda_2}{(1+q)^5}, \tag{4.8}$$

  with subleading terms introducing another tidal parameter that is generally not measurable, $\delta\tilde{\Lambda}$ (Wade et al., 2014). For BBH systems, $\Lambda_1 = \Lambda_2 = 0$ (Binnington and Poisson, 2009; Poisson, 2021).

- The CBC sampler further handles 3 extrinsic parameters either because they do not explicitly appear in the wavelet model (time of coalescence t_c, luminosity distance D_L) or because they are correlated with intrinsic parameters and need to be simultaneously sampled for efficiency (phase of coalescence ϕ_c). Despite the well-known correlation between the distance and the inclination (Cutler and Flanagan, 1994), we do *not* update the latter here as the inclination affects the wavelet model as well.

4.2.2.2 Priors

The priors on the binary intrinsic parameters depend on whether we consider BNSs or BBHs. We sample the detector frame chirp mass $\mathcal{M} \in [0.23, 2.6]\, M_\odot$ and total mass $M \in [1.0, 6.0]\, M_\odot$ for BNSs and $\mathcal{M} \in [0.23, 174]\, M_\odot$, $M \in [1.0, 400.0]\, M_\odot$ for BBHs and apply a Jacobian correction that makes the prior flat for the individual component masses m_1 and m_2 (Callister, 2021). We use flat priors for the dimensionless spin components $\chi_1, \chi_2 \in [-0.05, 0.05]$ for BNSs and $\chi_1, \chi_2 \in [-1, 1]$ for BBHs. In the BNS case we sample the effective dimensionless tidal parameters uniform $\tilde{\Lambda} \in [0, 1000]$ and $\delta\tilde{\Lambda} \in [-500, 500]$ with the constraint that the corresponding individual component tidal deformabilities can not be negative. We can optionally directly sample the individual component tidal parameters uniformly with the default prior range $\Lambda_{1,2} \in [0, 1000]$ though all results presented here use the former tidal prior. In all cases, the individual component tidal deformabilities are assumed to be independent of each other and not linked under any assumption about the NS EoS (Chatziioannou et al., 2018a; Zhao and Lattimer, 2018; Carson et al., 2019).

We sample the merger time relative to the GPS trigger time ($t_{\rm trig}$) with a default window of $t_c \in \left[t_{\rm trig} - 0.5{\rm s}, t_{\rm trig} + 1.5{\rm s}\right]$. Alternatively, it is possible to use a symmetric window of arbitrary size around the trigger time. We also sample in the logarithmic luminosity distance with an adjustment to the Jacobian such that the prior is uniform in volume (i.e., uniform in D_L^3) with $\mathrm{D}_L \in [1, 10000]$ Mpc. Finally, the coalescence phase prior is uniform in $\phi_c \in [0, 2\pi]$.

4.2.2.3 Data

Since the CBC and the wavelet models are active simultaneously, the overall likelihood has to take into account both CBCs and wavelets. The blocked nature of the code means that the wavelet model is not updated in the CBC sampler, so the wavelet-subtracted data $d^{\rm cbc} \equiv d - h^{\rm w}$ are constant as the CBC parameters are being updated. We compute $d^{\rm cbc} = d - h^{\rm w}$ once at the beginning of the CBC sampler block given the current wavelet parameters. The wavelet-subtracted data $d^{\rm cbc}$ then play the role of the data in the sense that $d^{\rm cbc} - h^{\rm cbc}$ is the residual during the likelihood computation.

4.2.2.4 Proposals

We employ the CBC sampler described in detail in Cornish (2021) within the `BayesWave` blocked Gibbs sampler, see Fig. 4.1. The sampler performs by default 300 iterations that update the CBC parameters described above using a mixture of proposals:

- *Fisher jumps*: Gaussian jumps along the eigenvectors of the Fisher information matrix. The Fisher information matrix is recomputed at the start of each CBC sampler block for each parallel chain and using each chain's current location.
- *Differential evolution (Ter Braak, 2006)*: Two samples from each chain's history propose a new location. The initial history array is drawn from the prior and then updated every 20 iterations. As each CBC sampler block does 300 iterations by default, the history array contains information about previous blocks.
- *Jiggle*: Small Gaussian jumps in individual parameter directions. The size of the proposed jump is scaled with the chain temperature such that high temperature chains have larger jumps and can explore a larger area of the parameter space. The jiggle proposal allows for exploring areas outside those covered by Fisher matrix jumps and are especially useful early on.
- *Mixed*: Uniform proposal for the tidal parameters and differential evolution (see above) for the other parameters. This proposal is employed 5% of the time and aids convergence for cases where the tidal parameters cannot be tightly constrained. The proposal facilitates taking larger steps for the tidal parameters whereas the other better constrained parameters evolve with smaller steps, thus leading to a higher acceptance rate than a fully uniform proposal.

4.2.2.5 Heterodyned likelihood

We integrate the heterodyned likelihood technique (Cornish, 2010, 2021) in the `BayesWave` CBC (and extrinsic, to be described later) sampler in order to speed up the likelihood evaluation. The technique is based on the fact that all CBC waveforms that contribute to the posterior (i.e. have a likelihood high enough that the MCMC proposal can be accepted) are similar. The likelihood of a model $h^{\rm cbc}$ can then be computed based on a fixed and precomputed reference model $h^{\rm cbc}_{\rm ref}$ that is similar to $h^{\rm cbc}$, and given the noise PSD and the data.

We can rewrite the residuals r using the reference heterodyne, omitting the detector subscript I for clarity, as:

$$r = d^{\rm cbc} - h^{\rm cbc} = d^{\rm cbc} - (h^{\rm cbc}_{\rm ref} - \Delta h^{\rm cbc}) = r_{\rm ref} + \Delta h^{\rm cbc}, \tag{4.9}$$

where $\Delta h^{\rm cbc}$ is the difference between the reference waveform and the exact waveform at a given set of CBC parameters, $d^{\rm cbc}$ is the wavelet-subtracted data as given to the CBC sampler, $d^{\rm cbc} = d - h^{\rm w}$, and $r_{\rm ref} \equiv d^{\rm cbc} - h^{\rm cbc}_{\rm ref}$. The likelihood from Eq. 4.3 then becomes

$$\log L \sim -\frac{1}{2}(r_{\rm ref} + \Delta h^{\rm cbc} | r_{\rm ref} + \Delta h^{\rm cbc}), \tag{4.10}$$

where we omit the second term of Eq. 4.3 as the PSD is held fixed during the CBC updates and the term does not affect the relative likelihoods. This form of the likelihood enables several computational advantages as we split the inner product in its components

$$\begin{aligned} -2\log L \sim & (r_{\rm ref}|r_{\rm ref}) + (\Delta h^{\rm cbc}|\Delta h^{\rm cbc}) \\ & + \left[(r_{\rm ref}|\Delta h^{\rm cbc}) + (\Delta h^{\rm cbc}|r_{\rm ref})\right]. \end{aligned} \tag{4.11}$$

A detailed description of how each components of the heterodyned likelihood above is calculated is available in Cornish (2010, 2021); we here briefly discuss how this is adapted and implemented within the `BayesWave` blocked Gibbs sampler. The first term of the inner product $(r_{\rm ref}|r_{\rm ref})$ is the most straightforward as it is computed once at the beginning of each CBC sampler block using the current wavelet-subtracted data and CBC parameters for each chain.

The key performance improvement of the heterodyne likelihood is achieved by recognizing that the latter two terms in Eq. 4.11 can be further split up into the product of two terms: a slowly-varying term that is computed at each iteration, and a rapidly-varying term that is computed once. Both terms are evaluated through a sum over frequencies with bin width Δf using a Legendre polynomial expansion (Cornish, 2021). The rapidly-varying term requires a small Δf but this is computed once. The slowly-varying term that needs to be calculated on each likelihood evaluation can instead make use of a coarse frequency grid of $\Delta f \leq 4$ Hz thus significantly reducing computational cost.

The term $(\Delta h^{\rm cbc}|\Delta h^{\rm cbc})$ in Eq. 4.11 introduces a slowly varying phase difference as a function of frequency between the reference waveform and the proposed waveform. However, in the noise weighted inner product, Eq. 4.4 the denominator might not be slowly varying if the PSD includes sharp spectral lines. Following the approach discussed above, we split the PSD in a smooth broadband component that is slowly varying as a function of frequency, and line features which are only evaluated at discrete intervals. If the PSD is held fixed throughout the analysis, the smooth part is computed once at the beginning. When using a variable PSD model instead, we update the smooth PSD component at the beginning of the CBC sampler block.

The final term of Eq. 4.11 combining $r_{\rm ref}$ and $\Delta h^{\rm cbc}$ is split up into the product of a slowly and rapidly varying part by introducing the smooth component of the PSD in the numerator and denominator of, respectively, the rapidly and slowly varying part. Due to this, the heterodyned difference in the waveforms $\Delta h^{\rm cbc}$ is effectively whitened using only the smooth PSD component on a coarse frequency grid. The rapidly varying part consisting of the heterodyned whitened residuals $r_{\rm ref}$ now includes an extra product with the smooth part of the PSD, but since this is only computed at the start of the CBC sampler block, it adds little computational cost.

The heterodyned likelihood speeds up the likelihood calculation by a factor of $\sim T_{\rm obs}\Delta f$, where $T_{\rm obs}$ is the duration of the analyzed data, thus making it especially useful for long duration signals such as BNSs. The heterodyne implementation in `BayesWave` is more complex than the one in Cornish (2021) as it has to account for a potentially changing PSD and the wavelet model. We thus expect more modest overall computational improvements, though the speed-up of each individual likelihood calculation is consistent with Cornish (2021).

4.2.3 Wavelet parameters: Wavelet sampler

`BayesWave` has a common RJMCMC sampler for its *signal* and *glitch* models where the non-Gaussian features (either a GW signal or a detector glitch) are modeled as a sum of a variable number of sine-Gaussian wavelets which can be expressed in the time-domain as

$$\psi(t; A, f_0, \tau, t_0, \phi_0) = Ae^{-(t-t_0)^2/\tau^2}\cos\left(2\pi f_0(t-t_0)+\phi_0\right), \tag{4.12}$$

where $\tau \equiv Q/(2\pi f_0)$ is the decay time, Q is the quality factor, A is the wavelet amplitude, f_0 is the wavelet central frequency, t_0 is the wavelet central time and ϕ_0 is the phase offset. The glitch model uses independent sums of wavelets for the different detectors, whereas the signal models use a set of wavelets at geocenter and projects them onto detectors through Eq. 4.5 with

$$h^{\rm w}_+(f) = \sum_{i=0}^{N_s}\psi(f; A_i, f_{0,i}, Q_i, t_{0,i}, \phi_{0,i}) \tag{4.13}$$

$$h^{\rm w}_\times(f) = \epsilon h^{\rm w}_+(f)e^{i\pi/2}, \tag{4.14}$$

where ϵ is the ellipticity parameter and N_s is the number of wavelets in the signal model. Details about the wavelet sampler and priors/proposals for each wavelet parameter and the number of wavelets are discussed in Cornish and Littenberg (2015); Cornish et al. (2021) and remain mostly unaltered in our analysis. We extend the standard settings by adding an option to limit the central frequency range of the wavelets to $f_0 \in [f^{\rm w}_{\rm min}, f^{\rm w}_{\rm max}]$ which defaults to the analysis bandwidth. Since the ellipticity ϵ is related to the remaining extrinsic parameters and to the binary inclination, it is not sampled by the wavelet sampler but the extrinsic sampler discussed next.

In Fig. 4.1 the wavelet sampler is depicted by the green box termed "wavelet RJMCMC" where $\theta^{\rm w}$ corresponds to the relevant parameters: the number of wavelets and the amplitude, quality factor frequency, time, and phase of each wavelet. Similarly to the CBC sampler, the wavelet sampler makes use of the CBC-subtracted

data $d^{\mathrm{w}} \equiv d - h^{\mathrm{cbc}}$, which are computed once at the beginning of each wavelet sampler block using the current CBC sample.

4.2.4 Common parameters: extrinsic sampler

Since both the CBC and wavelet models target the same GW source, they share a number of extrinsic source parameters that govern the signal projection onto the detector network, Eq. 4.5. These common parameters θ_{ext} are updated by the extrinsic sampler block, denoted as "extrinsic MCMC" in Fig. 4.1.

4.2.4.1 Parameters and Priors

The extrinsic parameters shared between the wavelet and CBC models are $\theta_{\mathrm{ext}} \equiv (\alpha, \delta, \psi, \epsilon, \varphi_0, A^{\mathrm{cbc}})$.

- The source sky location is given by the right ascension $\alpha \in [0, 2\pi]$ and the declination $\sin\delta \in [0, 1]$, and the corresponding prior is flat on the sphere.
- The polarization angle ψ is sampled with a uniform prior $\psi \in [0, 2\pi]$.
- The degree of elliptical polarization in the signal is encoded in $\epsilon \in [-1, 1]$, Eq. 4.14. For elliptically polarized signals, the ellipticity parameter can be related to the orbital inclination with respect to the line of sight to the binary ι as

$$\cos\iota = \frac{1 - \sqrt{1 - \epsilon^2}}{\epsilon}. \tag{4.15}$$

  We use a uniform prior in $\cos\iota \in [-1, 1]$, which is a deviation from previous works where `BayesWave` samples uniformly in ϵ instead.
- An overall phase shift φ_0 is sampled uniformly as $\varphi_0 \in [0, 2\pi]$ and it is equivalent to a phase shift in all wavelets simultaneously (see Eq. 4.12). Despite the degeneracy, introducing this parameter leads to more efficient convergence in the wavelet model.
- We also introduce an overall amplitude scaling parameter A^{cbc}, which is sampled uniformly and is degenerate with the CBC luminosity distance D_L. We restrict A^{cbc} to correspond to values for D_L within the D_L prior range and add a Jacobian factor such that the prior remains uniform in volume, i.e., $(D_L / A^{\mathrm{cbc}})^3$.

The extrinsic parameters can be efficiently sampled independently from the CBC intrinsic parameters by noting that a change in the extrinsic parameters induces an amplitude, time, and phase shift in the signal seen in each detector, while leaving its

phase evolution unaltered. We therefore compute the geocenter CBC waveform given the fixed intrinsic parameters once and then apply the shifts accordingly to evaluate the waveform at each new proposed set of extrinsic parameters. This has the advantage that we can sample the extrinsic parameters efficiently while not having to recompute the full, expensive CBC waveform phase at each iteration.

In practice we write the CBC response at a given detector I using the geocenter waveform $h_+^{\rm cbc}$ and applying the shifts as

$$h_I^{\rm cbc}(f) = h_+^{\rm cbc}(f) F_I^{\rm cbc} e^{i\lambda_I} e^{2\pi i f \Delta t_I}, \tag{4.16}$$

where Δt_I is the arrival time at the detector relative to the geocenter, the phase shift is

$$\lambda_I \equiv \arctan[\epsilon F^\times(\alpha,\delta,\psi)/F^+(\alpha,\delta,\psi)] + \phi_c, \tag{4.17}$$

and $F_I^{\rm cbc}$ is the magnitude of the detector response

$$F_I^{\rm cbc} \equiv A^{\rm cbc}\sqrt{\left[\frac{\cos\iota}{\epsilon}F^+(\alpha,\delta,\psi)\right]^2 + [\cos\iota\, F^\times(\alpha,\delta,\psi)]^2}, \tag{4.18}$$

Even though $A^{\rm cbc}$ is completely degenerate with D_L, we find that it increases sampling efficiency in the extrinsic sampler due to the distance-inclination degeneracy (Cutler and Flanagan, 1994).

At the end of the block of extrinsic updates, the scaling and shift parameters are reset and propagated to the corresponding CBC parameters before proceeding to the CBC sampler

$$D_L \to \frac{D_L}{A^{\rm cbc}} \tag{4.19}$$

$$t_c \to t_c + \Delta t_a \tag{4.20}$$

$$\phi_c \to \phi_c - \Delta\varphi_0 \tag{4.21}$$

In these equations Δt_a and $\Delta\varphi_0$ correspond to the overall time and phase shift induced on the waveform by the change in extrinsic parameters. The time of arrival of the merger at geocenter t_c changes due to the new sky location and this shift is encoded in Δt_a.

4.2.4.2 Proposals

To update the parameters $\theta_{\rm ext}$ the extrinsic sampler uses a mixture of proposals, which are described in detail in Cornish and Littenberg (2015); Cornish et al. (2021). In the following we focus on some CBC-specific updates made to the existing proposals.

The most commonly used proposal is the Fisher proposal, where jumps are proposed along the eigenvectors of the Fisher information matrix. The matrix elements, eigenvalues, and eigenvectors are computed once at the start of the extrinsic sampler block by finite differencing using the current chain position. We extend the existing `BayesWave` Fisher calculation by adding the CBC model such that the final Fisher matrix captures information from the full CBC+wavelet signal. When the sky location is known and fixed (for example for sources with an electromagnetic counterpart), the Fisher matrix is reduced to the parameters that are varied.

The `BayesWave` extrinsic sampler also employs a two-part proposal that exploits a degeneracy of the extrinsic parameters when multiple detectors are used (also proposed by Raymond and Farr (2014)). The first part consists of a sky-ring proposal that finds a new sky location $(\alpha, \sin\delta)$ such that the time delays between the detectors are preserved (Veitch et al., 2015; Cornish and Littenberg, 2015). The second part of the proposal updates the remaining parameters (ψ, $\cos\iota$, $A^{\rm CBC}$ and φ_0) either uniformly in their prior range or deterministically such that the waveform is identical at the new sky location compared to the waveform at the previous sky location, see Cornish (2021) for details. Since the waveform is conserved, so is the likelihood, though this does not guarantee that the proposed sky location will be accepted due the necessary Jacobian factor (Cornish, 2021).

Finally, the `BayesWave` extrinsic sampler includes a uniform proposal that draws from the prior of each parameter. These jumps are particularly useful for higher temperature chains to explore the entire space available as well as early on in the sampling if a signal has not been found.

4.2.4.3 Heterodyned likelihood

A change in the extrinsic parameters $\theta_{\rm ext}$ will affect both the CBC and the wavelet models, therefore both need to be taken into account when computing the likelihood. However, we find that the CBC part of the model dominates the computational cost so we again use the heterodyned likelihood to speed up the calculation. In this case, there are some differences compared to the implementation for the CBC sampler as discussed in Sec. 4.2.2.5 because we can no longer assume that the wavelet-subtracted data are fixed. Indeed, a new set of extrinsic parameters changes the full signal projection and the full residual $r = d - h^{\rm w} - h^{\rm cbc}$ has to be recomputed on each iteration.

Despite this, we can achieve a computational improvement by splitting up the heterodyne computation even further to avoid recomputing all components of the heterodyne when the residuals change, i.e., at every iteration. All subcomponents of the heterodyne that do not involve the residuals are computed and stored once at the

start of each extrinsic MCMC sampler block, whereas the remaining components are computed at each iteration. If the signal wavelet model is not employed (i.e. an analysis with only the CBC model), the heterodyne likelihood implementation falls back to the same, more efficient, version as for the CBC sampler.

Since the likelihood calculation is more involved and parts of the heterodyne procedure need to be recomputed in each iteration, we expect more modest computational improvements compared to those of Sec. 4.2.2.5.

4.2.5 Noise parameters: PSD sampler with `BayesLine`

The PSD of the gaussian noise is modeled as a combination of splines and Lorenzians targeting the broadband behavior and spectral lines respectively and marginalized over through an RJMCMC in `BayesLine` (Littenberg and Cornish, 2015; Chatziioannou et al., 2019; Cornish et al., 2021). The analysis presented here makes no updates on this noise marginalization procedure and can optionally include `BayesLine` as one of the blocks, corresponding to the "Noise RJMCMC" green box in Fig. 4.1. Alternatively, the analysis can skip the noise marginalization and use a predetermined fixed PSD. If the PSD is modeled by `BayesLine`, the heterodyne procedure needs to be repeated at the end of the `BayesLine` sampling block and before moving to CBC sampling with the latest PSD sample. However, this coincides with a recomputing of the heterodyne elements due to a change of the wavelet model that is subtracted from the data so no further computations are necessary. In the current implementation (Littenberg and Cornish, 2015), the wavelet model is subtracted from the data before entering the `BayesLine` sampling block and we extend this subtraction to include the CBC model.

4.3 Toy model: GW150914

To illustrate the wavelet+CBC model, we begin with a toy model by analyzing the first detected BBH merger, GW150914 (Abbott et al., 2016c). In reality, the full BBH signal can be modeled with existing CBC templates, something confirmed in the case of GW150914 through comparisons between different models and physical effects (Abbott et al., 2017c). In our case, and in order to test our hybrid analysis on what is probably the best studied signal, we pretend we lacked efficient modeling of a BBH merger and waveform templates terminated shortly before merger.

We analyze 4s of data from LIGO Hanford and LIGO Livingston available through the Gravitational Wave Open Science Center (GWOSC) (Gravitational Wave Open Science Center (GWOSC); Abbott et al., 2021) in the frequency band from 16Hz to 2048Hz. The noise PSD is marginalized over as discussed in Sec. 4.2.5. For the CBC model, we use

Label	GPS time (s)	Model	Waveform	Wavelet window (s)	Wavelet bandwidth (Hz)
CBC_{IMR}	1126259462.391	CBC/noise	PhenomD	N/A	N/A
CBC_{Insp}	1126259462.391	CBC/noise	PhenomDNRT	N/A	N/A
CBC_{Insp}+wavelet	1126259462.416	CBC/wavelet/noise	PhenomDNRT	(-0.025, 0.025)	(150, 2048)

TABLE 4.1: Settings for the runs of Sec. 4.3 on GW150914. All runs use a segment length of 4s, a low frequency cut off of 16Hz, and a sampling rate of 4096Hz. From left to right, columns correspond to the run label in Fig. 4.2, the GPS time (t_{trig}), the models active, the CBC waveform, the wavelet window around the trigger time, and the wavelet bandwidth.

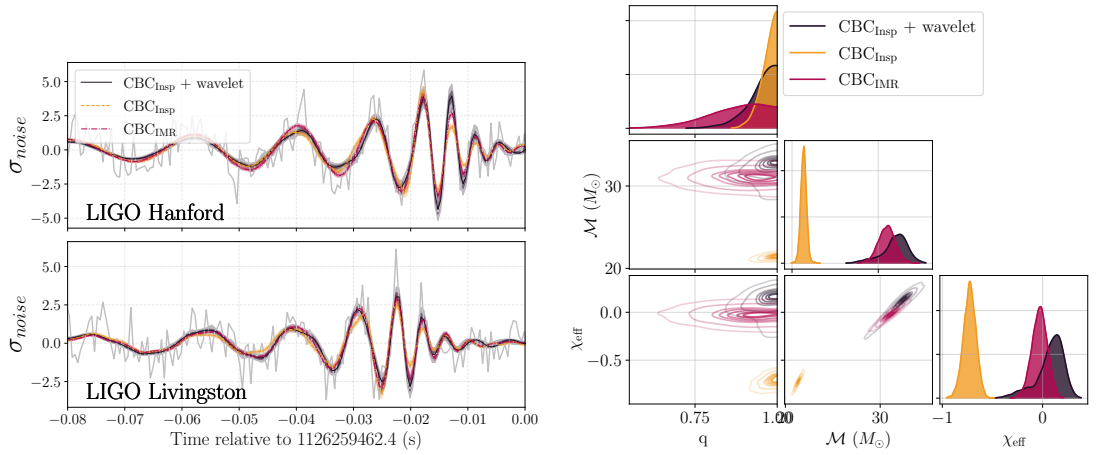

FIGURE 4.2: Results on the GW150914 data with different waveforms and analyses. The left panel presents 90% credible intervals for the whitened signal reconstructions in each detector while the right panel shows one- and two-dimensional marginalized posterior distributions for selected CBC parameters: mass ratio q, detector-frame chirp mass $\mathcal{M}$, and the effective spin $\chi_{\rm eff}$. In both panels red corresponds to an analysis with the CBC model and PhenomD (CBC_{IMR}), yellow corresponds again to a CBC model analysis but with PhenomDNRT (CBC_{Insp}), and black corresponds to an analysis with the CBC+wavelet model with PhenomDNRT (CBC_{Insp}+wavelet). In the left panel the gray dashed line indicates the data whitened with a fair draw PSD from the noise model posterior.

either the PhenomD waveform or the PhenomDNRT waveform. PhenomD is appropriate for BBH systems as it models the full IMR signal. On the other hand, we do not expect PhenomDNRT to accurately model the BBH data as it lacks an accurate merger and ringdown portion. Regardless, we use it to demonstrate how any leftover signal that is not captured by the CBC waveform can be captured by the wavelets.

We perform three analyses with different models and CBC templates:

1. A full "CBC_{IMR}" analysis where the data are described only by the CBC model with PhenomD.

2. A "CBC_{Insp}" analysis where the data are again described only by the CBC model but with PhenomDNRT.

3. A "CBC_{Insp}+wavelets" analysis where the data are described by a combination of the CBC model with PhenomDNRT and the wavelet model. We restrict the wavelet model to times $t_0 \in \left[t_{\text{trig}} - 0.025\text{s}, t_{\text{trig}} + 0.025\text{s}\right]$ and frequencies above 150Hz, targeting the merger portion of the signal that is missed by PhenomDNRT.

The run labels and relevant settings are given in Table 4.1.

In Fig. 4.2 we show the signal reconstructions and parameter posteriors for these three analyses. The left panel shows the whitened data and the 90% credible intervals for each reconstruction in LIGO Hanford (top) and LIGO Livingston (bottom). All analyses are able to identify the GW signal and they further lead to consistent signal reconstructions of the early portion of the signal. However, as the signal approaches the merger phase, the CBC_{Insp} analysis deviates from the reference CBC_{IMR} analysis as well as the hybrid CBC_{Insp}+wavelets analysis by underpredicting the strength of the signal. The hybrid CBC_{Insp}+wavelets agrees well with the full CBC_{IMR}, though the uncertainty of the former is larger in the merger phase where the signal is no longer modeled with CBC templates but with wavelets.

The recovered source parameters are given in the right panel of Fig. 4.2. The CBC_{IMR} analysis leads to consistent results with those reported in previous studies (Abbott et al., 2016b,a, 2018b), while the CBC_{Insp} analysis shows a significant bias away from the expected posteriors. The bias is most evident in the chirp mass: PhenomDNRT compensates for the lack of a merger and ringdown by decreasing the chirp mass, thus leading to a longer inspiral phase in an attempt to capture part of the missing merger cycles. Despite this bias, the CBC_{Insp} analysis is still not able to capture the full signal as seen on the left panel. The CBC_{Insp}+wavelets proof-of-concept model results in recovered parameter posteriors that are consistent with the full CBC_{IMR} analysis. The posteriors are not expected to be identical as now the CBC model has access only to the inspiral portion of the signal and therefore lacks information available to the CBC_{IMR} analysis. However, the wavelet model can efficiently capture the missing portion of the signal, allowing the CBC model to recover unbiased system parameters from the inspiral phase only.

Figure 4.3 examines the CBC_{Insp}+wavelets analysis in more detail and the complementary roles of the components in the hybrid model. We plot the signal reconstruction from each submodel as well as their sum. The sum of the individual model components is the same as the CBC_{Insp}+wavelets model in Fig. 4.2. Figure 4.3 shows that as expected the BBH inspiral phase is primarily reconstructed by the CBC PhenomDNRT model. As we move from the inspiral towards the merger time, the CBC model starts tapering off and deviating from the full GW signal. At the same time, the wavelet model captures the part of data that is not covered by the CBC model, such that the combination of CBC and wavelets accurately models the full GW

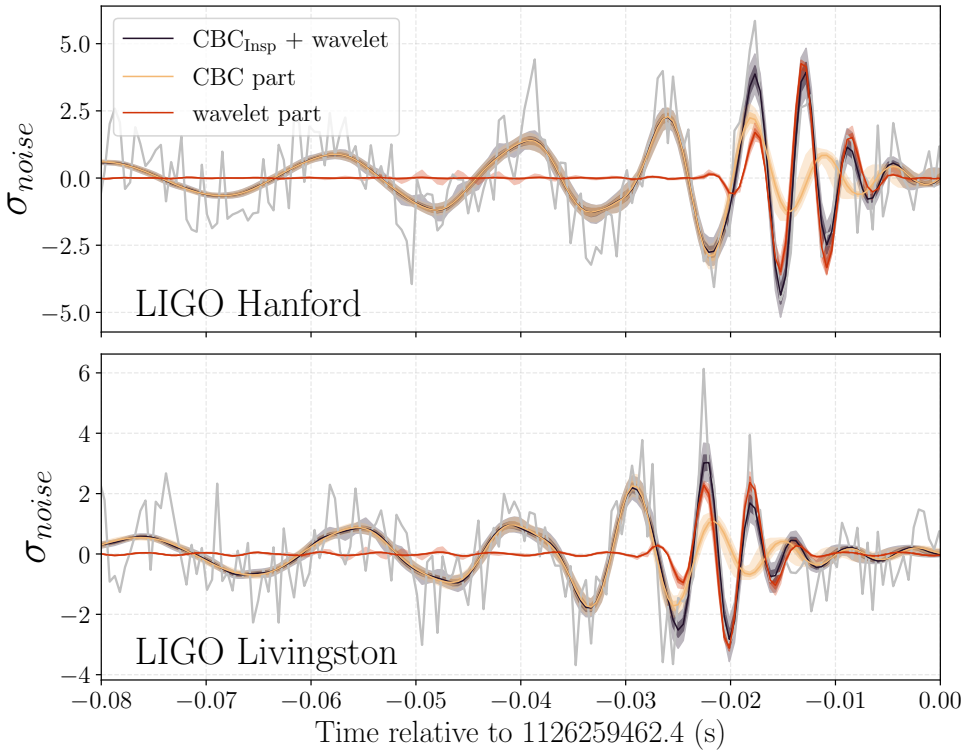

FIGURE 4.3: Breakdown of the various components of the hybrid CBC$_{\text{Insp}}$+wavelets analysis of GW150914. We present 90% credible intervals for the whitened signal reconstruction of the CBC model (yellow), the wavelet model (red), and the sum of the models (black). The sum is identical to the black lines from Fig. 4.2. The gray dashed line indicate the data whitened with a fair draw PSD from the noise model posterior.

signal. This toy model for GW150914 demonstrates the main concept of our hybrid analysis: the full GW signal is modeled with a well-understood component, i.e. the CBC waveform, and a model-agnostic component, i.e. the wavelet model, that covers any parts of the signal that are not included in the CBC model.

4.4 Constraints on the GW170817 post-merger

We turn our attention to GW170817, the first GW detection of a BNS coalescence (Abbott et al., 2017e, 2019b, 2018a). Both the low-frequency inspiral phase and a potential short-duration high-frequency post-merger phase of the signal have been studied separately in detail, with the latter remaining undetected (Abbott et al., 2017b, 2019b). Here we apply our hybrid CBC+wavelet model to study both parts of the GW170817 signal simultaneously and compare our results with separate analyses of the inspiral and post-merger signals.

We analyze 64s of data from LIGO Hanford and LIGO Livingston around GPS time 1187008882.446 (Gravitational Wave Open Science Center (GWOSC); Abbott et al., 2021) where the prominent glitch in LIGO Livingston has already been subtracted

Label	Segment length (s)	Sampling rate (Hz)	$f_{\rm low}$(Hz)	Model
$\rm CBC_{Insp}$	64	4096	32	CBC
$\rm CBC_{Insp}$+wavelet	64	8196	32	CBC/wavelet
wavelet-only	64	8196	1024	wavelet

Label	Waveform	Wavelet window (s)	Wavelet bandwidth (Hz)	$D_{\rm min}$
$\rm CBC_{Insp}$	PhenomDNRT	N/A	N/A	N/A
$\rm CBC_{Insp}$+wavelet	PhenomDNRT	(-0.125, 0.125)	(1024, 4096)	2
wavelet-only	N/A	(-0.125, 0.125)	(1024, 4096)	2

TABLE 4.2: Settings for the runs of Sec. 4.4 on GW170817. All runs are relative to GPS time 1187008882.446. From left to right, columns correspond to the run label in Fig. 4.4 and Fig. 4.5, the segment length, the sampling rate, the low frequency cut off, the models active, the CBC waveform, the wavelet window around the trigger time, the wavelet bandwidth, and the minimum number of wavelets.

(BayesWave Glitch Subtraction for GW170817; Pankow et al., 2018). While Virgo data are available for that time and aided in constraining the source sky location (Abbott et al., 2019b), we do not consider them due to the lower sensitivity and the fact that we fix the sky location to the known values of $\alpha = 3.446$ rad and $\delta = -0.408$ rad (Soares-Santos et al., 2017; Abbott et al., 2017a). We also use a fixed PSD rather than marginalize over the noise model for computational efficiency. We perform three analyses with different models and data:

1. A "$\rm CBC_{Insp}$" analysis with data in the frequency range (32,2048)Hz that are described only by the CBC model with the PhenomDRT waveform thus focusing on the inspiral signal.

2. A "$\rm CBC_{Insp}$+wavelet" analysis with data in the wider frequency range (32,4096)Hz that are described by a combination of the CBC model with PhenomDNRT and the wavelet model, thus targeting the full signal. The wavelets are restricted to (1024,4096)Hz targeting a potential high frequency post-merger signal.

3. A "wavelet-only" analysis where we restrict the frequency range to (1024Hz,4096)Hz and use only the wavelet model thus focusing only on the post-merger signal.

For both analyses with the wavelets model, we use a prior on the number of wavelets of $D \in [2, 100]$. The reason for selecting a minimum number of 2 wavelets is different for each analysis. In the case of the hybrid analysis, the CBC part of the model does not terminate at merger, but smoothly tapers into the post-merger phase for a few

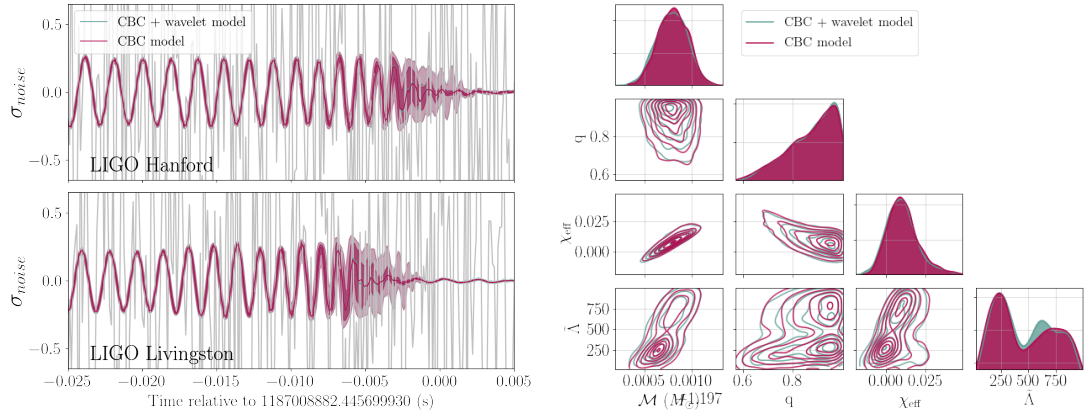

FIGURE 4.4: Results on the GW170817 data with different analyses. The left panel presents the median, 50% and 90% credible intervals for the whitened signal reconstructions in each detector while the right panel shows one- and two-dimensional marginalized posterior distributions for selected CBC parameters: mass ratio q, detector frame chirp mass $\mathcal{M}$, the effective spin $\chi_{\rm eff}$ and the tidal deformation parameter $\tilde{\Lambda}$. In both panels blue corresponds to the hybrid CBC_{Insp}+wavelet and magenta corresponds to the CBC_{Insp} analysis. In the left panel the gray dashed line gives the whitened data.

milliseconds, see Fig 4.9. We therefore need at least two wavelets, one to undo this effect of the CBC model and the other to capture the true post-merger signal. In the wavelet-only study we again need at least two wavelets such that one wavelet can capture the merger itself which extends into the analysis bandwidth and the other wavelet can capture the contribution from the post-merger signal. The full settings of our three analyses are detailed in Table 4.2.

We compare the two analyses that include the inspiral of the signal in Fig. 4.4 which shows the signal reconstructions focused on the end of the signal (left panel) and the marginalized posterior distribution for selected source parameters (right) of GW170817. Despite the high SNR of the signal, the amplitude is relatively weak compared to the noise level, however both analyses identify the signal and lead to consistent reconstructions for the late inspiral phase as shown on the left panel. As the binary approaches merger the reconstruction uncertainty increases due to the decreasing detector sensitivity at increasing frequencies. This behavior is evident in both analyses, suggesting that the GW170817 post-merger signal remains undetected as expected from the detector sensitivity at the time of GW170817 (Abbott et al., 2019b).

The right panel of Fig. 4.4 examines selected CBC parameters recovered from the inspiral signal. The recovered parameters are consistent with previous results using the same CBC model (Abbott et al., 2019b). The only difference is that our $\tilde{\Lambda}$ posterior has more support for the higher of the two modes present in the results of Abbott et al. (2019b). This can be attributed to different uses of prior: we use of a prior that is flat in $\tilde{\Lambda} - \delta\tilde{\Lambda}$ with the constraint that $\Lambda_1 > 0$ and $\Lambda_2 > 0$. This choice results in a

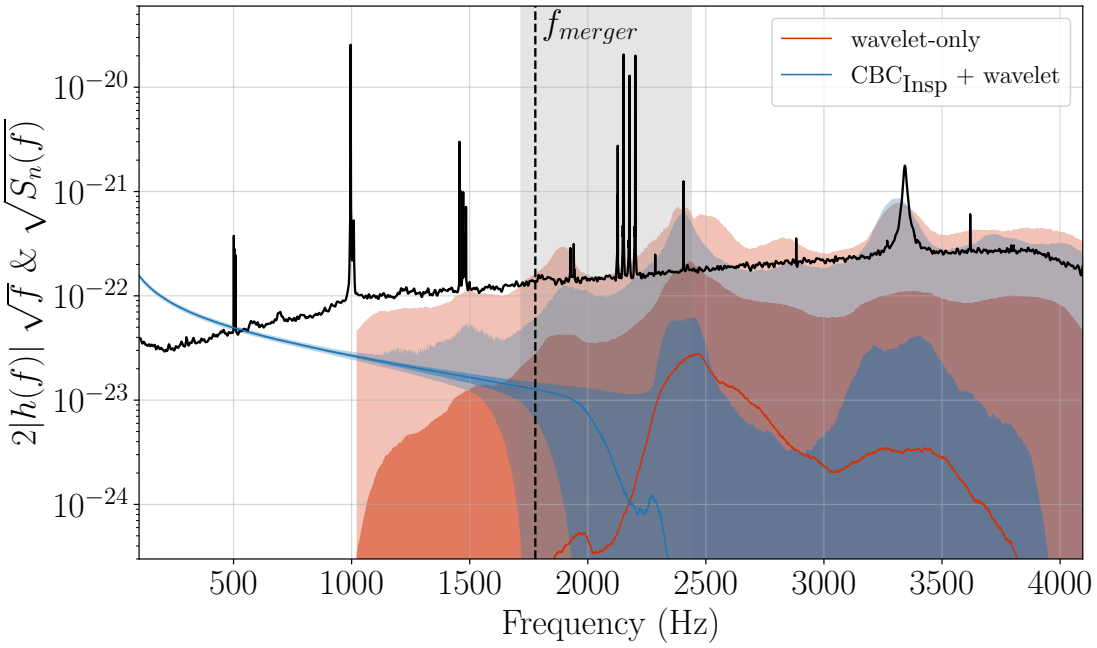

FIGURE 4.5: The median, 50% and 90% credible intervals of reconstructed GW spectrum in the LIGO Hanford detector in the CBC$_{\rm Insp}$+wavelet (blue) and the wavelet-only analysis (red). The noise amplitude spectral density is also overplotted (black dashed line). The grey band indicates the 90% credible interval of the merger frequency, with the dashed line indicating the peak of the merger frequency posterior at $f_{\rm merger} \sim 1778$Hz.

marginalized prior for $\tilde{\Lambda}$ that favors larger values. In contrast, Abbott et al. (2019b) uses a flat marginalized prior for $\tilde{\Lambda}$. We have verified that reweighting our posterior to a flat marginalized prior gives consistent results with Abbott et al. (2019b). Furthermore, the hybrid CBC$_{\rm Insp}$+wavelet analysis leads to essentially identical results for the CBC parameters as the traditional CBC$_{\rm Insp}$ analysis. This is consistent with Dudi et al. (2018) that finds that failure to account for the post-merger signal will not lead to biases in the CBC parameters for typical detector sensitivities, but also serves as a sanity check of the CBC$_{\rm Insp}$+wavelet analysis and the fact that sampling for the joint CBC+wavelet model has converged.

After confirming that the CBC$_{\rm Insp}$+wavelet analysis matches the CBC$_{\rm Insp}$ analysis as far as the inspiral portion of the signal is concerned, we switch to comparing the two analyses that include a wavelet model for a possible high-frequency post-merger signal. In Fig. 4.5 we show the reconstructed GW spectrum for CBC$_{\rm Insp}$+wavelet and wavelet-only. The vertical band gives the 90% credible interval for the merger frequency estimated using Eq. 11 in Dietrich et al. (2019) given the binary component masses and tidal parameters from the CBC model posteriors. This merger frequency is also the location at which the PhenomDNRT template starts tapering off, so it is related to the transition between the CBC and the wavelet models.

At frequencies below $\sim$ 1000Hz the CBC_{Insp}+wavelet reconstruction is dominated by the inspiral signal which is clearly detected as the spectrum lower limit is nonzero. With increasing frequency the detector sensitivity decreases, resulting in larger uncertainties in the reconstructed spectrum. Finally, at around the merger frequency, $\gtrsim$ 1600Hz, the reconstruction uncertainty is large and consistent with no detected signal. A similar picture is drawn by the wavelet-only analysis which again places only upper limits on the signal in frequencies above 1024Hz. The upper limit from the CBC_{Insp}+wavelet and the wavelet-only analyses is comparable, though the former is consistently lower across the frequency band, indicating more stringent constraints on the presence of a post-merger signal. Additionally, the wavelet-only analysis does not lead to a detection of the merger signal in the frequency range (1000,1500)Hz, unlike the CBC_{Insp}+wavelet case. This suggests that this high-frequency portion of the signal is not individually detectable, but only inferred coherently from the preceding inspiral signal.

4.5 Simulated signals

Going beyond the GW170817 upper limits, in this section we study simulated BNSs where the post-merger emission is detectable by future GW detectors. We assume a network of two detectors at the current location of LIGO Hanford and LIGO Livingston and a zero noise realization. Since all signals have high SNRs we assume that an electromagnetic counterpart has been identified and the sky location of the source is known. Following Torres-Rivas et al. (2019) we do not explicitly select any planned GW detector such as Cosmic Explorer (Abbott et al., 2017d; Reitze et al., 2019a,b) or the Einstein Telescope (Hild et al., 2011; Punturo et al., 2010) and their nominal sensitivity. We instead work with the Advanced LIGO design sensitivity (Barsotti et al., 2018) and gradually lower the noise PSD, emulating improving detector sensitivity and higher signal SNR.

Since no NR simulation of the full BNS signal as observed by ground-based detectors exists, all simulated signals are constructed in a hybrid fashion: the post-merger signal is obtained through NR simulations, while the premerger signal is computed with PhenomDRT using the same parameters as the NR simulation. We choose extrinsic parameters consistent with GW170817, namely distance $D_L = 40$Mpc, inclination $\iota = 2.635$, polarization angle $\psi = 0$, and the known sky location. All signals are analyzed with the full CBC+wavelets model with run settings equivalent to the CBC_{Insp}+wavelet analysis from Table 4.2.

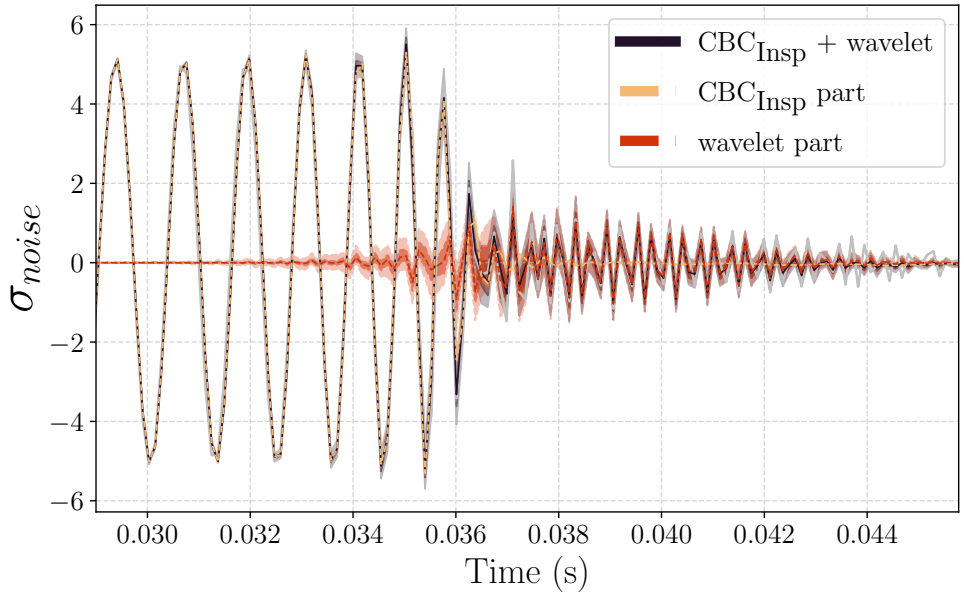

FIGURE 4.6: Whitened-time domain reconstruction of the simulated signal from Kastaun and Ohme (2021) at 6xDS. The gray dashed line gives the simulated data. The shaded regions give the 50% and 90% credible intervals for the full CBC+wavelets reconstruction (black), as well as the CBC (yellow) and the wavelets (red) component of the full analysis. The post-merger signal is reconstructed by the wavelets model.

4.5.1 Waveform from Kastaun and Ohme

The first full BNS waveform we consider was constructed and released by Kastaun and Ohme (2021). This hybrid waveform was constructed by combining the PhenomDNRT model for the inspiral and the results of NR simulations for the post-merger GW signal. The hydrodynamical simulations assume a GW170817-like initial system with the hadronic SFHO EoS (Hempel and Schaffner-Bielich, 2010), resulting in a short-lived hypermassive NS. We use the resulting hybrid waveform for a BNS system with mass ratio $q = 0.9$, detector frame component masses of $m_1 = 1.438 M_\odot$ and $m_2 = 1.294 M_\odot$, and dimensionless tidal parameters $\Lambda_1 = 280$ and $\Lambda_2 = 551$, giving a binary tidal parameter $\tilde{\Lambda} = 396$ and merger frequency $f_{\text{merger}} = 2110$ Hz.

We inject the signal in a detector network with sensitivities 2xDS, 4xDS, 6xDS which respectively denote 2, 4, and 6 time improved strain sensitivity compared to the LIGO design sensitivity (i.e., the design sensitivity divided by 2, 4 and 6). In terms of amplitude spectral density, they correspond to 8.23×10^{-24} Hz$^{-1/2}$, 4.12×10^{-24} Hz$^{-1/2}$, and 2.72×10^{-24} Hz$^{-1/2}$ at 3326 Hz, which is the frequency of the main post-merger mode of the signal, f_{peak}. The resulting network SNRs are 142, 384, and 425 for the pre-merger signal ($f < f_{\text{merger}}$), and 2.8, 5.6, and 8.4 for the post-merger signal ($f \geq f_{\text{merger}}$).

Figure 4.6 shows the whitened time-domain data and reconstruction for the injection at 6xDS, again focusing around the late stages of the signal. The full reconstruction

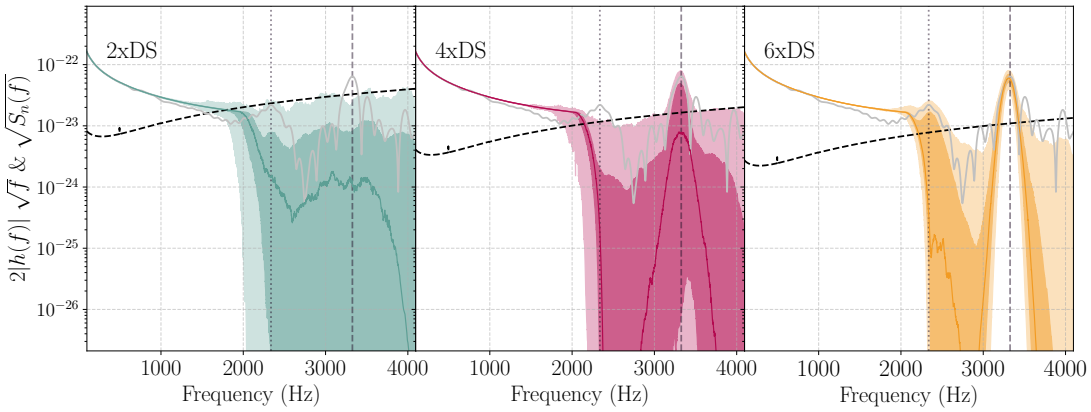

FIGURE 4.7: Median, 50%, and 90% credible intervals for the reconstructed spectrum from the Kastaun and Ohme (Kastaun and Ohme, 2021) waveform with different detector sensitivities; from left to right we assume a strain sensitivity two, four, and six times better than the advanced LIGO design sensitivity. The grey line gives the simulated data, which are the same in all cases, while the black dashed line gives the detection noise amplitude spectral density. The dashed and dotted vertical lines denote, respectively, the dominant and subdominant post-merger spectrum frequency.

with the combined CBC+wavelets model accurately captures the entire signal. We also plot the individual components of the full model, namely the CBC and the wavelets model separately. As expected, the CBC model captures everything up to merger and then tapers off. The wavelets overlap with the CBC model in the taper region around merger, effectively canceling out the ringdown-like oscillations of the CBC model that do not match the data. At later times, the wavelet model extends to the post-merger part of the signal, capturing its main oscillatory component.

Figure 4.7 shows the reconstructed spectrum from injections on different detectors sensitivities. As the detector sensitivity increases, the signal reconstruction becomes more accurate. The pre-merger signal is recovered in all analyses given its strength, and increasing sensitivity reduces the reconstruction uncertainty. The post-merger signal, on the other hand, is too weak to be detected in the 2xDS case, and therefore the reconstruction is essentially uninformative and similar to the GW170817 one. At increasing detector sensitivity, the post-merger signal emerges from the noise. At 4xDS the reconstructed spectrum starts identifying the main post-merger peak though only at the $\sim 50\%$ credible level. At 6xDS the reconstructed spectrum not only confidently identifies the main post-merger peak but also has evidence of a subdominant peak at around 2337 Hz.

Posteriors for select pre-merger and post-merger parameters are given in Fig. 4.8. The peak frequency of the post-merger signal $f_{\rm peak}$ is defined in the same way as Chatziioannou et al. (2017); Torres-Rivas et al. (2019): for each reconstruction posterior sample we compute the frequency at the maximum of the spectrum after the merger. If the spectrum possesses no maximum, then a sample is drawn from the $f_{\rm peak}$ prior. All posteriors are consistent with the injected values, and uncertainties decrease

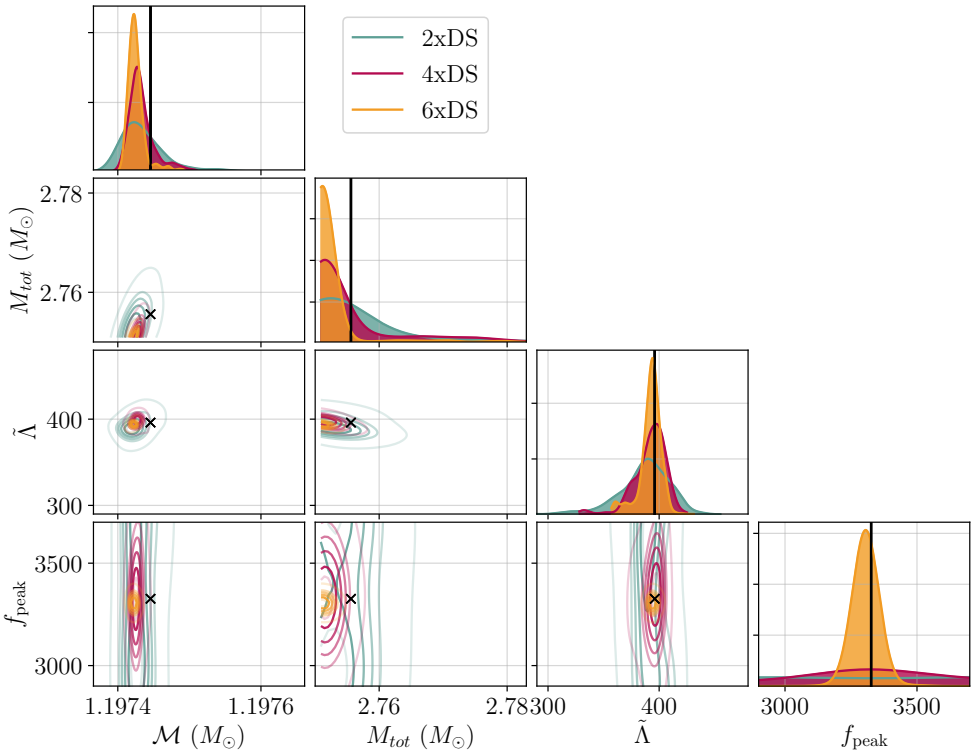

FIGURE 4.8: One- and two-dimensional marginalized posterior distributions for selected signal parameters from the analyses of Fig. 4.7: detector frame chirp mass $\mathcal{M}$, detector frame total mass M, tidal deformation parameter $\tilde{\Lambda}$, and post-merger peak frequency $f_{\rm peak}$. Vertical black lines or crosses denote the true values of each parameter.

with increasing detector sensitivity as expected. The $f_{\rm peak}$ posterior at 2xDS is essentially the prior, consistent with the fact that the post-merger signal was not detected. At 4xDS the $f_{\rm peak}$ posterior starts exhibiting a peak at the injected value, consistent with the partial identification of the post-merger signal in the middle panel of Fig. 4.7. At 6xDS the peak frequency is accurately measured.

4.5.2 Hadronic EoS

We construct further simulated signals based on the post-merger NR simulations from Torres-Rivas et al. (2019). All simulations in that study were at the time constructed to be consistent with GW170817, though in light of new data from NICER (Riley et al., 2019; Miller et al., 2019, 2021; Riley et al., 2021) some of the softest EoSs there are now disfavored. We work with EoS3 from Torres-Rivas et al. (2019), a hadronic EoS that is consistent with radii values inferred in Legred et al. (2021) and corresponds to $f_{\rm peak} = 2880$Hz and $R_{1.4} = 12.6$ km. The inspiral portion of the signal is again described by PhenomDNRT for a GW170817-like system with detector frame component masses $m_1 = m_2 = 1.362 M_\odot$, dimensionless tidal parameters $\Lambda_1 = \Lambda_2 = \tilde{\Lambda} = 587$, and zero spin.

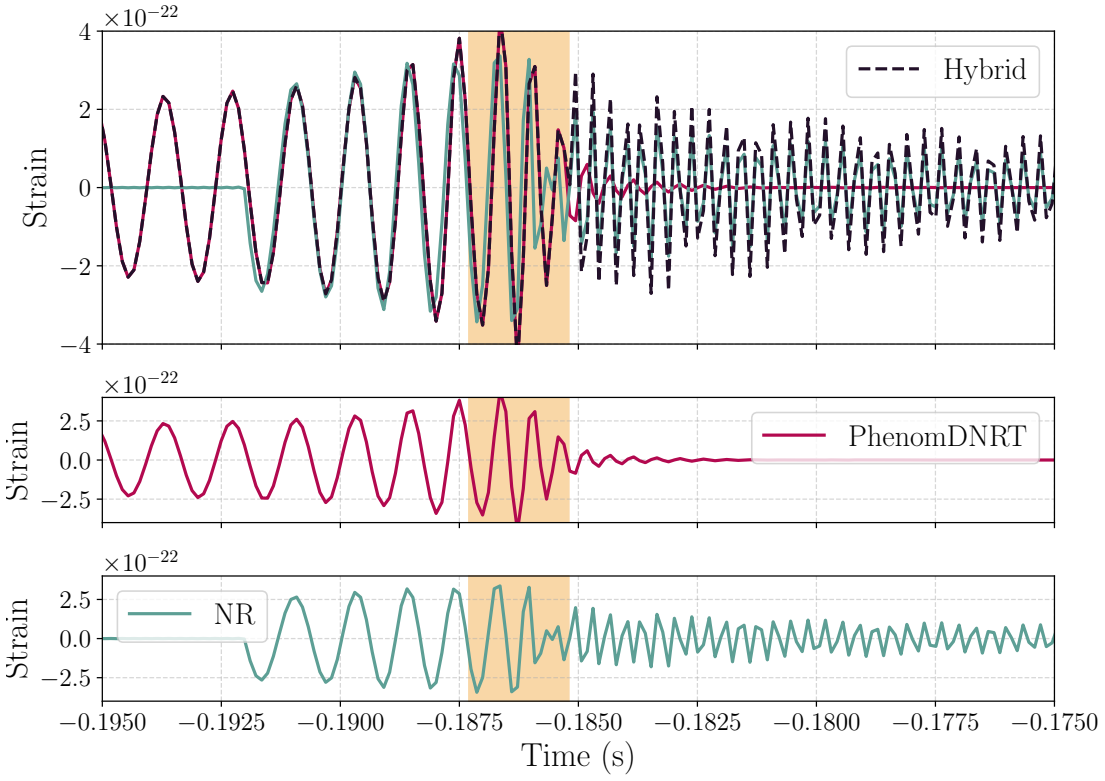

FIGURE 4.9: Comparison of the pre-merger (PhenomDNRT) and post-merger (NR simulation) waveform that are merged into a hybrid full waveform. The yellow band indicates the overlap region where the hybrid waveform transitions from being composed of the PhenomDNRT model to the post-merger NR simulation model. The maximum amplitude of the NR waveform used here is scaled to be 0.9 times the pre-merger peak amplitude.

The full waveform is constructed by aligning the projected pre-merger and post-merger waveforms in an overlap interval in the time domain. We use a transition window $[t_1, t_2]$ where the full waveform transitions from the PhenomDNRT template to the NR simulation data. The full waveform is then

$$h(t) = [1 - x(t)]\, h_{\text{inspiral}}(t) + x(t) A_{\text{scale}} h_{\text{NR}}(t), \tag{4.22}$$

where h_{inspiral} denotes the PhenomDNRT inspiral model, h_{NR} denotes the post-merger NR simulation, A_{scale} is a scale parameter that can be varied to account for the uncertainty in the post-merger amplitude, and x ensures a smooth transition through a Planck taper function

$$x(t) = \begin{cases} 0 & t \leq t_1 \\ \left[1 + \exp\left(\frac{t_2 - t}{t - t_1} + \frac{t_2 - t}{t - t_2}\right)\right]^{-1} & t_1 < t < t_2 \\ 1 & t \geq t_2. \end{cases} \tag{4.23}$$

Given uncertainties in the exact amplitude of the post-merger signal and possible

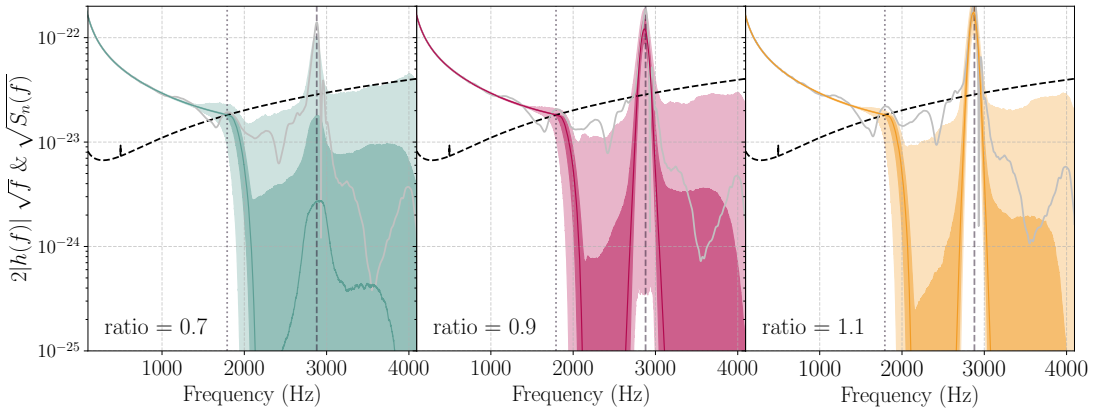

FIGURE 4.10: Median, 50% and 90% credible intervals of the reconstructed spectrum for the hadronic EoS3 signal from Sec. 4.5.2 and for pre-/post-merger peak amplitude ratios of 0.7 (left), 0.9 (middle), 1.1 (right) in the LIGO Hanford detector. The grey line gives the simulated data. The noise amplitude spectral density is also overplotted with a black dashed line.

leakage of power from the inspiral into the post-merger frequency regime we decide to fix the detector sensitivity to 2xDS and instead vary the amplitude of the post-merger signal through the scale parameter $A_{\rm scale}$. As a result all simulated signals have the same pre-merger SNR of 112, while the post-merger SNR varies and results are presented as a function of the ratio of the pre-merger to the post-merger peak amplitudes. We explore three values for the ratio of the peak post-merger to the peak pre-merger amplitudes: 0.7, 0.9 and 1.1 (corresponding to $A_{\rm scale}$ = 1.5, 2.0 and 2.5, respectively). We show an example of the above construction process in Fig. 4.9.

In Figure 4.10 we compare the reconstructed spectra from our analyses for the three cases of post-/pre-merger amplitude scaling. The inspiral portion of the signal and the corresponding reconstruction are similar in the three panels. As the amplitude of the post-merger signal increases from left to right the reconstructed spectrum includes more detailed features, progressing from hints of a post-merger peak on the left panel to increasingly more confident identification in the middle and right panel. Fig. 4.11 again shows select recovered pre-merger and post-merger parameters. The gray region overlapping with the $f_{\rm peak}$ posterior is the expected value for the post-merger frequency given the pre-merger signal. Specifically we use the fit of Chatziioannou et al. (2017) that holds for hadronic EoSs and the binary mass and tidal deformability as extracted from the pre-merger signal to compute the expected $f_{\rm peak}$. The expected and recovered posterior values for $f_{\rm peak}$ agree, showing that our analysis can correctly conclude that the pre-merger and post-merger signals are consistent with each other given expectations from NR simulations of hadronic EoSs.

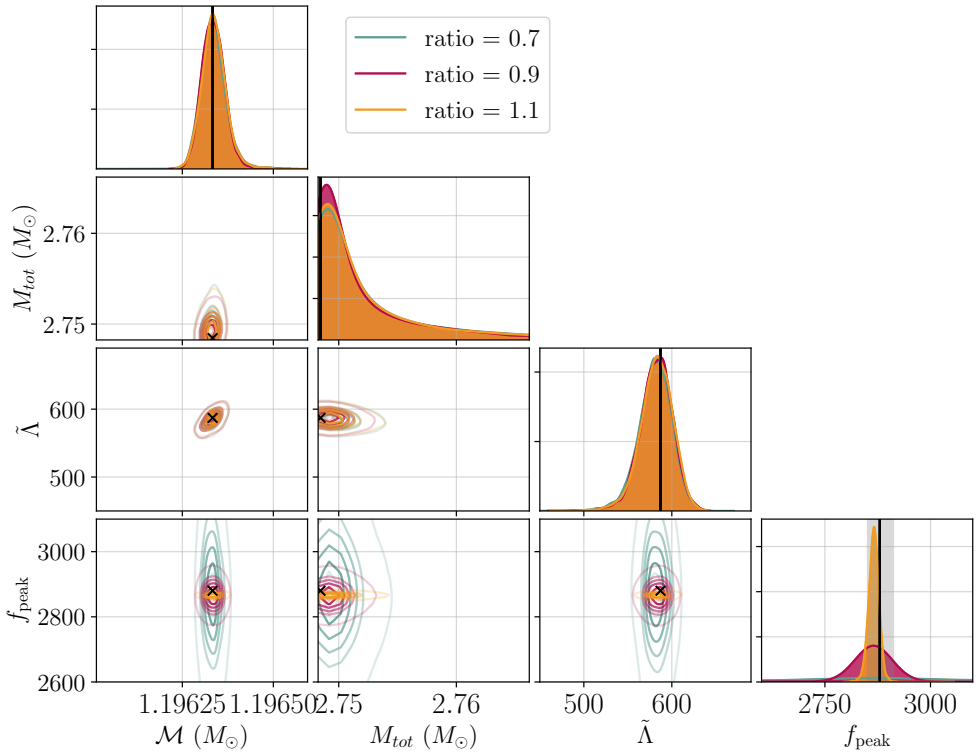

FIGURE 4.11: One- and two-dimensional marginalized posterior distributions for selected pre- and post-merger parameters for the analyses of Fig. 4.10. The parameters shown are the source-frame chirp mass $\mathcal{M}$ and total mass M, the tidal deformation parameter $\tilde{\Lambda}$ and the post-merger peak frequency $f_{\rm peak}$. Vertical black lines or crosses denote the true values of each parameter. The gray region is the expected value for $f_{\rm peak}$ given the pre-merger inferred parameters and the fit of Chatziioannou et al. (2017) that assumes hadronic EoSs. The pre-merger and post-merger results are consistent with expectations for hadronic EoSs.

4.5.3 EoS with phase transitions

Since the hypermassive NS that gives rise to the post-merger signal is characterized by higher core densities and temperatures than the pre-merger NSs, the post-merger signal has the potential to reveal new high-density physics. One possibility is a strong phase transition in the EoS toward quark degrees of freedom. Such phase transitions could occur at lower densities and thus be detectable with pre-merger data only (Del Pozzo et al., 2013; Agathos et al., 2015; Chatziioannou et al., 2015; Han and Steiner, 2019; Chen et al., 2020; Chatziioannou and Han, 2020; Han and Prakash, 2020; Zhang and Li, 2019; Pang et al., 2020; Drischler et al., 2021). However, as already mentioned above, the tidal deformability is a steeply decreasing function of the NS compactness, making pre-merger data less constraining about high-mass NSs that could contain quark cores. Post-merger data, on the other hand, can be used to probe higher densities and NR simulation-based studies have explored the potential signature of a high-density phase transition on the GW signal (Most et al., 2019;

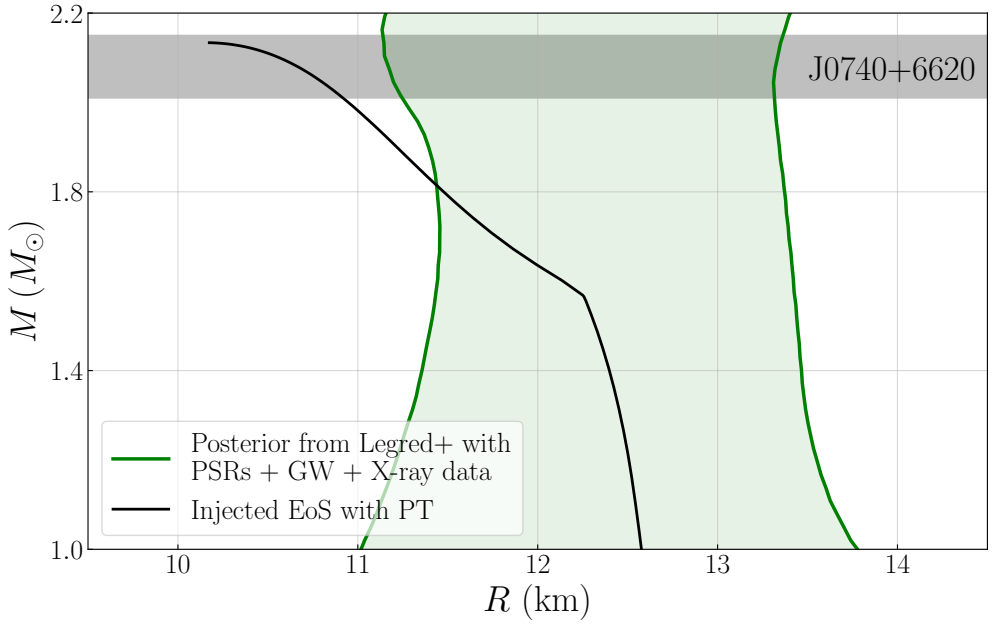

FIGURE 4.12: Mass-radius relation for EoS DD2-SF-4 used in our study of PTs (black line). The green shaded region gives the 90% credible level of the mass-radius posterior derived in Legred et al. (2021) using data from heavy pulsars, GWs, and the Miller et al. (2019, 2021) radius results using X-ray data. The adopted EoS is inconsistent with this posterior at the 90% level for masses above $\sim 1.8 M_\odot$, but it is consistent to the same level with the less constraining Riley et al. (2019, 2021) results.

Bauswein et al., 2019a,b; Weih et al., 2020; Bauswein and Blacker, 2020; Liebling et al., 2021; Prakash et al., 2021).

The most prominent signature of a sufficiently strong phase transition would be an increase in the post-merger peak frequency as the softening of the EoS would lead to a

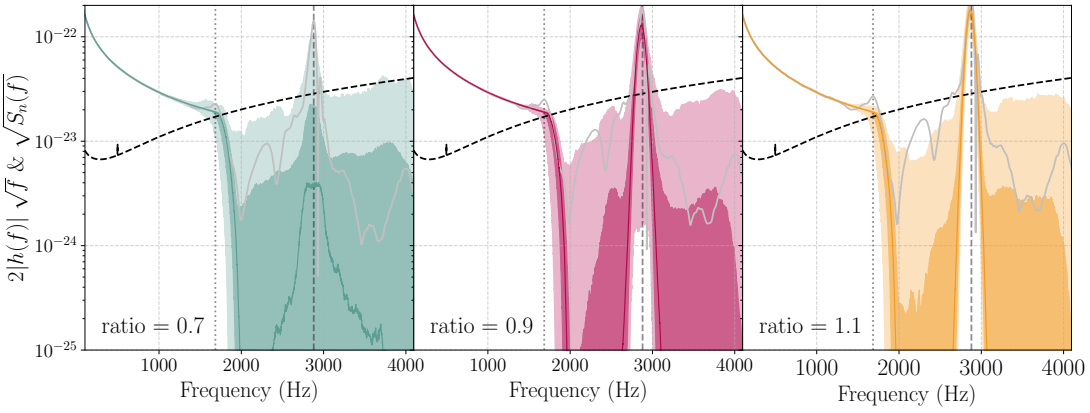

FIGURE 4.13: The median, 50% and 90% credible intervals of the reconstructed spectrum for a signal with EoS3 where the inspiral signal has been simulated with a stiffer EoS in order to mimic the phenomenology of a strong phase transitions from Sec. 4.5.3. We show pre-/post-merger peak amplitude ratios of 0.7 (left), 0.9 (middle), 1.1 (right). The grey line gives the simulated data. The noise amplitude spectral density is also overplotted with a black dashed line.

more compact merger remnant (Bauswein et al., 2019a). It has been proposed that such a frequency increase could be identified if one compares the pre-merger and post-merger data: for hadronic EoSs the tidal deformability and peak frequency follow the approximately EoS-insensitive relation shown in Figs. 4.11 and 4.17. This relation could be violated for EoSs with high-density phase transitions as the tidal deformability is determined by the hadronic part of the EoS alone, while the post-merger frequency is affected by the phase transition (Bauswein et al., 2019a; Bauswein and Blacker, 2020). The degree of deviation from the EoS-insensitive relation, and thus how detectable it is, depends on the strength of the transition (Bauswein et al., 2019a).

As a first example of such a signal, we work again with EoS3 from Torres-Rivas et al. (2019), but the pre-merger data are constructed with tidal parameters that are systematically shifted compared to their hadronic EoS value. This effectively results in a stiffer hadronic EoS that undergoes a phase transition which softens the post-merger signature to the level of EoS3. All intrinsic parameters remain to be the same as the simulations of Sec. 4.5.2 with the exception of $\tilde{\Lambda} = \Lambda_1 = \Lambda_2 = 800$.

Reconstructed spectra and parameter posteriors are shown in Figs. 4.13 and 4.14 respectively. We obtain qualitatively similar results to Figs. 4.10 and 4.11 with the main difference being that now the pre-merger and post-merger are now inconsistent as expected. Figure 4.14 shows the expected $f_{\rm peak}$ value given the pre-merger constraints on the binary properties under the assumption of a hadronic EoS. The recovered $f_{\rm peak}$ is inconsistent with this expectation to within its measurement uncertainty, signaling the presence of additional high-density effects in the EoS that affect the post-merger signal.

We further assess how well our hybrid analysis could detect a strong phase transition by simulating a signal with the DD2-SF-4 EoS (**??**) and corresponding post-merger simulation from Bauswein et al. (2019a). The mass-radius relation for this EoS is given in Fig. 4.12 and it exhibits the characteristic radius reduction due to strong phase transitions starting at $\sim 1.5 M_\odot$. Given our simulated masses, the inspiral signal is emitted by hadronic NSs, while the post-merger signal is affected by the onset of the phase transition (Bauswein et al., 2019a). For reference, we also show the 90% symmetric credible intervals of the mass-radius posterior for the EoS derived in Legred et al. (2021) using heavy pulsar, GW, and X-ray data. The DD2-SF-4 EoS is inconsistent with the posterior at the 90% level for large masses as it under-predicts the radius of a $2M_\odot$ NS, though it is consistent with current data at the 95% level. If we instead used the less constraining radius results from Riley et al. (2019, 2021), the EoS would be consistent with the posterior at the 90% level. Results are presented in Figs. 4.15 and 4.16, where again we find that the post-merger signal can be reconstructed for sufficiently loud signals, and the inconsistency between $\tilde{\Lambda}$ and $f_{\rm peak}$ can be identified.

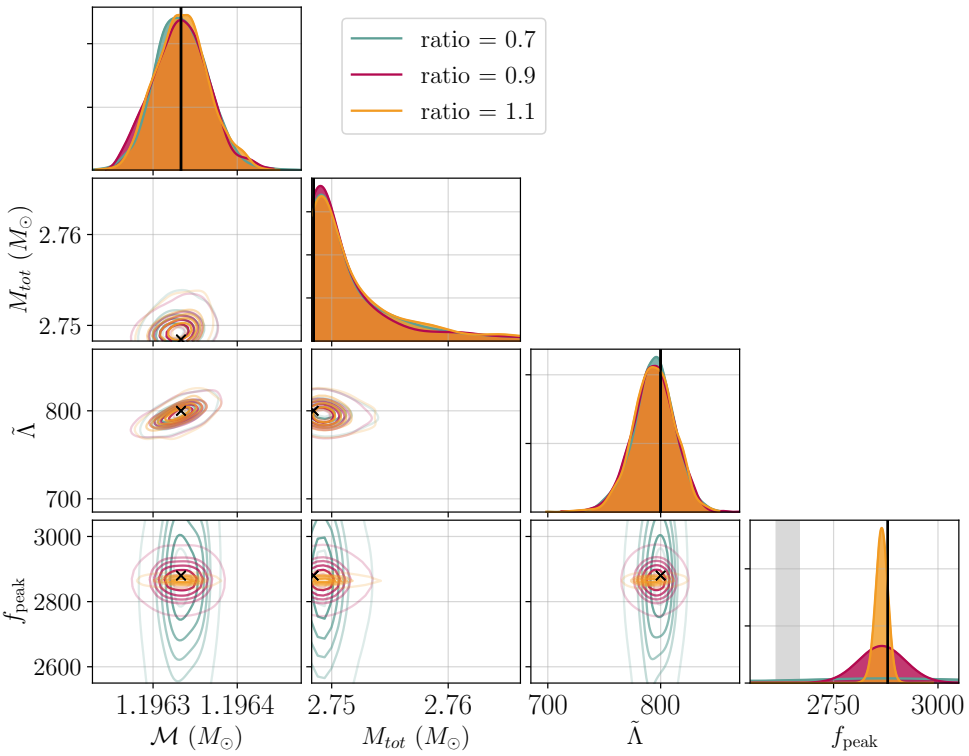

FIGURE 4.14: One- and two-dimensional marginalized posterior distributions for selected pre- and post-merger parameters for the analyses of Fig. 4.13. The parameters shown are the source-frame chirp mass $\mathcal{M}$ and total mass M, the tidal deformation parameter $\tilde{\Lambda}$ and the post-merger peak frequency $f_{\rm peak}$. Black vertical black lines or crosses denote the true values of each parameter. The gray region is the expected value for $f_{\rm peak}$ given the pre-merger inferred parameters and the fit of Chatziioannou et al. (2017) that assumes hadronic EoSs. The pre-merger and post-merger results are now inconsistent with expectations for hadronic EoSs.

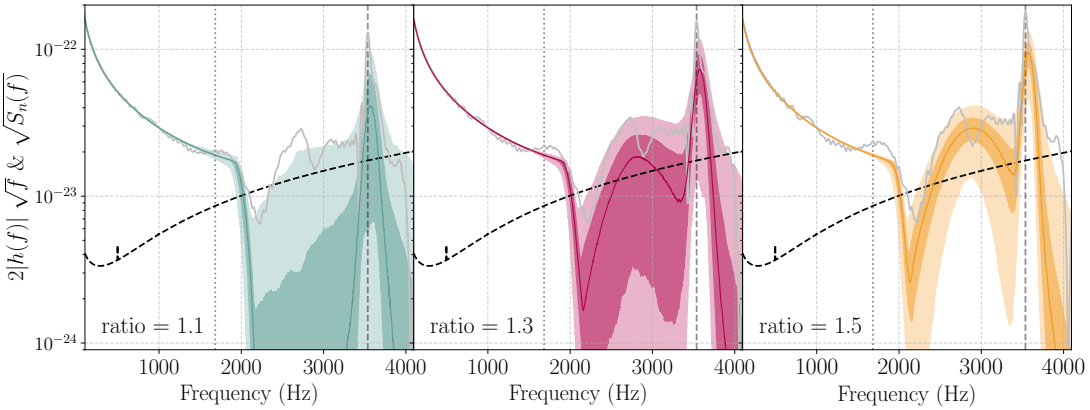

FIGURE 4.15: The median, 50% and 90% credible intervals of the reconstructed spectrum for a signal with the DD2-SF-4 EoS with phase transitions from Sec. 4.5.3 and for pre-/post-merger peak amplitude ratios of 1.1 (left), 1.3 (middle), 1.5 (right). Due to the increased $f_{\rm peak}$ value toward the less sensitive detector frequency range, such a post-merger signal would require a larger strain amplitude for detection, we therefore display results with larger values of the pre-/post-merger peak amplitude ratio compared to Fig. 4.13. The grey line gives the simulated data. The noise amplitude spectral density is also overplotted with a black dashed line.

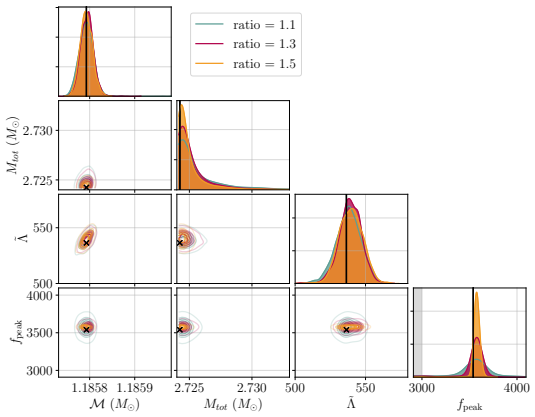

FIGURE 4.16: One- and two-dimensional marginalized posterior distributions for selected pre- and post-merger parameters for the analyses of Fig. 4.15. The parameters shown are the detector frame chirp mass $\mathcal{M}$ and total mass M, the tidal deformation parameter $\tilde{\Lambda}$ and the post-merger peak frequency $f_{\rm peak}$. Vertical black lines or crosses denote the true values of each parameter. The gray region is the expected value for $f_{\rm peak}$ given the pre-merger inferred parameters and the fit of Chatziioannou et al. (2017) that assumes hadronic EoSs. The pre-merger and post-merger data are now as expected inconsistent with expectations for hadronic EoSs.

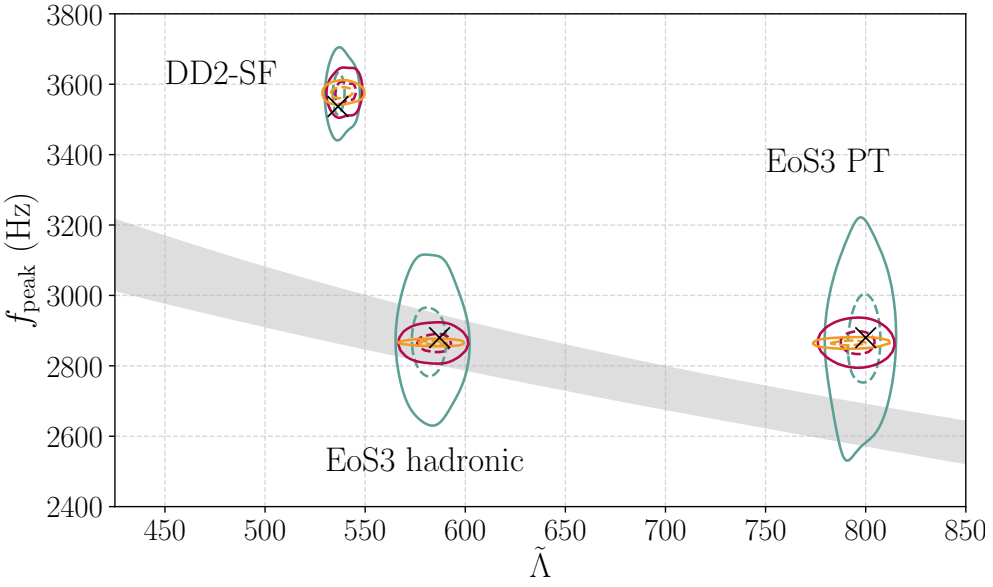

FIGURE 4.17: Comparison of the recovered $\tilde{\Lambda}$-$f_{\rm peak}$ posterior distribution from the EoS3 hadronic and PT data as well as the data with the DD2-SF-4 EoS. The grey bands corresponds to the EoS-independent fit relating the expected values of $\tilde{\Lambda}$-$f_{\rm peak}$ for hadronic NSs. We show 50% (dashed) and 90% (solid) contours, while colors correspond to those of Figs. 4.11, 4.14, 4.16. In all cases, we find the correct agreement or disagreement with the expected hadronic relation for sufficiently loud signals.

Finally, we elaborate on the pre-merger and post-merger consistency in Fig. 4.17. The grey band shows the expected relation between $\tilde{\Lambda}$ and $f_{\rm peak}$ assuming hadronic EoSs. We compute this relation by converting the $R - f_{\rm peak}$ EoS independent fit by Chatziioannou et al. (2017) (Fig. 4) to a $\tilde{\Lambda} - f_{\rm peak}$ relation. We convert R to Λ using the scaling relation calibrated for GW170817-like systems for our source mass parameters $M_{\rm tot}$ and $\mathcal{M}$ given by Eq. 17 in Zhao and Lattimer (2018). Overplotted are the two-dimensional posteriors for $\tilde{\Lambda} - f_{\rm peak}$ for the case of a hadronic EoS (EoS3), the phase-transition version of the EoS3-based signal, and the DD2-SF-4 EoS again with a strong phase transition. The injected values for the hadronic case is consistent with expectations and this is confirmed by the recovered values to within the statistical error. On the other hand, a strong phase transition leads to a $\tilde{\Lambda} - f_{\rm peak}$ combination that is inconsistent with hadronic expectations and the extracted posteriors are able to identify this behavior.

4.6 Conclusions

We present a hybrid approach to study the full GW signal emitted during a BNS coalescence, including a possible post-merger component. Our method models the inspiral part of the coalescence with waveform templates as implemented in `LALSimulation` and the post-merger part of the signal with a superposition of wavelets. We do not impose phase coherence in the transition between the template and the wavelets, however, the full signal is smooth and coherent as the template model used already extends coherently past merger and into the post-merger part of the signal. Applying our method to GW170817, we demonstrate that the inspiral parameters are consistent with previous results. We do not detect a post-merger signal for GW170817 but find that the high-frequency portion of the signal between 1000-1500 Hz is only detected using the full hybrid analysis in contrast with the traditional post-merger-only analyses. We apply our method to simulated signals with a detectable post-merger component and show that our full analysis simultaneously reconstructs both the inspiral tidal deformation and the post-merger dominant frequency peak when the SNR is sufficiently high.

In this work, we sample the inspiral tidal parameters and post-merger peak frequency independently, imposing no relation between them and no assumption on the nature of the EoS. This allows us to detect a possible signature of the hadron-quark phase transition in NS mergers, as it lead to a characteristic frequency shift of the post-merger signal. A possible extension of this analysis could use information extracted from a pre-merger signal to guide the analysis of the post-merger signal. If one restricts to a hadronic EoS, relations between the pre-merger tidal parameters and post-merger peak frequency could be used to predict the approximate location of the dominant post-merger peak from the inspiral tidal parameters. This could improve

the prospects of detecting the post-merger signal in weaker signals, at the expense of assuming the EoS is similar to a hadronic one. A further possible improvement concerns the use of "chirplets", sine-gaussian wavelets with an evolving frequency (Millhouse et al., 2018), which might be better suited for a time-evolving post-merger frequency mode.

Going beyond the main peak, numerical simulations of post-merger signals show subdominant peaks in the GW spectrum, which can also be used to characterize the properties of the merger remnant star (Stergioulas et al., 2011; Takami et al., 2014; Bauswein and Stergioulas, 2015). Our analysis can identify some of these peaks for sufficiently loud signals. Additionally, numerical simulations suggest that the main and the secondary peaks in the spectrum are related to each other (Bauswein and Stergioulas, 2015; Clark et al., 2016; Soultanis et al., 2021). In future work, this feature can be utilized in the form of a prior that links the two peaks and is parametrized in terms of the remnant compactness. Such a prior would enhance the sensitivity of the analysis to secondary peaks and offer additional information about the remnant star.

Next generation detectors are expected to detect thousands of BNSs (Regimbau et al., 2012; Sachdev et al., 2020; Adhikari et al., 2020), the majority of which will be weak with individually undetectable post-merger signals. While we study single loud sources here, our method can also be applied to the expected numerous weaker BNS sources which might dominate the overall constraints when combined (Haster et al., 2020). Additionally, our analysis concerns pre-merger SNRs in the hundreds, where systematic biases in the waveform models could be important and would have to be mitigated in the lead-up to such improved detectors (Dudi et al., 2018; Samajdar and Dietrich, 2019; Gamba et al., 2021; Chatziioannou, 2021; Kunert et al., 2021; Pratten et al., 2021). However, as we make minimal assumptions about the post-merger waveform and use sine-gaussian wavelets, we expect the reconstructed signal to be less susceptible to systematics. It is, however, possible that biases in the pre-merger waveform will lead to biased inferences about the post-merger when making use of EoS-independent relations connecting them. The extremely sensitive observations possible with next-generation detectors indeed require control over a wide range of potential systematic biases, and the flexible analysis presented here helps mitigate such biases from the post-merger signal.

Chapter 5

Conclusions

Neutron stars remain a complex but fascinating opportunity for the study of matter at extreme densities, allowing a glimpse into conditions and states of matter that cannot be accessed elsewhere in the universe. They are studied through a wide range of techniques using electromagnetic observations in the radio and X-ray spectrum and since 2017 also using gravitational waves from binary coalescences involving a neutron star. With ever-increasing detector capabilities in both fields and many open questions still remaining, the need grows for improved theoretical models in order to explain the more detailed observations as well as improved data-analysis methods in order to reveal any signals that could help constrain our theoretical models.

In this thesis I use two distinct approaches to constrain macroscopic properties of neutron stars in order to gain insight into the microscopic properties of their complex, rich interiors. In Part I we develop improved theoretical models of the neutron star microphysics that ultimately describe the thermal evolution of the neutron star surface which can be confronted with data. Theoretical models of cooling neutron stars allow for gaining insights into a broad range of microphysics such as possible neutron superfluidity (Page et al., 2011; Ho et al., 2015), exotic nucleonic matter formation (Alcain and Dorso, 2017; Caplan et al., 2018; Newton et al., 2021) and nuclear (burning) processes. Additionally, a measurement of the heat capacity of the core gives complementary information to a measurement of the equation of state and will help constrain the hypothesized (exotic) particle types and our understanding of particle interactions at supranuclear densities (Cumming et al., 2017).

In particular, we study diffusive nuclear burning (DNB) in the low density envelope and its effect on predictions of the neutron star thermal evolution. The envelope is a thin region still composed of familiar iron and hydrogen-like elements, but it is critical for our understanding of what occurs in the dense regions below. In the envelope, ionic matter isn't fully degenerate yet and therefore it acts as a thermal insulator between the interior and the surface and determines what we observe at the surface.

Its heat blanketing effects strongly depend on its composition, e.g., an envelope containing more light elements such as H and He is more transparent to the heat flux than when more heavier elements are present. Therefore, to infer the interior temperature from surface temperature observations, it is important to have a good model of the composition of the envelope.

Previously, the envelope was assumed to be static in cooling models, as its thermal timescale is much shorter than in the dense interior, and therefore it could be used as a boundary condition for numerical codes evolving the interior thermal state. In Chapter 2 we include the effect of diffusive nuclear burning in cooling models for the first time and find that it can lead to time-variable envelope compositions as the composition changes due to nuclear burning - as well as possible accretion from the interstellar medium. We provide new analytic temperature relations that include the effects of DNB. We show that including a changing composition can be important for the interpretation of cooling models and that it can explain the presence of a carbon atmosphere in young neutron stars such as Cassiopeia A, which is an important base for the first evidence of neutron superfluidity in the core.

Another way DNB could change the picture of neutron star cooling, is if it produces a non-negligible heat-flux. In Chapter 3 we focus on neutron star cooling in low mass X-ray binaries, where the crust cools down back to thermal equilibrium with the core over a timescale of years to decades after it was heated in a prolonged accretion outburst (Wijnands et al., 2017). In these systems, the conditions may be such that the temperature of the stratified envelope after the accretion outburst is high enough to drive high rates of diffusive nuclear burning. This could completely change the composition of the envelope after the outburst, as well as introduce a new source of heating in the envelope which would have been interpreted as originating from the interior. We calculate new temperature relations for the envelope that include DNB for H-C envelopes and calculate cooling models that include a time-variable envelope composition as well as a possible heat source in the envelope. Our most important conclusion from Part I is that the treatment of envelopes in cooling models as a static boundary condition may not be appropriate in all cases, and that it may be necessary to include more complex processes such as diffusive nuclear burning as they can alter the inferences made from cooling studies.

The study of cooling neutron star sources has many exciting prospects, and perhaps as many open issues to be resolved still. In our work, we have not included possible effects of a magnetic field on the thermal evolution and therefore limited ourselves to cases with low magnetic fields which are not expected to affect the heat flow. However, an exciting prospect is to extend these detailed cooling models to multiple dimensions and include the complicated effects of magnetic fields and their evolution (see Viganò et al. (2021)). This would open up cooling studies to be applied to many

more neutron star cooling sources, as a large number are measured to have extremely strong magnetic fields.

Confronting cooling models with observations of cooling neutron stars is one of several approaches to gain insight into their ultradense environments with interesting states of matter that would help us understand physics in extremes and one of the fundamental forces in the universe. It allows us to peek into the interiors while they are "cold" (i.e., 10^{8-10} K) compared to nuclear physics standards and matter is regarded to be in their lowest energy state (equivalent to when T=0 K). In Part II we move to the opposite regime in which neutron stars can be studied: The study of a possible heavy, hot neutron star remnant after two neutron stars merge that could be discovered using gravitational waves.

For the first and most famous detection of a binary neutron star (BNS) coalescence (GW170817) only the inspiral and merger was detected using gravitational waves, while its environment was observed in the electromagnetic spectrum until long after the merger. During the inspiral the neutron stars are tidally deformed as they orbit closer and closer to each other which causes them to loose more orbital energy as gravitational waves and merge faster than two black holes of the same mass would. The amount that an inspiraling neutron star deforms, is directly linked to its structure and equation of state: More compact stars are harder to deform than larger, less compact stars. Therefore, measurements of the tidal deformability of the neutron stars during the inspiral directly constrain the neutron star equation of state (Chatziioannou, 2020).

The inspiral dynamics of a compact binary is relatively well understood and its gravitational wave signal can be modeled using solutions to general relativity. The description of what happens during and after the merger is more complex and less well-understood. Relativistic simulations of this phase are numerically expensive as in these hot, dense and violent conditions, high temporal and spatial precision is necessary. Furthermore, complex physical processes need to be included such as thermal effects, turbulences and magnetohydrodynamical instabilities. Despite the complexity and divergence in the exact prediction of the creation and evolution of a post-merger remnant, numerical simulations provide a robust prediction for the primary peak in the frequency spectrum of the remnant's gravitational wave signal.

Detection of this post-merger remnant's peak frequency is an exciting prospect that is expected with future detector upgrades and more detections of BNS coalescences (Torres-Rivas et al., 2019). Information from the post-merger remnant is complimentary to that from the inspiral, as the heavier post-merger remnant with a higher maximum density allows access to the highest density regime of the equation of state. As no precise models exist for the description of the post-merger signal, extracting a post-merger signal from data requires a different approach than for the

inspiral. `BayesWave` has been used for detecting simulated post-merger signals, as it can model any signal morphology without making assumptions about its physical origin using a superposition of sine-Gaussian wavelet functions. So far, all studies of and searches for the BNS post-merger signal were done separately from the inspiral and vice versa. While no post-merger signal has been detected, this has a significant computational advantage, as the inspiral requires long duration ($\sim$ 128s) but relatively low frequency ($\leq$2048Hz) data while a study of only the post-merger could use short duration data ($\sim$1s) while requiring high frequencies (up to 4096Hz). However, as detections of a post-merger remnant are anticipated, the need for an analysis method that includes the full BNS signal grows.

In Part II of this thesis, we present an analysis method to coherently analyze the full BNS gravitational wave signal including both the inspiral and post-merger phase. Our method uses template waveform solutions to general relativity for the inspiral, and sine-Gaussian wavelets to simultaneously model any signal component that is not captured by the inspiral waveforms. We extend `BayesWave`'s RJMCMC with a CBC inspiral sampler and include several techniques to reduce memory and computation times as the full analysis requires a large amount of data. We interface our code with the `LALSimulation` (LIGO Scientific Collaboration, Virgo Collaboration, 2018) waveforms suite such that in principle any non-precessing waveform could be used, although for our application we use `PhenomD` and `PhenomDNRT` to model, respectively, binary black hole and binary neutron star inspirals.

We create full waveforms that combine the simulations of Torres-Rivas et al. (2019) to model the post-merger and the `LALSimulation` waveforms to model the inspiral and also use the full waveforms computed by Kastaun and Ohme (2021). We perform full waveform analyses using our hybrid model that includes both the CBC inspiral model and the wavelet model. Our method is able to accurately reconstruct the inspiral parameters such as chirp mass, spin, mass ratio and tidal parameters. When our injected post-merger signal-to-noise is high enough, the code simultaneously extracts the post-merger dominant peak. We also use the code to reproduce the primary post-merger peak for injections that include a phase transition between the premerger and the post-merger equation of state and successfully extract and find evidence for a phase transition.

We also reanalyze GW170817 and perform a simultaneous search for the full signal. Consistent with previous analyses, we do not detect a post-merger signal but are able to provide upper limits on the post-merger gravitational wave strain. We compare our full hybrid analyses with analyses of the inspiral and the post-merger frequency range separately. We find that the separate analysis does not lead to a detection of the merger signal in the frequency range 1000-1500 Hz, unlike our hybrid analysis. This suggests that this high-frequency portion of the signal is not individually detectable, but only inferred coherently from the preceding inspiral signal using a hybrid analysis.

While we apply the hybrid analysis code to studies of the full BNS signal, it can be used for other applications too. For example, to use the CBC model for the inspiral signal and the wavelet model to capture any detector glitches as we show in Chatziioannou et al. (2021a). In principle, the hybrid analysis has wide applications for compact binary inspirals where any unmodelled signal accompanying the CBC inspiral model will be caught by the wavelet model. We hope it will be a useful tool for the next era of gravitational wave detections that will undoubtedly bring new predicted and unexpected discoveries from compact binary coalescences.

This thesis has been divided into two parts that contribute to distinct approaches aiming to improve our understanding of neutron stars and the mysteries they hold within. It contains a part that is more theoretical in nature: Building neutron stars in computer models based on the current understanding of their microphysics and evolving them over time to make predictions for properties that can be observed. The other part focuses more on the practical aspect of going from data to discoveries: Building a tool that can detect and extract physical properties of full binary neutron star merger signals that are expected to be observed in the not too distant future. Both aspects are important pieces of the puzzle that only requires more attention in the future: we need physical models that can make sense of what we see, as well as data analysis methods that can help extract the signals to constrain those models. Both parts have a highly computational nature and highlight the increasingly important role of software development and numerical education, resources and tools in (astro)physics research. With improved detectors and computational resources, it will be exciting to see what the next era of neutron star research will bring and its impact on the understanding of astrophysical phenomena as well as fundamental physics.

6. Publications and author contributions

Thesis chapters

The chapters in this thesis are based on the following papers of which I am the first author. The co-authors were involved with supervising my work and helped with the interpretation of the results unless otherwise indicated. Wynn Ho is the main supervisor of my PhD work. Philip Chang provided a diffusive nuclear burning code to start from and Dany Page is the main developer of *NSCool* used in the first two papers. For the final paper, I developed code using the existing `BayesWave` and `QuickCBC` (by Neil Cornish) codes. In the chapters it is noted in detail where my work is based on other codes and other work.

- **Chapter 2**: *Diffusive nuclear burning in cooling simulations and application to new temperature data of the Cassiopeia A neutron star*
  Wijngaarden M. J. P., Ho W. C. G., Chang P., Heinke C. O., Page D., Beznogov M., Patnaude D. J., 2019, MNRAS, 484, 974
  Wynn Ho performed the Cas A Chandra data analysis and led the writing of subsections 2.6.1 and 2.6.2 and Appendix A.2.

- **Chapter 3**: *The effect of diffusive nuclear burning in neutron star envelopes on cooling in accreting systems*
  Wijngaarden M. J. P., Ho W. C. G., Chang P., Heinke C. O., Page D., Beznogov M., Patnaude D. J., 2019, MNRAS, 493, 4936

- **Chapter 4**: *Analyzing of the full gravitational wave signal from neutron star binary coalescences*
  Wijngaarden M. J. P., K. Chatziioannou, A. Bauswein, N. Cornish, et. al.
  Paper in preparation to be submitted to Phys. Rev. D. Katerina Chatziioannou is the main scientific supervisor of this work.

Other publications

As part of my PhD work I contributed to the following publications:

- *Continued cooling of the accretion-heated neutron star crust in the X-ray transient IGR J17480-2446 located in the globular cluster Terzan 5*
  L. S. Ootes, S. Vats, D. Page, R. Wijnands, A. S. Parikh, N. Degenaar, **M. J. P. Wijngaarden**, D. Altamirano, A. Bahramian, E. M. Cackett, C. O. Heinke, J. Homan, J. M. Miller, 2019, MNRAS, 487, 1447 (I)

- *Consistent accretion-induced heating of the neutron-star crust in MXB 1659-29 during two different outbursts*
  A. S. Parikh, R. Wijnands, L. S. Ootes, D. Page, N. Degenaar, A. Bahramian, E. F. Brown, E. M. Cackett, A. Cumming, C. Heinke, J. Homan, A. Rouco Escorial, **M. J. P. Wijngaarden**, 2019, *A&A*, 624, A84 (II)

- *Cooling of the Cassiopeia A neutron star and the effect of diffusive nuclear burning*
  Wynn C. G. Ho, **M. J. P. Wijngaarden**, Philip Chang, Craig O. Heinke, Dany Page, Mikhail Beznogov, Daniel J. Patnaude, 2019, AIP Conference Proceedings, 2127, 020007 (III)

- *Early neutron star evolution in high-mass X-ray binaries*
  Wynn C. G. Ho, **M. J. P. Wijngaarden**, Nils Andersson, Thomas M. Tauris, F. Haberl, 2020, MNRAS, 494, 44 (IV)

- *Modeling compact binary signals and instrumental glitches in gravitational wave data*
  Katerina Chatziioannou, Neil Cornish, **Marcella Wijngaarden**, Tyson B. Littenberg, 2021, Phys. Rev. D, 103, 044013 (V)

- *Model-independent constraints on superfluidity from the cooling neutron star in Cassiopeia A*
  Peter S. Shternin, Dmitry D. Ofengeim, Wynn C. G. Ho, Craig O. Heinke, **M. J. P. Wijngaarden**, Daniel J. Patnaude, 2021, MNRAS, 506, 709 (VI)

- *X-ray bounds on cooling, composition, and magnetic field of the Cassiopeia A neutron star and young central compact objects*
  Wynn C. G. Ho, Yue Zhao, Craig O. Heinke, D. L. Kaplan, Peter S. Shternin, **M. J. P. Wijngaarden**, 2021, MNRAS, 506, 5015 (VII)

Appendix A

Cas A contamination model + T_s-T_b relations

A.1 Analytic T_s-T_b relations

In this section, we present accurate fits to the computed T_s-T_b-ρ^* data which can be used in cooling simulations. All the fits are obtained for an envelope with surface gravity $g_{s,0}$ = 2.4271 $\times$ 10^{14} cm s^{-2} but can be scaled for any g_s using $Y = (T_s/1\text{MK})(g_{s,0}/g_s)^{1/4}$ (Gudmundsson et al., 1983). Firstly, we obtain $T_b(T_s, \rho^*)$ fits and use an adapted version of the shape of the analytic functions (Equation B.1) presented by Beznogov et al. (2016).

$$T_b(Y,\rho*) = 10^7 K \times \left(f_4(Y) + [f_1(Y) - f_4(Y)] \times \left[f_5(Y,\rho*) + \left(\frac{\rho*}{f_2(Y)} \right)^{f_3(Y)} \right]^{p_{12}} \right) \tag{A.1}$$

A.1.1 H-He mixture

The analytic results for the H-He mixture are applicable to an envelope with a lower boundary of ρ_b = 10^8 g cm^{-3}. While the H-He temperature relations can be fit with a simple function needing less fit parameters, we fit the same function as Beznogov et al. (2016) for consistency. The best fit parameters correspond to a maximum relative error of 0.0054 and a rms of the relative error of 0.0028.

TABLE A.1: Best fit parameters for the H-He mixture with ρ_b = 10^8 g cm^{-3} (Equations A.2).

p_1	p_2	p_3	p_4	p_5	p_6	p_7	p_8
0.8254	2.086	18.70	-1.025	3.346	1.589	0.03848	1.572

p_9	p_{10}	p_{11}	p_{12}	p_{13}	p_{14}	p_{15}
21.03	3.268	0.2416	0.1923	-1.355	-0.07620	-9.864 $\times$ 10^3

$$
\begin{aligned}
f_1(Y) &= p_1 Y^{p_2} \sqrt{1 + p_3 Y^{p_4}} \\
f_2(Y) &= p_5 Y^{p_6} \sqrt{1 + p_7 Y^{p_8}} \\
f_2(Y) &= \frac{p_9 Y^{p_1 0}}{(1 - p_{11} Y + p_{12} Y^2)^2} \\
f_3(Y) &= p_{13} Y^{-p_{14}} \\
f_5(Y) &= 1
\end{aligned}
\tag{A.2}
$$

A.2 Analysis of Cassiopeia A data using previous *Chandra* contamination model N0010

Here we show results of our analysis of the same data as that in Section 2.6 but using a previous CALDB (4.7.8 in CIAO 4.9), with ACIS-S contamination model N0010, and XSPEC (12.9.1). The results presented in this Appendix are to illustrate the small changes that the ACIS-S contamination model N0012 (in CALDB 4.8.1) introduces, as well as to facilitate comparisons with the results of Posselt and Pavlov (2018) who use CALDB 4.7.3 and N0010. The analysis follows the same procedure as that described in Section 2.6, and the final fit results are given in Table A.2. Comparison with the fit results in Table 2.1 show differences in temperature and flux well within uncertainties, keeping in mind that some of the small changes are due to differing best-fit radius (13.0 km versus 12.9 km) and hence gravitational redshift.

TABLE A.2: Surface temperature $T_{s,6}$ (10^6 K), absorbed 0.5–7 keV flux F^{abs}_{-13} (10^{-13} erg cm^{-2} s^{-1}), and absorption column $N_{H,22}$ (10^{22} cm^{-2}) determined from model fits using CALDB 4.7.8 (in comparison to Table 2.1 which results from using CALDB 4.8.1) to *Chandra* ACIS-S GRADED spectra of the neutron star in the Cassiopeia A supernova remnant. Three spectral fit results are shown, two with constant N_H and one with changing N_H. Each set of 13 temperatures and fluxes are fit to a linear decline, with decline rate and fit statistic as shown. For merged ObsIDs, the MJD listed is that of the first ObsID. Number in parentheses is 1σ uncertainty in last digit.

		$N_{H,22}$	1.73		1.68		see below		
		M (M_{Sun})	1.65		1.65		1.65		
		R (km)	10.3		13.0		13.0		
		χ^2/dof	1550/1450		1512/1450		1486/1437		
ObsID	Date	MJD	$T_{s,6}$	F^{abs}_{-13}	$T_{s,6}$	F^{abs}_{-13}	$T_{s,6}$	F^{abs}_{-13}	$N_{H,22}$
114	2000 Jan 30	51573.4	2.127(10)	7.4(2)	1.862(8)	7.3(2)	1.870(10)	7.3(2)	1.73(3)
1952	2002 Feb 6	52311.3	2.125(10)	7.4(2)	1.861(8)	7.3(2)	1.869(10)	7.3(2)	1.73(3)
5196	2004 Feb 8	53043.7	2.106(10)	7.1(1)	1.845(8)	7.0(2)	1.842(10)	7.0(1)	1.66(3)
9117/9773	2007 Dec 5/8	54439.9	2.097(9)	7.0(2)	1.835(8)	6.9(2)	1.842(10)	6.9(2)	1.72(4)
10935/12020	2009 Nov 2/3	55137.9	2.089(10)	6.9(2)	1.828(8)	6.8(2)	1.832(11)	6.8(2)	1.71(4)
10936/13177	2010 Oct 31/Nov 2	55500.2	2.080(10)	6.7(1)	1.820(8)	6.6(2)	1.812(11)	6.6(1)	1.64(4)
14229	2012 May 15	56062.4	2.044(9)	6.4(2)	1.791(7)	6.3(1)	1.797(11)	6.3(2)	1.72(4)
14480	2013 May 20	56432.6	2.062(9)	6.5(1)	1.806(8)	6.4(2)	1.810(10)	6.5(1)	1.70(4)
14481	2014 May 12	56789.1	2.045(9)	6.3(1)	1.792(7)	6.2(2)	1.799(10)	6.2(2)	1.72(4)
14482	2015 Apr 30	57142.5	2.064(9)	6.6(2)	1.809(7)	6.6(2)	1.795(10)	6.5(2)	1.60(4)
19903/18344	2016 Oct 20/21	57681.2	2.059(9)	6.6(2)	1.803(8)	6.5(1)	1.789(11)	6.4(1)	1.60(4)
19604	2017 May 16	57889.7	2.050(9)	6.5(2)	1.798(7)	6.4(2)	1.795(10)	6.4(1)	1.67(4)
19605	2018 May 15	58253.7	2.050(9)	6.5(2)	1.798(7)	6.5(2)	1.777(10)	6.4(1)	1.53(4)
		10-year decline rate	$2.2 \pm 0.2\%$	$7.3 \pm 1.0\%$	$2.1 \pm 0.2\%$	$7.2 \pm 1.0\%$	$2.6 \pm 0.3\%$	$7.6 \pm 0.9\%$	
		χ^2/dof	15.8/11	12.9/11	16.8/11	15.4/11	6.3/11	11.1/11	

Appendix B

H-C envelopes: T_s-T_b relations

B.1 Analytic T_s-T_b relations

In this section, we present accurate fits to the computed T_s-T_b-y_H data which can be used in cooling simulations. All the fits are obtained for an envelope with surface gravity $g_{s,0} = 2.4271 \times 10^{14}$ cm s^{-2} but can be scaled for any g_s using $Y = (T_s/1\text{MK})(g_{s,0}/g_s)^{1/4}$ (Gudmundsson et al., 1983). We use an adapted version of the analytic functions (Equation B.1) presented by Beznogov et al. (2016).

$$T_b(Y, y_H) = 10^7 K \times \left(f_4(Y) + [f_1(Y) - f_4(Y)] \times \left[1 + \left(\frac{y_H}{f_2(Y)} \right)^{f_3(Y)} \right]^{-0.2} \right) \tag{B.1}$$

The analytic relation fitted for the H-C mixture for an envelope with a lower boundary of $\rho_b = 10^{10}$ g cm^{-3} is given by equation B.1 which consists of the functions B.2. Here, the functions $f_1 - f_4$ are identical to those for the He-C mixture in Beznogov et al.

TABLE B.1: Best fit parameters for the H-C mixture with $\rho_b = 10^{10}$ g cm^{-3} (Equations B.2).

p_1	p_2	p_3	p_4	p_5	p_6
5.708	0.09519	1.667	4.676	0.05793	1.705

p_7	p_8	p_9	p_{10}	p_{11}
16710	3.178	4.681	1.719	0.9215

(2016). We obtain a maximum relative error of 0.016. The root mean square of the relative error is 0.0037.

$$
\begin{aligned}
f_1(Y) &= p_1 Y^{p_2 \log_{10} Y + p_3} \\
f_2(Y) &= p_7 Y^{p_8 (\log_{10} Y)^2 + p_9} \\
f_3(Y) &= p_{10} \sqrt{\frac{Y}{Y^2 + p_{11}^2}} \\
f_4(Y) &= p_4 Y^{p_5 \log_{10} Y + p_6}
\end{aligned}
\tag{B.2}
$$

References

J. Aasi et al. Advanced LIGO. *Class. Quant. Grav.*, 32:074001, 2015. .

B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya, and et al. Multi-messenger observations of a binary neutron star merger. *The Astrophysical Journal*, 848(2):L12, Oct 2017a. ISSN 2041-8213. . URL http://dx.doi.org/10.3847/2041-8213/aa91c9.

B. P. Abbott et al. Binary Black Hole Mergers in the first Advanced LIGO Observing Run. *Phys. Rev. X*, 6(4):041015, 2016a. . [Erratum: Phys.Rev.X 8, 039903 (2018)].

B. P. Abbott et al. Properties of the Binary Black Hole Merger GW150914. *Phys. Rev. Lett.*, 116(24):241102, 2016b. .

B. P. Abbott et al. Search for Post-merger Gravitational Waves from the Remnant of the Binary Neutron Star Merger GW170817. *Astrophys. J. Lett.*, 851(1):L16, 2017b. .

B. P. Abbott et al. GW170817: Measurements of neutron star radii and equation of state. *Phys. Rev. Lett.*, 121(16):161101, 2018a. .

B. P. Abbott et al. GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs. 2018b.

B. P. Abbott et al. Search for gravitational waves from a long-lived remnant of the binary neutron star merger GW170817. *Astrophys. J.*, 875(2):160, 2019a. .

B. P. Abbott et al. GW190425: Observation of a Compact Binary Coalescence with Total Mass $\sim 3.4 M_{\odot}$. *Astrophys. J. Lett.*, 892(1):L3, 2020. .

Benjamin P. Abbott et al. Prospects for Observing and Localizing Gravitational-Wave Transients with Advanced LIGO, Advanced Virgo and KAGRA. *Living Rev. Rel.*, 19: 1, 2013. .

Benjamin P. Abbott et al. Effects of waveform model systematics on the interpretation of GW150914. *Class. Quant. Grav.*, 34(10):104002, 2017c. .

Benjamin P Abbott et al. Exploring the Sensitivity of Next Generation Gravitational Wave Detectors. *Class. Quant. Grav.*, 34(4):044001, 2017d. .

Benjamin P. Abbott et al. GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Phys. Rev. Lett.*, 119(16):161101, 2017e. .

B.P. Abbott et al. Observation of Gravitational Waves from a Binary Black Hole Merger. *Phys. Rev. Lett.*, 116(6):061102, 2016c. .

B.P. Abbott et al. Properties of the binary neutron star merger GW170817. *Phys. Rev. X*, 9(1):011001, 2019b. .

Rich Abbott et al. Open data from the first and second observing runs of Advanced LIGO and Advanced Virgo. *SoftwareX*, 13:100658, 2021. .

F. Acernese et al. Advanced Virgo: a second-generation interferometric gravitational wave detector. *Class. Quant. Grav.*, 32(2):024001, 2015. .

R. X. Adhikari et al. A cryogenic silicon interferometer for gravitational-wave detection. *Class. Quant. Grav.*, 37(16):165003, 2020. .

Michalis Agathos, Jeroen Meidam, Walter Del Pozzo, Tjonnie G. F. Li, Marco Tompitak, John Veitch, Salvatore Vitale, and Chris Van Den Broeck. Constraining the neutron star equation of state with gravitational wave signals from coalescing binary neutron stars. *Phys. Rev. D*, 92(2):023012, 2015. .

A. Akmal, V. R. Pandharipande, and D. G. Ravenhall. Equation of state of nucleon matter and neutron star structure. *Phys. Rev. C*, 58:1804–1828, September 1998. .

Mohammad Al-Mamun, Andrew W. Steiner, Joonas Nättilä, Jacob Lange, Richard O'Shaughnessy, Ingo Tews, Stefano Gandolfi, Craig Heinke, and Sophia Han. Combining Electromagnetic and Gravitational-Wave Constraints on Neutron-Star Masses and Radii. *Phys. Rev. Lett.*, 126(6):061101, 2021. .

P. Alcain and C. Dorso. Dynamics of Fragment Formation in Neutron Rich Matter. *ArXiv e-prints*, May 2017.

D. Alp, J. Larsson, C. Fransson, M. Gabler, A. Wongwathanarat, and H.-T. Janka. X-Ray Absorption in Young Core-collapse Supernova Remnants. *ApJ*, 864:175, September 2018. .

Nils Andersson. A Superfluid Perspective on Neutron Star Dynamics. *Universe*, 7(1): 17, January 2021. .

C. Angulo, M. Arnould, M. Rayet, P. Descouvemont, D. Baye, C. Leclercq-Willain, A. Coc, S. Barhoumi, P. Aguer, C. Rolfs, R. Kunz, J. W. Hammer, A. Mayer, T. Paradellis, S. Kossionides, C. Chronidou, K. Spyrou, S. degl'Innocenti, G. Fiorentini, B. Ricci, S. Zavatarelli, C. Providencia, H. Wolters, J. Soares, C. Grama,

J. Rahighi, A. Shotter, and M. Lamehi Rachti. A compilation of charged-particle induced thermonuclear reaction rates. *Nuclear Phys. A*, 656:3–183, August 1999. .

John Antoniadis, Paulo C.C. Freire, Norbert Wex, Thomas M. Tauris, Ryan S. Lynch, et al. A Massive Pulsar in a Compact Relativistic Binary. *Science*, 340(6131):1233232, 2013. .

K. A. Arnaud. XSPEC: The First Ten Years. In G. H. Jacoby and J. Barnes, editors, *Astronomical Data Analysis Software and Systems V*, volume 101 of *Astronomical Society of the Pacific Conference Series*, page 17, 1996.

M. Baldo, G. F. Burgio, H.-J. Schulze, and G. Taranto. Nucleon effective masses within the Brueckner-Hartree-Fock theory: Impact on stellar neutrino emission. *Phys. Rev. C*, 89(4):048801, April 2014. .

Lisa Barsotti, Peter Fritschel, Matthew Evans, and Slawomir Gras. Updated advanced ligo sensitivity design curve. Technical report, 2018. URL https://dcc.ligo.org/LIGO-T1800044/public.

Thomas W. Baumgarte, Stuart L. Shapiro, and Masaru Shibata. On the Maximum Mass of Differentially Rotating Neutron Stars. *ApJ*, 528(1):L29–L32, January 2000. .

Thomas W. Baumgarte, Stuart L. Shapiro, and Masaru Shibata. On the maximum mass of differentially rotating neutron stars. *Astrophys. J. Lett.*, 528:L29, 2000. .

A. Bauswein and H. T. Janka. Measuring Neutron-Star Properties via Gravitational Waves from Neutron-Star Mergers. *Phys. Rev. Lett.*, 108(1):011101, January 2012. .

A. Bauswein and N. Stergioulas. Unified picture of the post-merger dynamics and gravitational wave emission in neutron star mergers. *Phys. Rev. D*, 91:124056, Jun 2015. . URL http://link.aps.org/doi/10.1103/PhysRevD.91.124056.

A. Bauswein, H. T. Janka, K. Hebeler, and A. Schwenk. Equation-of-state dependence of the gravitational-wave signal from the ring-down phase of neutron-star mergers. *Phys. Rev. D*, 86(6):063001, September 2012. .

A. Bauswein, H.T. Janka, K. Hebeler, and A. Schwenk. Equation-of-state dependence of the gravitational-wave signal from the ring-down phase of neutron-star mergers. *Phys. Rev. D*, 86:063001, 2012. .

Andreas Bauswein and Sebastian Blacker. Impact of quark deconfinement in neutron star mergers and hybrid star mergers. *Eur. Phys. J. ST*, 229(22-23):3595–3604, 2020. .

Andreas Bauswein, Nikolaos Stergioulas, and Hans-Thomas Janka. Exploring properties of high-density matter through remnants of neutron-star mergers. *Eur. Phys. J. A*, 52(3):56, 2016. .

Andreas Bauswein, Niels-Uwe F. Bastian, David B. Blaschke, Katerina Chatziioannou, James A. Clark, Tobias Fischer, and Micaela Oertel. Identifying a first-order phase transition in neutron star mergers through gravitational waves. *Phys. Rev. Lett.*, 122 (6):061102, 2019a. .

Andreas Bauswein, Niels-Uwe Friedrich Bastian, David Blaschke, Katerina Chatziioannou, James Alexander Clark, Tobias Fischer, Hans-Thomas Janka, Oliver Just, Micaela Oertel, and Nikolaos Stergioulas. Equation-of-state Constraints and the QCD Phase Transition in the Era of Gravitational-Wave Astronomy. *AIP Conf. Proc.*, 2127(1):020013, 2019b. .

Andreas Bauswein, Sebastian Blacker, Vimal Vijayan, Nikolaos Stergioulas, Katerina Chatziioannou, James A. Clark, Niels-Uwe F. Bastian, David B. Blaschke, Mateusz Cierniak, and Tobias Fischer. Equation of state constraints from the threshold binary mass for prompt collapse of neutron star mergers. *Phys. Rev. Lett.*, 125(14):141103, 2020. .

BayesWave Glitch Subtraction for GW170817. https://dcc.ligo.org/LIGO-T1700406/public.

Gordon Baym, Tetsuo Hatsuda, Toru Kojo, Philip D. Powell, Yifan Song, and Tatsuyuki Takatsuka. From hadrons to quarks in neutron stars: a review. *Rept. Prog. Phys.*, 81(5):056902, 2018. .

Bence Bécsy, Peter Raffai, Neil J. Cornish, Reed Essick, Jonah Kanner, Erik Katsavounidis, Tyson B. Littenberg, Margaret Millhouse, and Salvatore Vitale. Parameter Estimation for Gravitational-wave Bursts with the BayesWave Pipeline. *ApJ*, 839(1):15, April 2017. .

S. Bernuzzi, T. Dietrich, and A. Nagar. Modeling the Complete Gravitational Wave Spectrum of Neutron Star Mergers. *Physical Review Letters*, 115(9):091101, August 2015. .

Sebastiano Bernuzzi. Neutron Star Merger Remnants. *Gen. Rel. Grav.*, 52(11):108, 2020. .

M. V. Beznogov, A. Y. Potekhin, and D. G. Yakovlev. Diffusive heat blanketing envelopes of neutron stars. *MNRAS*, 459:1569–1579, June 2016. .

M. V. Beznogov, A. Y. Potekhin, and D. G. Yakovlev. Heat blanketing envelopes of neutron stars. *Phys. Rep.*, 919:1–68, July 2021. .

Mikhail V. Beznogov, Dany Page, and Enrico Ramirez-Ruiz. Thermal Evolution of Neo-neutron Stars. I. Envelopes, Eddington Luminosity Phase, and Implications for GW170817. *ApJ*, 888(2):97, January 2020. .

Taylor Binnington and Eric Poisson. Relativistic theory of tidal Love numbers. *Phys. Rev.*, D80:084018, 2009. .

Bhaskar Biswas. Impact of PREX-II, the revised mass measurement of PSRJ0740+6620, and possible NICER observation on the dense matter equation of state. 5 2021.

O. Blaes and P. Madau. Can we observe accreting, isolated neutron stars? *ApJ*, 403: 690–705, February 1993. .

Luc Blanchet. Gravitational Radiation from Post-Newtonian Sources and Inspiralling Compact Binaries. *Living Reviews in Relativity*, 17(1):2, December 2014. .

Luc Blanchet. Gravitational Radiation from Post-Newtonian Sources and Inspiralling Compact Binaries. *Living Rev. Rel.*, 17:2, 2014. .

D. Blaschke, H. Grigorian, D. N. Voskresensky, and F. Weber. Cooling of the neutron star in Cassiopeia A. *Phys. Rev. C*, 85(2):022802, February 2012. .

D. Blaschke, H. Grigorian, and D. N. Voskresensky. Nuclear medium cooling scenario in light of new Cas A cooling data and the 2M$_?$ pulsar mass measurements. *Phys. Rev. C*, 88(6):065805, December 2013. .

S. Bogdanov. Modeling the X-Rays from the Central Compact Object PSR J1852+0040 in Kesteven 79: Evidence for a Strongly Magnetized Neutron Star. *ApJ*, 790:94, August 2014. .

A. Bonanno, M. Baldo, G. F. Burgio, and V. Urpin. The neutron star in Cassiopeia A: equation of state, superfluidity, and Joule heating. *A&A*, 561:L5, January 2014. .

Sukanta Bose, Kabir Chakravarti, Luciano Rezzolla, B. S. Sathyaprakash, and Kentaro Takami. Neutron-star Radius from a Population of Binary Neutron Star Mergers. *Phys. Rev. Lett.*, 120(3):031102, 2018. .

Matteo Breschi, Sebastiano Bernuzzi, Francesco Zappa, Michalis Agathos, Albino Perego, David Radice, and Alessandro Nagar. kiloHertz gravitational waves from binary neutron star remnants: time-domain model and constraints on extreme matter. *Phys. Rev. D*, 100(10):104029, 2019. .

E. F. Brown, L. Bildsten, and P. Chang. Variability in the Thermal Emission from Accreting Neutron Star Transients. *ApJ*, 574:920–929, August 2002. .

Edward F. Brown and Andrew Cumming. Mapping Crustal Heating with the Cooling Light Curves of Quasi-Persistent Transients. *ApJ*, 698:1020–1032, Jun 2009. .

Edward F. Brown, Lars Bildsten, and Robert E. Rutledge. Crustal Heating and Quiescent Emission from Transiently Accreting Neutron Stars. *ApJ*, 504:L95–L98, September 1998. .

Edward F. Brown, Andrew Cumming, Farrukh J. Fattoyev, C. J. Horowitz, Dany Page, and Sanjay Reddy. Rapid Neutrino Cooling in the Neutron Star MXB 1659-29. *Phys. Rev. Lett.*, 120:182701, May 2018. .

T. Callister. A Thesaurus for Common Priors in Gravitational-Wave Astronomy. 4 2021.

M. E. Caplan, A. S. Schneider, and C. J. Horowitz. Elasticity of Nuclear Pasta. *Phys. Rev. Lett.*, 121(13):132701, September 2018. .

Zack Carson, Katerina Chatziioannou, Carl-Johan Haster, Kent Yagi, and Nicolás Yunes. Equation-of-state insensitive relations after GW170817. *Phys. Rev. D*, 99(8): 083016, 2019. .

N. Chamel. Entrainment in Superfluid Neutron-Star Crusts: Hydrodynamic Description and Microscopic Origin. *Journal of Low Temperature Physics*, 189(5-6): 328–360, December 2017. .

Nicolas Chamel and Pawel Haensel. Physics of Neutron Star Crusts. *Living Reviews in Relativity*, 11(1):10, Dec 2008. .

P. Chang and L. Bildsten. Diffusive Nuclear Burning in Neutron Star Envelopes. *ApJ*, 585:464–474, March 2003. .

P. Chang and L. Bildsten. Evolution of Young Neutron Star Envelopes. *ApJ*, 605: 830–839, April 2004. .

P. Chang, L. Bildsten, and P. Arras. Diffusive Nuclear Burning of Helium on Neutron Stars. *ApJ*, 723:719–728, November 2010. .

Katerina Chatziioannou. Neutron star tidal deformability and equation of state constraints. *Gen. Rel. Grav.*, 52(11):109, 2020. .

Katerina Chatziioannou. Uncertainty limits on neutron star radius measurements with gravitational waves. 8 2021.

Katerina Chatziioannou and Sophia Han. Studying strong phase transitions in neutron stars with gravitational waves. *Phys. Rev. D*, 101(4):044019, 2020. .

Katerina Chatziioannou, Kent Yagi, Antoine Klein, Neil Cornish, and Nicolas Yunes. Probing the Internal Composition of Neutron Stars with Gravitational Waves. *Phys. Rev. D*, 92(10):104008, 2015. .

Katerina Chatziioannou, James Alexander Clark, Andreas Bauswein, Margaret Millhouse, Tyson B. Littenberg, and Neil Cornish. Inferring the post-merger gravitational wave emission from binary neutron star coalescences. *Phys. Rev. D*, 96 (12):124035, 2017. .

Katerina Chatziioannou, Carl-Johan Haster, and Aaron Zimmerman. Measuring the neutron star tidal deformability with equation-of-state-independent relations and gravitational waves. *Phys. Rev.*, D97(10):104036, 2018a. .

Katerina Chatziioannou, Geoffrey Lovelace, Michael Boyle, Matthew Giesler, Daniel A. Hemberger, Reza Katebi, Lawrence E. Kidder, Harald P. Pfeiffer, Mark A. Scheel, and Béla Szilágyi. Measuring the properties of nearly extremal black holes with gravitational waves. *Phys. Rev. D*, 98(4):044028, 2018b. .

Katerina Chatziioannou, Carl-Johan Haster, Tyson B. Littenberg, Will M. Farr, Sudarshan Ghonge, Margaret Millhouse, James A. Clark, and Neil Cornish. Noise spectral estimation methods and their impact on gravitational wave measurement of compact binary mergers. *Phys. Rev. D*, 100:104004, Nov 2019. . URL https://link.aps.org/doi/10.1103/PhysRevD.100.104004.

Katerina Chatziioannou, Neil Cornish, Marcella Wijngaarden, and Tyson B. Littenberg. Modeling compact binary signals and instrumental glitches in gravitational wave data. *Phys. Rev. D*, 103(4):044013, 2021a. .

Katerina Chatziioannou, Maximiliano Isi, Carl-Johan Haster, and Tyson B. Littenberg. Morphology-independent test of the mixed polarization content of transient gravitational wave signals. *Phys. Rev. D*, 104(4):044005, 2021b. .

Hsin-Yu Chen, Paul M. Chesler, and Abraham Loeb. Searching for exotic cores with binary neutron star inspirals. *Astrophys. J. Lett.*, 893(1):L4, 2020. .

J. M. C. Chen, J. W. Clark, R. D. Davé, and V. V. Khodel. Pairing gaps in nucleonic superfluids. *Nuclear Phys. A*, 555:59–89, April 1993. .

J. Clark, A. Bauswein, L. Cadonati, H. T. Janka, C. Pankow, and N. Stergioulas. Prospects For High Frequency Burst Searches Following Binary Neutron Star Coalescence With Advanced Gravitational Wave Detectors. *Phys. Rev. D*, 90(6): 062004, 2014. .

James Alexander Clark, Andreas Bauswein, Nikolaos Stergioulas, and Deirdre Shoemaker. Observing Gravitational Waves From The Post-Merger Phase Of Binary Neutron Star Coalescence. *Class. Quant. Grav.*, 33(8):085003, 2016. .

Neil J. Cornish. Fast Fisher Matrices and Lazy Likelihoods. *arXiv e-prints*, art. arXiv:1007.4820, July 2010.

Neil J. Cornish. Rapid and Robust Parameter Inference for Binary Mergers. *Phys. Rev. D*, 103(10):104057, 2021. .

Neil J. Cornish and Tyson B. Littenberg. BayesWave: Bayesian Inference for Gravitational Wave Bursts and Instrument Glitches. *Class. Quant. Grav.*, 32(13): 135012, 2015. .

Neil J. Cornish, Tyson B. Littenberg, Bence Bécsy, Katerina Chatziioannou, James A. Clark, Sudarshan Ghonge, and Margaret Millhouse. BayesWave analysis pipeline in the era of gravitational wave observations. *Phys. Rev. D*, 103(4):044006, 2021. .

H. Thankful Cromartie et al. Relativistic Shapiro delay measurements of an extremely massive millisecond pulsar. *Nature Astron.*, 4(1):72–76, 2019. .

Andrew Cumming, Edward F. Brown, Farrukh J. Fattoyev, C. J. Horowitz, Dany Page, and Sanjay Reddy. Lower limit on the heat capacity of the neutron star core. *Phys. Rev. C*, 95:025806, Feb 2017. .

Curt Cutler and Eanna E. Flanagan. Gravitational waves from merging compact binaries: How accurately can one extract the binary's parameters from the inspiral wave form? *Phys. Rev. D*, 49:2658–2697, 1994. .

Gergely Dálya, Peter Raffai, and Bence Bécsy. Bayesian Reconstruction of Gravitational-wave Signals from Binary Black Holes with Nonzero Eccentricities. *Class. Quant. Grav.*, 38(6):065002, 2021. .

J. E. Davis. Event Pileup in Charge-coupled Devices. *ApJ*, 562:575–582, November 2001. .

A. De Luca. Central compact objects in supernova remnants. In *Journal of Physics Conference Series*, volume 932 of *Journal of Physics Conference Series*, page 012006, December 2017. .

N. Degenaar, Z. Medin, A. Cumming, R. Wijnands, M. T. Wolff, E. M. Cackett, J. M. Miller, P. G. Jonker, J. Homan, and E. F. Brown. Probing the Crust of the Neutron Star in EXO 0748-676. *ApJ*, 791:47, Aug 2014. .

N. Degenaar, L. S. Ootes, D. Page, R. Wijnands, A. S. Parikh, J. Homan, E. M. Cackett, J. M. Miller, D. Altamirano, and M. Linares. Crust cooling of the neutron star in Aql X-1: Different depth and magnitude of shallow heating during similar accretion outbursts. *arXiv e-prints*, art. arXiv:1907.07205, Jul 2019.

Alex Deibel, Andrew Cumming, Edward F. Brown, and Dany Page. A Strong Shallow Heat Source in the Accreting Neutron Star MAXI J0556-332. *ApJ*, 809:L31, Aug 2015. .

Walter Del Pozzo, Tjonnie G. F. Li, Michalis Agathos, Chris Van Den Broeck, and Salvatore Vitale. Demonstrating the feasibility of probing the neutron star equation of state with second-generation gravitational wave detectors. *Phys. Rev. Lett.*, 111(7): 071101, 2013. .

T. Dietrich, M. Ujevic, W. Tichy, S. Bernuzzi, and B. Brügmann. Gravitational waves and mass ejecta from binary neutron star mergers: Effect of the mass ratio. *Phys. Rev. D*, 95(2):024029, January 2017. .

Tim Dietrich, Sebastiano Bernuzzi, and Wolfgang Tichy. Closed-form tidal approximants for binary neutron star gravitational waveforms constructed from high-resolution numerical relativity simulations. *Phys. Rev. D*, 96(12):121501, 2017. .

Tim Dietrich, Sebastian Khan, Reetika Dudi, Shasvath J. Kapadia, Prayush Kumar, Alessandro Nagar, Frank Ohme, Francesco Pannarale, Anuradha Samajdar, Sebastiano Bernuzzi, Gregorio Carullo, Walter Del Pozzo, Maria Haney, Charalampos Markakis, Michael Pürrer, Gunnar Riemenschneider, Yoshinta Eka Setyawati, Ka Wa Tsang, and Chris Van Den Broeck. Matter imprints in waveform models for neutron star binaries: Tidal and self-spin effects. *Phys. Rev. D*, 99(2): 024029, January 2019. .

Tim Dietrich, Anuradha Samajdar, Sebastian Khan, Nathan K. Johnson-McDaniel, Reetika Dudi, and Wolfgang Tichy. Improving the NRTidal model for binary neutron star systems. *Phys. Rev. D*, 100(4):044003, 2019a. .

Tim Dietrich, Michael W. Coughlin, Peter T. H. Pang, Mattia Bulla, Jack Heinzel, Lina Issa, Ingo Tews, and Sarah Antier. Multimessenger constraints on the neutron-star equation of state and the Hubble constant. *Science*, 370(6523):1450–1453, 2020. .

Tim Dietrich, Tanja Hinderer, and Anuradha Samajdar. Interpreting Binary Neutron Star Mergers: Describing the Binary Neutron Star Dynamics, Modelling Gravitational Waveforms, and Analyzing Detections. *Gen. Rel. Grav.*, 53(3):27, 2021. .

Tim Dietrich et al. Matter imprints in waveform models for neutron star binaries: Tidal and self-spin effects. *Phys. Rev.*, D99(2):024029, 2019b. .

V. Doroshenko, V. Suleimanov, and A. Santangelo. CXOU J160103.1-513353: another central compact object with a carbon atmosphere? *A&A*, 618:A76, October 2018. .

F. Douchin and P. Haensel. A unified equation of state of dense matter and neutron star structure. *A&A*, 380:151–167, December 2001. .

Christian Drischler, Sophia Han, James M. Lattimer, Madappa Prakash, Sanjay Reddy, and Tianqi Zhao. Limiting masses and radii of neutron stars and their implications. *Phys. Rev. C*, 103(4):045808, 2021. .

Reetika Dudi, Francesco Pannarale, Tim Dietrich, Mark Hannam, Sebastiano Bernuzzi, Frank Ohme, and Bernd Brügmann. Relevance of tidal effects and post-merger dynamics for binary neutron star parameter estimation. *Phys. Rev. D*, 98(8):084061, 2018. .

W. E. East, V. Paschalidis, and F. Pretorius. Equation of state effects and one-arm spiral instability in hypermassive neutron stars formed in eccentric neutron star mergers. *Classical and Quantum Gravity*, 33(24):244004, December 2016. .

Paul J. Easter, Sudarshan Ghonge, Paul D. Lasky, Andrew R. Casey, James A. Clark, Francisco Hernandez Vivanco, and Katerina Chatziioannou. Detection and parameter estimation of binary neutron star merger remnants. *Phys. Rev. D*, 102(4): 043011, 2020. .

K. G. Elshamouty, C. O. Heinke, G. R. Sivakoff, W. C. G. Ho, P. S. Shternin, D. G. Yakovlev, D. J. Patnaude, and L. David. Measuring the Cooling of the Neutron Star in Cassiopeia A with all Chandra X-Ray Observatory Detectors. *ApJ*, 777:22, November 2013. .

Reed Essick, Ingo Tews, Philippe Landry, Sanjay Reddy, and Daniel E. Holz. Direct Astrophysical Tests of Chiral Effective Field Theory at Supranuclear Densities. *Phys. Rev. C*, 102(5):055803, 2020. .

Reed Essick, Philippe Landry, Achim Schwenk, and Ingo Tews. A Detailed Examination of Astrophysical Constraints on the Symmetry Energy and the Neutron Skin of ^{208}Pb with Minimal Modeling Assumptions. 7 2021a.

Reed Essick, Ingo Tews, Philippe Landry, and Achim Schwenk. Astrophysical Constraints on the Symmetry Energy and the Neutron Skin of ^{208}Pb with Minimal Modeling Assumptions. 2 2021b.

Marc Favata. Systematic parameter errors in inspiraling neutron star binaries. *Phys. Rev. Lett.*, 112:101101, 2014. .

Robert A. Fesen, Molly C. Hammell, Jon Morse, Roger A. Chevalier, Kazimierz J. Borkowski, Michael A. Dopita, Christopher L. Gerardy, Stephen S. Lawrence, John C. Raymond, and Sidney van den Bergh. The Expansion Asymmetry and Age of the Cassiopeia A Supernova Remnant. *ApJ*, 645:283–292, July 2006. .

E. Fonseca et al. Refined Mass and Geometric Measurements of the High-mass PSR J0740+6620. *Astrophys. J. Lett.*, 915(1):L12, 2021. .

F. Foucart, R. Haas, M. D. Duez, E. O'Connor, C. D. Ott, L. Roberts, L. E. Kidder, J. Lippuner, H. P. Pfeiffer, and M. A. Scheel. Low mass binary neutron star mergers: Gravitational waves and neutrino emission. *Phys. Rev. D*, 93(4):044019, February 2016. .

Rossella Gamba, Matteo Breschi, Sebastiano Bernuzzi, Michalis Agathos, and Alessandro Nagar. Waveform systematics in the gravitational-wave inference of tidal parameters and equation of state from binary neutron star signals. *Phys. Rev. D*, 103(12):124015, 2021. .

Sudarshan Ghonge, Katerina Chatziioannou, James A. Clark, Tyson Littenberg, Margaret Millhouse, Laura Cadonati, and Neil Cornish. Reconstructing gravitational wave signals from binary black hole mergers with minimal assumptions. 3 2020.

E. V. Gotthelf, R. Perna, and J. P. Halpern. Modeling the Surface X-ray Emission and Viewing Geometry of PSR J0821-4300 in Puppis A. *ApJ*, 724:1316–1324, December 2010. .

E. V. Gotthelf, J. P. Halpern, and J. Alford. The Spin-down of PSR J0821-4300 and PSR J1210-5226: Confirmation of Central Compact Objects as Anti-magnetars. *ApJ*, 765: 58, March 2013. .

Gravitational Wave Open Science Center (GWOSC). https://www.gw-openscience.org/.

Peter J. Green. Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*, 82(4):711–732, 12 1995. ISSN 0006-3444. . URL https://dx.doi.org/10.1093/biomet/82.4.711.

E. H. Gudmundsson, C. J. Pethick, and R. I. Epstein. Structure of neutron star envelopes. *ApJ*, 272:286–300, September 1983. .

M. E. Gusakov, A. D. Kaminker, D. G. Yakovlev, and O. Y. Gnedin. Enhanced cooling of neutron stars via Cooper-pairing neutrino emission. *A&A*, 423:1063–1071, September 2004. .

P. Haensel and J. L. Zdunik. Non-equilibrium processes in the crust of an accreting neutron star. *A&A*, 227:431–436, January 1990.

P. Haensel and J. L. Zdunik. Nuclear composition and heating in accreting neutron-star crusts. *A&A*, 404:L33–L36, June 2003. .

J. P. Halpern and E. V. Gotthelf. Spin-Down Measurement of PSR J1852+0040 in Kesteven 79: Central Compact Objects as Anti-Magnetars. *ApJ*, 709:436–446, January 2010. .

K. Hamaguchi, N. Nagata, K. Yanagi, and J. Zheng. Limit on the axion decay constant from the cooling neutron star in Cassiopeia A. *Phys. Rev. D*, 98(10):103015, November 2018. .

Peter Hammond, Ian Hawke, and Nils Andersson. Thermal aspects of neutron star mergers. 8 2021.

Sophia Han and Madappa Prakash. On the Minimum Radius of Very Massive Neutron Stars. *Astrophys. J.*, 899(2):164, 2020. .

Sophia Han and Andrew W. Steiner. Tidal deformability with sharp phase transitions in (binary) neutron stars. *Phys. Rev. D*, 99(8):083014, 2019. .

Carl-Johan Haster, Katerina Chatziioannou, Andreas Bauswein, and James Alexander Clark. Inference of the neutron star equation of state from cosmological distances. *Phys. Rev. Lett.*, 125(26):261101, 2020. .

C. O. Heinke and W. C. G. Ho. Direct Observation of the Cooling of the Cassiopeia A Neutron Star. *ApJ*, 719:L167–L171, August 2010. .

Matthias Hempel and Jürgen Schaffner-Bielich. A statistical model for a complete supernova equation of state. *Nuclear Physics A*, 837(3-4):210–254, June 2010. .

S. Hild et al. Sensitivity studies for third-generation gravitational wave observatories. *Classical and Quantum Gravity*, 28(9):094013, May 2011. .

Tanja Hinderer, Benjamin D. Lackey, Ryan N. Lang, and Jocelyn S. Read. Tidal deformability of neutron stars with realistic equations of state and their gravitational wave signatures in binary inspiral. *Phys. Rev. D*, 81(12):123016, June 2010. .

W. C. G. Ho. Central compact objects and their magnetic fields. In J. van Leeuwen, editor, *Neutron Stars and Pulsars: Challenges and Opportunities after 80 years*, volume 291 of *IAU Symposium*, pages 101–106, March 2013. .

Wynn C. G. Ho and Craig O. Heinke. A neutron star with a carbon atmosphere in the Cassiopeia A supernova remnant. *Nature*, 462:71–73, November 2009. .

Wynn C. G. Ho, Khaled G. Elshamouty, Craig O. Heinke, and Alexander Y. Potekhin. Tests of the nuclear equation of state and superfluid and superconducting gaps using the Cassiopeia A neutron star. *Phys. Rev. C*, 91:015806, January 2015. .

Wynn C. G. Ho, M. J. P. Wijngaarden, Philip Chang, Craig O. Heinke, Dany Page, Mikhail Beznogov, and Daniel J. Patnaude. Cooling of the Cassiopeia A neutron star and the effect of diffusive nuclear burning. In *Xiamen-CUSTIPEN Workshop on the Equation of State of Dense Neutron-Rich Matter in the Era of Gravitational Wave Astronomy*, volume 2127 of *American Institute of Physics Conference Series*, page 020007, July 2019. .

Wynn C. G. Ho, M. J. P. Wijngaarden, Nils Andersson, Thomas M. Tauris, and F. Haberl. Early neutron star evolution in high-mass X-ray binaries. *MNRAS*, 494(1): 44–49, May 2020. .

Wynn C. G. Ho, Yue Zhao, Craig O. Heinke, D. L. Kaplan, Peter S. Shternin, and M. J. P. Wijngaarden. X-ray bounds on cooling, composition, and magnetic field of the Cassiopeia A neutron star and young central compact objects. *MNRAS*, 506(4): 5015–5029, October 2021. .

Jeroen Homan, Joel K. Fridriksson, Rudy Wijnands, Edward M. Cackett, Nathalie Degenaar, Manuel Linares, Dacheng Lin, and Ronald A. Remillard. A Strongly Heated Neutron Star in the Transient Z Source MAXI J0556-332. *ApJ*, 795(2):131, Nov 2014. .

K. Hotokezaka, K. Kyutoku, H. Okawa, M. Shibata, and K. Kiuchi. Binary neutron star mergers: Dependence on the nuclear equation of state. *Phys. Rev. D*, 83(12):124008, June 2011. .

K. Hotokezaka, K. Kiuchi, K. Kyutoku, H. Okawa, Y.-i. Sekiguchi, M. Shibata, and K. Taniguchi. Mass ejection from the merger of binary neutron stars. *Phys. Rev. D*, 87 (2):024001, January 2013. .

Kenta Hotokezaka, Kenta Kiuchi, Koutarou Kyutoku, Takayuki Muranushi, Yu-ichiro Sekiguchi, Masaru Shibata, and Keisuke Taniguchi. Remnant massive neutron stars of binary neutron star mergers: Evolution process and gravitational waveform. *Phys. Rev. D*, 88:044026, 2013. .

R. Hunt. A fluid dynamical study of the accretion process. *MNRAS*, 154:141, January 1971. .

Sascha Husa, Sebastian Khan, Mark Hannam, Michael Pürrer, Frank Ohme, Xisco Jiménez Forteza, and Alejandro Bohé. Frequency-domain gravitational waves from nonprecessing black-hole binaries. I. New numerical waveforms and anatomy of the signal. *Phys. Rev. D*, 93(4):044006, 2016. .

A. F. Illarionov and R. A. Sunyaev. Why the Number of Galactic X-ray Stars Is so Small? *A&A*, 39:185, February 1975.

Jonah B. Kanner, Tyson B. Littenberg, Neil Cornish, Meg Millhouse, Enia Xhakaj, Francesco Salemi, Marco Drago, Gabriele Vedovato, and Sergey Klimenko. Leveraging waveform complexity for confident detection of gravitational waves. *Phys. Rev. D*, 93:022002, Jan 2016. . URL https://link.aps.org/doi/10.1103/PhysRevD.93.022002.

W. Kastaun and F. Galeazzi. Properties of hypermassive neutron stars formed in mergers of spinning binaries. *Phys. Rev. D*, 91(6):064027, March 2015. .

Wolfgang Kastaun and Frank Ohme. Numerical inside view of hypermassive remnant models for GW170817. *Phys. Rev. D*, 104(2):023001, 2021. .

Sebastian Khan, Sascha Husa, Mark Hannam, Frank Ohme, Michael Pürrer, Xisco Jiménez Forteza, and Alejandro Bohé. Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era. *Phys. Rev. D*, 93(4):044007, 2016. .

D. Klochkov, V. Suleimanov, M. Sasaki, and A. Santangelo. Study of a new central compact object: The neutron star in the supernova remnant G15.9+0.2. *A&A*, 592: L12, August 2016. .

M. R. Krumholz, C. F. McKee, and R. I. Klein. Bondi Accretion in the Presence of Vorticity. *ApJ*, 618:757–768, January 2005. .

Nina Kunert, Peter T. H. Pang, Ingo Tews, Michael W. Coughlin, and Tim Dietrich. Quantifying modelling uncertainties when combining multiple gravitational-wave detections from binary neutron star sources. 10 2021.

Philippe Landry, Reed Essick, and Katerina Chatziioannou. Nonparametric constraints on neutron star matter with existing and upcoming gravitational wave and pulsar observations. *Phys. Rev. D*, 101(12):123007, 2020. .

Nicolas Lartillot and Hervé Philippe. Computing Bayes Factors Using Thermodynamic Integration. *Systematic Biology*, 55(2):195–207, 04 2006. ISSN 1063-5157. . URL https://doi.org/10.1080/10635150500433722.

J. M. Lattimer and M. Prakash. Neutron Star Structure and the Equation of State. *ApJ*, 550:426–442, March 2001. .

James M. Lattimer and Madappa Prakash. The Equation of State of Hot, Dense Matter and Neutron Stars. *Phys. Rept.*, 621:127–164, 2016. .

Isaac Legred, Katerina Chatziioannou, Reed Essick, Sophia Han, and Philippe Landry. Impact of the PSR J0740+6620 radius constraint on the properties of high-density matter. *Phys. Rev. D*, 104(6):063003, 2021. .

Luis Lehner, Steven L. Liebling, Carlos Palenzuela, O. L. Caballero, Evan O'Connor, Matthew Anderson, and David Neilsen. Unequal mass binary neutron star mergers and multimessenger signals. *Class. Quant. Grav.*, 33(18):184002, 2016. .

L. B. Leinson. Axion mass limit from observations of the neutron star in Cassiopeia A. *J. Cosmology Astropart. Phys.*, 8:031, August 2014. .

Steven L. Liebling, Carlos Palenzuela, and Luis Lehner. Effects of High Density Phase Transitions on Neutron Star Dynamics. *Class. Quant. Grav.*, 38(11):115007, 2021. .

LIGO Scientific Collaboration and Virgo Collaboration. BayesWave, https://git.ligo.org/lscsoft/bayeswave, 2018. URL https://git.ligo.org/lscsoft/bayeswave.

Tyson B. Littenberg and Neil J. Cornish. Bayesian inference for spectral estimation of gravitational wave detector noise. *Phys. Rev.*, D91(8):084034, 2015. .

Tyson B. Littenberg, Jonah B. Kanner, Neil J. Cornish, and Margaret Millhouse. Enabling high confidence detections of gravitational-wave bursts. *Phys. Rev. D*, 94: 044050, Aug 2016. . URL https://link.aps.org/doi/10.1103/PhysRevD.94.044050.

Helei Liu, Yasuhide Matsuo, Masa-aki Hashimoto, Tsuneo Noda, and Masayuki Y. Fujimoto. Quiescent Light Curve of Accreting Neutron Star MAXI J0556-332. *Journal of the Physical Society of Japan*, 86:123901, Dec 2017. .

L. A. Lopez and R. A. Fesen. The Morphologies and Kinematics of Supernova Remnants. *Space Sci. Rev.*, 214:44, February 2018. .

Ben Margalit and Brian D. Metzger. Constraining the Maximum Mass of Neutron Stars from Multi-messenger Observations of GW170817. *ApJ*, 850(2):L19, December 2017. .

Ben Margalit and Brian D. Metzger. The Multi-Messenger Matrix: the Future of Neutron Star Merger Constraints on the Nuclear Equation of State. *Astrophys. J. Lett.*, 880(1):L15, 2019. .

LIGO Scientific Collaboration, Virgo Collaboration. LALSuite. https://git.ligo.org/lscsoft/lalsuite, 2018. URL https://doi.org/10.7935/GT1W-FZ16.

Rachael L. Merritt, Edward M. Cackett, Edward F. Brown, Dany Page, Andrew Cumming, Nathalie Degenaar, Alex Deibel, Jeroen Homan, Jon M. Miller, and Rudy Wijnand s. The Thermal State of KS 1731-260 after 14.5 years in Quiescence. *ApJ*, 833:186, Dec 2016. .

M. C. Miller et al. PSR J0030+0451 Mass and Radius from *NICER* Data and Implications for the Properties of Neutron Star Matter. *Astrophys. J. Lett.*, 887(1):L24, 2019. .

M. C. Miller et al. The Radius of PSR J0740+6620 from NICER and XMM-Newton Data. 5 2021.

Margaret Millhouse, Neil J. Cornish, and Tyson Littenberg. Bayesian reconstruction of gravitational wave bursts using chirplets. *Phys. Rev.*, D97(10):104057, 2018. .

Elias R. Most, L. Jens Papenfort, Veronica Dexheimer, Matthias Hanauske, Stefan Schramm, Horst Stöcker, and Luciano Rezzolla. Signatures of quark-hadron phase transitions in general-relativistic neutron-star mergers. *Phys. Rev. Lett.*, 122(6): 061101, 2019. .

S. S. Murray, S. M. Ransom, M. Juda, U. Hwang, and S. S. Holt. Is the Compact Source at the Center of Cassiopeia A Pulsed? *ApJ*, 566:1039–1044, February 2002. .

R. Negreiros, S. Schramm, and F. Weber. Impact of rotation-driven particle repopulation on the thermal evolution of pulsars. *Physics Letters B*, 718:1176–1180, January 2013. .

William G. Newton, Mark Alexander Kaltenborn, Sarah Cantu, Shuxi Wang, Amber Stinson, and Jirina Rikovska Stone. Glassy quantum nuclear pasta in neutron star crusts. *arXiv e-prints*, art. arXiv:2104.11835, April 2021.

T. Noda, M.-a. Hashimoto, N. Yasutake, T. Maruyama, T. Tatsumi, and M. Fujimoto. Cooling of Compact Stars with Color Superconducting Phase in Quark-hadron Mixed Phase. *ApJ*, 765:1, March 2013. .

Roland Oechslin and H.-T. Janka. Gravitational waves from relativistic neutron star mergers with nonzero-temperature equations of state. *Phys. Rev. Lett.*, 99:121102, 2007. .

M. Oertel, M. Hempel, T. Klahn, and S. Typel. Equations of state for supernovae and compact stars. *Rev. Mod. Phys.*, 89(1):015007, 2017. .

L. S. Ootes, S. Vats, D. Page, R. Wijnands, A. S. Parikh, N. Degenaar, M. J. P. Wijngaarden, D. Altamirano, A. Bahramian, and E. M. Cackett. Continued cooling of the accretion-heated neutron star crust in the X-ray transient IGR J17480-2446 located in the globular cluster Terzan 5. *MNRAS*, 487(1):1447–1461, Jul 2019. .

Laura S. Ootes, Dany Page, Rudy Wijnands, and Nathalie Degenaar. Neutron star crust cooling in KS 1731-260: the influence of accretion outburst variability on the crustal temperature evolution. *MNRAS*, 461:4400–4405, Oct 2016. .

Laura S. Ootes, Rudy Wijnands, Dany Page, and Nathalie Degenaar. A cooling neutron star crust after recurrent outbursts: modelling the accretion outburst history of Aql X-1. *MNRAS*, 477(3):2900–2916, Jul 2018. .

F. Özel and P. Freire. Masses, Radii, and the Equation of State of Neutron Stars. *Ann. Rev. Astron. Astrophys.*, 54:401–440, 2016. .

D. Page. NSCool: Neutron star cooling code. Astrophysics Source Code Library, September 2016.

Dany Page and Sanjay Reddy. Forecasting Neutron Star Temperatures: Predictability and Variability. *Phys. Rev. Lett.*, 111:241102, Dec 2013. .

Dany Page, James M. Lattimer, Madappa Prakash, and Andrew W. Steiner. Minimal Cooling of Neutron Stars: A New Paradigm. *The Astrophysical Journal Supplement Series*, 155:623–650, December 2004. .

Dany Page, Madappa Prakash, James M. Lattimer, and Andrew W. Steiner. Rapid Cooling of the Neutron Star in Cassiopeia A Triggered by Neutron Superfluidity in Dense Matter. *Phys. Rev. Lett.*, 106:081101, February 2011. .

Peter T. H. Pang, Tim Dietrich, Ingo Tews, and Chris Van Den Broeck. Parameter estimation for strong phase transitions in supranuclear matter using gravitational-wave astronomy. *Phys. Rev. Res.*, 2(3):033514, 2020. .

Peter T. H. Pang, Ingo Tews, Michael W. Coughlin, Mattia Bulla, Chris Van Den Broeck, and Tim Dietrich. Nuclear-Physics Multi-Messenger Astrophysics

Constraints on the Neutron-Star Equation of State: Adding NICER's PSR J0740+6620 Measurement. 5 2021.

Chris Pankow et al. Mitigation of the instrumental noise transient in gravitational-wave data surrounding GW170817. *Phys. Rev.*, D98(8):084016, 2018. .

A. S. Parikh, R. Wijnands, L. S. Ootes, D. Page, N. Degenaar, A. Bahramian, E. F. Brown, E. M. Cackett, A. Cumming, C. Heinke, J. Homan, A. Rouco Escorial, and M. J. P. Wijngaarden. Consistent accretion-induced heating of the neutron-star crust in MXB 1659-29 during two different outbursts. *A&A*, 624:A84, April 2019. .

Aastha S. Parikh, Jeroen Homan, Rudy Wijnands, Laura Ootes, Dany Page, Diego Altamirano, Nathalie Degenaar, Edward F. Brown, Edward Cackett, Andrew Cumming, Alex Deibel, Joel K. Fridriksson, Dacheng Lin, Manuel Linares, and Jon M. Miller. Different Accretion Heating of the Neutron Star Crust during Multiple Outbursts in MAXI J0556-332. *ApJ*, 851:L28, Dec 2017. .

R. Perna, R. Narayan, G. Rybicki, L. Stella, and A. Treves. Bondi Accretion and the Problem of the Missing Isolated Neutron Stars. *ApJ*, 594:936–942, September 2003. .

P. P. Plucinsky, A. Bogdan, G. Germain, and H. L. Marshall. The evolution of the ACIS contamination layer over the 16-year mission of the Chandra X-ray Observatory. In *Space Telescopes and Instrumentation 2016: Ultraviolet to Gamma Ray*, volume 9905 of *Proc. SPIE*, page 990544, July 2016. .

P. P. Plucinsky, A. Bogdan, H. L. Marshall, and N. W. Tice. The complicated evolution of the ACIS contamination layer over the mission life of the Chandra X-ray Observatory. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 10699 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 106996B, July 2018. .

Eric Poisson. Compact body in a tidal environment: New types of relativistic Love numbers, and a post-Newtonian operational definition for tidally induced multipole moments. *Phys. Rev. D*, 103(6):064023, 2021. .

B. Posselt and G. G. Pavlov. Upper Limits on the Rapid Cooling of the Central Compact Object in Cas A. *ApJ*, 864:135, September 2018. .

B. Posselt, G. G. Pavlov, V. Suleimanov, and O. Kargaltsev. New Constraints on the Cooling of the Central Compact Object in Cas A. *ApJ*, 779:186, December 2013. .

A. Y. Potekhin and D. G. Yakovlev. Thermal structure and cooling of neutron stars with magnetized envelopes. *A&A*, 374:213–226, July 2001. .

A. Y. Potekhin, G. Chabrier, and D. G. Yakovlev. Internal temperatures and cooling of neutron stars with accreted envelopes. *A&A*, 323:415–428, July 1997.

A. Y. Potekhin, D. A. Baiko, P. Haensel, and D. G. Yakovlev. Transport properties of degenerate electrons in neutron star envelopes and white dwarf cores. *A&A*, 346: 345–353, June 1999.

A. Y. Potekhin, A. F. Fantina, N. Chamel, J. M. Pearson, and S. Goriely. Analytical representations of unified equations of state for neutron-star matter. *A&A*, 560:A48, December 2013. .

A. Y. Potekhin, G. Chabrier, and W. C. G. Ho. Opacities and spectra of hydrogen atmospheres of moderately magnetized neutron stars. *A&A*, 572:A69, December 2014. .

A. Yu Potekhin. Atmospheres and radiating surfaces of neutron stars. *Physics Uspekhi*, 57:735-770, August 2014. .

Alexander Y. Potekhin, Dmitry G. Yakovlev, Gilles Chabrier, and Oleg Y. Gnedin. Thermal Structure and Cooling of Superfluid Neutron Stars with Accreted Magnetized Envelopes. *ApJ*, 594:404–418, September 2003. .

Alexander Y. Potekhin, José A. Pons, and Dany Page. Neutron Stars—Cooling and Transport. *Space Sci. Rev.*, 191:239–291, October 2015. .

Aviral Prakash, David Radice, Domenico Logoteta, Albino Perego, Vsevolod Nedora, Ignazio Bombaci, Rahul Kashyap, Sebastiano Bernuzzi, and Andrea Endrizzi. Signatures of deconfined quark phases in binary neutron star mergers. 6 2021.

Geraint Pratten, Patricia Schmidt, and Natalie Williams. Impact of Dynamical Tides on the Reconstruction of the Neutron Star Equation of State. 9 2021.

P. Predehl, E. Costantini, G. Hasinger, and Y. Tanaka. XMM-Newton observation of the galactic centre - evidence against the X-ray reflection nebulae model? *Astronomische Nachrichten*, 324:73–76, 2003. .

M. Punturo et al. The third generation of gravitational wave observatories and their science reach. *Classical and Quantum Gravity*, 27(8):084007, April 2010. .

Michael Pürrer, Mark Hannam, and Frank Ohme. Can we measure individual black-hole spins from gravitational-wave observations? *Phys. Rev. D*, 93(8):084042, 2016. .

G. Raaijmakers, S. K. Greif, K. Hebeler, T. Hinderer, S. Nissanke, A. Schwenk, T. E. Riley, A. L. Watts, J. M. Lattimer, and W. C. G. Ho. Constraints on the dense matter equation of state and neutron star properties from NICER's mass-radius estimate of PSR J0740+6620 and multimessenger observations. 5 2021.

G. Raaijmakers et al. Constraining the dense matter equation of state with joint analysis of NICER and LIGO/Virgo measurements. *Astrophys. J. Lett.*, 893(1):L21, 2020. .

Etienne Racine. Analysis of spin precession in binary black hole systems including quadrupole-monopole interaction. *Phys. Rev.*, D78:044021, 2008. .

David Radice, Sebastiano Bernuzzi, Walter Del Pozzo, Luke F. Roberts, and Christian D. Ott. Probing Extreme-Density Matter with Gravitational Wave Observations of Binary Neutron Star Merger Remnants. *Astrophys. J. Lett.*, 842(2): L10, 2017. .

Carolyn A. Raithel, Vasileios Paschalidis, and Feryal Özel. Realistic Finite-Temperature Effects in Neutron Star Merger Simulations. 4 2021.

V. Raymond and W. M. Farr. Physically motivated exploration of the extrinsic parameter space in ground-based gravitational-wave astronomy. 2 2014.

Brendan T. Reed, F. J. Fattoyev, C. J. Horowitz, and J. Piekarewicz. Implications of PREX-2 on the Equation of State of Neutron-Rich Matter. *Phys. Rev. Lett.*, 126(17): 172503, 2021. .

J. E. Reed, J. J. Hester, A. C. Fabian, and P. F. Winkler. The Three-dimensional Structure of the Cassiopeia A Supernova Remnant. I. The Spherical Shell. *ApJ*, 440: 706, February 1995. .

Tania Regimbau et al. A Mock Data Challenge for the Einstein Gravitational-Wave Telescope. *Phys. Rev. D*, 86:122001, 2012. .

David Reitze et al. The US Program in Ground-Based Gravitational Wave Science: Contribution from the LIGO Laboratory. *Bull. Am. Astron. Soc.*, 51:141, 2019a.

David Reitze et al. Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO. *Bull. Am. Astron. Soc.*, 51:035, 2019b.

Thomas E. Riley et al. A *NICER* View of PSR J0030+0451: Millisecond Pulsar Parameter Estimation. *Astrophys. J. Lett.*, 887(1):L21, 2019. .

Thomas E. Riley et al. A NICER View of the Massive Pulsar PSR J0740+6620 Informed by Radio Timing and XMM-Newton Spectroscopy. 5 2021.

M. Ruffert and H. Th. Janka. Coalescing neutron stars - A step towards physical models. III. Improved numerics and different neutron star masses and spins. *A&A*, 380:544–577, December 2001. .

M. Ruffert, H. Th. Ruffert, and H. Th. Janka. Coalescing neutron stars - a step towards physical models. 3. Improved numerics and different neutron star masses and spins. *Astron. Astrophys.*, 380:544, 2001. .

Surabhi Sachdev, Tania Regimbau, and B. S. Sathyaprakash. Subtracting compact binary foreground sources to reveal primordial gravitational-wave backgrounds. *Phys. Rev. D*, 102(2):024051, 2020. .

Anuradha Samajdar and Tim Dietrich. Waveform systematics for binary neutron star gravitational wave signals: Effects of spin, precession, and the observation of electromagnetic counterparts. *Phys. Rev. D*, 100(2):024046, 2019. .

J. A. Sauls, N. Chamel, and M. A. Alpar. Superfluidity in Disordered Neutron Stars Crusts. *arXiv e-prints*, art. arXiv:2001.09959, January 2020.

Achim Schwenk, Bengt Friman, and Gerald E. Brown. Renormalization group approach to neutron matter: quasiparticle interactions, superfluid gaps and the equation of state. *Nuclear Phys. A*, 713:191–216, January 2003. .

A. Sedrakian. Rapid cooling of Cassiopeia A as a phase transition in dense QCD. *A&A*, 555:L10, July 2013. .

Stuart L. Shapiro and Saul A. Teukolsky. *Black holes, white dwarfs, and neutron stars : the physics of compact objects*. 1983.

M. Shibata and K. Taniguchi. Merger of binary neutron stars to a black hole: Disk mass, short gamma-ray bursts, and quasinormal mode ringing. *Phys. Rev. D*, 73: 064027, March 2006.

Masaru Shibata. Constraining Nuclear Equations of State Using Gravitational Waves from Hypermassive Neutron Stars. *Phys. Rev. Lett.*, 94(20):201101, May 2005. .

Masaru Shibata, Keisuke Taniguchi, and Kōji Uryū. Merger of binary neutron stars with realistic equations of state in full general relativity. *Phys. Rev. D*, 71(8):084021, April 2005. .

Masaru Shibata, Keisuke Taniguchi, and Koji Uryu. Merger of binary neutron stars with realistic equations of state in full general relativity. *Phys. Rev. D*, 71:084021, 2005. .

Peter S. Shternin, Dmitry G. Yakovlev, Craig O. Heinke, Wynn C. G. Ho, and Daniel J. Patnaude. Cooling neutron star in the Cassiopeia A supernova remnant: evidence for superfluidity in the core. *MNRAS*, 412:L108–L112, March 2011. .

Peter S. Shternin, Dmitry D. Ofengeim, Wynn C. G. Ho, Craig O. Heinke, M. J. P. Wijngaarden, and Daniel J. Patnaude. Model-independent constraints on superfluidity from the cooling neutron star in Cassiopeia A. *MNRAS*, 506(1): 709–726, September 2021. .

M. Soares-Santos, et al., Dark Energy Survey, and Dark Energy Camera GW-EM Collaboration. The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/Virgo GW170817. I. Discovery of the Optical Counterpart Using the Dark Energy Camera. *Astrophysical Journal, Letters*, 848(2):L16, October 2017. .

Theodoros Soultanis, Andreas Bauswein, and Nikolaos Stergioulas. Analytic model of the spectral properties of gravitational waves from neutron star merger remnants. 11 2021.

Nikolaos Stergioulas, Andreas Bauswein, Kimon Zagkouris, and Hans-Thomas Janka. Gravitational waves and non-axisymmetric oscillation modes in mergers of compact object binaries. *Monthly Notices of the Royal Astronomical Society*, 418(1):427–436, November 2011. .

V. F. Suleimanov, D. Klochkov, J. Poutanen, and K. Werner. Probing the possibility of hotspots on the central neutron star in HESS J1731-347. *A&A*, 600:A43, April 2017. .

K. Takami, L. Rezzolla, and L. Baiotti. Constraining the Equation of State of Neutron Stars from Binary Mergers. *Physical Review Letters*, 113(9):091104, August 2014. .

K. Takami, L. Rezzolla, and L. Baiotti. Spectral properties of the post-merger gravitational-wave signal from binary neutron stars. *Phys. Rev. D*, 91(6):064001, March 2015. .

G. Taranto, G. F. Burgio, and H.-J. Schulze. Cassiopeia A and direct Urca cooling. *MNRAS*, 456:1451–1458, February 2016. .

Cajo J. F. Ter Braak. A Markov Chain Monte Carlo version of the genetic algorithm Differential Evolution: easy Bayesian computing for real parameter spaces. *Statistics and Computing*, 16(3):239–249, September 2006. .

Andoni Torres-Rivas, Katerina Chatziioannou, Andreas Bauswein, and James Alexander Clark. Observing the post-merger signal of GW170817-like events with improved gravitational-wave detectors. *Phys. Rev. D*, 99(4):044014, 2019. .

Ka Wa Tsang, Michiel Rollier, Archisman Ghosh, Anuradha Samajdar, Michalis Agathos, Katerina Chatziioannou, Vitor Cardoso, Gaurav Khanna, and Chris Van Den Broeck. A morphology-independent data analysis method for detecting and characterizing gravitational wave echoes. *Phys. Rev. D*, 98(2):024023, 2018. .

Ka Wa Tsang, Tim Dietrich, and Chris Van Den Broeck. Modeling the postmerger gravitational wave signal and extracting binary properties from future binary neutron star detections. *Phys. Rev. D*, 100(4):044047, 2019. .

Ka Wa Tsang, Archisman Ghosh, Anuradha Samajdar, Katerina Chatziioannou, Simone Mastrogiovanni, Michalis Agathos, and Chris Van Den Broeck. A morphology-independent search for gravitational wave echoes in data from the first and second observing runs of Advanced LIGO and Advanced Virgo. *Phys. Rev. D*, 101(6):064012, 2020. .

J. Veitch et al. Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library. *Phys. Rev. D*, 91(4):042003, 2015. .

Daniele Viganò, Alberto Garcia-Garcia, José A. Pons, Clara Dehman, and Vanessa Graber. Magneto-thermal evolution of neutron stars with coupled Ohmic, Hall and ambipolar effects via accurate finite-volume simulations. *Computer Physics Communications*, 265:108001, August 2021. .

Jacco Vink. Supernova remnants: the X-ray perspective. *Astronomy and Astrophysics Review*, 20:49, December 2012. .

Leslie Wade, Jolien D.E. Creighton, Evan Ochsner, Benjamin D. Lackey, Benjamin F. Farr, et al. Systematic and statistical errors in a bayesian approach to the estimation of the neutron-star equation of state using advanced gravitational wave detectors. *Phys. Rev. D*, 89(10):103012, 2014. .

A. C. Waterhouse, N. Degenaar, R. Wijnands, E. F. Brown, J. M. Miller, D. Altamirano, and M. Linares. Constraining the properties of neutron star crusts with the transient low-mass X-ray binary Aql X-1. *MNRAS*, 456(4):4001–4014, Mar 2016. .

Lukas R. Weih, Matthias Hanauske, and Luciano Rezzolla. Postmerger Gravitational-Wave Signatures of Phase Transitions in Binary Mergers. *Phys. Rev. Lett.*, 124(17):171103, 2020. .

Rudy Wijnands, Nathalie Degenaar, and Dany Page. Cooling of Accretion-Heated Neutron Stars. *Journal of Astrophysics and Astronomy*, 38, September 2017. .

M. J. P. Wijngaarden, Wynn C. G. Ho, Philip Chang, Craig O. Heinke, Dany Page, Mikhail Beznogov, and Daniel J. Patnaude. Diffusive nuclear burning in cooling simulations and application to new temperature data of the Cassiopeia A neutron star. *MNRAS*, 484:974–988, March 2019. .

M. J. P. Wijngaarden, Wynn C. G. Ho, Philip Chang, Dany Page, Rudy Wijnands, Laura S. Ootes, Andrew Cumming, Nathalie Degenaar, and Mikhail Beznogov. The effect of diffusive nuclear burning in neutron star envelopes on cooling in accreting systems. *MNRAS*, 493(4):4936–4944, April 2020. .

Marcella Wijngaarden, Katerina Chatziioannou, Andreas Bauswein, and Neil Cornish et. al. Analyzing of the full gravitational wave signal from neutron star binary coalescences. *to be submitted to Phys. Rev. D.*, 2021.

J. Wilms, A. Allen, and R. McCray. On the Absorption of X-Rays in the Interstellar Medium. *ApJ*, 542:914–924, October 2000. .

Zhu-Ge Xing, Joan M. Centrella, and Stephen L. W. McMillan. Gravitational radiation from coalescing binary neutron stars. *Phys. Rev.*, D50:6247–6261, 1994. .

D. G. Yakovlev, A. D. Kaminker, and O. Y. Gnedin. ¡SUP¿1¡/SUP¿S¡SUB¿0¡/SUB¿ neutron pairing vs. observations of cooling neutron stars. *A&A*, 379:L5–L8, November 2001. .

D. G. Yakovlev, W. C. G. Ho, P. S. Shternin, C. O. Heinke, and A. Y. Potekhin. Cooling rates of neutron stars and the young neutron star in the Cassiopeia A supernova remnant. *MNRAS*, 411:1977–1988, March 2011. .

S.-H. Yang, C.-M. Pi, and X.-P. Zheng. Rapid Cooling of the Neutron Star in Cassiopeia A and r-mode Damping in the Core. *ApJ*, 735:L29, July 2011. .

Nai-Bo Zhang and Bao-An Li. Implications of the Mass $M = 2.17^{+0.11}_{-0.10}\mathrm{M}_{\odot}$ of PSR J0740+6620 on the Equation of State of Super-dense Neutron-rich Nuclear Matter. *Astrophys. J.*, 879(2):99, 2019. .

Tianqi Zhao and James M. Lattimer. Tidal Deformabilities and Neutron Star Mergers. *Phys. Rev. D*, 98(6):063020, 2018. .