Numerical simulations of accretion onto neutron stars

David Anthony Abarca

A thesis submitted for the degree of
Doctor of Philosophy

from the

Nicolaus Copernicus Astronomical Center
Polish Academy of Sciences

Supervisor: Prof. Włodek Kluźniak
March 2022

Contents

Abstract

This thesis is comprised of three papers concerning accretion related phenomena around neutron stars and an introduction describing general relativistic radiation magnetohydrodynamics (GRRMHD), numerical techniques for solving the GRRMHD equations on a computer, accretion disks, neutron stars, and some unpublished results.

The motivation for much of the work in this thesis concerns pulsating ultraluminous X-ray sources (PULXs). These are extra-bright ($L < 10^{39}$ erg s^{-1}) extra-galactic X-ray point sources observed outside the centers of galaxies which also show (~ 1 s) coherent pulsations. The pulsations indicate the presence of neutron stars so that the observed luminosities must be many times the Eddington limit ($L_{\rm Edd}$). Much of this thesis is dedicated to understanding the physical processes which could allow accreting neutron stars to produce such large luminosities using numerical simulations and to a small extend analytical calculations.

In Chapter 2, which consists of the publication Abarca and Kluźniak (2016), I extended the results in Wielgus et al. (2015) to include a first order, linear perturbation analysis to study oscillations about the equilibrium configuration of a radiation supported atmosphere. While not directly concerned with PULXs, this paper serves as an illustration of the necessity of including radiation in relativistic hydrodynamics to properly encompass the physics around accreting neutron stars and so it helps to set the stage for the next two chapters. The main result of the publication was that the lowest frequency eigenmode is consistent with the 300-600 Hz quasi periodic oscillations (QPOs) seen in several X-ray bursting low-mas X-ray binaries. However, when the full effects of radiation drag were included in the calculation, the oscillations were found to be over-damped.

The second paper, Abarca et al. (2018) included in Chapter 3, involves simulations of super-Eddington accretion onto a black hole and a non-magnetized neutron star, the later of which obtained a hard-surface implemented as a sticky reflective inner radial boundary condition. The simulations showed that gas collected on the surface of the neutron star and filled the domain with so much material that any radiation released by the accretion disk hitting the surface was trapped. The radiation which was able to escape was found to be nearly isotropic and around one Eddington luminosity which does not resemble a ULX in any sense. This paper made it possible to disentangle the effects of including a magnetic field and a hard-surface into an accreting NS simulation.

The last chapter concerns a letter (Abarca et al. 2021) in which I run a 2D axisymmetric GRRMHD simulation of super-Eddington accretion onto a neutron star with a 2×10^{10} G dipolar magnetic field. In order to handle the large magnetizations present in the magnetosphere I implement the method from Parfrey and Tchekhovskoy (2017) and adapt it to work with GRRMHD simulations. I also use a boundary condition which is meant to model gas hitting the surface, becoming shocked and releasing a fraction (in this case 0.75) of its kinetic energy as outflowing radiation. The disk formed in the simulation is truncated by the magnetic field and the flow is driven along field lines forming accretion columns. A large amount of radiation is released at the base of the column and this radiation becomes collimated by the outflowing gas so that when it reaches the observer it appears to be originating from a source many times brighter ($\sim 140\ L_{\rm Edd}$). The actual amount of released radiation is much smaller, showing that the system is able to beam the radiation to a large degree. This shows that weakly-magnetized accretion neutron stars could be considered to be candidates for PULXs.

Streszczenie

(Abstract in Polish)

Niniejsza praca składa się z trzech artykułów dotyczących zjawisk związanych z akrecją na gwiazdy neutronowe oraz wprowadzenia opisującego ogólnorelatywistyczną magnetohydrodynamikę promienistą (GRRMHD), metody numeryczne do rozwiązywania równań GRRMHD na komputerze, dyski akrecyjne, gwiazdy neutronowe i pewne niepublikowane wyniki.

Motywacją dla dużej części pracy w tej rozprawie są pulsujące ultrajasne źródła promieniowania rentgenowskiego (PULX-y). Są to bardzo jasne ($L < 10^{39}$ erg s^{-1}) pozagalaktyczne punktowe źródła promieniowania rentgenowskiego obserwowane poza centrami galaktyk, które wykazują również spójne pulsacje o okresie 1 s. Pulsacje te wskazują na obecność gwiazd neutronowych, tak więc obserwowane jasności muszą być wielokrotnie większe od granicy Eddingtona (L_{Edd}). Znaczna część tej rozprawy poświęcona jest zrozumieniu procesów fizycznych, które mogłyby pozwolić akreującym gwiazdom neutronowym na uzyskanie tak dużych jasności, przy wykorzystaniu symulacji numerycznych i w niewielkim stopniu obliczeń analitycznych.

W rozdziale 2, który składa się z publikacji Abarca and Kluźniak (2016), rozszerzyłem wyniki z Wielgus et al. (2015) o analizę perturbacji liniowych pierwszego rzędu, aby zbadać oscylacje wokół konfiguracji równowagowej atmosfery podtrzymywanej przez promieniowanie. Chociaż praca ta nie dotyczy bezpośrednio PULX-ów, służy jako ilustracja konieczności uwzględnienia promieniowania w hydrodynamice relatywistycznej, aby właściwie opisać fizykę akrecji na gwiazdę neutronową i w ten sposób pomaga przygotować grunt pod następne dwa rozdziały. Głównym wynikiem tej publikacji było stwierdzenie, że tryb własny o najniższej częstotliwości jest zgodny z oscylacjami kwazi-okresowymi (QPO-y) o częstotliwości 300-600 Hz, obserwowanymi w kilku małomasywnych rentgenowskich układach podwójnych z rozbłyskami rentgenowskimi. Jednakże, gdy w obliczeniach uwzględniono pełne efekty oporu promieniowania, okazało się, że oscylacje są przetłumione.

Druga praca, Abarca et al. (2018) zamieszczona w rozdziale 3, dotyczy symulacji super-Eddingtonowskiej akrecji na czarną dziurę i na nienamagnetyzowaną gwiazdę neutronową, przy czym twarda powierzchnia tej drugiej została zaimplementowana przez warunek brzegowy na wewnętrznym promieniu, który absorbuje moment pędu lecz odbija składową radialną. Symulacje wykazały, że na powierzchni gwiazdy neutronowej gromadzi się gaz i wypełnia obszar tak dużą ilością materiału, że promieniowanie uwalniane przez dysk akrecyjny uderzający w powierzchnię zostaje uwięzione. Promieniowanie, które zdołało się wydostać, okazało się być prawie izotropowe i o jasności około jednej jasności Eddingtona, co w żadnym wypadku nie przypomina ULX. Praca ta pozwoliła rozdzielić efekty włączenia pola magnetycznego i twardej powierzchni do symulacji akrecji na gwiazdy neutronowe.

Ostatni rozdział dotyczy pracy (Abarca et al. 2021), w której przeprowadziłem dwuwymiarową osiowo-symetryczną symulację GRRMHD super-Eddingtonowskiej akrecji na gwiazdę neutronową z dipolowym polem magnetycznym 2×10^{10} G. Aby poradzić sobie z dużymi magnetyzacjami występującymi w magnetosferze, zaimplementowałem metodę z Parfrey and Tchekhovskoy (2017) i przystosowałem ją do pracy z symulacjami GRRMHD. Użyłem także warunku brzegowego, który ma modelować gaz uderzający w powierzchnię, ulegający szokowi i uwalniający część (w tym przypadku 0.75) swojej energii kinetycznej jako wypływające promieniowanie. Dysk uformowany w symulacji jest obcięty przez pole magnetyczne, a przepływ jest skierowany wzdłuż linii pola, tworząc kolumny akrecyjne. U podstawy kolumny uwalniana jest duża ilość promieniowania, które zostaje skolimowane przez wypływający gaz, zatem gdy dociera do obserwatora, wydaje się pochodzić ze źródła wielokrotnie jaśniejszego ($\sim 140\ L_{\mathrm{Edd}}$) niż w rzeczywistości. To oznacza, że układ jest w stanie w znacznym stopniu skupiać wiązkę promieniowania. Wynika z tego, że słabo namagnetyzowane akrecyjne gwiazdy neutronowe mogą być uważane za kandydatów na PULX-y.

Acknowledgments

It goes without saying that I could not have produced this work alone. I may have written the words, but in writing them I drew from the knowledge, experience, encouragement, companionship, and inspiration of others. This thesis is the result of not only years of scientific work, but a multitude of experiences shared with the people around me, only of subset of whom I have the room to mention below, and so I can only hope that those not mentioned by name also realise their impact on my work and my life.

I begin by thanking my supervisor, Prof. Włodek Kluźniak, without whom this thesis would absolutely not be possible. I thank him for his wealth of scientific knowledge, keen eye for details, patience, understand, motivation, and odd jokes. I learned from him a great deal not only applicable to science, but to life as well.

I also thank Dr. Aleksander (Olek) Sądowski. Olek was the bridge that brought me across the Atlantic. Without him, I would not have had the opportunity or the bravery to come to Poland to begin my PhD studies. I see him as an advisor, a mentor, and a friend and I am immensely grateful for all that he taught me.

Next, I thank Prof. Kyle Parfrey. Our collaboration was absolutely essential during the last years of my PhD and without his knowledge and advice, I am honestly not sure if I would have been able to finish.

I must additionally thank all of my professors and colleagues who contributed to my education and scientific output, including, but not limited to Prof. Ramesh Narayan, Prof. Marek Abramowicz, Dr. Maciek Wielgus, Debora Lančová, Prof. Omer Blaes, Dr. Miljenko Cemeljic, and Dr. Deepika Bollimpalli.

I thank my friends at CAMK both old and new, Abbas, Ago, Andrés, Bhupe, Deepika, Dominik, Fatima, Filip, Gonzalo, JB, José, Kasia, Kasia, Maitrayee, Nicolas, Ruchi, Stefa, Sylvain, et al. You helped me adjust to life in Warsaw, kept me sane during the Pandemic, supported me, celebrated with me, and improved my PhD experience a thousand-fold.

Finally, I thank my family, whose support is immeasurable and for whom my gratitude cannot be expressed in words. I thank my parents, Manuel and Roberta, for raising me, motivating me, and fostering my curiosity for the world. I thank my step-parents, Shawna and Duane, for bringing more happiness to my family than I could have hoped for. I thank my sister, Anamaria, for putting up with me and being a positive force in the world. I thank my mother and father-in-law, Ewa and Tadeusz, for accepting me into their family and providing me a home away from home. I thank my dog, Batat,

for the unwavering loyalty and perpetual joy that only animals are capable of.

And above all, I thank my wife, *Moje Kochanie na zawsze,* Katarzyna Rusinek-Abarca, for her genuine spirit, endless compassion, remarkable tolerance, and enduring love. This PhD thesis is just one small thing out of many made possible by our shared joy and mutual understanding of and adoration for one another. *Dziękuję z całego serca.*

1 Introduction

The subject of this thesis is a collection of extremes. From the fields of physics and astrophysics are united the extremely hot, extremely bright, extremely heavy, extremely fast, and extremely magnetized. The equations describing the physical processes are so complicated that they must be solved numerically on a super-computer using hundreds to thousands of processors.

The astrophysical motivation is based on the phenomena of pulsating ultraluminous X-ray sources (PULXs), extra-galactic objects emitting copious amounts of intense X-rays alongside coherent pulsations. The most widely accepted explanation for these sources involves super-Eddington accretion onto neutron stars.

Neutron stars are highly compact, and understanding them requires the invocation of Einstein's theory of general relativity (GR). In almost all astrophysical situations, the emission of X-rays can be understood as a consequences of the accretion of plasma onto a compact object. Plasma flows are governed by the equations of magnetohydrodynamics (MHD). Furthermore, the amount of radiation produced is enough to have a dynamical effect on the accreting plasma, and so the evolution of radiation must be included along with plasma using radiation hydrodynamics (RHD). Stitching it all together we have the relevant physical laws used in this thesis, general relativistic radiation magnetohydrodynamics (GRRMHD).

In this introduction, I will describe the equations of GRRMHD (Sec. 1.1) and show how they are solved using advanced numerical methods (Sec. 1.2). I will then lay the astrophysical foundation by discussing accretion disks (Sec. 1.3), the primary laboratories of GRRMHD, and finally, I will finish the discussion by focusing on accreting neutron stars and their magnetic fields (Sec. 1.4).

1.1 General relativistic radiation magnetohydrodynamics

1.1.1 General relativistic magnetohydrodynamics (GRMHD)

The strong gravity around black holes and neutron stars requires a fully relativistic formulation of hydrodynamics. A modern and complete reference on the subject is provided in Rezzolla and Zanotti (2013) which contains a proper derivation of the following equations which describe the evolution of a perfect fluid starting with a relativistic version of the Boltzmann equation.

The evolution equations for a perfect fluid in general relativity are given by,

$$\nabla_\mu(\rho u^\mu) = 0, \tag{1.1}$$

$$\nabla_\mu T^{\mu\nu} = 0, \tag{1.2}$$

where ∇_μ is the covariant derivative. These equations describe conservation of mass (Eqn. 1.1) and conservation of energy-momentum (Eqn. 1.2). The mass conservation equation is comprised of the fluid rest mass density, ρ, and the fluid four-velocity, u^μ. The latter equation makes use of the stress-energy tensor of a perfect fluid which can be written in component form as

$$T^{\mu\nu} = (\rho + p + u_{\rm int})u^\mu u^\nu + p g^{\mu\nu}, \tag{1.3}$$

and which also includes the fluid internal energy $u_{\rm int}$ and pressure $p = (\gamma - 1)u_{\rm int}$, which is related to the internal energy by the adiabatic index γ. It is also possible to relate the pressure to the density through the ideal gas equation of state,

$$p = \frac{k_{\rm B}}{m_{\rm p}}\rho T, \tag{1.4}$$

where $k_{\rm B}$ is the Boltzmann constant, m_p is the proton mass, and T is the temperature which comes from assuming a Maxwell-Boltzmann distribution for the individual particle velocities.

A physical interpretation of the stress-energy tensor, $T^{\mu\nu}$, is the flux of the "μ" component of momentum in the "ν" direction. T^{tt} would then be the fluid energy density, T^{it} would be the fluid momentum density, and T^{ij} gives the fluid stress-tensor where the diagonal components T^{ii} correspond to the fluid pressure, all these quantities can be read off from Eqn. 1.3 ($i, j = 1, 2, 3$).

I have also introduced the metric, $g_{\mu\nu}$ whose components dictate how to measure distances in four dimensional space-time. In this thesis, I am only considering at most slowly-rotating objects, so in every instance it is sufficient to use the Schwarzschild metric whose components in spherical coordinates are given by,

$$\mathrm{d}s^2 = g_{\mu\nu}\mathrm{d}x^\mu \mathrm{d}x^\nu = -\left(1 - \frac{2GM}{c^2 r}\right) c\,\mathrm{d}t^2 + \left(1 - \frac{2GM}{c^2 r}\right)^{-1} \mathrm{d}r^2 + r^2(\mathrm{d}\theta^2 + \sin\theta\,\mathrm{d}\phi^2). \tag{1.5}$$

The Schwarzschild metric becomes singular at $r = 2GM/c^2 \equiv 2r_g$, where r_g is the gravitational or mass radius. This radius, the Schwarzschild radius $r_s = 2r_g$, gives the location of the event horizon, the surface from below which no information can escape. Keep in mind that this is only a coordinate singularity which can be remedied by a clever change of coordinates, while the true singularity at $r = 0$ is always present.

From this point on, as is conventional, I will switch to geometrical units where $c = G = 1$, which simplifies many of the equations. Thanks to Birkhoff's theorem, we can also use the Schwarzschild metric to describe the space-time outside of a non-rotating neutron star. In this case, we do not need to worry about singularities at all since the typical radii of neutron stars have $R \gtrsim 5r_g$.

Astrophysical fluids are mainly comprised of hydrogen gas, and at sufficiently high temperatures, $T \gtrsim 10^4$ K, the hydrogen atoms becomes ionized forming a plasma, a collection of protons surrounded by a sea of electrons which are free to flow and produce currents which generate magnetic fields, which in turn affect the motion of the plasma producing more currents and so on. This self-interaction causes the equation of motions for plasma to be highly nonlinear.

A relativistic formulation of electrodynamics can be found in Misner et al. (1973) and for relativistic magnetohydrodynamics (MHD) one can find an extensive reference in Anile (1989) but useful summaries are given in Komissarov (1999); Gammie et al. (2003).

The most basic way to formulate a theory of electromagnetism, is to take a charge and move it around and measure the force on the charge. If a test charge is in an electromagnetic field, then one will measure one force when the charge is stationary, the source of which is called the 'electric' field, $\mathbf{E}$, and one force which is only present when the charge is in motion necessitating another field called the 'magnetic' field, $\mathbf{B}$. After enough measurements, it might be possible to notice a pattern and to write down the Lorentz force as a function of these two fields, the particle velocity, and the charge, q,

$$\mathbf{F} = q(\mathbf{E} + \mathbf{u} \times \mathbf{B}). \tag{1.6}$$

In relativity, one finds that the usual electric and magnetic fields, $\mathbf{E}$ and $\mathbf{B}$ respectively, are rather two sides of the same coin, where one field can be transformed into another via Lorentz boost i.e., the fields will change when measured in a different frame. Thus, it is necessary to describe them as one object, an antisymmetric rank-two tensor $F^{\mu\nu}$, the electromagnetic or "Faraday" tensor. Using the same process, one might measure the six independent components of the tensor in some (locally inertial[1]) frame and write down the Lorentz force as a function of the Faraday tensor and particle 4-velocity,

$$m\hat{a}^{\mu} = q\hat{F}^{\mu\nu}\hat{u}_{\nu}, \tag{1.7}$$

where $\hat{a}^{\mu}$ is the four-acceleration of the particle, finding the components of $\hat{F}^{\mu\nu}$ to be,

$$\hat{F}^{\mu\nu} = \begin{pmatrix} 0 & \hat{E}_1 & \hat{E}_2 & \hat{E}_3 \\ -\hat{E}_1 & 0 & \hat{B}_3 & -\hat{B}_2 \\ -\hat{E}_2 & -\hat{B}_3 & 0 & \hat{B}_1 \\ -\hat{E}_3 & \hat{B}_2 & -\hat{B}_1 & 0 \end{pmatrix}, \tag{1.8}$$

where $\hat{\mathbf{E}} \equiv \hat{E}^i$ and $\hat{\mathbf{B}} \equiv \hat{B}^i$ are the usual electric and magnetic fields that can go into Eqn. 1.6. Of course Eqn. 1.7 is true in any frame, and one can recover a different set of $\mathbf{E}$ and $\mathbf{B}$ in another frame from the Faraday tensor as,

$$E^i = F^{0i}, \tag{1.9}$$

$$B^i = \frac{1}{2}\epsilon^{ijk}F_{jk}, \tag{1.10}$$

where ϵ^{ijk} is the Levi-Civita symbol and the metric signature is $(-,+,+,+)$. It is then possible to construct the stress-energy tensor for electromagnetism,

$$T^{\mu\nu}_{(\mathrm{EM})} = F^{\mu\alpha}F^{\nu}{}_{\alpha} - \frac{1}{4}g^{\mu\nu}F_{\alpha\beta}F^{\alpha\beta}, \tag{1.11}$$

which behaves as you might expect, i.e., in an inertial frame $\hat{T}^{00}_{(\mathrm{EM})} = (\hat{\mathbf{E}}^2 + \hat{\mathbf{B}}^2)/2$ gives the electromagnetic energy density, $\hat{T}^{0i}_{(\mathrm{EM})} = \hat{\mathbf{E}} \times \hat{\mathbf{B}}$ gives the Poynting flux and $\hat{T}^{ij}_{(\mathrm{EM})}$ gives the Maxwell stress tensor. When we take a covariant derivative we get

$$\nabla_{\nu}T^{\mu\nu}_{(\mathrm{EM})} = -F^{\mu\alpha}J_{\alpha}, \tag{1.12}$$

---

[1] A vector whose components are given by v^{μ} in the lab frame will have components in an inertial frame indicated by $\hat{v}^{\mu}$. In GR it is always possible to identify a locally inertial frame in which (to first order) $g^{\mu\nu} = \eta^{\mu\nu} \equiv \mathrm{diag}(-1, 1, 1, 1)$.

where J^μ is the four-current density, analogous to qu^μ which means that Eqn. 1.12 is just minus the Lorentz force on a fluid with current J^μ, which would evolve according to,

$$\nabla_\nu T^{\mu\nu} = F^{\mu\alpha} J_\alpha. \tag{1.13}$$

Then, if we simply add the two stress-energy tensors together while renaming the fluid contribution as $T^{\mu\nu}_{(\mathrm{Rey})}$ (Reynolds stress), we get the full conservation equations for GRMHD (Dixon 1978),

$$\nabla_\mu \left(T^{\mu\nu}_{(\mathrm{Rey})} + T^{\mu\nu}_{(\mathrm{EM})} \right) = 0. \tag{1.14}$$

The form of the equations I have just presented are not always very useful in practice, so I will apply a simplifying assumption to reformulate stress-energy tensor. I will assume that the fluid is sufficiently conductive so that I can invoke the ideal MHD approximation in which electric fields are generated purely from Faraday's law, $\mathbf{E} = -\mathbf{v} \times \mathbf{B}$, which, in relativistic language, implies that the Lorentz force in the fluid frame[2] vanishes giving,

$$\widehat{u}_\mu \widehat{F}^{\mu\nu} = 0 \implies u_\mu F^{\mu\nu} = 0. \tag{1.15}$$

We can then introduce a new four-vector,

$$b^\mu = \frac{1}{2} \epsilon^{\mu\nu\kappa\lambda} u_\nu F_{\lambda\kappa}, \tag{1.16}$$

using the Levi-Civita tensor, $\epsilon^{\mu\nu\kappa\lambda}$ (a fully covariant analog to the Levi-Civita symbol). The magnetic four-vector, b^μ, has the property it is orthogonal to the four-velocity, $u_\mu b^\mu = 0$. Eqn. 1.16 can be inverted to get,

$$F^{\mu\nu} = \epsilon^{\mu\nu\kappa\lambda} u_\kappa b_\lambda, \tag{1.17}$$

which then leads to a much simpler form of the stress-energy tensor (Komissarov 1999; Gammie et al. 2003),

$$T^{\mu\nu}_{(\mathrm{EM})} = b^2 u^\mu u^\nu + \frac{1}{2} b^2 g^{\mu\nu} - b^\mu b_\nu. \tag{1.18}$$

So far, I have shown how the plasma responds to the influence of an electromagnetic field, but to complete the theory, we need to know how the electromagnetic field evolves as well. To do this, one can then take the Hodge dual of $F^{\mu\nu}$,

$$*F_{\mu\nu} = \frac{1}{2} F^{\alpha\beta} \epsilon_{\alpha\beta\mu\nu}, \tag{1.19}$$

which has a much simpler form in terms of b^μ,

$$*F^{\mu\nu} = b^\mu u^\nu - b^\nu u^\mu. \tag{1.20}$$

This allows us to write the conservative form of Maxwell's equations,

$$\nabla_\mu (*F^{\mu\nu}) = 0, \tag{1.21}$$

$$\nabla_\mu F^{\mu\nu} = -J^\nu, \tag{1.22}$$

---

[2] Here, I have introduced a special inertial frame, the fluid frame (denoted by wide hats $\widehat{v}$) which has the additional property that in this frame, the fluid velocity vanishes $\widehat{u}^\mu = (1, 0, 0, 0)$.

in their simplified forms. It is possible to check that in the orthonormal fluid frame, the more familiar version of Maxwell's equations are recovered. The second equation is used to compute the four-current density, J^ν, although in ideal GRMHD, the current is not dynamically relevant.

The first equation is the most important for the evolution of the system and encodes the induction equation and the zero divergence condition of the magnetic field. If we bring back the magnetic 3-field following Komissarov (1999) we have,

$$b^t = B^i u^\mu g_{i\mu}, \tag{1.23}$$

$$b^i = \frac{B^i + b^t u^i}{u^t}, \tag{1.24}$$

where products with repeated instances of i correspond to sums over spatial components. Including the magnetic 3-field gives simple forms of the induction equation and zero-divergence conditions,

$$\partial_t(\sqrt{-g}B^i) = -\partial_j\left[\sqrt{-g}(b^j u^i - b^i u^j)\right], \tag{1.25}$$

$$\partial_i\left(\sqrt{-g}B^i\right) = 0, \tag{1.26}$$

where $g \equiv \det g_{\mu\nu}$. Then combining Eqns. 1.14 and 1.25 along with an equation of state for the gas pressure give a set of closed equations for the full GRMHD system.

1.1.2 General relativistic radiative magnetohydrodynamics (GRRMHD)

If we move to even higher temperatures, the thermal radiation of the ionized plasma begins to have a dynamical effect (as opposed to a purely thermodynamic effect, i.e., cooling the plasma) on the fluid motion. A full reference of radiation hydrodynamics is provided by Mihalas and Mihalas (1984), but I will be mostly paraphrasing Sądowski et al. (2013), who provides a streamlined exposition of how to include radiation into the GRMHD equations.

One usually begins their study of radiative transfer by consulting Rybicki and Lightman (1986) who start by introducing the most fundamental radiation quantity, the specific radiative intensity, I_ν which is a function of position $\mathbf{x}$, direction $\hat{\mathbf{r}}$, and frequency ν. There are a number of situations in which it is important to consider the intensity as a function of frequency but I do not discuss any of them in this thesis, therefore we make the 'grey' approximation and assume all radiation quantities are integrated over frequency, $I = \int I_\nu \mathrm{d}\nu$.

We can then proceed to take various moments of the intensity by integrating over all directions, this is equivalent to integrating over the unit sphere, $\int \mathrm{d}\Omega$. Again, we simplify our lives by first working in the orthonormal[3] fluid frame, where the various radiative quantities have simple definitions. The zeroth moment corresponds to the radiation energy density,

$$\widehat{E} = \int \widehat{I}\mathrm{d}\Omega. \tag{1.27}$$

We can take higher moments by considering integrals of the intensity along the three unit vector directions,

$$\widehat{F}^i = \int \widehat{I}N^i\mathrm{d}\Omega, \tag{1.28}$$

---

[3] The fluid frame can be constructed so that its basis vectors form an orthonormal tetrad.

where N^i can be thought of as $\mathbf{\hat{r}} \cdot \hat{x}^i$ where $\hat{x}^i$ are the unit vectors along each dimension of the considered space. This gives the three components of the radiative flux. The second moments form a matrix which is computed via,

$$\widehat{P}^{ij} = \int \widehat{I} N^i N^j \mathrm{d}\Omega, \tag{1.29}$$

which corresponds to the radiation pressure tensor.

We can now construct the full radiation stress-energy tensor which is constructed from the various moments we just computed,

$$\widehat{R}^{\mu\nu} = \begin{pmatrix} \widehat{E} & \widehat{F}^i \\ \widehat{F}^j & \widehat{P}^{ij} \end{pmatrix}. \tag{1.30}$$

We then need two transformations to get back to the lab frame quantities. Up until now, I have not described the transformations necessary to move between a coordinate basis and an orthonormal tetrad basis. If we consider a fluid which has four-velocity as measured by the orthonormal zero-angular momentum observer (ZAMO[4]) , $\tilde{u}^\mu$ (Bardeen et al. 1972), then we can move between the ZAMO frame and the fluid frame using a Lorentz boost defined by,

$$\Lambda^\mu{}_\nu(\tilde{u}) = \begin{pmatrix} \gamma & \gamma\tilde{v}^i \\ \gamma\tilde{v}^j & \delta^{ij} + \frac{\tilde{v}^i\tilde{v}^j(\gamma-1)}{\tilde{v}^2}, \end{pmatrix} \tag{1.31}$$

where $\tilde{v}^i = \tilde{u}^i/\tilde{u}^t$ and $\gamma = 1/\sqrt{1-\tilde{v}^2}$. Then, to move between the lab-frame and the orthonormal ZAMO frame, we can use the tetrads defined in Bardeen et al. (1972), $e^\mu{}_\nu, \tilde{e}^\mu{}_\nu$, so that a vector in the lab-frame x^μ can be transformed to the ZAMO frame and vice versa via,

$$\tilde{x}^\mu = x^\nu \tilde{e}^\mu{}_\nu, \tag{1.32}$$

$$x^\mu = \tilde{x}^\nu e^\mu{}_\nu. \tag{1.33}$$

The ZAMO frame is useful for many calculations because the metric is locally Minkowski, but the transformation is slightly cumbersome. Luckily, as mentioned previously, I am only concerned with the Schwarzschild geometry in this thesis, and so one can take, for example, the orthonormal tetrad used in Abramowicz et al. (1990),

$$\tilde{e}^t{}_t = \sqrt{1-2M/r}, \quad \tilde{e}^r{}_r = \frac{1}{\sqrt{1-2M/r}}, \quad \tilde{e}^\theta{}_\theta = r, \quad \tilde{e}^\phi{}_\phi = r\sin\theta, \tag{1.34}$$

and all other components equal to 0, with the inverse transformation $e^\mu{}_\nu$ given by the reciprocals. Finally we can perform the full transformation between the Boyer-Lindquist (essentially spherical coordinates) lab frame and the orthonormal fluid frame where each transformation gets one matrix for each index,

$$R^{\mu\nu} = e^\mu{}_\sigma e^\nu{}_\kappa \tilde{R}^{\sigma\kappa}, \tag{1.35}$$

$$\tilde{R}^{\mu\nu} = \Lambda^\mu{}_\sigma(\tilde{u})\Lambda^\nu{}_\kappa(\tilde{u})\widehat{R}^{\sigma\kappa}, \tag{1.36}$$

again, with the inverse transformations for the other direction.

Now that we have a prescription to calculate $R^{\mu\nu}$ we can write down the conservation law of GRRMHD,

$$\nabla_\mu(T^{\mu\nu} + R^{\mu\nu}) = 0. \tag{1.37}$$

---

[4] Quantities in the ZAMO frame are denoted by tildes $\tilde{v}^\mu$.

This is similar to what was done for GRMHD, however, in this case we don't have the Lorentz force and Maxwell's equations to close the system. Instead we need to split this equation into two,

$$\nabla_\mu T^{\mu\nu} = G^\nu, \tag{1.38}$$

$$\nabla_\mu R^{\mu\nu} = -G^\nu. \tag{1.39}$$

Now we can evolve gas and radiation separately as fluids, with their secular evolution provided by the homogeneous version of the above equations, and the coupling provided by the radiation four-force density, G^ν, provided in Mihalas and Mihalas (1984) as

$$G^\mu = \int (\chi_\nu I_\nu - \eta_\nu) N^i \mathrm{d}\nu \mathrm{d}\Omega, \tag{1.40}$$

where χ_ν, I_ν, and η_ν correspond to the frequency dependant opacity, intensity, and emissivity. This equation has a much simpler form in the fluid frame,

$$\widehat{G}^\mu = \begin{pmatrix} \kappa(\widehat{E} - 4\pi\widehat{B}) \\ \chi\widehat{F}^i \end{pmatrix}, \tag{1.41}$$

where $\widehat{B} = \sigma T^4/\pi$ is the frequency integrated Planck function for a black body with temperature, T, and σ is the Stefan-Boltzmann constant. The total opacity coefficient $\chi = \kappa + \kappa_{\mathrm{es}}$ is the sum of the frequency integrated absorption and electron scattering opacities. Depending on the state of the gas, there are several radiative processes which can be important, and different ways to compute the frequency integrated quantities, but the basic structure of $\widehat{G}^\nu$ is the same. An addition to $\widehat{G}^\nu$ is necessary when temperatures reach such levels that Comptonization becomes important. There are a number of ways to do this, but the general idea is compute a fitting function which approximates the energy exchange between gas and radiation as a function of their temperature difference with functions that drive that difference to zero in the fluid frame. As this is purely a thermodynamic process, it is only introduced by a factor $\widehat{G}^t_{\mathrm{Comp}}$. A more advanced method involves evolving the number density of photons, as well as the radiation energy density, so that the energy per photon becomes important. This is explained in detail in Sądowski and Narayan (2015), but I will not repeat the discussion here because I rely on the simpler method in this thesis. The total radiation four-force is then computed via $\widehat{G}^\mu = \widehat{G}^t_{\mathrm{Comp}}\widehat{u}^\mu + \widehat{G}^\mu_0$, where now $\widehat{G}^\mu_0$ is computed from Eqn. 1.41 and relabelled with a subscript. The new expression for $\widehat{G}^\mu$ is then transformed to the lab frame to get G^ν.

Now that we have an expression for G^ν, it is possible to solve Eqns. 1.38 and 1.39 along with the equation for mass conservation and the induction equation in the coordinate basis by expanding the covariant derivatives,

$$\partial_t(\sqrt{-g}T^t{}_\nu) + \partial_i(\sqrt{-g}T^i{}_\nu) = \sqrt{-g}T^\kappa{}_\lambda\Gamma^\lambda_{\nu\kappa} + \sqrt{-g}G_\nu, \tag{1.42}$$

$$\partial_t(\sqrt{-g}R^t{}_\nu) + \partial_i(\sqrt{-g}R^i{}_\nu) = \sqrt{-g}R^\kappa{}_\lambda\Gamma^\lambda_{\nu\kappa} - \sqrt{-g}G_\nu. \tag{1.43}$$

Here, for the first time in this thesis, I have used the Christoffel symbols, $\Gamma^\lambda_{\nu\kappa}$, coordinate dependent objects which are computed from derivatives of the metric.

The M_1 closure scheme

Eqn. 1.43 evolves $R^{t\nu}$ in time and the symmetry of the stress-energy tensor provides $R^{\mu t}$. What then of the remaining components, R^{ij}? For this, it is necessary to use a closure scheme, something to fill in the rest of the tensor when only the first row and column are known. A natural choice for GRRMHD is the M_1 closure scheme (Levermore 1984). The premise assumes that there exists a frame[5], not necessarily corresponding to the fluid or lab frames, in which the radiation pressure tensor, with radiation energy density $\bar{E}$, is isotropic,

$$\bar{P}^{ij} = \begin{pmatrix} \bar{E}/3 & 0 & 0 \\ 0 & \bar{E}/3 & 0 \\ 0 & 0 & \bar{E}/3 \end{pmatrix}, \tag{1.44}$$

and in which the radiation flux vanishes, $\bar{F}^i = 0$. If we define the four-velocity of such a frame as $u^\mu_{\rm R}$, where $\bar{u}^\mu_{\rm R} = (1,0,0,0)$, then we can write the radiation stress-energy tensor in the compact form given by,

$$\bar{R}^{\mu\nu} = \frac{4}{3}\bar{E}\bar{u}^\mu_{\rm R}\bar{u}^\nu_{\rm R} + \frac{1}{3}\bar{E}g^{\mu\nu}. \tag{1.45}$$

This is a covariant equation, and so it is true in any frame. If we transform the vector and tensor quantities to the lab frame then we simply have,

$$R^{\mu\nu} = \frac{4}{3}\bar{E}u^\mu_{\rm R}u^\nu_{\rm R} + \frac{1}{3}\bar{E}g^{\mu\nu}. \tag{1.46}$$

If we then assume that we know $R^{t\mu}$ from Eqn. 1.39, it is possible to solve for $\bar{E}$ and $u^\mu_{\rm R}$, and so the rest of the tensor can be constructed. M_1 allows the radiation to flow along geodesics in the optically thin regime, and to diffuse through the gas in the optically thick regime.

The M_1 closure scheme causes the gas and radiation to be nearly symmetrical in their formulation. The microscopic behavior of gas is described by the Boltzmann equation, and the radiative transfer equation describes the evolution of radiation. Then by taking moments up to second order of both equations, we get the conservation Equations 1.38 and 1.39. M_1 introduces the radiation velocity which obeys the same rules in GR as the fluid velocity, even though photons travel along null geodesics.

Since, in this formulation, radiation behaves like a fluid, it is subject to drawbacks in situations where fluid behavior would not adequately describe the radiation field. The most obvious is when two beams of radiation from different sources intersect. The evolution of the intensity should be independent along different angles, and so the two beams are expected to pass through each-other without interacting. In M1, however they collide and interact. It is possible to alleviate this problem to some degree by trying to solve the problem along several directions using a variable Eddington tensor scheme, but this introduces a large degree of complexity to the problem.

1.1.3 Levitating atmospheres

The equations of motion for GRRMHD are highly non-linear and require the use of a computer to solve in nearly all cases, even when using the simplest of closure schemes. However there is one situation where even an analytical treatment of radiation in GR provides an interesting result. A peculiar consequence of radiative transfer in the context of general relativity is the existence of the Eddington capture sphere. I will, for a moment, revert to the Newtonian regime to better illustrate the concept. Consider an isotropically radiating sphere or point source. The luminosity of such a source is simply

---

[5] Again, I introduce yet another frame, the radiation rest frame, this time denoted with bars $\bar{v}$.

given by the integral of radiative flux over a sphere centered on the source,

$$L = 4\pi r^2 F. \tag{1.47}$$

It can be easily shown from the definition of intensity that the flux should decrease proportionally to r^{-2} making luminosity constant with radius. If one measures the luminosity at one radius, then one knows the luminosity at all radii. If the luminosity of the source is known a priori, then it is simple to compute the flux as a function of radius. Let us now consider the force a plasma parcel feels due to radiation pressure around an isotropically radiating sphere. The radiation force, $F_{\rm R}$, acts on electrons,

$$F_{\rm R} = \frac{\sigma_{\rm ES}}{c} F = \frac{\sigma_{\rm ES}}{c} \frac{L}{4\pi r^2}, \tag{1.48}$$

where $\sigma_{\rm ES}$ is the electron scattering cross section and I have temporarily restored the c. If the source has mass, M, then we can also figure out the gravitational force on the plasma. The gravitational force on electrons is negligible compared to protons, and so we get,

$$F_g = -\frac{GMm_p}{r^2}. \tag{1.49}$$

Electrons and protons are tied together due to Coulomb interactions and so the total radiation force would be the sum of the radiation pressure on each electron, and the total gravitational force would be the sum of the forces on each proton. If we consider ionized hydrogen, then the number densities of protons and neutrons are equal, which we can write as n. The total force density,

$$\mathcal{F} = n(F_{\rm R} + F_g) = n\left(\frac{\sigma_{\rm ES} L}{4\pi c} - GMm_p\right)\frac{1}{r^2}. \tag{1.50}$$

This shows that if the two forces are balanced at one radius, then they should be balanced at all radii, and the luminosity must necessarily be

$$L_{\rm Edd} \equiv \frac{4\pi GMm_p c}{\sigma_{\rm ES}}, \tag{1.51}$$

hereby referred to as the Eddington luminosity. The Eddington luminosity is an important unit of measurement in astrophysics and represents the maximum luminosity a spherical source (of pure ionized hydrogen) should have before it is blown apart by its own radiation pressure. One should note that the Eddington luminosity is purely a function of mass and physical constants.

Once general relativity is introduced, the picture changes. In the Schwarzschild metric, due to gravitational redshift, the luminosity of a source decreases with distance,

$$L(r) = \frac{L_\infty}{1 - 2M/r}, \tag{1.52}$$

where L_∞ is the luminosity observed at infinity. It can then be shown that there is only one surface where the radiation and gravitational forces balance (Abramowicz et al. 1990),

$$r_{\rm ECS} = \frac{2M}{1 - (L_\infty/L_{\rm Edd})^2}, \tag{1.53}$$

referred to as the Eddington Capture Sphere (ECS; Stahl et al. 2012; Wielgus et al. 2012). Below this surface, radiation pressure dominates over gravity, and above, gravity dominates causing $r_{\rm ECS}$ to be surface of stable equilibrium. It has also been shown that plasma traveling around Eddington or even near-Eddington sources should have its angular momentum removed due to the Poynting-Robertson radiation drag (Bini et al. 2009; Sok Oh et al. 2011). A basic explanation is that a particle moving transversely through uniform radiation field does not actually see a uniform radiation field in its frame. Instead the radiation should be relativistically beamed so that there is a component of the flux along the direction of motion, slowing the particle down. Wielgus et al. (2015) pointed out that if enough particles are collected, then an optically thin, hydrostatic atmosphere should form with its pressure maximum located at r_{ECS}. Such an atmosphere should be completely detached from the radiating object, levitating above the surface and have so been deemed, *levitating atmospheres*. Wielgus et al. (2016) extended the calculation to show that optically thick levitating atmospheres could also be constructed. Such a result is only possible when considering radiation in a GR framework and provides a great example for why GRRMHD is important in the first place.

1.2 Numerical methods

The equations of GRRMHD comprise twelve coupled highly nonlinear partial differential equations. In almost all cases it is necessary to solve them on a computer. Over the next section I will explain how this is done in general and how it is done in the GRRMHD code `Koral`, my primary tool throughout this thesis.

1.2.1 Godunov, a finite-volume method for conservation laws

While all of the equations of motion look rather complicated, they all have the same basic structure,

$$\frac{\partial}{\partial t}\mathbf{U} + \nabla \cdot \mathbf{F}(\mathbf{U}) = \mathbf{S}(\mathbf{U}). \tag{1.54}$$

In words, the rate of change of some vector of quantities, $\mathbf{U}$, plus the divergence of the flux of those quantities (which is in general a function of those quantities), $\mathbf{F}(\mathbf{U})$, is equal to the source of those quantities (again, in general a function of those quantities), $\mathbf{S}(\mathbf{U})$. This is the basic differential form of a conservation law and there are many well developed methods to solve them. The method I use can more easily be seen from the integral form of Eqn. 1.54 were we take an integral over the volume defined by Σ and use the divergence theorem to perform a surface integral over its boundary, $\partial\Sigma$,

$$\frac{\partial}{\partial t}\int_{\Sigma} \mathbf{U} \mathrm{d}V + \oint_{\partial\Sigma} \mathbf{F} \cdot \hat{\mathbf{n}}\, \mathrm{d}A = \int_{\Sigma} \mathbf{S}\, \mathrm{d}V. \tag{1.55}$$

This is called the conservative form of the equation and shows us that the rate of change of a conserved quantity in some volume of space plus the total flux of that quantity leaving the volume equals the source of that quantity throughout the volume. The conservative form is able to handle discontinuities due to the lack of spatial derivatives, which would correspond to shocks for the GRRMHD equations. The conservative form also shows us how we would discretize the equations. A complete reference is given in LeVeque (2002) but I will outline the general idea.

Consider dividing space into many tiny volume units, $\Delta\Sigma_i$. Then we care only about the average[6] amount of the conserved quantity, $\bar{\mathbf{U}}_i$, inside the volume (now the horizontal bar refers to the average

---

[6] The formulation is the same if we consider the total amount of the conserved quantity, we would just have to multiply through by $\Delta\Sigma_i$.

of the quantity, not the radiation rest frame). It is convenient but not necessary[7] to use simple shapes for $\Delta\Sigma_i$ such as cubes or cuboids in Cartesian coordinates, x, y, z. The individual volume elements (grid cells) can be indexed in the conventional way, with i, j, and k indexing the x, y, and z coordinates, respectively. This gives, $\Delta\Sigma_{ijk} = \Delta x \Delta y \Delta z$[8] leaving us with,

$$\int_{\Delta\Sigma_{ijk}} \mathbf{U} \mathrm{d}V = \bar{\mathbf{U}}_{ijk} \Delta x \Delta y \Delta z. \tag{1.56}$$

Now we can consider the fluxes on each of the six faces. It is typical to refer to a cell face using a 'half' index, for example $\mathbf{F}_{i+1/2,j,k}$ would be the vector of fluxes across the interface between cell i, j, k and cell $i+1, j, k$ with commas added to the indices for clarity. The dot product $\mathbf{F} \cdot \hat{\mathbf{n}}$ introduces an orientation so that when we sum the fluxes (performing the integral) the three '$+1/2$' fluxes receive a positive sign and the three '$-1/2$' fluxes receive a negative sign. The projection of the flux vector onto the normal vector is done in the calculation of the flux components, i.e., $\mathbf{F}_{i+1/2,j,k}$ is calculated differently than $\mathbf{F}_{i,j+1/2,k}$ using the formulas for the x components of the fluxes instead of the y components, which I notate with an upper index. The differential, $\mathrm{d}A$, requires multiplying the fluxes by the areas of the faces so the x fluxes get factors of $\Delta y \Delta Z$ and so on. We can then approximate the integral,

$$\begin{aligned}\oint_{\partial\Delta\Sigma_{ijk}} \mathbf{F} \cdot \hat{\mathbf{n}} \, \mathrm{d}A \approx & \left(\mathbf{F}^x_{i+1/2,j,k} - \mathbf{F}^x_{i-1/2,j,k}\right) \Delta y \Delta z \\ & + \left(\mathbf{F}^y_{i,j+1/2,k} - \mathbf{F}^y_{i,j-1/2,k}\right) \Delta x \Delta z \\ & + \left(\mathbf{F}^z_{i,j,k+1/2} - \mathbf{F}^z_{i,j,k-1/2}\right) \Delta x \Delta y \end{aligned} \tag{1.57}$$

Next, we need to approximate the time derivative. The simplest approximation is to use a first order difference to get,

$$\frac{\partial \bar{\mathbf{U}}_{ijk}}{\partial t} \approx \frac{\bar{\mathbf{U}}^{n+1}_{ijk} - \bar{\mathbf{U}}^n_{ijk}}{\Delta t}, \tag{1.58}$$

where we have introduced the time index as a superscript. The flux term and the source term do not contain time derivatives so we are free to use their values anywhere from time n to time $n+1$, or to try to compute some estimate for the average along the time interval. For now I will avoid making the choice and drop the time indices for the flux and source terms altogether. It is now possible to integrate Eqn. 1.55 numerically by introducing the approximations above and solving for $\bar{\mathbf{U}}^{n+1}_{ijk}$.

$$\begin{aligned}\bar{\mathbf{U}}^{n+1}_{ijk} = \bar{\mathbf{U}}^n_{ijk} & - \frac{\Delta t}{\Delta x}\left(\mathbf{F}^x_{i+1/2,j,k} - \mathbf{F}^x_{i-1/2,j,k}\right) \\ & - \frac{\Delta t}{\Delta y}\left(\mathbf{F}^y_{i,j+1/2,k} - \mathbf{F}^y_{i,j-1/2,k}\right) \\ & - \frac{\Delta t}{\Delta z}\left(\mathbf{F}^z_{i,j,k+1/2} - \mathbf{F}^z_{i,j,k-1/2}\right) \\ & + \bar{\mathbf{S}}_{ijk}. \end{aligned} \tag{1.59}$$

---

[7] It is usually always possible to perform the calculation on a Cartesian grid but to include a coordinate transformation so that more complicated geometries can be considered (such as spherical). In grid-based general relativistic hydrodynamics (GRHD), the geometrical information is encoded in the metric and all its derived quantities so it is rather straight-forward to consider more complicated geometries. Furthermore, such schemes exist which use a moving mesh where the grid points are advected with the fluid, and the grid constructed at each time step using the Voronoi tessellation of an unstructured mesh (Springel 2010; Weinberger et al. 2020).

[8] The grid spacings, $\Delta x, \Delta y, \Delta z$, do not have to be equal to each other. They do not even have to be constant, then they would be indexed as $\Delta x_i, \Delta y_j, \Delta z_k$, but I will assume they are constant for simplicity.

To evolve the equation forward in time, it is then only necessary to know the current state at a specific cell ijk, the source term at cell ijk, and the six fluxes along with the lengths of cells in each direction. To evolve the entire system, it is necessary to repeat this process for each cell. Much of the complexity involves calculating the fluxes on the faces of the cells, the only respite is that the value of the flux at one interface can be used twice for each of the two neighboring cells. This makes sense because the value of the flux should depend on the states of both sides of the interface. Choosing a suitable method to compute the interface flux is the last step in the process. The numerical method I described here is called the Godunov scheme (Godunov and Bohachevsky 1959), a type of finite-volume method for conservation laws .

The Riemann problem

The Godunov scheme approximates the continuous function $\mathbf{U}(\mathbf{x})$ as $\bar{\mathbf{U}}_{ijk}$. This is the piecewise constant approximation where the whole cell is assumed to have constant values for each variable in the vector of conserved quantities. At the interface between each cell (for example between cells ijk and $i+1, j, k$) we then have a discontinuity, a left and right state, $\mathbf{U}_R = \bar{\mathbf{U}}_{ijk}$ and $\mathbf{U}_L = \bar{\mathbf{U}}_{i+1,j,k}$, with corresponding left and right fluxes, $\mathbf{F}_L = \mathbf{F}(\mathbf{U}_L)$ and $\mathbf{F}_R = \mathbf{F}(\mathbf{U}_R)$, from which we need to produce a single vector of fluxes. This is what is referred to as the Riemann problem. In one dimension, it consists of two fluids with different but constant states separated by an imaginary membrane. The membrane is then removed and the system evolves in time. We are then interested in what the fluxes are at the location of this membrane at some later time after the system begins to evolve.

The Riemann problem can be solved exactly in some cases. One should consult Toro (2009) for an in depth treatment of the subject. A common test problem for computational fluid codes is the Sod shock tube problem (Sod 1978), a specific case of the Riemann problem which has a semi-analytical solution. However, such methods are usually time consuming, and if they have to be computed at every cell interface for every time step, it becomes impractical. It is much easier, and usually nearly as effective, to use an approximate solution.

One strategy for coming up with approximate Riemann solvers is to approximate the solution by the superposition of a number of waves moving left and right after the system begins to evolve in time. A two-wave solution which uses the maximum and minimum velocity signals was proposed by Harten et al. (1983) hereby referred to as the HLL Riemann solver. If the maximum and minimum signal velocities (or maximum left and right moving signal speeds), S_R and S_L, respectively, can be computed (there are a number of ways to do this), then the HLL Riemann solver gives, for left and right states, $\mathbf{U}_L$ and $\mathbf{U}_R$, with corresponding left and right fluxes, $\mathbf{F}_L$ and $\mathbf{F}_R$,

$$\mathbf{F}_{\mathrm{HLL}} = \begin{cases} \mathbf{F}_L & 0 < S_L \\ \dfrac{S_R\mathbf{F}_L - S_L\mathbf{F}_R + S_L S_R(\mathbf{U}_R - \mathbf{U}_L)}{S_R - S_L} & S_L \leq 0 \leq S_R \,. \\ \mathbf{F}_R & S_R < 0 \end{cases} \tag{1.60}$$

If the minimum signal velocity, S_L is positive, this means all waves should be moving to the right and so one can use the flux corresponding to the left side of the interface. If the maximum signal velocity is negative, then the opposite is true and all waves move to the left, so we take the flux from the right. The intermediate case takes a combination of fluxes weighted by the wavespeeds with the addition of a term that computes the overlap of the left and right conserved quantities.

HLL is used extensively due to its simplicity while still producing meaningful results, but many other Riemann solvers exists. There is a third type of wave in ordinary hydrodynamics which HLL does not consider, the contact discontinuity corresponding to a jump in density without a corresponding jump in pressure. HLL is unable to provide satisfactory solutions when the contact discontinuity becomes important and a three wave Riemann solver was proposed by Toro (2009) (HLLC) which includes it. In MHD and GRMHD one also has to consider Alfvén and magnetosonic waves. These wave speeds can be used in the HLL or HLLC Riemann solvers but more accuracy is achieved by considering even more intermediate states as in the HLLD Riemann solver (Miyoshi and Kusano 2005) which assumes a five wave solution at the interface.

Wavespeeds and the CFL condition

There are a number of ways to compute the minimum and maximum signal velocities, S_L and S_R. The general strategy usually involves writing the conservation laws into a linear or quasi-linear form. In one dimension with no source terms we have,

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial}{\partial x}\mathbf{F}(\mathbf{U}) = 0 \implies \frac{\partial \mathbf{U}}{\partial t} + \boldsymbol{\mathcal{A}}(\mathbf{U})\frac{\partial \mathbf{U}}{\partial x} = 0, \tag{1.61}$$

where

$$\boldsymbol{\mathcal{A}}(\mathbf{U}) = \frac{\partial \mathbf{F}}{\partial \mathbf{U}},$$

is the Jacobian of the linear system. The signal velocities then correspond to Eigenvalues of the Jacobian. In ordinary HD, if c_s is the sound speed then the three Eigenvalues correspond to,

$$\lambda_1 = u + c_s,\ \lambda_2 = u,\ \lambda_3 = u - c_s, \tag{1.62}$$

which would correspond to the shock wave, contact discontinuity and rarefaction wave respectively. The HLL Riemann solver considers only the shock and rarefaction waves for both the left and right states.

In magnetohydrodynamics the functional form of $\boldsymbol{\mathcal{A}}$ is rather complicated, and only in the co-moving fluid-frame does it become (relatively) easy to find the eigenvalues analytically (Komissarov 1999). The resulting eigenvalues correspond to three additional waves, the fast/slow magnetosonic waves and the Alfvén wave, and serve as their phase velocities in the fluid frame. The phase velocity is given by $\widehat{\mu}^2 = \widehat{\omega}^2/\widehat{k}^2$, where $\widehat{\omega}$ is the wave frequency and $\widehat{k}^2 = \widehat{k}^i\widehat{k}^j\eta_{ij}$ (a reminder that $\eta_{ij} = \delta_{ij}$ in our metric signature) is the magnitude of the wave three-vector. We are only concerned with wave vectors oriented along the coordinate axes (orthogonal to the cell interfaces) which makes it easy to construct the wave four-vector $\widehat{k}^\mu = (\widehat{\omega}, \widehat{k}^i)$ in the fluid frame. Then one needs only to transform the wave four-vector to the lab frame and recompute the various phase velocities $\mu^2 = \omega^2/k^2$ (see Komissarov 1999, for more details). An almost identical approach involves plugging wave solutions into the linearized GRMHD equations to get a dispersion relation. One can then plug in the three k^i along the three coordinate axes, solve for ω and again transform to the lab frame (Gammie et al. 2003).

In relativistic and radiative hydrodynamics the speed of light becomes important as well. In relativity, the sound speed is given by,

$$c_\mathrm{s}^2 = \left(\frac{\partial}{\partial p}(\rho + u_\mathrm{int})\right)^{-1}, \tag{1.63}$$

assuming an ideal gas with an adiabatic equation of state $p \propto \rho^\Gamma$. As $\Gamma \to 4/3$, $c_\mathrm{s} \to c/\sqrt{3}$. A Γ of $4/3$ well describes a relativistic or optically-thick radiation-dominated fluid.

Speaking of radiation, it is also necessary to identify the radiation wavespeeds as radiation can be evolved using an identical scheme. Luckily, the wavespeeds for gas and radiation can be decoupled, as they are only relevant for the advective (homogeneous) part of the evolution equation. It may be tempting to set the radiation wavespeeds to c but if one computes the eigenvalues for $\boldsymbol{\mathcal{A}}$ for the radiation quantities, one will find they can drop to $c/\sqrt{3}$. A problem occurs in the regime that absorption κ is very small, but total opacity χ is very large due to the gas being in the scattering-dominated regime. In this limit, the photons are not absorbed by the gas but merely bounce around, scattering off electrons and diffusing through the gas via Brownian motion. However, nowhere in our evolution equations did we include a diffusion equation. In face, if we use the "correct" wavespeeds of $\pm c/\sqrt{3}$, then we would introduce large amount of numerical diffusion, which is unphysical. It is possible, however, to limit the wavespeeds using a function of the optical depth of the cell, $\rho\chi\Delta x_i$, in such a way that the resulting numerical diffusion matches the expected physical diffusion (for a slightly more detailed explanation, see Sądowski et al. 2013).

No matter what the values or the type of waves the maximum and minimum signal velocities correspond to, it is important that no signal travels the entire length of the cell in one time step,

$$\Delta x_i > a_i \Delta t, \tag{1.64}$$

where a_i is the maximum signal speed in cell i. This condition must be satisfied for every cell in the simulation. Otherwise the computed fluxes from the Riemann solver would not be accurate as there would be more waves and more states mixed into the solution. Therefore, one must limit the time step so that,

$$\Delta t = C \left(\frac{\Delta x_i}{a_i} \right)_\mathrm{min}, \tag{1.65}$$

where $C < 1$ and that Δt is still less than the maximum signal crossing time for every cell in the domain. This is known as the Courant–Friedrichs–Lewy (CFL) condition (Courant et al. 1928). One can pick any value for C less than unity and lower values of C give more accurate and more stable solutions however they also require longer to run so a balance needs to be found but in most cases $C \sim 0.5$ is sufficient.

1.2.2 Higher order methods

Reconstruction methods

The original Godunov method is technically zero'th order accurate in space. We can achieve more accuracy by moving from a piecewise constant approximation for $\mathbf{U}$ to a piecewise linear approximation. To do this, first I will introduce the concept of primitive variables. While I never actually mentioned it earlier, usually it is very difficult to compute the fluxes as a function of the conserved quantities. Instead, it is easier to use a different set of variables referred to as the primitives, $\mathbf{P}$. The primitives are usually naturally well-defined quantities such as density, velocity, and internal energy for which the computation of the fluxes and the conserved quantities is straight forward, i.e., $\mathbf{F}(\mathbf{P})$ and $\mathbf{U}(\mathbf{P})$ have analytical expressions. The trade-off is that, we must also perform the inversion from conserved to primitive variables, $\mathbf{P}(\mathbf{U})$, which, while simple for in ordinary HD, becomes much more complicated once additional physics is involved. In GRMHD is it necessary to perform the inversion numerically such as with a Newton-Rhapson method. A good scheme for GRMHD is discussed in Noble et al. (2006) who show that the system can be reduced to one equation by a clever change of variables.

Assuming that $\mathbf{F}(\mathbf{P})$, $\mathbf{U}(\mathbf{P})$, and $\mathbf{P}(\mathbf{U})$ can all be computed, then we are able to reconstruct the primitive variables at the cell interfaces. If we can find slopes $\frac{\Delta \mathbf{P}_i}{\Delta x}$ and $\frac{\Delta \mathbf{P}_{i+1}}{\Delta x}$ then the left and right primitives at the cell interface $\mathbf{P}^R_{i+1/2}$ and $\mathbf{P}^L_{i+1/2}$ can be computed as follows,

$$\mathbf{P}^R_{i+1/2} = \mathbf{P}_i + \frac{1}{2}\frac{\Delta \mathbf{P}_i}{\Delta x}, \tag{1.66}$$

$$\mathbf{P}^L_{i+1/2} = \mathbf{P}_{i+1} - \frac{1}{2}\frac{\Delta \mathbf{P}_{i+1}}{\Delta x}. \tag{1.67}$$

Then, one only needs to compute the left and right fluxes and conserved quantities and use a Riemann solver and described above. And of course, I dropped the indices j, k for simplicity but it is still necessary to include them in a full 3D scheme and so reconstruction must be performed three times for every cell, once in each direction.

One must use caution when obtaining the slopes $\frac{\Delta \mathbf{P}_i}{\Delta x}$, as it is possible to introduce new maxima and minima into the approximate solution of $\mathbf{U}$ (and therefore $\mathbf{P}$). This causes the numerical scheme to become unstable producing spurious oscillations in the solution which grow with time. To avoid this problem, it is necessary to use a total variation diminishing (TVD) scheme. A common method is to use a slope limiter which adjusts the computed slopes to prevent the formation of new maxima and minima. One widely used scheme is the generalized minmod (van Leer 1979) slope limiter which produces a modified slope $\frac{\Delta \mathbf{P}_i}{\Delta x}$ as a function of $\Delta \mathbf{P}_- \equiv \mathbf{P}_i - \mathbf{P}_{i-1}$ and $\Delta \mathbf{P}_+ \equiv \mathbf{P}_{i+1} - \mathbf{P}_i$,

$$\frac{\Delta \mathbf{P}_i}{\Delta x} = \frac{1}{\Delta x}\begin{cases} 0 & \Delta \mathbf{P}_- \Delta \mathbf{P}_+ \leq 0 \\ \min\left(\theta \Delta \mathbf{P}_-, 0.5\left(\Delta \mathbf{P}_- + \Delta \mathbf{P}_+\right), \theta \Delta \mathbf{P}_+\right) & \Delta \mathbf{P}_\pm > 0 \\ \max\left(\theta \Delta \mathbf{P}_-, 0.5\left(\Delta \mathbf{P}_- + \Delta \mathbf{P}_+\right), \theta \Delta \mathbf{P}_+\right) & \Delta \mathbf{P}_\pm < 0 \end{cases}. \tag{1.68}$$

The minmod slope limiter is actually a one-parameter mix between the more diffusive van Leer scheme (van Leer 1974) when $\theta = 1$ and the more accurate (but less stable) monotonized central (MC) scheme (van Leer 1977) when $\theta = 2$. A common choice is $\theta = 1.5$.

Slope-limiters have the additional function of preserving large discontinuities in the fluid across cell interfaces. Numerical schemes which include these limiters are refereed to as shock-capturing schemes.

Runge-Kutta

Probably one of first topics studied in an introductory course on numerical analysis is the numerical solution to an ordinary differential equation (ODE) using the forward Euler method. What proceeds immediately is usually a demonstration of how terribly inaccurate it is. It then follows that the next topic would be methods to improve it. One such method is the midpoint method. For a differential equation,

$$\frac{dx}{dt} = f(x, t),$$

if one knows the value of x at time t, one can approximate the solution after a time h by first performing an Euler step to time, $t + h/2$, and then using the value of the derivative computed from the approximation at time $t + h/2$ to jump all the way from time t to time $t + h$. The formula for the approximation at time $t + h$ is then,

$$x(t+h) \approx x(t) + hf\left(x(t) + f(x,t)\frac{h}{2}, t + \frac{h}{2}\right). \tag{1.69}$$

It is not necessary to use the value at the midpoint to approximate the derivative for second-order methods. Other values along the interval from t to $t + h$ introduce trade-offs between accuracy and stability. It is also not necessary to restrict oneself to only two steps or to a constant step size. Many methods use a combination of nth and $(n + 1)$th order methods to judge the accuracy of the numerical approximation in order to increase or decrease the step size when necessary, but this is not used in computational hydrodynamics because the step size is controlled by the CFL condition.

1.2.3 Stiff ODEs and the semi-implicit method

Not every differential equation can be integrated using explicit methods. If we again consider the most general first order differential equation,

$$\frac{dx}{dt} = f(x,t),$$

depending on the form of $f(x,t)$ any explicit method no matter the order or size of the time-step the integration becomes unstable and the equation is said to be stiff. A well known simple fix to this problem is to use an implicit integration scheme. Consider the following finite difference where instead of evaluating $f(x,t)$ at time n as before we evaluate it at time $n + 1$,

$$x^{n+1} = x^n + \Delta t f(x^{n+1}, t^{n+1}). \tag{1.70}$$

This gives an equation which is implicit in x^{n+1} and for all but the simplest cases, one must use a numerical root finding algorithm such as the Newton-Rhapson method. Using a numerical root finding algorithm every time step for every cell in a fluid simulation can be time consuming but in the case of GRRMHD, it is the only reliable way to include the time evolution of the source terms, G_ν.

1.2.4 Magnetic fields

The $\nabla \cdot \mathbf{B} = 0$ condition in MHD

Much of our discussion so far has focused on the numerical solution of conservation laws of ordinary fluids. We cannot however use *exactly* the same methods when Maxwell's equations are involved. While Eqn. 1.25, the induction equation, may look like an ordinary conservation equation, it is subject to the constraint given by Eqn. 1.26, the zero divergence condition. Nothing in the previous methods I have described is designed to obey such a constraint and so if one were to use them to evolve the magnetic field one would find that the accumulated divergence of the magnetic field would quickly cause the fluid to evolve in unphysical ways. There are many methods to alleviate this issue, several of which are described in Tóth (2000), two of which I will discuss below.

The first method is to simply remove the unphysical component of the magnetic field. Consider a vector field, $\mathbf{B}^*$. From Helmholtz's theorem, we know that we can write this field as the sum of a curl and a gradient of two other fields, $\mathbf{A}$ and ϕ respectively,

$$\mathbf{B}^* = \nabla \times \mathbf{A} + \nabla\phi. \tag{1.71}$$

Taking the gradient of both sides gives us an equation (the Poisson equation) which can be solved for ϕ,

$$\nabla \cdot \mathbf{B}^* = \nabla^2 \phi. \tag{1.72}$$

Then, one can get a divergence-free field from the gradient of ϕ,

$$\mathbf{B} = \mathbf{B}^* - \nabla \phi = \nabla \times \mathbf{A}. \tag{1.73}$$

The drawback is that one needs to solve the Poisson equation fairly often, which is expensive and non-local so it is difficult to parallelize. There are similar methods which do not remove the divergence completely but they follow the same idea. This family of methods are called divergence cleaning algorithms but there is a another family of methods whose advantages and disadvantages make them more suitable for GRMHD codes.

The second method involves evolving the field in such a way that no new divergence is introduced. Then, one only needs to start with a divergence-free field and it will be maintained for the duration of the simulation. The basic idea is to compute the induced electric field, $\Omega = \mathbf{v} \times \mathbf{B}$, at four locations for each component of $\mathbf{B}$ surrounding the point where $\mathbf{B}$ is to be evolved, and then to evolve $\mathbf{B}$ using a finite difference in Ω. This makes it so that any round-off errors which introduce divergence into the field are cancelled out. This family of methods constitutes *constrained transport* (CT) methods (Evans and Hawley 1988). The original CT methods uses a staggered grid, where the magnetic field is stored and evolved, not in the cell centers, but in the center of the cell faces, then the four values of Ω are computed at the cell corners. It is also possible to use the values of $\mathbf{B}$ and $\mathbf{v}$ at cell centers which are then interpolated to the cell faces where the field is evolved and then interpolated back to the cell centers. This is called *flux-interpolated constrained transport* (flux-CT) (Tóth 2000). A simple description of how it works is described in Gammie et al. (2003) which I will repeat here.

In 2D Cartesian coordinates in special relativity, the induction equation looks like

$$\partial_t B^i = -\partial_j \left(b^j u^i - b^i u^j \right) = -\partial_j F^{ji}, \tag{1.74}$$

where

$$\begin{aligned} F^{xx} &= 0, \\ F^{yy} &= 0, \\ F^{xy} &= b^x u^y - b^y u^x, \\ F^{yx} &= b^y u^x - b^x u^y. \end{aligned}$$

The Riemann solver will compute F^{xx}, F^{xy} at the location $i - 1/2, j$ (the interface between cell i, j and cell $i - 1, j$) and F^{yy}, F^{yx} at the location $i, j - 1/2$ which I will denote with lower indices $F^{xy}_{i-1/2,j}$, $F^{yx}_{i,j-1/2}$, etc. The flux-CT method will replace these fluxes using a combination of neighboring fluxes in the following way,

$$\tilde{F}^{xx}_{i+1/2,j} = 0, \tag{1.75}$$

$$\tilde{F}^{yy}_{i,j+1/2} = 0, \tag{1.76}$$

$$\tilde{F}^{xy}_{i-1/2,j} = \frac{1}{8}\left(2F^{xy}_{i-1/2,j} + F^{xy}_{i-1/2,j+1} + F^{xy}_{i-1/2,j-1} - F^{yx}_{i,j-1/2} - F^{yx}_{i,j+1/2} - F^{yx}_{i-1,j-1/2} - F^{yx}_{i-1,j+1/2}\right), \tag{1.77}$$

$$\tilde{F}^{yx}_{i,j-1/2} = \frac{1}{8}\left(2F^{yx}_{i,j-1/2} + F^{yx}_{i+1,j-1/2} + F^{yx}_{i-1,j-1/2} - F^{xy}_{i-1/2,j} - F^{xy}_{i+1/2,j} - F^{xy}_{i-1,j-1/2} - F^{xy}_{i-1,j+1/2}\right). \tag{1.78}$$

It is then straight-forward, but tedious to extend the scheme to 3D curved space.

Large magnetizations and force-free electrodynamics

There are many situations in GRMHD which require large magnetic fields such as the interior of magnetically launched jets or in the magnetospheres of neutron stars. One measure of the strength of the magnetic field is the ratio of the magnetic to rest mass energy densities, $\sigma = b^2/(2\rho)$, referred to as the magnetization. When the magnetization becomes high the numerical schemes of GRMHD start to break down and become prone to error. There are a few reasons for this. First, in the stress-energy tensor of a magnetized fluid, b^2 and p are summed together. When the conserved to primitive inversion $\mathbf{P}(\mathbf{U})$ is performed, small errors in the magnetic field translate to large errors in the pressure and thus internal energy of the fluid. Such errors could come from the difference in magnetic field that would be computed from the CT scheme which updates $T^{\mu\nu}_{\rm EM}$ and the normal Godunov method which updates $T^{\mu\nu}$ leading to errors in the evolution of $T^{\mu\nu}_{\rm (Rey)} = T^{\mu\nu} - T^{\mu\nu}_{\rm EM}$. A further problem is that it is impossible for a conservative numerical scheme to guarantee that the acceleration due to the Lorentz force is orthogonal to the magnetic field (Tóth 2000). If the magnetic field is small enough this effect is negligible but when $\sigma \gg 1$, gas starts to become accelerated along magnetic field lines. The third problem is that at large enough magnetizations, the equations of MHD are no longer a good approximation for the plasma flow and it can be shown that the current must be zero or parallel to the magnetic field. In this limit, known as force-free electrodynamics the plasma momentum can be effectively ignored. One must then either devise a numerical scheme which is appropriate for force-free electrodynamics (McKinney 2006; Mahlmann et al. 2021) or find a way to let a regular GRMHD scheme handle large magnetizations (Parfrey and Tchekhovskoy 2017).

1.2.5 Koral

At this point I have discussed at length nearly all of the numerical techniques required to understand the primary tool I use in this thesis, the `Koral` code which solves the equations of GRRMHD on a stationary 3D mesh for a static spacetime of an arbitrary, analytical metric. The code was written by Dr. Aleksander Sądowski. At the time of writing this thesis it is currently being maintained by Dr. Andrew Chael and contributed to by Prof. Ramesh Narayan. It has received smaller contributions from a number of other scientists and is used by a small number of groups internationally. I contributed a number of routines, many of which were adapted from Parfrey and Tchekhovskoy (2017), that allow the code to handle the highly magnetized environment of a neutron star magnetosphere which I will discuss in Chapter 4.

The main implementation notes of the code can be found in Sądowski et al. (2013) and Sądowski et al. (2015). There have been a number of features added including but not limited to photon-counting Comptonization (Sądowski and Narayan 2015), two-temperature electron and ion heating and cooling (Sądowski et al. 2016), and the evolution of non-thermal electrons (Chael et al. 2017). I will describe here the main features relevant to this thesis.

`Koral` evolves the following vector of conserved quantities, $\mathbf{U} = (\rho u_t, T^t{}_t, T^t{}_i, B^i, R^t{}_t, R^t{}_i)$, where $i = 1, 2, 3$ correspond to the three spatial indices in the chosen coordinate system. The corresponding primitive vector is $\mathbf{P} = (\rho, u_{\text{int}}, v^i, B^i, \bar{E}, v^i_R)$ where $v_i = u^i/u^t$ and $v^i_R = u^i_R/u^t_R$ are the 3-velocities of the gas and radiation rest frame, respectively, and $\bar{E}$ is the radiation energy density in the radiation rest frame.

The evolution of the gas and radiation fields are given by Eqns. 1.42 and 1.43. To evolve the equations forward in time a semi-implicit/explicit method is used. The explicit step is comprised of the advection and geometric source terms while the implicit step takes into account the radiative source terms, $\pm\sqrt{-g}G_\nu$.

The explicit evolution contains the following steps:

1. The primitives, $\mathbf{P}$, are reconstructed on the cell interfaces using the minmod or MC slope limiters.
2. The left and right biased fluxes and conserved quantities are computed from the reconstructed primitives.
3. The left and right moving wavespeeds are calculated and the Riemann problem solved using the HLL Riemann solver to get the intercell flux.
4. The magnetic fluxes are recomputed using the flux-CT method.
5. The CFL condition is used to compute the length of the time step Δt by searching for the minimum signal crossing time, $\Delta x/a$.
6. The conserved quantities are stepped forward in time by summing the fluxes and the geometric source terms.
7. The cell centered primitives are then recovered using the new conserved quantities via the $\mathbf{U}(\mathbf{P})$ inversion from (Noble et al. 2006). During the inversion, numerical errors are flagged, problem cells are fixed up, and numerical floors are applied.

The implicit step while simpler is more numerically demanding:

1. The radiative and gas energies are compared to see which fluid is more energetically dominant in each cell.
2. The less dominant fluid is selected and its primitives iterated over in a 4D Newton-Rhapson root finder to solve the set of equations either $T^{t,(n+1)}{}_\mu = T^{t,(n)}{}_\mu + G^{(n+1)}_\nu \Delta t$, or $R^{t,(n+1)}{}_\mu = R^{t,(n)}{}_\mu - G^{(n+1)}_\nu \Delta t$ where the Jacobian is estimated numerically and the primitives from the previous step used as the initial guess.

3. Cells which failed to find a solution are flagged and fixed up using neighboring cells.

The explicit and implicit steps are combined together and performed twice per time step using a 2nd order Runge-Kutta method for the time evolution.

It is also necessary to set the initial and boundary conditions, but these vary from simulation to simulation. The boundary conditions are set by extending the grid by a number of couple cells in each direction corresponding to the order of the spatial reconstruction scheme. These cells, referred to as *ghost cells* are not evolved in time but set manually and used to compute the intercell flux at the boundary.

The grid also varies from simulation to simulation, but all of the simulations presented in this thesis are run on grids in spherical coordinates, r, θ, ϕ. As mentioned before, the geometry of the grid is contained in the metric and so what the code actually works on is a cubical grid with numerical coordinates, x^1, x^2, x^3 which are then transformed to r, θ, ϕ only when necessary. This allows the grid to be stretched and warped and a common transformation is to make the grid exponentially spaced in r with the following transformation,

$$r = \exp x^1 + R_0, \tag{1.79}$$

where R_0 is a parameter which controls how quickly the grid grows with radius. Other such transformations can be used to concentrate grid cells towards the equatorial plane when studying accretion disks, or towards the polar region when studying jets for example.

For simulations in which an event horizon is present, it is possible to transform to Kerr-Schild coordinates. In these coordinates the coordinate singularity at the event horizon is removed, making it possible to perform simulations around black holes.

In D dimensions, as the resolution ($n_{xyz...}$) increases, the computational time required increases proportionally to $n_{xyz...}^{D+1}$. In two or three dimensions, this rapidly becomes to computationally expensive for a single machine. It becomes necessary to use a supercomputer and run the simulation on thousands of CPUs where each CPU works on a different section of the grid and information is exchanged between them using a standard message passing interface (MPI) implementation. All of the simulations presented in these thesis were run on the Prometheus supercomputer using around a thousand cores at a time running for three of four days each.

1.3 Accretion disks

At this point, I have discussed the necessary physics and numerical techniques needed to understand the subsequent chapters in this thesis, so now I can begin to lay out the astrophysical foundations. I am concerned with the study of accretion disks. In a binary star system, when a low mass star star begins to overfill its Roche lobe, or a high mass star undergoes significant mass loss through winds, a significant portion of the matter can be transferred to the second star, of mass M. The low-mass systems are referred to as a low-mass X-ray binaries (LMXBs) while the high-mass systems are called high-mass X-ray binaries (HMXBs) due to their emitted spectra peaking in the X-ray regime (Reig 2011).

In LMXBs, the transferred matter stream has some angular momentum, ℓ, and enters an elliptical orbit before self-interactions from the self-intersecting stream causes the matter to circularize into a ring at a radius,

$$r_{\text{circ}} = \frac{\ell^2}{GM}. \tag{1.80}$$

Viscous processes cause the ring to spread out into a disk through which gas is transported to eventually accrete onto the central object. The study of these processes constitute the field of accretion physics. One should consult Frank et al. (1992); Kato et al. (1998); Lasota (2016) for detailed references on the subject.

The rate at which the gas moves through the disk is $\dot{M}$, and in order for the gas to fall from radius r_1 to radius r_2 then the entire disk needs to release energy at a rate of $[\epsilon(r_1) - \epsilon(r_2)]\,\dot{M}c^2$, where $\epsilon(r)$ is the specific energy of an orbit at radius, r. If we assume that $\epsilon(r_1) \sim 0$ which is true for most gravitational potentials when r_1 is large compared to r_2, then the rate of energy release, L, is just a function of the accretion rate and specific energy at the inner edge of the disk. If we make the naive assumption that the gravitational potential is Newtonian and the disk ends at 6 r_g, which is the case for accretion disks around static black holes, then we have,

$$L = \frac{1}{12}\dot{M}c^2. \tag{1.81}$$

This factor of $1/12$ is called the efficiency, η, of the disk, and if we assume all of the energy is lost in the form of radiation (which is not always a good assumption) then this would be the radiative efficiency. More detailed models can give $\eta = 1/16$ for a pseudo-Newtonian potential, and $\eta = 0.057$ for a relativistic estimate around a non-spinning black hole which can rise all the way up to $\eta \sim 0.4$ for a maximally rotating black hole (Shakura and Sunyaev 1973). Regardless of the specific value used, we can see that the efficiency of energy extraction for accretion processes around compact objects is orders of magnitude more efficient that any other process besides matter-antimatter annihilation.

We will assume $\eta \sim 0.1$, which is a reasonable assumption for the models presented in this thesis. If we have an accretion disk which is accreting enough matter so that the total rate of energy release (luminosity) is equal to the Eddington luminosity L_{Edd}, then we can define an equivalent accretion rate,

$$\dot{M}_{\text{Edd}} = \frac{L_{\text{Edd}}}{\eta_0 c^2}. \tag{1.82}$$

The Eddington accretion rate $\dot{M}_{\text{Edd}}$ will prove to be a useful unit of measurement for $\dot{M}$ although one must keep in mind that it includes η and that many authors use many different values for η, even if the particular η chosen for the unit does not correspond to the actually η corresponding to the model. In this instance it is useful to differentiate the $\dot{M}_{\text{Edd}}$ efficiency as η_0 and the true or measured efficiency as η, with $L = \eta\dot{M}c^2$.

1.3.1 Mass in, angular momentum out

The basic mechanic in an accretion disk is the transfer of mass and angular momentum. A simple argument considers two particles of different masses and angular momenta. It can be shown that the minimum energy configuration is when a particle with infinitesimal mass and all the angular momentum is carried to infinity while the rest of the mass settles at the center (Lynden-Bell and Pringle 1974; Pringle 1981; Kato et al. 1998). In the accretion disk this translates to the flow of angular momentum outwards and mass inward.

The physical mechanism responsible for this transfer is viscosity. In a differentially rotating accretion flow, the vertically integrated viscous stress is given by,

$$T_{r\phi} = \nu \Sigma r \frac{\mathrm{d}\Omega}{\mathrm{d}r}, \tag{1.83}$$

where

$$\Sigma = \int_{-\infty}^{\infty} \rho \, \mathrm{d}z, \tag{1.84}$$

$$T_{r\phi} = \int_{-\infty}^{\infty} t_{r\phi} \, \mathrm{d}z, \tag{1.85}$$

where $t_{r\phi}$ is the viscous force per unit area excerpted in the ϕ direction at the interface between two annuli, Σ or surface density corresponds to the vertically integrated density, and Ω is the angular velocity. The rate of angular momentum transfer, and thus the timescale for accretion is decided by the kinematic viscosity, ν, where $t_{\text{vis}} \sim R^2/\nu$ where R is the typical length scale of the accretion disk. Typical values of the molecular viscosity, $\nu_{\text{mol}} \sim 10^3$ cm^2 s^{-1} give $t_{\text{vis}} \sim 10^{11}$ years which exceeds the age of the universe by far, therefore, some other processes must be responsible for the source of the viscosity, i.e., turbulence and magnetic fields, and it is necessary to choose a value for ν to construct a full analytical model of an accretion disk. A mechanism which leads to a much larger value of viscosity was proposed by Shakura and Sunyaev (1973) in which magnetic turbulence leads to the stress being proportional to the total pressure. Such a relation was shown for the first time in simulations by Hirose et al. (2009) where turbulence is supplied by the magnetorotational instability (Balbus and Hawley 1991).

1.3.2 Fundamental models

As mentioned above, a major step forward in accretion disk theory came from Shakura and Sunyaev (1973) who suggested the shear, $t_{r\phi}$, should be proportional to the pressure, $t_{r\phi} = -\alpha p$ where $\alpha < 1$. Typical estimates for α assume a constant value between 0.01 and 0.1, although Penna et al. (2013b) measured its radial dependence from GRMHD simulations of accretion disks.

With the α-viscosity prescription, a fully analytic model of accretion disks can be constructed with the following assumptions (Shakura and Sunyaev 1973; Kato et al. 1998):

1. The gravitational field is due only to the central object with self-gravity of the disk being ignored.
2. The disk is steady $\partial/\partial t = 0$.
3. The disk is axisymmetric, $\partial/\partial\phi = 0$.
4. The disk is geometrically thin, $h/r \ll 1$.
5. The disk is (nearly) Keplerian, $v_r \ll v_\phi = \sqrt{GM/r} \equiv r\Omega$.
6. The disk is in vertical hydrostatic equilibrium.
7. The disk is optically thick (and therefore radiatively efficient) in the vertical direction.
8. The viscous stress is proportional to the pressure.

Then, by combining mass conservation, momentum conservation, angular-momentum conservation, hydrostatic equilibrium, thermal equilibrium, an equation of state, a model for the opacity, and the viscosity prescription, one can produce a model for the scale height h/r, the surface density Σ, the radial velocity v_r, the central temperature T_c, and the optical depth τ, as a function of $\alpha, M, \dot{M}$. Almost immediately after Shakura and Sunyaev published their model, a relativistic version was published by Novikov and Thorne (1973).

These and similar standard disk models have been wildly successful in reproducing the observed properties of X-ray binaries, however it suffers from a rather catastrophic instability when

$$\left(\frac{\dot{M}}{\dot{M}_{\rm Edd}}\right) \geq \frac{1}{170}\left(\alpha\frac{M}{M_\odot}\right)^{-1/8}. \tag{1.86}$$

At these accretion rates, the inner region becomes radiation pressure dominated and the pressure's new dependence on temperature becomes unstable. So far no analytical model has been able to resolve the discrepancy between the overwhelming observational evidence that these disks must exist and the lack of a stable analytic solution. Numerical simulations are only recently beginning to probe the parameter space where the instability should manifest itself (Mishra et al. 2016; Sądowski 2016; Lančová et al. 2019).

1.3.3 The magnetorotational instability

Another major step forward in our understand of accretion processes came from the advent of the magnetorotational instability (MRI) as the primary mechanism of angular momentum transport in accretion disks (Balbus and Hawley 1998; Balbus 2009). A common simple (but wrong) interpretation of the instability involves considering two weakly coupled fluid elements at different radii in the disk. The element at the lower radius has a larger angular velocity and so it will tend to "pull" on the outer element transferring its angular momentum and causing it to fall to even lower orbits, increasing the force and thus the momentum transfer between the two elements, thus leading to instability. The MRI can only take place when the magnetic field is weak. This is usually measured by $\beta = p_{\rm gas}/p_{\rm mag}$ and MHD simulations of MRI have shown that β tends to saturate around 10 in the bulk of the disk. Simulations are necessary to study MRI because it is a turbulent process although growth rates can be predicted using perturbation analysis. The most important result is that the MRI can develop even from very weak seed-fields and so it is an excellent candidate to explain angular momentum transfer in accretion disks.

1.3.4 Super-Eddington accretion and slim disks

Aside from the problem with instability, standard thin disk models are also problematic in that their underlying assumptions are not valid at all accretion rates. When the accretion rate becomes very small, $\dot{M}/\dot{M}_{\rm Edd} \lesssim 10^{-4}$, the assumption of optical thickness no longer applies. At these accretion rates, the gas becomes so thin that it can no longer cool efficiently. The disk becomes hot and pressure gradients cause the disk to expand. Another consequence is that the radial velocity increases and the advection of thermal energy becomes important. Self-similar models of such flows called advection dominated accretion flows (ADAFs) where derived by Narayan and Yi (1994). Due to their geometrical thickness, low optical depths, and larger inflow velocities, ADAFs are much easier to simulate numerically and there have been many such publications on the subject, e.g., Igumenshchev et al. (2000, 2003); Narayan et al. (2012); Sądowski et al. (2016); Chael et al. (2017).

On the other end of the spectrum, for $\dot{M} \gtrsim \dot{M}_{\rm Edd}$, so-called super-Eddington accretion disks, the optical depth becomes so large that the assumption that the viscous heat is radiated locally ($Q^+ = Q^-$) no longer applies. One can compute the advection of heat (Shakura and Sunyaev 1973; Kato et al. 1998; Sadowski 2011)

$$Q^{\rm adv} \approx \frac{\dot{M}}{2\pi r^2}\frac{P}{\Sigma}, \tag{1.87}$$

and, along with the assumptions of radiation pressure domination, can then write the ratio of advection to radiative cooling as,

$$\frac{Q^{\rm adv}}{Q^-} = \frac{\dot{M}\kappa_{es}H}{32\pi r^2}, \tag{1.88}$$

which gives the accretion rate at which advection becomes important. Analytical models which include advection at larger accretion rates constitute the slim disk models (Abramowicz et al. 1988; Szuszkiewicz 1990; Sadowski 2011). Another often missed model of super-Eddington accretion can be found near the end of Shakura and Sunyaev (1973). At the radius where the total radiation emitted becomes super-Eddington called the spherization radius, i.e., $L(r_{\rm sph}) = L_{\rm Edd}$, an outflow is launched, driven by radiation pressure. The accretion rate then decreases linearly with radius so that at the inner edge of the disk it is back to near $\dot{M}_{\rm Edd}$.

Another phenomenon at large optical depths is that photons generated in the inner parts of the disk are advected significant distances before they can diffuse out of the flow. A large fraction of these photons never actually escape the disk before they are accreted onto the central object. The excess radiation pressure causes these disks to puff up as well and so, for similar reasons as ADAFs, they are also much easier to simulate (Ohsuga et al. 2002, 2005; Sądowski et al. 2013, 2015; Jiang et al. 2014; Sądowski and Narayan 2015, 2016). Super-Eddington flows like these are the main subject of this thesis.

1.4 Accretion onto neutron stars

Up until this point, I have been mostly concerned with the physics of the accretion disk without regards to the central object, from which all that was necessary to know was the gravitational potential it was responsible for. I will now turn our attention to the case where the central object is something specific, which has a profound effect on the structure of the inner disk and on the resulting outflows and radiation output.

1.4.1 Neutron Stars

I will consider accretion onto neutron stars (NSs). These objects are the end-states of massive stars whose cores exceed the $1.4M_\odot$ Chandrasekhar mass limit for electron degeneracy pressure-supported stars but which are not so massive as to go on to form black holes. Neutron stars are supported by neutron degeneracy pressure and they constitute the densest known non-singular objects in the universe with typical masses between 1-2 $M_\odot$ and typical radii between 5-15 km. Their central densities, while difficult to estimate because their equations of state cannot be precisely constrained, are often estimated to be several times the nuclear density, $\rho_s \sim 2.7 \times 10^{14}$ g cm^{-3} (Lattimer 2012). The surfaces of neutron stars are comprised of ionized iron nuclei reaching densities of up to 10^{11} g cm^{-3} in the lower surface layer which transitions to a crystal lattice of neutral iron atoms closer to 10 g cm^{-3} (Chamel and Haensel 2008). This is several orders of magnitude larger than the expected plasma densities of even highly super-Eddington accretion flows and so the inflowing gas is expected to slow down and release its kinetic energy through electron ion collisions (Syunyaev and Shakura 1986; Kluzniak and Wilson 1991; Narayan and Yi 1995; Narayan et al. 1997; Inogamov and Sunyaev 1999; Sibgatullin and Sunyaev 2000; Popham and Sunyaev 2001; Mukhopadhyay and Fang 2002) while depositing its angular momentum and spinning up the star (Kluzniak and Wagoner 1985) releasing large amounts of X-rays and driving gas outflows.

The magnetic fields of neutron stars are expected to be quite high. If we take the Sun (even though it is too small to end up as a neutron star) which has a typical magnetic field ~ 1 G and magnetic moment of $\sim 10^{32}$ G cm^3, and squeeze it down to the radius of a neutron star (conserving the magnetic moment), we would expect surface magnetic fields around 10^{14} G. The actual magnetothermal evolution of neutron stars is quite complex (Pons and Viganò 2019), but the range of magnetic fields strengths has been inferred from observations to span from 10^8 to 10^{15} G.

For accreting neutron stars, the magnetic field can be strong enough to truncate the inner region of the disk, forcing the gas to flow along magnetic field lines towards the magnetic poles forming accretion columns. Elsner and Lamb (1977) computed the location at which this occurs, the Alfvén radius. A basic estimate can be done by assuming the simplest configuration of spherical accretion free-falling onto a dipolar magnetic field which gives,

$$\begin{aligned} r_A = \left(\frac{\mu^4}{2GM\dot{M}^2}\right)^{1/7} &\approx 3.2 \times 10^8 \times \mu_{30}^{4/7}\, m^{-1/7}\, \dot{M}_{17}^{-2/7}\, \mathrm{cm} \\ &\approx 1.0 \times 10^3 \times \mu_{30}^{4/7}\, m^{-10/7}\, \dot{m}^{-2/7}\, r_g, \end{aligned} \tag{1.89}$$

where μ is the neutron star dipole moment, $m = M/M_\odot$, and $\dot{m} = \dot{M}/\dot{M}_{\rm Edd}$. At high enough accretion rates with low enough magnetic moments the field can be entirely buried by the accretion flow, while at the opposite end, the field can be truncated at thousands of gravitational radii. In Chapter 3, I study the former case while in Chapter 4, I study the case where the Alfvén radius sits at about 10 r_g.

1.4.2 Pulsars

X-ray Pulsars

There are about 200 LMXBs observed in the Milky way (Liu et al. 2006, 2007) and while a small fraction of them have been confirmed to be accreting black holes (Corral-Santana et al. 2016), the rest are thought to be accreting neutron stars. As discussed above, when the magnetic field of the neutron far is sufficiently strong, the accretion disk is truncated at the Alfvén radius where it flows along magnetic field lines forming accretion columns near the magnetic poles. At lower accretion rates ($\dot{M} \lesssim 0.1\dot{M}_{\rm Edd}$), the gas in the accretion column is not expected to be radiatively efficient, so it flows along the magnetic field lines until it is deposited near the poles, at which case it releases its thermal and kinetic energy in the form of X-rays. If the dipole field axis is misaligned with the rotation axis, then these hot-spots which form at the base of the column will also rotate, resulting in what are observed as X-ray (accretion powered) pulsars (Caballero and Wilms 2012). The X-ray spectra produced are very difficult to explain, but generally consist of a power law in the $(5 \sim 20)$ keV range with an exponential cut off at higher energies (Coburn et al. 2002). The spectra are further complicated by iron Kα fluorescent and cyclotron absorption lines.

If we consider higher accretion rates $(0.1 \sim 10\,\dot{M}_{\rm Edd})$, the column starts to become optically thick, and so photons are not able to escape so easily from the base. This forms a radiation shock which rises above the surface of the neutron star, confined by the magnetic field. Gas sinks slowly through the dense radiation field before settling on the NS surface, meanwhile photons diffuse out the side of the column. The classical picture is well described by Basko and Sunyaev (1976). More recent studies were conducted by Mushtukov et al. (2015a,b); Revnivtsev and Mereghetti (2015), explaining the complicated pulse profiles expected from the accretion column when the rotation and magnetic axes are misaligned. The spectra are expected to be dominated by bulk and thermal Comptonization of bremsstrahlung, cyclotron, and blackbody emission leading to a flat continuum with a high-energy quasi-exponential cutoff (Becker and Wolff 2007). Single accretion columns in such sources were investigated using local radiative hydrodynamic simulations by Kawashima et al. (2016); Kawashima and Ohsuga (2020) who considered both the low and high magnetic field cases by allowing or restricting the gas motion in the θ direction. A radiative shock above the NS surface along with a large flux of photons escaping out the sides of the column were observed in the low magnetic field case. In the high magnetic field case, several "chimneys" in the interior of the column appeared allowing radiation to escape radially as well as tangentially.

At even higher accretion rates $(\dot{M} \geq 10\dot{M}_{\rm Edd})$ it has been theorized that the accretion column spreads out to envelop a closed region around the neutron star surface which is filled with radiation. Most photons which originate from near the NS surface are not observed directly but possibly emitted thermally from the accretion 'curtain' leading to smooth pulsed profiles (Mushtukov et al. 2017).

Accretion can deposit angular momentum onto the spinning neutron star if its Alfvén radius is smaller than its corotation radius (the radius at which Keplerian and NS rotation period are the same). This spins up the star and is thought to be how millisecond X-ray pulsars, accreting pulsars with periods less than 0.1 s, are formed (Patruno and Watts 2021). If the pulsar gains too much angular momentum so that its corotation radius falls below the Alfvén radius then the pulsar enters the propeller regime (Illarionov and Sunyaev 1975), where the accreting gas is stopped and blown away and X-ray pulses are no longer produced.

Radio Pulsars

Another type of pulsar is formed when the NS survives the evolution of its binary companions or when an isolated neutron star maintains a sufficient amount of angular momentum and magnetic field during its formation. In this case the rapid rotation of the magnetic field powers radio emission which are observed as radio pulses and the sources are called radio pulsars[9] (Beskin et al. 2015). They are worth mentioning as they are the most numerous source of observed neutron stars, and the first radio pulsar, which was discovered in 1967 by Jocelyn Bell (Hewish et al. 1968), is widely credited as being

---

[9] It may be more correct to refer to these objects as 'rotation-powered pulsars' as they are also capable of producing X-rays and gamma rays and some have even been observed to produce no radio emission at all (Kaspi et al. 2006).

the first observationally confirmed neutron star in existence. However, Shklovsky (1967) correctly concluded based off observations of Scorpius X-1 that it was powered by an accretion neutron star mere months before Bell's observation.

1.4.3 PULXs

Recently, a new class of objects has been observed which has opened up a whole new set of questions regarding accretion. For a long time, ultra-luminous X-ray sources (ULXs) had been observed in other galaxies. These are bright X-ray sources with luminosities greater than 10^{39} erg s^{-1} which are located outside the central regions of galaxies (thus they are not related to super-massive black holes). As 10^{39} erg s^{-1} exceeds the Eddington luminosity of a neutron star by a factor of 10, early explanations attributed the exceptionally large luminosities to sub-Eddington accretion by intermediate mass black holes with masses $\sim 100\, M_\odot$ (Colbert and Mushotzky 1999). A competing explanation claimed that it was possible for solar mass objects to produce these luminosities from super-Eddington accretion via the beaming of the radiation from the accretion disk outflows (King et al. 2001).

The question for at least a few of these objects was answered when a ULX was discovered emitting ~ 1 s coherent pulsations (Bachetti et al. 2014), showing that the central accretor had to be a neutron star accreting above the Eddington limit, $\dot{M}_{\mathrm{Edd}}$, and producing luminosities from 10-1000 L_{Edd}. There are now a significant number of known pulsating ULXs (PULXs) (Trudolyubov 2008; Bachetti et al. 2014; Motch et al. 2014; Fürst et al. 2016; Israel et al. 2017a,b; Townsend et al. 2017; Tsygankov et al. 2017; Brightman et al. 2018; Carpano et al. 2018; Doroshenko et al. 2018; Fürst et al. 2018; Heida et al. 2019; Chandra et al. 2020), and further evidence that many other ULXs may indeed be powered by neutron stars (Kluźniak and Lasota 2015; King et al. 2017; Wiktorowicz et al. 2017; Pintore et al. 2017).

The question now remains, how are these accreting neutron stars able to produce such large luminosities. I already discussed super-Eddington accretion, however I neglected to mention that the expected luminosities rises logarithmically with the accretion rate when the spherization radius is larger than the ISCO radius (Shakura and Sunyaev 1973),

$$L \sim L_{\mathrm{Edd}} \left[1 + \ln\left(\dot{M}/\dot{M}_{\mathrm{Edd}}\right)\right], \tag{1.90}$$

so that large total luminosities, at least in the classical picture, would be unrealistic.

One answer is to invoke the presence of extremely large magnetic fields ($B \gtrsim 10^{13}$ G). At these field strengths electrons are confined to move along magnetic field lines, lowering their scattering cross sections in the perpendicular direction. This would allow large amounts of radiation to escape but it should do so nearly isotropically (Ekşi et al. 2015; Mushtukov et al. 2015b).

It is also possible to consider a more detailed description of the physics of the accretion column. Although Basko and Sunyaev (1976) lay the foundation for much of our understanding of accretion column physics, they do not discuss the photon bubble instability (Arons 1992), which could very well be very important for understanding the luminosity and variable of bright accretion columns. In a highly radiative magnetized plasma, if a compression wave is present with some component transverse to the magnetic field, then it is possible for photons to enter the under-dense regions of the wave due to the restriction of matter to move along magnetic field lines. This could cause oscillations in the column (Klein et al. 1996), and the structure could possibly resemble long elongated bubbles like the fingers observed in Kawashima and Ohsuga (2020). Radiative MHD simulations of photon bubbles confirm the linear theory and lead to instability and in static radiative atmospheres (Zhang et al. 2021). It is uncertain how the results will be affected when considering an accretion column subject to accreting matter, but it is possible that a large amount of photons may be able to escape either tangentially or radially, increasing the maximum theoretical luminosity of the source.

An alternative explanation I have already mentioned involves beaming via the outflows from the super-Eddington disk (King et al. 2001). If the total luminosity of a source is equal to L, but the radiation only escapes through some fraction, b, of solid angle of the sphere, then an observer will overestimate the total luminosity as,

$$L_{\mathrm{iso}} = L/b, \tag{1.91}$$

if they assume the source is isotropic. The lower the beaming factor, b, the more highly concentrated the emerging radiation is, and the larger the inferred isotropic luminosity, L_{iso}, will be. Such an effect has been included into a model (King et al. 2017; King and Lasota 2019, 2020, hereafter referred to as the KLK model) which also incorporates period, spin-up, and luminosity to predict the magnetic field strength and degree of beaming. The model assumes the accretion disk is truncated at the Alfvén or magnetospheric radius, R_M. At this radius angular momentum is transferred to the star until the spin period equals the Keplerian period at R_M. Furthermore, if $R_M < R_{\mathrm{sph}}$, then the disk should launch an outflow which collimates the radiation leading to a beaming factor as a function of the mass accretion rate,

$$b \approx \frac{73}{(\dot{M}/\dot{M}_{\mathrm{Edd}})^2}. \tag{1.92}$$

Such an outflow is necessary for the KLK model to be valid, although one could argue that the geometry of the accretion column could significantly beam the resulting emission, but this would come with its own problems. Strong and ultra-fast outflows have been observed in ULXs (Pinto et al. 2016) and even in pulsating ULXs (Kosec et al. 2018), showing that beaming in these objects is a very real possibility.

The KLK model predicts magnetic fields strengths ranging from 10^9 to 10^{13} G although most of their results lay between 10^{10} and 10^{11} G. Lower magnetic fields are also supported by the presence of cyclotron lines in NGX300 ULX1 (Walton et al. 2018) and M51 ULX-8 (Brightman et al. 2018; Middleton et al. 2019). Further evidence is provided by additional observations of strong outflows in some ULXs (Mushtukov et al. 2019), although this seems to be complicated by the implication that strong outflows/and thus strong beaming should be incompatible with large pulse fractions (Mushtukov et al. 2021) which would require large opening angles. This is perhaps why in nearly all cases the KLK model predicts $R_M \sim R_{\mathrm{sph}}$, due to the selection effect of only observing pulsations in ULXs with the lowest possible values of the spherization radius to sufficiently beam the emission but not to obscure the pulsations.

PULXs and ULXs in general contain all the right ingredients that make them suitable for studying using GRRMHD simulations. If they are indeed powered by neutron stars accreting above the Eddington limit, then the radiation should have a dynamic role in their evolution and the strong gravity around the neutron star requires the use of general relativity. Both the magnetized and unmagnetized cases are interesting to study. First, to see what effect the magnetic field has, if any, on the observable and structural properties of super-Eddington accretion. And second, because there should be many ULXs powered by neutron stars which are non-pulsating, and may therefore have very weak magnetic fields.

Until recently only one such GRRMHD simulation of accretion onto a magnetized neutron star has been performed (Takahashi and Ohsuga 2017), and one onto a non-magnetized neutron star. (Takahashi et al. 2018). The magnetized case showed promising results. The measured luminosity was an order of magnitude larger than the Eddington limit and the radiation was tightly beamed around the axis. The non-magnetized simulation showed all of the inflowing gas to be blown off the surface of the neutron star leading to strong outflows which obscured the inner radii of the accretion flow. Both simulations are important works but they have several shortcomings. The most striking are the short duration of the simulations, and the lack of proper treatment of the highly magnetized zones. Another issue is that it is not entirely clear how the inner boundary was treated. I aimed to overcome these issue in the papers presented in Chapters 3 and 4 in which I performed simulations of super-Eddington accretion onto a non-magnetized and weakly-magnetized neutron star, respectively.

1.4.4 Transient behavior in accreting neutron stars

Turbulence is by definition chaotic, and since accretion is powered by turbulent viscosity, it must also contain an element of chaos. This is evident in the variability of their light curves, measurements of the luminosity as it changes with time, and while it is impossible to account for every minor blip, some patterns can be found in the chaos and attributed to physical processes.

Probably the simplest of these patterns would be epicyclic oscillations. In Newtonian physics, if you take a particle on a circular orbit and perturb it in the radial direction, it will enter a different, slightly elliptical orbit that is very close to the original circular one. At some locations, the object will be outside its original orbit, and at others it will be inside. An observer orbiting at the original orbit would see the object appear to oscillate radially about its original position with a period equal to the orbital period. The same is true if the original displacement is perpendicular to the orbital plane. These oscillations are referred to as radial and vertical epicycles, and in general relativity, their frequencies differ from each other, and they usually contribute strong peaks in the periodograms of accretion disks.

A periodogram (the power spectrum when considering a set of periodic basis functions) measures the amount of ’power’ or the contribution that an individual frequency provides to the total variability of the light curve. The most common way to compute a periodogram/power spectrum involves a discrete Fourier transform, which provides the recipe for reconstructing the original light curve using a linear combination of sine and cosine waves whose frequencies are integer multiples of each other.

Accreting neutron stars have been observed to occasionally erupt in thermonuclear (Type I) X-ray bursts. The accreting gas thermalizes and is compressed to the point where thermonuclear fusion can occur igniting an eruption which spreads over the surface of the star lasting anywhere from 10-100 seconds (e.g. Bildsten 2000) . These X-ray burst occur repeatedly as fresh fuel is replaced by the accretion flow. The power spectra of the light curves during these bursts contain peaks referred to as quasi periodic oscillations (QPOs), as they are not delta functions in frequency space but rather have some width, meaning they are comprised of a collection of similar frequencies. In other words, the oscillations are not coherent. QPOs present in X-ray bursts are termed bust oscillations and have been observed to span frequencies between 300 and 600 Hz (Strohmayer and Bildsten 2006; Watts 2012). Type 1 X-ray bursts can be divided into three phases: rise, peak, and tail. Burst oscillations during the rise and peak have been shown to be related to the neutron star spin (Chakrabarty et al. 2003; Strohmayer and Markwardt 2002), but oscillations during the tail phase have yet to be properly explained.

Type 1 X-ray bursts are associated with large luminosities with some bursts reaching near their Eddington limits (Lewin et al. 1984; Tawara et al. 1984; Lewin et al. 1993). As was previously discussed, if the redshifted luminosity exceeds the Eddington limit then the ECS should be present and there should be a surface where plasma can accumulate. Since this surface is located at a stable equilibrium, it is reasonable to consider oscillations of an atmosphere at this surface and ask if they are in any way related to the oscillations observed in the tail phase of Type 1 X-ray bursts. The next chapter is comprised of a publication (Abarca and Kluźniak 2016) which addresses this very question, taking into account the effects of radiation drag to see if the lowest order possible mode could explain these oscillations. The results in Abarca and Kluźniak (2016) set the stage for a series of other works exploring the oscillations of Eddington supported atmospheres, Bollimpalli and Kluźniak (2017) and Bollimpalli et al. (2019), the later of which identified a new way to perform simultaneous mass and radius measurements if such oscillations are present around a neutron star.

1.5 Simulation details

The simulations described in Chapters 3 and Chapters 4 were presented in sufficient detail to be accepted for publication, but for the sake of clarity in this thesis, I would like to elaborate on some of the specifics.

1.5.1 Initial conditions

Each simulation required an initial condition. In the `Koral` simulations, this correspond to setting the primitives $\mathbf{P} = (\rho, u_{\rm int}, v^i, B^i, \bar{E}, v^i_R)$ to some initial state. In accretion disk simulations, it is common to initialize the accretion flow with a "Polish donut", or equilibrium torus. Such models were originally developed in Warsaw, Poland (Abramowicz et al. 1978; Paczyński and Wiita 1980) as alternative models to explain super-Eddington luminosities. The general structure involves an axisymmetric torus with a non-Keplerian angular momentum profile which is confined by the effective potential in a frame orbiting with Keplerian angular momentum at $r_{\rm max}$, which describes the location of the pressure maximum. For all the simulations presented in this thesis, I used the equilibrium torus solution from Penna et al. (2013a) which has a few advantages over more traditional solutions such as Fishbone and Moncrief (1976), namely finer control over the initial Bernoulli parameter.

The initial torus has an adiabatic equation of state which means it has a well defined edge. Outside of the torus, the simulation is set to a low density "background" atmospheric profile with densities and internal energies several orders of magnitude lower than inside the torus. The temperature of the background is usually much hotter than the torus and typically decays linearly with radius.

The magnetic field needs to be initialized so that it has zero divergence. This is done by setting the vector potential in such a way that it gives the desired magnetic field after a numerical curl operator is applied. For the simulation described in Chapter 3, it is only necessary to add magnetic field lines to the initial torus. Many configurations are possible. I chose a multi-loop configuration which does not allow significant magnetic flux to accumulate at the inner boundary for Chapter 3 and a single loop configuration in Chapter 4 which allows the disk field to reconnect with the stellar dipole. Once the field is computed from the vector potential, its magnitude is scaled so that the maximum plasma $\beta = p_{\rm mag}/p_{\rm gas} = 20$, as MRI does not require a very strong initial magnetic field. The dipole field for the neutron star described in Chapter 4 was initialized from the relativistic potential described in Wasserman and Shapiro (1983), but with the outer field lines deformed to warp around the initial torus.

To initialize the radiation field it is necessary to redistribute pressure between gas and radiation as the solution from Penna et al. (2013a) is non-radiative. In local thermal equilibrium, the total pressure is just the sum of the gas and radiation pressures,

$$p_{\rm tot} = p_{\rm gas} + p_{\rm rad} = \frac{k_{\rm B}\rho}{\mu m_p} T + \frac{4\sigma_{\rm SB}}{3} T^4, \tag{1.93}$$

where μ is the mean-molecular weight (I assumed pure ionized hydrogen so $\mu = 1/2$ for each simulation described), and $\sigma_{\rm SB}$ is the Stefan-Boltzmann constant. This provides a fourth order polynomial which can be solved for T, which is then used to compute $u_{\rm int}$ and $\bar{E}$. The radiation velocity is then set to the fluid velocity and a low density radiative atmosphere which follows a power law is initialized outside of the torus.

1.5.2 Force-free flooring scheme for high magnetizations

One consequence of including a stellar magnetic field is that from the beginning the simulation will be subject to large magnetizations. In Sec. 1.2.4 I explained why this is a problem. Here, I elaborate on the method I used to overcome this issue. The scheme was introduced in Parfrey and Tchekhovskoy (2017) and I adapted it to work with radiation.

What is normally done when large magnetizations appear in the simulation is to inject gas so that the magnetization never rises above a predefined value, usually between 20 and 100. This introduces a relative floor in the density. The basic idea of the new scheme is to, in a sense, keep track of this extra gas from the flooring scheme and then to adjust the fluid quantities based on the amount of extra gas present in a cell. But, instead of using a relative floor on the density which fluctuates based of the magnitude of the magnetic field, we use a static floor, ρ_f which follows a function of radius. The floor is chosen as a broken power-law so that initially the magnetization has a maximum on the neutron star surface, usually around $\sigma \sim 10^4$, and then drops with radius to about unity at the break of the power law. At larger radii, the density floor decays with $\rho_f \propto r^{-6}$, keeping the magnetization constant.

To keep track of the contribution of the floor to the total density, an additional scalar field, $\mathcal{F}$ is evolved with the simulation, according to,

$$\nabla_\mu(\mathcal{F}\rho u^\mu) = 0. \tag{1.94}$$

The scalar field $\mathcal{F}$ represents the fraction of the density which is due to the density floor, initially $\mathcal{F}$ is set to 0 inside the torus and 1 outside. As $\mathcal{F}$ is advected with the density it is also used to adjust the density floor, ρ_f. When $\rho < \rho_f$, $\mathcal{F} \to (\rho_f - \rho)/\rho_f$ and $\rho \to \rho_f$, so ρ is brought up to the value of the density floor, and $\mathcal{F}$ keeps track of the amount of density added. When $\mathcal{F}\rho > \rho_f$, $\rho \to \rho - \mathcal{F}(\rho - \rho_f)$ and $\mathcal{F} \to \rho_f/\rho$, so that if the density contribution from the floor is too large it can be removed and $\mathcal{F}$

lowered appropriately. In this way the total density can be adjusted to make the GRRMHD equations behave while still keeping track of the "true" density.

The quantity $\mathcal{F}$ can then be used to adjust the other fluid quantities which are affected by large magnetization. Analogous to ρ_f a ceiling on the internal energy is also used $u_{\rm int,f}$ which has nearly the same radial dependence as ρ_f. If $u_{\rm int} > u_{\rm int,f}$ then $u_{\rm int} \to u_{\rm int,f} + (1-\mathcal{F})(u_{\rm int} - u_{\rm int,f})$. Since the internal energy cannot be split in the same way as density, it is linearly interpolated between its value from the GRRMHD equations when $\mathcal{F} \sim 0$ and $u_{\rm int,f}$ when $\mathcal{F} \sim 1$.

Finally, the fluid velocity is adjusted by reducing its component parallel to the magnetic field in the frame of the stationary observer, $u^\mu b_{\mu,s}$. At first glance, this seems strange because the orthogonality of the 4-velocity and 4-magnetic field is part of the definition of b^μ. However, b^μ, even though it is defined in the lab frame, represents the magnetic field seen by the fluid. This can be seen by transforming Eqn. 1.24 to the orthonormal fluid frame, where one finds $\widehat{b}^i = \widehat{B}^i$. One can easily recover the corresponding b^μ seen by an observer with four-velocity u'^μ as $b'^\mu = -{*}F^{\mu\nu}u'_\mu$, which is evident when $u'^\mu = u^\mu$. Then to get the field as observed by the stationary observer[10] with four-velocity $\eta^\mu = (\eta^t, 0, 0, 0)$ we can simply do

$$b_s^\nu = -{*}F^{\mu\nu}\eta_\mu = (u^\mu\eta_\mu)b^\nu - (b^\mu\eta_\mu)u^\nu. \tag{1.95}$$

From this point, it can be seen that u^μ and b_s^μ are not orthogonal and $u_\mu b_s^\mu = b^\mu\eta_n u$. This may seem like a lot of machinery for what, at least in the Schwarzschild geometry, is simply proportional to $B^i u_i$ as expected, but for more complicated geometries, especially for metrics with off-diagonal terms, mixing of the coordinates between frames makes the analysis much less straightforward and the above formalism becomes much more useful. In this thesis, I stick to the Schwarzschild geometry but extending the method to Kerr for example should not be difficult.

To actually reduce $k \equiv u_\mu b_s^\mu$, we split the four-velocity into parallel and orthogonal components,

$$u_\parallel^\mu = k\frac{b_s^\mu}{b_s^2}, \tag{1.96}$$

$$u_\perp^\mu = u^\mu - k\frac{b_s^\mu}{b_s^2}. \tag{1.97}$$

From this, one can easily see that $u^\mu = u_\parallel^\mu + u_\perp^\mu$, $u_\parallel^\mu u_{\perp,\mu} = 0$, and $u^2 = u_\parallel^2 + u_\perp^2$. In the flooring scheme, we use the value $\mathcal{F}$ to transform $u^\mu \to A(u_\perp^\mu + \mathcal{F}u_\parallel^\mu)$ where $A = \sqrt{-1/(u_\perp^2 + \mathcal{F}^2 u_\parallel^2)}$ is the normalization factor of the new four-velocity.

This takes inspiration from force-free electrodynamics (FFE) in which the important dynamic objects are the electric and magnetic fields. In FFE, the plasma momentum is unimportant, and so in this flooring scheme, we only preserve the orthogonal component of the velocity which is necessary to evolve the magnetic field in the GRMHD equations. This provides stability to the code as the parallel component tends to grow artificially large at high magnetizations.

The above described scheme is applied after every explicit step after the conserved-to-primitive inversion occurs.

---

[10] The magnetic field as observed by the stationary observer is defined in the lab-frame so it does not receive a head-piece (hat, tilde, bar, etc.). Instead I denote the components of the 4-vector b_s^μ, where s is not meant to be thought of as an index but merely a label.

It is also possible to gradually turn off the flooring scheme at larger radii by multiplying all instances of $\mathcal{F}$ by a smooth function of r which goes to zero at some distance away from the star. In this thesis I set the maximum radius of the flooring scheme to be $r = 20\, r_g$.

1.5.3 Boundary conditions

The last numerical concept to elaborate on before moving to the main results of the thesis are the boundary conditions. In the absence of infinite grids, boundary conditions are necessary for all grid-based numerical schemes. All of the simulations presented in this thesis use spherical coordinate systems which means that we have six boundary conditions to consider, a lower and upper boundary for each coordinate. All but the lower radial boundary condition are the same for every simulation described here.

First, I will mention the azimuthal boundary conditions as they are the simplest. In 2D axisymmetry, we have only one grid cell in the ϕ direction and so we simply copy the entire simulation into the ghost cells in either direction.

Next are the upper and lower boundary conditions on the polar angle. In 2D axisymmetry it is impossible for any vector to have a nonzero θ component at $\theta = 0$ and $\theta = \pi$. Furthermore, the factor of $\sin(\theta)$ originating from the metric but used in many different places in the code can causes numerical problems as $\sin(\theta) \to 0$ near the poles. For this reason, the first two cells in the domain around the poles are overwritten. This makes the boundary condition largely irrelevant. All of the scalar quantities and r and ϕ components of the vector quantities are copied from the 3rd closest cell in the domain and the θ components are interpolated to zero.

The upper and lower radial boundary conditions are the most complicated. First we discuss the outer radial boundary condition. The idea here is for gas, radiation, and magnetic field to flow out of the simulation. This is done by extrapolating all of the scalar quantities and radial components of the vector quantities such that $r^2 P^i$ is conserved across the boundary where P^i is the primitive quantity of interest. The polar and azimuthal quantities are extrapolated such that rP^i is conserved. If either the gas or radiation radial velocity are negative (inflowing), then they are set to zero in the ghost cells.

Inner radial boundary for the non-magnetized neutron star

The inner radial boundary condition is the most important for these simulations as we are studying accretion onto compact objects, this is where the properties of those compact objects must be encoded. Chapter 3 concerns accretion onto a black hole and a non-magnetized neutron star. In the black hole case, the boundary condition is simple. As long as there are enough cells below the event horizon (5), you can put almost anything as the inner boundary and nothing will escape the black hole. This is purely a consequence of the geometry of space-time. One has to keep in mind that it is not possible to use the Schwarzschild metric here, it is necessary to change to horizon-penetrating coordinates such as Kerr-Schild coordinates.

For the non-magnetized neutron star we model the hard-surface created by the crystal lattice as a reflective boundary. From a numerical point of view, the idea is for the radial fluxes excluding radial momentum to be zero at the inner boundary. This is usually achieved by copying all quantities symmetrically across the boundary, and multiplying the radial velocity by negative one. This causes the reconstructed left and right radial fluxes of mass and energy at the boundary to be antisymmetric about zero so that the Riemann solver produces zero flux for these quantities. Momentum flux is not zero because momentum needs to be transferred from the gas to the star. Because the grid used in Chapter 3 is not constant in radius, I extrapolated the radial velocity so that linear interpolation between the cell center would produce zero at the boundary. Otherwise the fluxes would not line up properly and there was leakage of gas into the domain.

In order to make the stellar surface "sticky" i.e., to give it some friction so that the gas does not just slide along the surface indefinitely, I simply set the tangential velocities to zero. This allows some angular momentum to flow through the boundary, but it is not really physical as the rate of momentum transfer is related to not only the difference in momentum flux between cells, but also to numerical properties such as the resolution, the reconstruction scheme, and the Riemann solver. I believe however, that this is still the most reasonable choice, as it would be very difficult to predict the actual rate of angular momentum transfer without detailed knowledge of the physics of the neutron star crust and its interaction with accreting plasma. The results of this simulation are described in detail in Chapter 3.

Inner radial boundary for the magnetized neutron star

In Chapter 4, I move onto magnetized neutron stars, and this requires me to include some more physics into the boundary condition. For gas quantities, I use the same inner boundary condition as described in Parfrey and Tchekhovskoy (2017), which allows gas to accrete smoothly through the inner boundary while sliding along magnetic field lines.

The main idea is to extrapolate quantities smoothly into the ghost cells using slope-limited linear extrapolation in cells where gas is accreting onto the star, and to set the gas quantities to their force-free background states otherwise.

For accreting cells (defined by cells with negative u^r and $\mathcal{F} < 0.5$), the gas density and internal energy are directly extrapolated, but again the velocity is treated differently. However, before the velocity can be extrapolated, it is necessary to set the magnetic field in the ghost cells. This is done by first setting the radial component to the dipole solution (Wasserman and Shapiro 1983) and then smoothly extrapolating the tangential components. This allows field lines to smoothly pierce the stellar surface but also to return to the equilibrium configuration if the whole field is deformed (such as by accreting gas). It is also necessary to set the tangential electromotive forces used in the flux-CT algorithm to prevent field lines from sliding along the surface.

Again, the component of the velocity parallel to the magnetic field as viewed by the stationary observer[11], $u^\mu b_{s,\mu}$ is measured and this quantity extrapolated into the boundary, and the velocity reconstructed using the value of the magnetic field in the ghost cells.

In the absence of radiation, the gas would accrete onto the neutron star without experiencing any effects from the surface as seen in Parfrey and Tchekhovskoy (2017). In order to restore the surface I implement what I call an "energy-reflective" boundary. First I measure the inflowing flux of kinetic, thermal, and radiative energy, then I return a fraction of that energy as outflowing radiation. The point is to mimic the effect of gas hitting the surface and releasing its energy as radiation. Because it is

---

[11] In Parfrey and Tchekhovskoy (2017) the magnetic field of the corotating observer is used for the extrapolation, but as our neutron stars are not rotating, the corotating and stationary observer are the same.

Name	Albedo (β)
Accr	0
Enre060	0.6
Enre065	0.65
Enre070	0.7
Enre075	0.75
Enre085	0.85
Enre095	0.95
Enre100	1

Table 1.1: Simulation name and corresponding albedo.

unknown exactly how much energy is reflected and how much is buried underneath the star by the accretion flow, I study a range of this reflective fraction, or albedo from 0% to 100%.

I will conclude this introduction by presenting the results of a boundary condition study which allow me to be comfortable with my choice of boundary condition in Chapter 4.

1.5.4 Results of the boundary condition study

In order to study the effects of the albedo, β, I run a series of eight simulations with albedos described in Table 1.1.

All the runs are initialized using the setup described in Section 1.5.1. The initial torus is set so that the accretion rate is approximately, $\dot{M} \approx 20\dot{M}_{\rm Edd}$ and the initial stellar dipole magnetic field set to have a maximum value on the stellar surface of $B_{\rm max} = 2 \times 10^{10}$ G. The simulations are then run for $t_{\rm max} = 40\,000\,t_g$.

Accretion disk and column structure

Time averages from $t = 20\,000 t_g - 40\,000\,t_g$ are used for most of the analysis. Colormaps of the gas rest mass and radiation energy densities are shown in Fig. 1.1. I will start by describing run `Accr`, the zero albedo set-up.

As the simulation starts, MRI builds up in the initial torus causing an accretion disk to form and matter to flow towards the star. As the gas flows inwards, it compresses and reconnects with field lines from the neutron star until it reaches the point where magnetic pressure from the dipole field matches the inflowing ram pressure. At this location, called the Alfvén surface which is located at a radius of $r_a \approx 10\,r_g$, the gas is forced to slide along magnetic field lines, forming accretion columns which deposit the inflowing matter near the poles of the star. In the case of `Accr`, the gas simply flows into the star without interruption. This is evident from Fig. 1.1 since no shocks are visible at the base of the accretion column. There is a little bit of radiation which is able to diffuse out the sides of the accretion column and escape along the poles, which are completely devoid of gas.

As the albedo increases, the gas experiences a shock at the base of the accretion column as the inflowing gas is slowed down by outflowing radiation. The height of this shock increases with albedo. It is also possible to notice an increasing amount of radiation which is escaping along the poles. This is due to the increasing shock height. As the gas spends more time sinking through the dense radiation field, more photons are able to diffuse through the accretion column and escape along the poles.

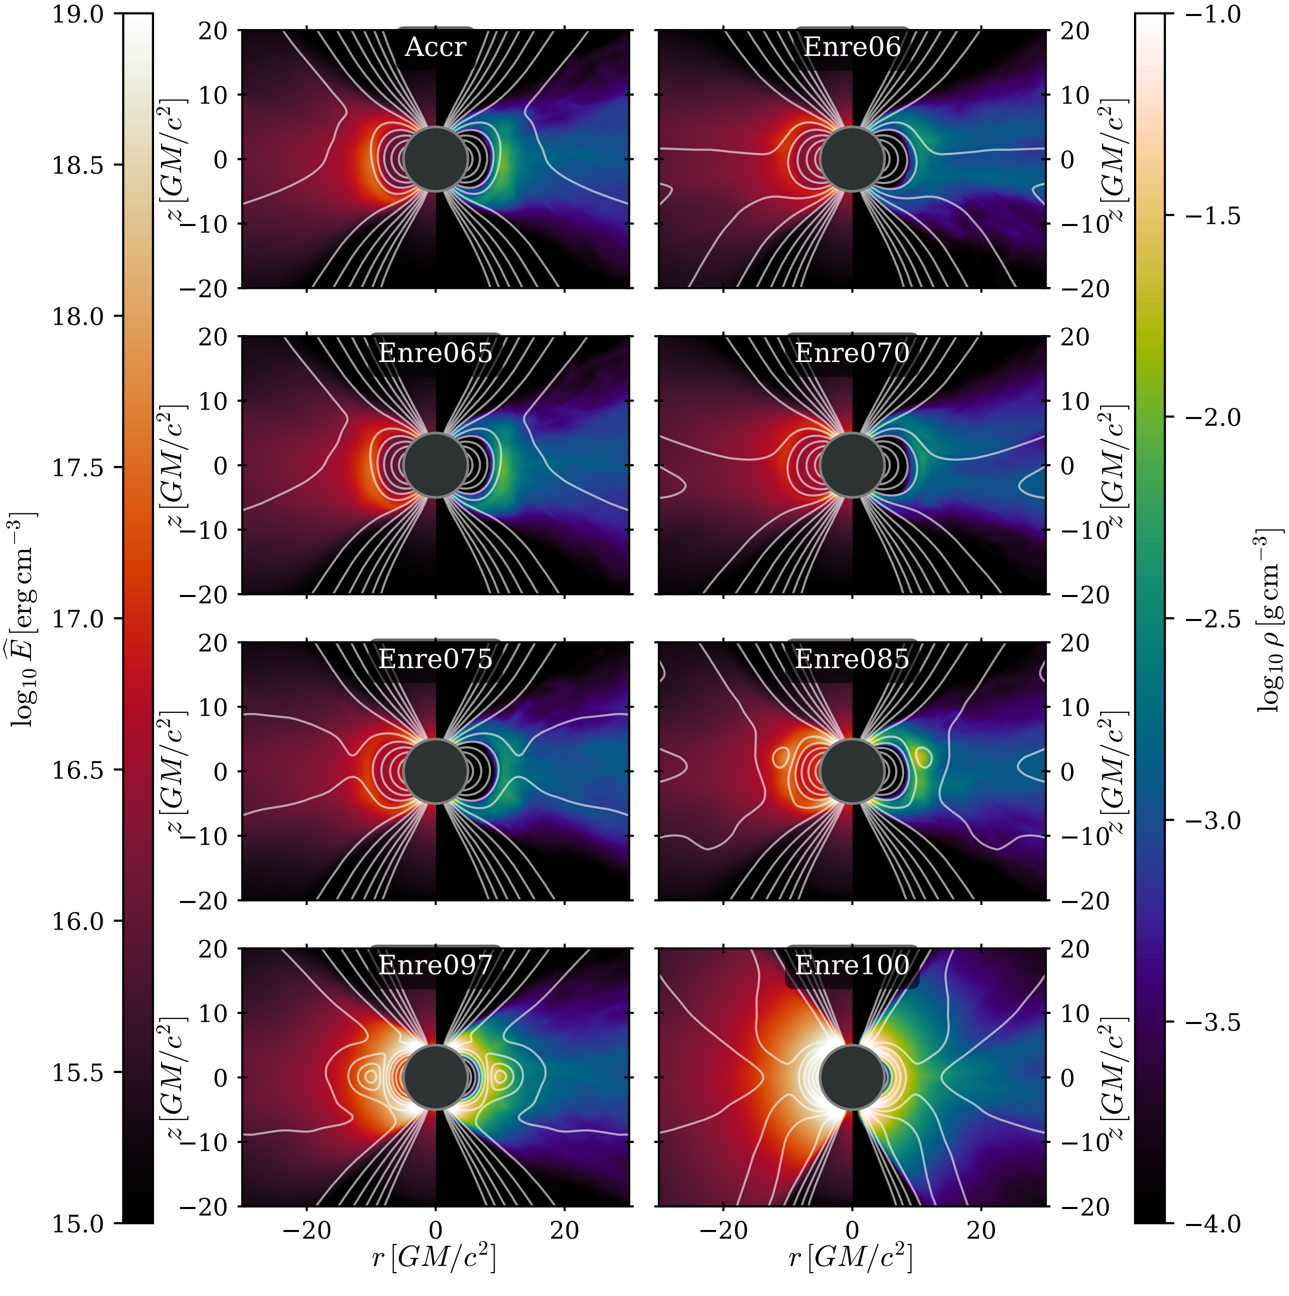

Figure 1.1: Colormaps of radiative energy density and rest mass density for the eight simulations in the boundary condition study. The colormaps are produced from time-averaged data spanning a length of 20 000 t_g. Also plotted are poloidal magnetic field lines.

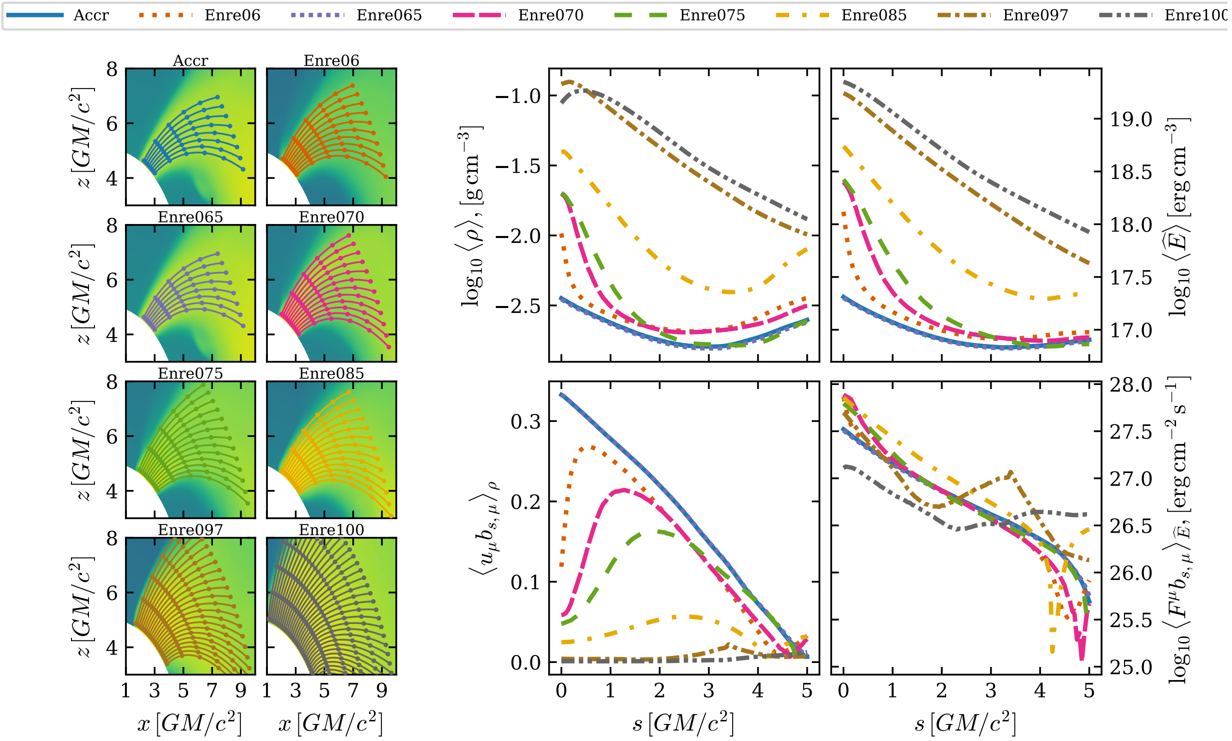

Figure 1.2: Quantities along magnetic fields are shown in the four right panels. The quantities are interpolated onto paths which follow magnetic field lines and then averaged together. The paths are shown in the eight panels on the left. From upper-left to lower-right the four quantities plotted are the gas density, radiation energy density, the component to of four-velocity along the magnetic field line, the component of the radiative flux along the magnetic field line.

Somewhere between an albedo of 85% and 95%, the radiative shock propagates all the way into the disk and a quasi-spherical atmosphere starts to form, somewhat similar to the simulation I will present in Chapter 3, although the magnetic field keeps the gas from filling in the polar regions.

To get a better understanding of the structure of the accretion columns, I plot quantities along magnetic field lines shown in Fig. 1.2. Starting at the surface of the star, for each cell in the θ direction which is inside the column, I integrate a path which follows the magnetic field. I then interpolate relevant scalar quantities onto the paths and then average them together. The paths for each simulation are shown in the eight left panels of Fig. 1.2. Units of $1\,r_g$ are marked along the paths and those points correspond to the tick marks for the horizontal axes of the right four panels. I plot four quantities against s, the distance along the path from the surface of the star, $\langle\rho\rangle$, $\langle\widehat{E}\rangle$, $\langle u^\mu b_{s,\mu}\rangle_\rho$, and $\langle F^\mu b_s\mu\rangle_{\widehat{E}}$. The gas and radiation energy densities are computed via direct average, the velocity along the field line is computed with a density-weighted average, and the radiative flux along the field line is computed with a radiation-energy-density-weighted average.

The plot of density in Fig. 1.2 for `Accr` starts with an initial decrease at the top of the column as the gas begins to accelerate towards the star, and then gently rises again as it is compressed by the magnetic field. There is no indication of an interaction with the inner boundary. As the albedo increases the behavior is similar except for a sharp rise in density near the surface of the star. This corresponds to the radiative shock caused by gas hitting the surface of the star and slowing down. The shock rises with albedo and reaches farther up the column, which was also visible in Fig. 1.1. The

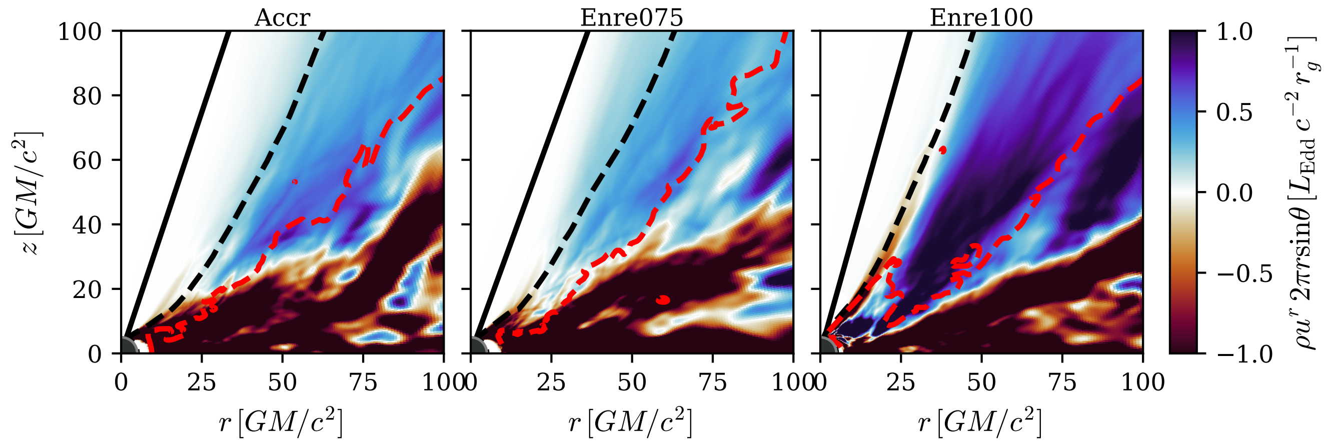

Figure 1.3: Colormap showing the radial mass flux density integrated into the poloidal plane. Blue colors show outward moving gas, and red colors show inflowing gas. The solid lines represent the photosphere as integrated radially from the outer boundary. The dashed lines represent the photosphere as integrated along θ from the pole. The dashed red line is the zero Bernoulli surface.

slight dips in density at the surface for the two highest albedo runs are numerical in nature but they do not affect the results significantly.

The same trend is also visible in the plots of radiation energy density, this is just a consequence of local thermal equilibrium as the accretion columns are optically thick.

The shock is also visible in the plot of the component of velocity along the magnetic field line. For `Accr`, the velocity increases smoothly towards surface of the neutron star. As albedo increases, the velocity as the surface is pushed towards zero and the rest of the column has to transition between free-fall to the shock region.

In general, the shocks present themselves as being rather smooth. One thing to keep in mind is that I am showing the data from time-averages. In actuality, the accretion rate is quite variable causing the extent of the shock to move up and down so it gets smeared out in the time-average.

The last panel shows the component of the radiative flux along the magnetic field. It is harder to make sense of this plot. Initially we can see the flux rise near the surface for lower albedos, but at some point it begins to drop again. This is because there is a significant amount of outgoing flux at higher albedos which cancels out with the inflowing flux.

Outflows

As discussed previously, outflows play a large role in super-Eddington accretion and the appearance of its emerging radiation. I pick three simulations ($\beta = 0, 0.75, 1$) and plot the azimuthally integrated mass flux which is shown in Fig. 1.3. The blue part of the colormap shows outflowing gas and the red shows inflowing gas. Also plotted are three surfaces. The two black lines correspond to the photosphere as measured radially (solid) or along θ from the axis (dashed). These surfaces give us a sense of how the radiation should be collimated. All the outflowing radiation which escapes the gas should be collimated above the dashed line, however, not all of that radiation is expected to reach the observer. Some of it will be absorbed by the very thin outflowing gas. All the radiation above the solid black line however should manage to reach the observer. Also shown is the zero specific energy

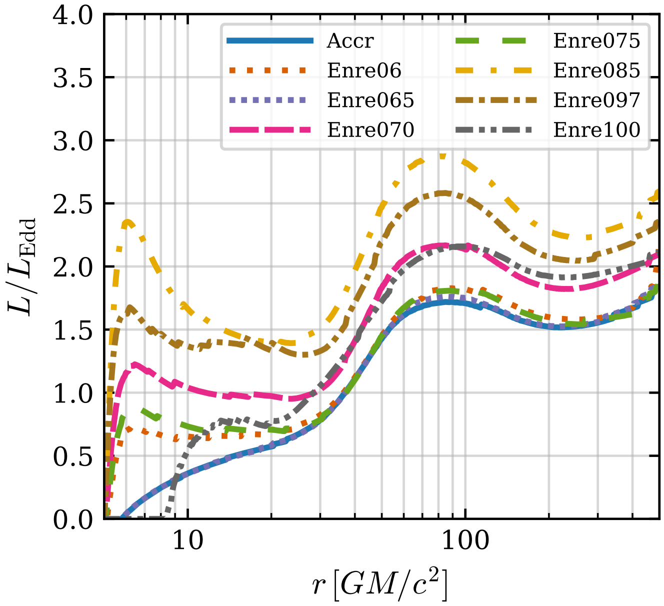

Figure 1.4: Luminosity as a function of radius.

surface, or the surface where the relativistic Bernoulli parameter is equal to zero. How these surfaces relate to the luminosity is explained in more detail in Chapter 4.

The two lower albedo runs look similar. There is a bit more gas flux between the two photospheres for the 75% albedo run but it does not appear that there would be a significant difference in collimation. The 100% run is drastically different. Very little gas can escape through the boundary condition and so when it hits the NS surface it produces lots of photons which drive a much stronger outflow. There is a noticeable difference in the volume confined by the photosphere and so it is expected that the 100% albedo simulation would collimate the outflowing radiation more effectively.

1.5.5 Radiative output

The question that is most important to answer is can weakly magnetized neutron stars produce luminosities comparable to those observed in pulsating ULXs. In order to answer this question, it is necessary to know the effect the boundary condition has on the radiative output. Once this information is known, it is possible to chose one albedo to study in more detail, which is what I have presented in Chapter 4.

The first quantity to measure is the total radiative output or luminosity of the source. This is not so interesting from an observational point of view but more so for understanding the physics of the problem as it depends not only on how much radiation is generated from the accretion flow, but also, how much can escape to infinity. This is very difficult to measure as these simulations are only performed for the inner most part of the system. If I had infinite computing power I could probably simulate nearly the entire system, then a measure of the luminosity would be trivial. Instead, I have to make some approximations.

As the simulation begins, the accretion disk is formed near the center and with time it grows outwards as the simulation reaches its quasi-steady or equilibrium state. This is closely related to inflow equilibrium in which the time-averaged plot of $\dot{M}(r)$ becomes constant. The speed at which the disk grows is roughly proportional to the viscous time which increases exponentially with radius, so usually I am only able to simulate the disk out to between $50-100\,r_g$ reliably, so the first approximation is that the majority of the radiation is produced from this range.

The collimation of the radiation is achieved through the outflows, so it is also necessary to ensure that they are converged to their steady state solution as well. Since the outflowing gas has much higher velocities, the outflows are usually trustworthy throughout most of the domain as long as they can be traced back to portions of the disk which are converged.

The last approximation involves the decision of which photons will reach the observer or not. There are several ways to approach the problem and I discuss them in detail in Chapter 4. The method I use here corresponds to integrating the radiative flux over values of θ where the optical depth measured radially to the outer edge of the simulation is less than one. This corresponds to all of the flux which is traveling above the solid black lines in Fig. 1.3.

The result of this measurement is shown in Fig. 1.4. In steady state and in the absence of gas (such as the exterior of a star) the luminosity should be constant with radius. Its obvious from the figure that this is not the case. There are two main contributions to spatial variations here. One is that the outflows are absorbing and emitting radiation so they can decrease and increase the luminosity with radius. The second, is that the surface defining the photosphere has some curvature so that radiation is either included or excluded from the total luminosity. If I were to instead integrate the flux around the poles up to an angle to $10°$, the luminosity plots would appear nearly flat.

From Fig. 1.4 we can see that the luminosity varies the most near stellar surface. This makes sense because the only differences between the simulations are from the boundary conditions, so farther away from the center those are masked to a certain extent by the disk contribution. All the plots show a peak around $r = 50\,r_g$. It is at this point that the outflows stop emitting radiation and start absorbing radiation so the luminosity decreases until around $r = 250\,r_g$. Here the outflows start to become optically thin and so photons which were earlier excluded are now included in the luminosity. At $r = 500\,r_g$, which is the largest radius I include, the luminosities vary between 1.75 $L_{\rm Edd}$ and 2.5 $L_{\rm Edd}$. This is encouraging that difference in luminosities is less than a factor of two and likely means any boundary condition should get close to correct answer.

The last quantity to measure is much simpler to compute, and arguably much simpler to understand from an observational perspective and that is the inferred isotropic luminosity, $L_{\rm iso}$. As explained in Section 1.4.3, $L_{\rm iso}$ is the luminosity inferred from a flux measurement by assuming that the source emits isotropically. If the source is not isotropic then $L_{\rm iso}$ overestimates the true luminosity, L, by a factor for $1/b$ where b is the beaming factor. To measure $L_{\rm iso}$ we simply have to multiply the observed flux by the surface area of the sphere with radius corresponding to the location the measurement is occurring. In real observations, this corresponds to the distance between the source and the Earth, but for simulations we can perform this measurement at any radius.

I perform this measurement at $r = 500\,r_g$ and the result as a function of the observation angle is shown in Fig. 1.5. We can see that almost all of the emission is concentrated between 0 and 20°. The largest isotropic luminosities are seen when viewing the source along the pole. The peak luminosity is about 140 $L_{\rm Edd}$ which is many times the measured value of L implying very strong beaming. This is encouraging as these luminosities are consistent with those observed in ULXs. These results show that it is likely possible for weakly magnetized NSs to power ULXs even when the interaction between the

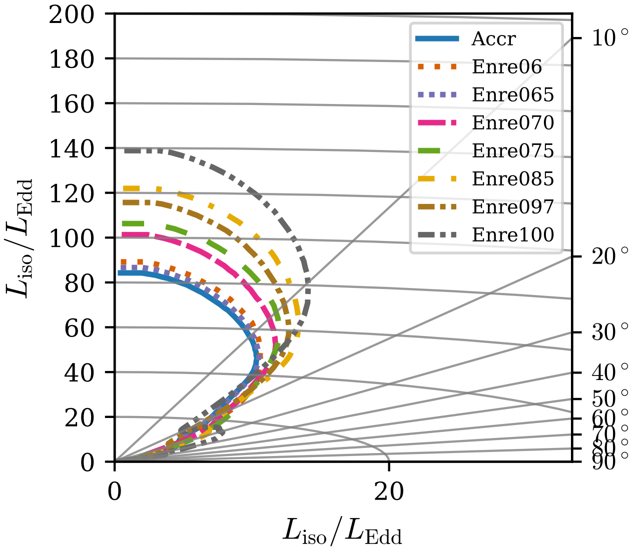

Figure 1.5: Inferred isotropic flux as a function of viewing angle measuring at radius $r = 500\,r_g$. The plot is presented in a polar fashion so that the polar coordinate present the viewing angle and the radial coordinate represents the luminosity.

gas and the neutron star surface is unknown. For this reason we pick an albedo of $\beta = 0.75$ to study in more detail in Chapter 4.

1.6 Summary

Thus far, in Section 1.1 I laid the foundation to understand the physics necessary to study accreting neutron stars, namely GRRMHD. In addition, I described a special case in which radiation and general relativity give an interesting result not possible in the Newtonian formulation, the existence of the Eddington capture sphere. In Section 1.2, I presented an overview of the numerical techniques used in the rest of the thesis in order to solve the equations of GRRMHD on a computer. I also described some of the specifics of the GRRMHD code `Koral`, the code used to run the simulations in this work. Later, in Section 1.3, I moved the discussion to astrophysics. Here I described both classical and contemporary models of accretion disks including super-Eddington accretion. Accretion disks can occur around a number of objects, such as black holes and neutron stars in X-ray binaries. In Section 1.4 I focused mainly on accreting neutron stars and ULXs, the main astrophysical object I am concerned about. In Section 1.5, I described the set-up of the simulations in a level of detail which is usually not present in publications so that the reader could have a deeper understanding of what was done to obtain the results. I also presented some unpublished results which help to provide context for the simulation described in Chapter 4.

The three remaining chapters consist of three publications, the first of which describes an analytical study and the last two involve GRRMHD simulations. The analytical study in Chapter 2 was included because it described an interesting phenomenon which can only be present around a near-Eddington luminosity neutron star. Abarca and Kluźniak (2016) was written with Type I X-ray bursts in mind but it is not unreasonable to consider the situation could be applied to ULXs with neutron star accretors.

The subject of the paper considered oscillations of an atmosphere around a neutron star which is supported against gravity by radiation, completely disconnected from the stellar surface. The main result of the paper showed that the lowest order mode was overdamped by radiation drag. While we were unable to show that QPOs in Type-1 X-ray bursts could be explained by these oscillations, we set the stage for additional publications on the topic which produced interesting results.

The second paper, Abarca et al. (2018) included in Chapter 3, involves simulations of super-Eddington accretion onto a black hole and a non-magnetized neutron star. While much of the debate around pulsating ULXs relates to the strength of the magnetic field, the goal of this publication was to investigate the possibility of having no magnetic field. There was a hard surface implemented for the neutron star case, however. What the simulations showed was that gas collected on the surface of the neutron star and filled the domain so that any radiation released by the accretion disk hitting the surface was trapped. The radiation which was able to escape was nearly isotropic and around one Eddington luminosity which does not resemble a ULX in any sense, but possibly a super soft source. Including this paper makes it possible to disentangle the effects of including a magnetic field and a hard-surface into an accreting NS simulation.

The third and last paper is probably the most interesting result, the most difficult to obtain, and represents the culmination of my PhD studies. In Abarca et al. (2021) I implemented the technique described in Parfrey and Tchekhovskoy (2017) to work in GRRMHD so that we could perform simulations of super-Eddington accretion onto magnetized neutron stars. The resulting simulations showed that accreting weakly-magnetized ($\sim 10^{10}$ G) neutron stars could produce highly beamed emission which would appear to the observer to be in excess of 10^{40} erg s^{-1}. And while it does not answer the question of exactly how PULXs operate, it clearly supports the beaming scenario and shows that super-high magnetic fields are not necessary for producing large luminosities.

2 Paper 1: Radial oscillations of a radiation-supported levitating shell in Eddington luminosity neutron stars

In this paper (Abarca and Kluźniak 2016), I extended the results in Wielgus et al. (2015) to include a first order, linear perturbation analysis to study oscillations about the equilibrium configuration. This paper serves as an illustration of the necessity of including radiation in GRHD to properly encompass the physics around accreting neutron stars and so it helps to set the stage for the next two chapters. They main result I found was that the eigenfunction of the lowest frequency corresponded to a radial, incompressible mode. The frequency of these oscillations is consistent with the 300-600 Hz QPOs seen in several X-ray bursting low-mas X-ray binaries (Strohmayer et al. 1996). However, when the full effects of radiation drag were included in the calculation, the oscillations were found to be over-damped. Nevertheless, this work laid the foundation for future studies which considered higher order modes in both in Newtownian gravity (Bollimpalli and Kluźniak 2017) and full radiative GRHD (Bollimpalli et al. 2019) which showed that the underdamped second and third order modes could provide simultaneous mass and radius measurements for neutron stars.

Monthly Notices
of the
ROYAL ASTRONOMICAL SOCIETY

MNRAS **461,** 3233–3238 (2016)
Advance Access publication 2016 June 15

doi:10.1093/mnras/stw1432

Radial oscillations of a radiation-supported levitating shell in Eddington luminosity neutron stars

David Abarca★ and Włodek Kluźniak★
Nicolaus Copernicus Astronomical Center, Bartycka 18, Warsaw PL-00-716, Poland

Accepted 2016 June 13. Received 2016 June 13; in original form 2016 April 15

ABSTRACT

In general relativity, it has been shown that radiation-supported atmospheres exist well outside the surface of a radiating spherical body close to a radius where the gravitational and radiative forces balance each other. We calculate the frequency of oscillation of the incompressible radial mode of such a thin atmospheric shell and show that in the optically thin case, this particular mode is overdamped by radiation drag.

Key words: stars: atmospheres – stars: neutron – X-rays: bursts.

1 INTRODUCTION

Neutron stars have been shown to erupt in thermonuclear (Type I) X-ray bursts. In addition, pulsations of an ultraluminous X-ray source have been explained by accretion onto a neutron star (Bachetti et al. 2014). At such large (∼100 times Eddington) luminosities, it is easy to imagine that neutron star systems may produce luminosities above the Eddington limit.

In Newtonian physics, the radiative force is proportional to the flux, which falls off as $1/r^2$ for a spherically symmetric source. If the radiative force exceeds the gravitational force at one radius, then it will exceed it at all radii. In general relativity, the radiative force can be shown to increase faster than the gravitational force at smaller radii (Phinney 1987; Abramowicz, Ellis & Lanza 1990). It turns out that for a given luminosity, there exists a radius in the Schwarzschild metric where the gravitational and radiative forces are equal, forming an imaginary, spherically symmetric surface referred to as the Eddington Capture Sphere (ECS; Stahl et al. 2012; Wielgus et al. 2012) onto which particles are captured by radiation drag (Bini, Jantzen & Stella 2009; Oh, Kim & Lee 2010). Wielgus et al. (2015) have shown that it is possible to create an optically thin atmosphere at this radius which levitates above the surface of the star, supported entirely by radiation. Wielgus et al. (2016) have extended the analysis to include optically thick atmospheres as well.

We are interested if oscillations of these atmospheres can explain the still unresolved problem of the source of quasi-periodic oscillations (QPOs), the transient peaks observed in the power spectrum of highly compact sources (Remillard & McClintock 2006), including neutron stars (van der Klis 2006). Specifically in X-ray bursting neutron stars, there have been observations of hectoHertz QPOs; Strohmayer (2001) report on several QPOs from X-ray bursts from low-mass X-ray binaries all with frequencies between 300 and 600 Hz. Moreover, most sources show an increase of frequency with time during the decay phase of the X-ray burst (Strohmayer 2001; Strohmayer & Bildsten 2006).

We investigate the possibility of oscillations of atmospheres around the ECS which could possibly provide an explanation for X-ray burst hectoHertz QPOs. We begin by finding an incompressible radial oscillation mode and calculating its oscillation frequency. We then compute the effects of radiation drag and discuss the viability of such a mode as a model of a QPO.

2 EIGENMODE OF AN INCOMPRESSIBLE THIN SHELL

First we demonstrate the equations that are used to construct the atmospheres from Wielgus et al. (2015). This is essentially the relativistic equation for hydrostatic equilibrium for an optically thin fluid subject to the radiative force from a spherical source. We also explicitly include the derivation of the ECS from Stahl et al. (2012), because the equations involved are also used to calculate the frequency of oscillations of a thin shell about the ECS, as well as the contribution of radiation drag.

2.1 Relativistic hydrostatic equilibrium

For this relativistic calculation, we use the Schwarzschild metric,

$$ds^2 = -B(r)dt + B(r)^{-1}dr + r^2 d\Omega^2, \quad (1)$$

where $B(r) = 1 - 2r_g/r$ for $r_g = GM/c^2$. We find it convenient to use units where $G = c = 1$.

Let us consider the equation for the conservation of stress-energy given by

$$\nabla_\mu T^{\mu\nu} = 0, \quad (2)$$

where $T^{\mu\nu}$ is the stress-energy tensor for a perfect fluid given by, $T^{\mu\nu} = (p + \rho + u)u^\mu u^\nu + pg^{\mu\nu}$, for pressure, p, rest mass density, ρ, and internal energy density, u. We can project the conservation equation on to the space orthogonal to the four-velocity using the

---

★ E-mail: dabarca@camk.edu.pl (DA); wlodek@camk.edu.pl (WK)

Published by Oxford University Press on behalf of the Royal Astronomical Society

projection tensor, $h^{\mu\nu} = u^{\mu}u^{\nu} + g^{\mu\nu}$, to get the relativistic Euler equation,

$$u^{\mu}\nabla_{\mu}u^{\nu} + \frac{h^{\mu\nu}\nabla_{\mu}p}{p+\rho+u} = f^{\mu}, \tag{3}$$

where we have also added a four-force, f^{μ}, which corresponds to the radiation force due to Thomson scattering with cross-section, σ, in an optically thin fluid around a luminous star. As in Wielgus et al. (2015), we assume that the fluid is at a sufficiently low temperature such that $\rho \gg u + p$. This allows us to write the four-force in terms of the radiative flux, F^{μ}, as

$$f^{\mu} = \frac{\sigma}{m}F^{\mu}. \tag{4}$$

Here, m is the proton mass.

One can construct an atmosphere with its pressure maximum located at, r_0, the radius of the ECS. These atmospheres obey the equation of hydrostatic equilibrium which can be calculated by substituting $u^{\mu} = u^{t}(1, 0, 0, 0)$ into the relativistic Euler equation, which gives

$$\frac{1}{\rho}\frac{\partial p}{\partial r} = -\frac{M}{1-2M/r}\left(\frac{1}{r^2} - \frac{f^r}{M}\right). \tag{5}$$

2.2 Derivation of the equilibrium surface

Let us now demonstrate a quick derivation of, r_0, the radius of the ECS. See Stahl et al. (2012, 2013), Phinney (1987) and Abramowicz et al. (1990) for original work on the ECS. We start with the expression in parentheses in equation (5), which we will eventually set to zero.

For convenience, we name it $\mathcal{F}(r, 0) = 1/r^2 - f^r/M$, the reason for which will be explained in the next section. In terms of the flux, F^r we have

$$f^r = \frac{\sigma}{m}F^r. \tag{6}$$

Expressions for the flux can be found in Stahl et al. (2013). For a stationary particle, we have

$$F^r = T^{tr}u_t. \tag{7}$$

There T^{tr} is the rt component of the radiation stress-energy tensor, $T^{\mu\nu}$, outside a luminous star, first derived in Abramowicz et al. (1990). We have

$$T^{rt} = \pi I \sin^2\alpha, \tag{8}$$

for intensity, I, and viewing angle, α, defined as

$$\alpha = \arcsin\frac{R}{r}\frac{\sqrt{1-2M/r}}{\sqrt{1-2M/R}}, \tag{9}$$

where R is the radius of the star.

We can write the intensity in terms of the luminosity at infinity, L_{∞},

$$I = \frac{L_{\infty}}{4\pi^2R^2}\frac{1-2M/R}{(1-2M/r)^2}. \tag{10}$$

Putting all of this back into the expression for the acceleration and using the usual expression for the Eddington luminosity, $L_{\mathrm{Edd}} = 4\pi M m/\sigma$, we get

$$f^r = \frac{L_{\infty}}{L_{\mathrm{Edd}}}\frac{M}{r^2\sqrt{1-2M/r}}. \tag{11}$$

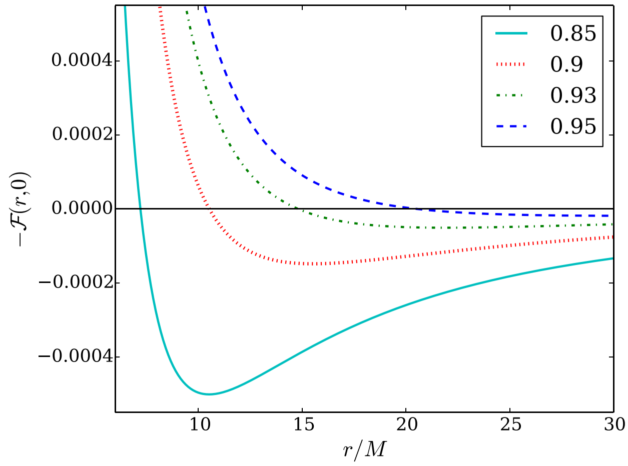

Figure 1. Plots of $-\mathcal{F}(r, 0)$ for various values of λ shown in the legend. The location where $\mathcal{F}(r, 0) = 0$ shows the radius of the ECS. $-\mathcal{F}(r, 0)$ shows the forces acting on a stationary fluid, and so one can see how atmospheres can be constructed around the ECS.

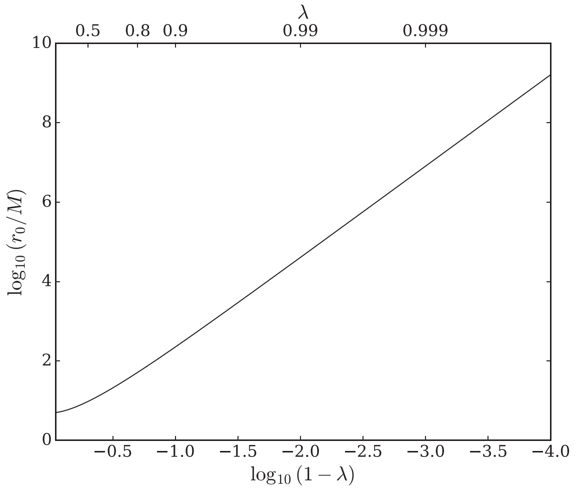

Figure 2. We have plotted r_0 as a function of λ. To illustrate the sensitivity of r_0 to λ, we have plotted against two different axis scales. The upper is λ and the lower shows $\log(1-\lambda)$.

For brevity we define, $\lambda = L_{\infty}/L_{\mathrm{Edd}}$. We can write down the equation for $\mathcal{F}(r, 0)$ for a stationary fluid,

$$\mathcal{F}(r,0) = \frac{1}{r^2}\left(1 - \frac{\lambda}{\sqrt{1-2M/r}}\right). \tag{12}$$

Plots of $\mathcal{F}(r, 0)$ for different λ are shown in Fig. 1. We note the location of the ECS to be where $\mathcal{F}(r_0, 0) = 0$, this gives

$$r_0 = \frac{2M}{1-\lambda^2}. \tag{13}$$

It is important to note that r_0 is extremely sensitive to λ as $\lambda \to 1$. This is further illustrated in Fig. 2.

2.3 Atmospheric solution

The equation for hydrostatic equilibrium, equation (5) was solved in Wielgus et al. (2015), who took an optically thin fluid (again, at a sufficiently low temperature, such that $\rho \gg p + u$), and derived atmospheric solutions with their pressure maximum located at radius r_0. They have provided a set of polytropic atmosphere solutions with polytropic index, Γ, given by

$$\rho(r) = \rho_0 \left[\ln(1 - 2M/r)^{-1/2} - \lambda(1 - 2M/r)^{-1/2} + \lambda\right]^{1/\Gamma - 1}. \tag{14}$$

Such an atmospheric shell is entirely supported by radiation and there are regions between the atmosphere and the surface of the star with no gas at all.

3 RADIAL EIGENMODE OF AN INCOMPRESSIBLE THIN SHELL

Now we turn from hydrostatic equilibrium to the time-dependent equation, the relativistic Euler equation, with which we will apply perturbations to derive the oscillation frequency. While we are in the relativistic regime of strong gravity, we will make the assumption that our velocities are small and so non-relativistic. This allows us to write the four velocity, u^μ, to first order in $u^r = \mathrm{d}r/\mathrm{d}\tau$, or in $v = \mathrm{d}r/\mathrm{d}t$.

$$u^\mu = (u^t, u^r, 0, 0) = u^t(1, v, 0, 0) \approx \frac{1}{\sqrt{1 - 2M/r}}(1, v, 0, 0). \tag{15}$$

This allows us to simplify equation (3) to get,

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}\tau}u^r + \frac{1}{\rho}\left(g^{rr}\frac{\partial p}{\partial r} + u^r\frac{\mathrm{d}p}{\mathrm{d}\tau}\right) &= -M\left(\frac{1}{r^2} - \frac{f^r}{M}\right) \\ &= -M\mathcal{F}(r, u^r),\end{aligned} \tag{16}$$

We have combined the two non-fluid forces into one expression, $\mathcal{F}$, which is in general a function of r and u^r. We will demonstrate that $\mathcal{F}$ can be divided into the sum of two terms, one of which is only a function of r, the other a function of both r and u^r. The former corresponds to the radiation force, and the latter to radiation drag, both of which will play an important role in our analysis.

We let equation (5) serve as the background over which we consider spherically symmetric radial perturbations. In this work, we will consider an incompressible mode, where the whole atmosphere is transported by a small radial distance, ξ, while preserving its pressure and density profiles, $p(r) \to p_\mathrm{b}(r - \xi)$, and $\rho(r) \to \rho_\mathrm{b}(r - \xi)$, where the b index indicates the background solutions. We also have, $u^r = \mathrm{d}\xi/\mathrm{d}\tau$. We will demonstrate that for a sufficiently thin atmosphere, we will recover a radial eigenmode that oscillates with the same frequency as a test particle about r_0. This gives us the following two equations to solve,

$$\frac{\mathrm{d}^2\xi}{\mathrm{d}\tau^2} + \frac{g^{rr}(r)}{\rho_\mathrm{b}(r - \xi)}\left(\frac{\partial}{\partial r}p_\mathrm{b}(r - \xi)\right) = -M\mathcal{F}(r, u^r), \tag{17}$$

$$\frac{g^{rr}(r)}{\rho_\mathrm{b}(r)}\frac{\partial p_\mathrm{b}(r)}{\partial r} = -M\mathcal{F}(r, 0), \tag{18}$$

where we have been explicit with the r dependence for clarity. The system can be simplified to one equation after expanding in terms of ξ, and substituting equation (18) into equation (17) to get,

$$\frac{1}{M}\frac{\mathrm{d}^2\xi}{\mathrm{d}\tau^2} - \mathcal{F}(r, 0) + \frac{\partial \mathcal{F}(r, 0)}{\partial r}\xi - \mathcal{F}(r, 0)\frac{\partial \ln g^{rr}(r)}{\partial r}\xi = -\mathcal{F}(r, u^r). \tag{19}$$

At this point we invoke the thin shell limit. We assume that we have a thin atmosphere concentrated around, r_0. This allows us to expand $\mathcal{F}$ around $r = r_0$ and $u^r = 0$ to get,

$$\mathcal{F}(r, u^r) \approx (r - r_0)\left.\frac{\partial \mathcal{F}}{\partial r}\right|_{r_0,0} + u^r\left.\frac{\partial \mathcal{F}}{\partial u^r}\right|_{r_0,0}, \tag{20}$$

where we have used, $\mathcal{F}(r_0, 0) = 0$, which we recall from Section 2.2. Substituting in the expansion and noting that $(r - r_0)\xi$ is second order in small quantities, we are left with the following equation,

$$\frac{1}{M}\frac{\mathrm{d}^2\xi}{\mathrm{d}\tau^2} + \left.\frac{\partial \mathcal{F}}{\partial u^r}\right|_{r_0,0}\frac{\mathrm{d}\xi}{\mathrm{d}\tau} + \left.\frac{\partial \mathcal{F}}{\partial r}\right|_{r_0,0}\xi = 0. \tag{21}$$

This is the equation for a damped harmonic oscillator (assuming the appropriate signs for the derivatives of $\mathcal{F}$). The equation of motion contains no fluid terms and so we expect this particular incompressible mode to oscillate about the equilibrium position with the same frequency as a test particle, in fact, the trajectories should be identical in the limit of small perturbations.

4 OSCILLATIONS

4.1 Undamped oscillation frequency

Our equation of motion in the thin shell limit is given by equation (21). In order to calculate the frequency of oscillations, first we will neglect the damping (second) term in equation (21). If

$$\left.\frac{\partial \mathcal{F}}{\partial r}\right|_{r_0,0} > 0, \tag{22}$$

then we can expect the harmonic oscillator solution, with the angular frequency in gravitational units as

$$\tilde{\omega} = \sqrt{M\left.\frac{\partial \mathcal{F}}{\partial r}\right|_{r_0,0}}. \tag{23}$$

We can tell that the above inequality holds from Fig. 1, but we will calculate it explicitly to find the frequency that the atmosphere would oscillate at if the radiation drag were negligible.

Taking the derivative and substituting for the radius of the equilibrium position in terms of λ, we get

$$\tilde{\omega} = \frac{(1 - \lambda^2)^2}{4M\lambda}. \tag{24}$$

To put our angular frequency in terms of s^{-1}, we restore c and G to get $\omega = c^3\tilde{\omega}/G$.

We have calculated our angular frequency with respect to the proper time experienced by the shell, τ, so we multiply by a factor of g_{tt} to redshift the angular frequency into the coordinate time, t. $\omega' = \omega\sqrt{1 - 2M/r_0} = \lambda\omega$. The frequency, $\nu = \omega'/2\pi$, is then

$$\nu = \frac{c^3}{2\pi G}\lambda\tilde{\omega} \approx 8.08(1 - \lambda^2)^2\frac{\mathrm{M}_\odot}{M}\ \mathrm{kHz}. \tag{25}$$

In Fig. 3 we can see a plot of ν as a function of λ. An example value of $\lambda = 0.8$ with a mass of $1.4\,\mathrm{M}_\odot$ gives, $\nu = 750$ Hz.

4.2 Damped oscillations

At this point, we calculate the strength of the radiation drag for small velocities. If the drag coefficient is small, then we should further explore this and other oscillation modes as a possible explanation for X-ray burst QPOs. If we find that the drag coefficient is large enough to overdamp the oscillations, then we should rule out this particular mode of the optically thin atmospheres. Although, one should still study the oscillatory behaviour of the optically thick atmospheres which may be different.

Our equation of motion is now given by

$$\frac{1}{M}\frac{\mathrm{d}^2\xi}{\mathrm{d}\tau^2} + \left.\frac{\partial \mathcal{F}}{\partial u^r}\right|_{r_0,0}\frac{\mathrm{d}\xi}{\mathrm{d}\tau} + \frac{(1 - \lambda^2)^4}{16M^3\lambda^2}\xi = 0. \tag{26}$$

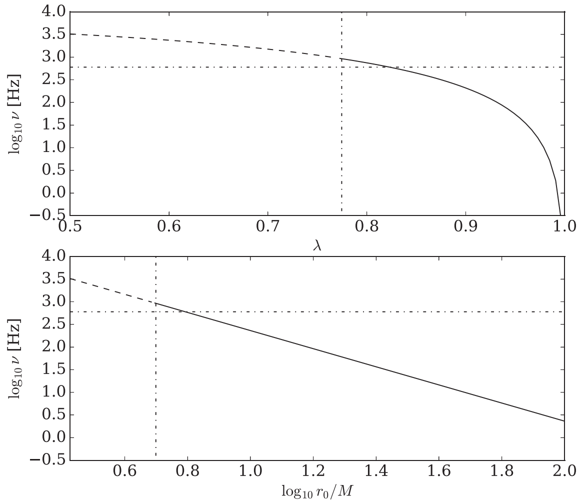

Figure 3. A plot of frequency of oscillations, ν, about r_0 with respect to the Eddington ratio, λ, above, and with respect to r_0 below. The neutron star mass is chosen to be $M = 1.4\,\mathrm{M}_\odot$. Also shown is the interesting frequency, 600 Hz which corresponds to $\lambda \sim 0.82$ and the λ for which $r_0 = 5M$, below which the frequency is not physical for a neutron star of that radius.

To calculate the drag coefficient for our linearized equation of motion, we just need to evaluate the u^r derivative of $\mathcal{F}$.

To calculate the u^r derivative, we return to our equation for the flux, where we have already simplified the first term. Substituting for $I(r)$ gives

$$F^r = \frac{L_\infty}{4\pi r^2} - \frac{L_\infty}{4\pi R^2}\frac{1-2M/R}{(1-2M/r)^2}\left[\tilde{T}^{(r)(r)} + \tilde{\epsilon}\right] u^r. \qquad (27)$$

We take the definition of $\tilde{\epsilon}$ from Stahl et al. (2013),

$$\tilde{\epsilon} = \frac{T^{\mu\nu}u_\mu u_\nu}{\pi I(r)}, \qquad (28)$$

and the relevant, dimensionless tetrad components of $T^{\mu\nu}$ written as, $\tilde{T}^{(t)(t)}$ and $\tilde{T}^{(r)(r)}$, derived in Abramowicz et al. (1990), given by

$$\tilde{T}^{(t)(t)} = 2(1 - \cos\alpha), \qquad (29)$$

$$\tilde{T}^{(r)(r)} = \frac{2}{3}(1 - \cos^3\alpha), \qquad (30)$$

also in terms of the viewing angle, $\alpha(r)$. Substituting back into the equation of motion keeping terms to first order in u^r, we now have

$$\mathcal{F}(r, u^r) = \frac{1}{r^2} - \frac{\lambda}{r^2\sqrt{1-2M/r}} + \frac{1}{R^2}\frac{1-2M/R}{(1-2M/r)^2}\left[\tilde{T}^{(r)(r)} + \tilde{T}^{(t)(t)}\right] u^r, \qquad (31)$$

which makes evaluating the derivative very simple,

$$\frac{\partial\mathcal{F}}{\partial u^r} = \frac{\lambda}{R^2}\frac{1-2M/R}{(1-2M/r)^2}\left[\tilde{T}^{(r)(r)} + \tilde{T}^{(t)(t)}\right]. \qquad (32)$$

At the point, we only need to evaluate the derivative at r_0,

$$\left.\frac{\partial\mathcal{F}}{\partial u^r}\right|_{r_0,0} = \frac{1-2M/R}{\lambda^3 R^2}\left[\frac{2}{3}(1-\cos\alpha_0)\left(\cos^2\alpha_0 + \cos\alpha_0 + 4\right)\right], \qquad (33)$$

where

$$\sin^2\alpha(r_0) = (1-\lambda^2)^2\frac{\lambda^2}{4}\frac{(R/M)^2}{1-2M/R}. \qquad (34)$$

From here, it becomes much simpler to switch to dimensionless variables scaled by the neutron star mass,

$$\frac{R}{M} \to X$$
$$\frac{\xi}{M} \to \tilde{\xi}$$
$$\frac{\tau}{M} \to \tilde{\tau}$$

Then the equation of motion is

$$\frac{\mathrm{d}^2\tilde{\xi}}{\mathrm{d}\tilde{\tau}^2} + M^2\left.\frac{\partial\mathcal{F}}{\partial u^r}\right|_{r_0,0}\frac{\mathrm{d}\tilde{\xi}}{\mathrm{d}\tilde{\tau}} + \left.\frac{\partial\mathcal{F}}{\partial r}\right|_{r_0,0} M^3\tilde{\xi} = 0. \qquad (35)$$

If, for convenience, we set the dimensionless quantities

$$f = M^3\left.\frac{\partial\mathcal{F}}{\partial r}\right|_{r_0,0} = \frac{(1-\lambda^2)^4}{16\lambda^2}, \qquad (36)$$

$$g = M^2\left.\frac{\partial\mathcal{F}}{\partial u^r}\right|_{r_0,0} = \frac{1-2/X}{\lambda^3 X^2} \times\left[\frac{2}{3}(1-\cos\alpha_0)\left(\cos^2\alpha_0 + \cos\alpha_0 + 4\right)\right], \qquad (37)$$

then we can classify the damping from the sign of

$$\tilde{\omega}_d^2 = \frac{g^2}{4} - f. \qquad (38)$$

If $\omega_\mathrm{d}^2 < 0$, then we have an underdamped oscillation, $\omega_\mathrm{d}^2 > 0$ is overdamped and $\omega_\mathrm{d} = 0$ is critically damped.

If we take some example values as before, $\lambda = 0.8$ and $R = 5M$, then we get

$f \approx 0.0016$, $g \approx 0.089$,

$\omega_\mathrm{d} \approx 0.06 = 885\,\mathrm{s}^{-1}$.

Because ω_d is real, we have an overdamped solution with an exponential decay without oscillation. For this reason, we do not label non-oscillating quantities of dimension $(\mathrm{s})^{-1}$ with Hz, so that they are not to be thought of as oscillation frequencies, but rather inverse decay time-scales. In fact, over the whole reasonable parameter space of λ and R, the eigenmode is overdamped. This can be seen in Fig. 4, a plot of the ratio of $\tilde{\omega}_\mathrm{d}^2$ to $\tilde{\omega}^2$. We can see that for small λ, we are strongly overdamped and as λ increases, the solution approaches critical damping. Our damped harmonic oscillator equation yields two exponential solutions of the form $C_\pm\exp(\gamma_\pm\tau)$. where

$$\gamma_\pm = -\frac{g}{2} \pm \sqrt{\frac{g^2}{4} - f}. \qquad (39)$$

Both of which correspond to decay constants of $\gamma_+ = -0.025$, $\gamma_- = -0.064$. If we convert to seconds and redshift for the observer at infinity, we have

$$\gamma'_\pm = \frac{c\sqrt{B(r_0)}}{2\pi r_g}\gamma_\pm = 32.31\times 10^3\,\mathrm{s}^{-1}\left(\frac{M}{\mathrm{M}_\odot}\right)^{-1}\lambda\gamma_\pm, \qquad (40)$$

$\gamma'_+ = 369\,\mathrm{s}^{-1}$, $\gamma_- = 945\,\mathrm{s}^{-1}$.

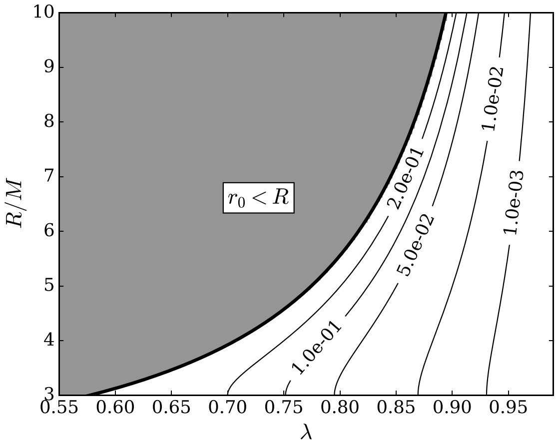

Figure 4. Contours of $\tilde{\omega}_{\rm d}^2/\tilde{\omega}^2$ over a reasonable parameter space for neutron star radius, R, and Eddington ratio, λ. The grey region on the left-hand side of the plot corresponds to $r_0 < R$ which is unphysical. The thick line marks the boundary where $r_0 = R$. There is no region on the plot where $\tilde{\omega_d}^2$ is less than zero which means that the thin shell eigenmode is overdamped over the parameter space we consider.

The scales of decay for each solution are then $1/\gamma_\pm$ so

$$\tau_+ = 0.0027\,\mathrm{s},\ \tau_- = 0.0011\,\mathrm{s}.$$

5 VALIDITY OF THE LINEAR REGIME

A key assumption mentioned at the beginning of the work relied on the atmosphere's thickness and velocities being small enough to be able to linearize $\mathcal{F}(r, u^r)$. This allows us to obtain analytical trajectories for the fluid for the thin shell mode, but we are interested to see the extent to which this linearization is valid, both in position and velocity. While we do not expect oscillatory behaviour to exist at larger velocities and displacements, we can still study how the time-scale for decay varies from that predicted by the analytical treatment.

To explore the degree to which the linear regime of $\mathcal{F}$ is valid, we compare the trajectories of the thin shell mode to those of test particles which obey the equation of motion,

$$\frac{\mathrm{d}^2\tilde{\xi}}{\mathrm{d}\tilde{\tau}^2} = -M^2\mathcal{F}(r, u^t). \qquad (41)$$

We expect the trajectories from this equation, given by $\xi_e(t)$ to be similar to the trajectories given by equation (21), which we will denote by $\xi_a(t)$.

To integrate this equation, we use the ODEINT routine from the SCIPY package which relies on the LSODA routine from the FORTRAN library ODEPACK. We test a variety of initial conditions. We plot these against the analytical trajectories given by solving the linear equation of motion, of which the trajectories are just linear combinations of exponentials. The plots, shown in Fig. 5, show both trajectories from the linear equation of motion (blue) and from the full equation of motion (red dashed).

The first plot in Fig. 5 shows trajectories starting with zero initial velocity. The initial positions of the test particle are readable from the axis where the y-axis represents the fractional displacement from the equilibrium position $f = r(0)/r_0 - 1$. We can see from the plot that deviations from the test particle solution start to occur between $f = 0.1$ and 0.2. The deviation is small, however, and the atmosphere settles into the equilibrium position at about the same time as the test particle. We expect that for $v(0) = 0$, the atmosphere trajectory is valid up to initial displacements of 10 per cent from the equilibrium solution, and a reasonable approximation up to 20 per cent away.

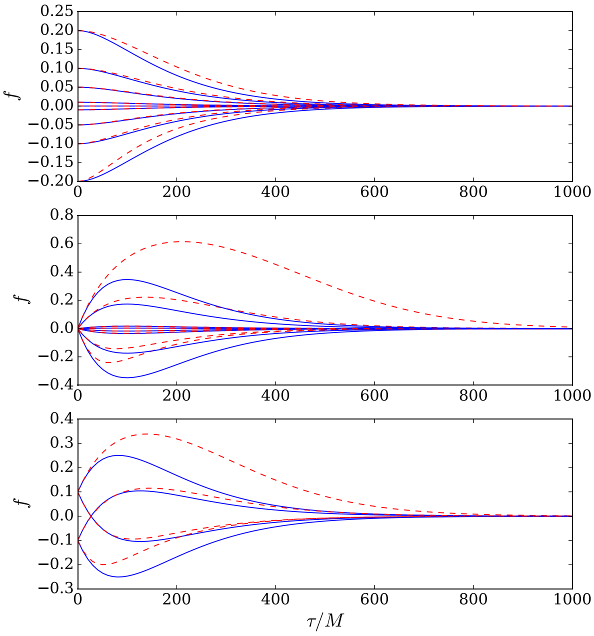

Figure 5. Trajectories from a variety of initial conditions analytically integrated from the linearized equation of motion, equation (21) (blue), and numerically integrated from the test particle equation of motion, equation (41) (red dashed). The vertical axes show, f, the fraction the initial position is from the ECS. Upper: the initial velocity is zero and the initial position can be read from the axis. Middle: the initial position is at the ECS and the initial velocities are given by, $v(0)/c = \pm 0.005, \pm 0.01, \pm 0.05, \pm 0.1$. Lower: Initial conditions are given by $(f(0), v(0)/c) = (0.1, 0.05), (0.1, -0.05), (-0.1, 0.05), (-.1, -0.05)$.

The second plot in Fig. 5 shows trajectories with an initial position at the equilibrium position, but with different initial velocities of $v(0)/c = \pm 0.005, \pm 0.01, \pm 0.05, \pm 0.1$. We can immediately see that $v(0) = \pm 0.1c$ shows vastly different trajectories, which are also highly asymmetrical with respect to the equilibrium position, indicating that the atmospheric approximation is no longer valid and even at $v(0)/c = \pm 0.05$, there are significant deviations. We expect the linear regime in this case to hold up to $v(0)/c = \pm 0.01$.

The third plot in Fig. 5 shows trajectories corresponding to the following initial conditions:

$$(f(0), v(0)/c) = (0.1, 0.05), (0.1, -0.05),$$
$$(-0.1, 0.05), (-.1, -0.05).$$

The two inner trajectories show much better agreement with the test particle. This is because their initial conditions keep them closer to the equilibrium position where the atmospheric equation of motion more accurately reflects that of the test particle.

5.1 Convergence

We also find it necessary to numerically confirm the convergence of the linearized fluid equation to the test particle equation. In essence, we want to show that the trajectories from the equation (21) approach those from equation (41) as $\xi(0) \to 0$, $u^r(0) \to 0$. We measure the deviation in the trajectory by summing up the fractional difference along $n = 100$ points, denoted by τ_i, which

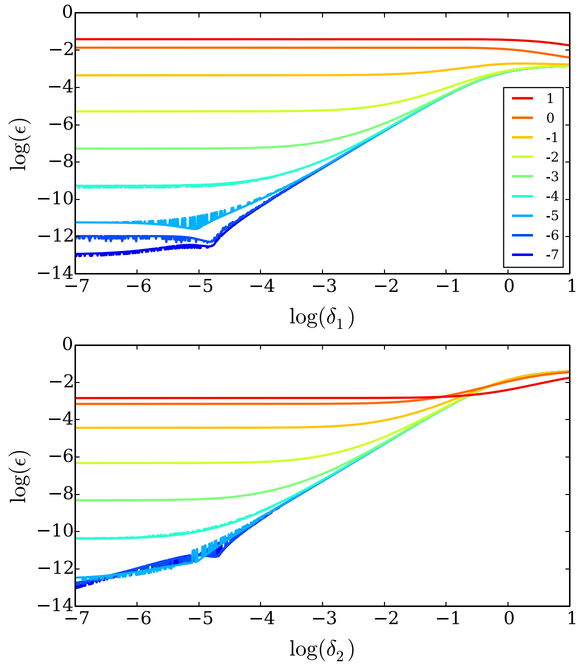

Figure 6. Here we show the results of the convergence test for the linearization of the equation of motion. For initial conditions, $(\xi(0), u^r(0)) = (\delta_1, \delta_2)$. The upper panel shows different curves corresponding to different δ_2 as indicated by the legend. The curves match the legend as read from top to bottom. The δ_1s are shown on the horizontal axis. The lower plot shows the same except for different curves of δ_1 with δ_2 on the horizontal axis.

are evenly sampled along the trajectories shown in Fig. 5. The error per timestep is then calculated by,

$$\epsilon = \frac{1}{n}\sum_{i=0}^{n}\frac{|\xi_e(\tau_i) - \xi_a(\tau_i)|}{\xi_e(\tau_i)}, \tag{42}$$

which we calculate for different values of the initial conditions, $(\xi(0), u^r(0)) = (\delta_1, \delta_2)$. A plot of the convergence is shown in Fig. 6. We can see that ϵ shrinks to the level of machine precision as δ_1, δ_2 approach zero.

6 DISCUSSION AND CONCLUSIONS

6.1 Consequences for the stability of atmospheres

We have shown that for optically thin levitating atmospheres, as in Wielgus et al. (2015), the radial, incompressible, thin shell modes are stable against radial oscillations owing to the strength of the radiation drag term. The natural frequency of these oscillations is of the order of $\sim 10^2$ Hz if they were not overdamped. Moreover, the frequency increases as the luminosity decreases. This would lead to an increase of frequency with time if the luminosity were to decay with time, such as during the decay phase of an X-ray burst. The trajectories of the fluid particles in these oscillations are exactly the same as for those of the test particles, and the incompressible mode is constructed in a way that does not allow for extra forces due to pressure gradients. It is possible to conceive of a mode where the pressure terms become important, such as a breathing mode where the shell expands and contracts in opposite directions about the pressure maximum so that the shell becomes thinner and thicker. It is also possible to construct pressure corrections to the thin shell incompressible mode to extend its validity to larger thicknesses. If the eigenfrequency of such a mode is large enough, one can imagine that underdamped oscillations may occur, and so it would be worth exploring such other modes.

6.2 Optically thick oscillations

In this work, we have only considered optically thin solutions. Wielgus et al. (2016) have extended their work to include optically thick atmospheres. These atmospheres are constructed using numerical techniques, so it is difficult to calculate analytical oscillation modes. In future work, we plan to extend this analysis to include numerical simulations of both optically thin and thick Eddington supported atmospheres. Since photons diffuse through the optically thick atmospheres, as opposed to free streaming through the thin ones, it is possible that the radiation drag is less efficient at damping oscillations. We expect the drag to only be effective in the optically thin edges of the atmospheres.

ACKNOWLEDGEMENTS

The authors thank Maciek Wielgus and Omer Blaes for useful discussions during the work. The research was supported by the Polish NCN grant 2013/08/A/ST9/00795.

REFERENCES

Abramowicz M. A., Ellis G. F. R., Lanza A., 1990, ApJ, 361, 470
Bachetti M. et al., 2014, Nat, 514, 202
Bini D., Jantzen R. T., Stella L., 2009, Class. Quantum Gravity, 26, 055009
Oh J. S., Kim H., Lee H. M., 2010, Phys. Rev. D, 81, 084005
Phinney E. S., 1987, in Zensus J. A., Pearson T. J., eds, Superluminal Radio Sources. Cambridge Univ. Press, Cambridge, p. 301
Remillard R. A., McClintock J. E., 2006, ARA&A, 44, 49
Stahl A., Wielgus M., Abramowicz M., Kluźniak W., Yu W., 2012, A&A, 546, A54
Stahl A., Kluźniak W., Wielgus M., Abramowicz M., 2013, A&A, 555, A114
Strohmayer T. E., 2001, AIP Conf. Ser. Vol. 599, X-ray Astronomy: Stellar Endpoints, AGN, and the Diffuse X-ray Background. Am. Inst. Phys., New York, p. 377
Strohmayer T., Bildsten L., 2006, in Lewin W., van der Klis M., eds, Compact Stellar X-ray Sources. Cambridge Univ. Press, Cambridge, p. 113
van der Klis M., 2006, Lew in W. van der Klis M., eds, Compact Stellar X-ray Sources. Cambridge Univ. Press, Cambridge, p. 39
Wielgus M., Stahl A., Abramowicz M., Kluźniak W., 2012, A&A, 545, A123
Wielgus M., Kluźniak W., Sądowski A., Narayan R., Abramowicz M., 2015, MNRAS, 454, 3766
Wielgus M., Sądowski A., Kluźniak W., Abramowicz M., Narayan R., 2016, MNRAS, 458, 3420

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

3 Paper 2: Radiative GRMHD simulations of accretion and outflow in non-magnetized neutron stars and ultraluminous X-ray sources

In this paper (Abarca et al. 2018), I performed two 2D axisymmetric GRRMHD simulations of accretion disks with accretion rates around $15\dot{M}_{\rm Edd}$. One simulation contains a black hole as the central object, and the other an unmagnetized neutron star, both with masses equal to $1.4\,M_\odot$ so that they could be directly compared to one another. I ran them for long durations (80 000 GMc^{-3} for the black hole, and 160 000 GMc^{-3} for the neutron star) in order for the solution to converge at large radii providing more realistic assessment of the nature of the gas outflows.

The black hole simulation was consistent with previous simulations in the literature (although no other black hole was considered with such a low mass) showing strong outflows with collimated emission in the funnel region with inferred isotropic luminosities peaking around $10\,\dot{M}_{\rm Edd}$ making it a reasonable candidate for the some of the lowest luminosity ULXs.

I modeled the neutron star using a radially reflective, angular momentum absorbing boundary condition. As gas hit the inner boundary, it slowed down and started to accumulate on the surface forming an atmosphere which was highly radiation pressure dominated. The outer layers of the atmosphere launched a nearly spherical outflow obscuring the entire inner region from view and radiating nearly isotropically at around $1\,L_{\rm Edd}$. I concluded that non-magnetized neutron stars would not be good candidates for ULXs because it is necessary to have something to clear the funnel region so that radiation can escape and become collimated by the outflows.

Monthly Notices
of the
ROYAL ASTRONOMICAL SOCIETY

MNRAS **479,** 3936–3951 (2018)
Advance Access publication 2018 June 15

doi:10.1093/mnras/sty1602

Radiative GRMHD simulations of accretion and outflow in non-magnetized neutron stars and ultraluminous X-ray sources

David Abarca,[1]* Włodek Kluźniak[1] and Aleksander Sądowski[2]

[1]*Nicolaus Copernicus Astronomical Center, Warsaw, Poland*
[2]*Akuna Capital, 585 Massachusetts Avenue, Cambridge, MA 02139, USA*

Accepted 2018 June 13. Received 2018 June 6; in original form 2017 December 21

ABSTRACT

We run two general relativistic radiation magnetohydrodynamic simulations of super-Eddington accretion discs around a black hole and a non-magnetized, non-rotating neutron star (NS). The NS was modelled using a reflective inner boundary condition. We observe the formation of a transition layer in the inner region of the disc in the NS simulation which leads to a larger mass outflow rate and a lower radiative luminosity over the black hole case. Sphereization of the flow leads to an observable luminosity at infinity around the Eddington value when viewed from all directions for the NS case, contrasting to the black hole case where collimation of the emission leads to observable luminosities about an order of magnitude higher when observed along the disc axis. We find the outflow to be optically thick to scattering, which would lead to the obscuring of any NS pulsations observed in corresponding ultraluminous X-ray sources.

Key words: accretion, accretion discs – MHD – stars: neutron.

1 INTRODUCTION

The transfer of matter with angular momentum onto a compact object occurs via an accretion disc through which mass and angular momentum are transported in opposite directions through viscous processes (Shakura & Sunyaev 1973). The viscous process is now believed to be magnetic turbulence induced by the magnetorotational instability (Balbus & Hawley 1991).

One particular class of accreting objects which has gained interest in recent years are ultraluminous X-ray sources (ULXs). These are X-ray bright objects observed outside the centres of galaxies with luminosities from 10^{39} to 10^{41} ergs s^{-1}. The first explanations for such bright X-ray objects favoured the elusive intermediate mass black holes radiating at or below the Eddington luminosity (Colbert & Mushotzky 1999). Up until the discovery of ULXs, the most luminous stellar-mass, persistent X-ray source was known to be Sco X-1, a neutron star (NS) radiating at its Eddington limit at around 10^{38} erg s^{-1} (Shklovsky 1967; Bradshaw, Fomalont & Geldzahler 1999). Currently, the leading explanation for ULXs is beamed emission from accretion in an X-ray binary (King et al. 2001), implying that super-Eddington accretion is responsible for the large observed luminosities. In particular, a set of three such objects were observed which reveal X-ray pulsations with a period on the order of 1 s (Bachetti et al. 2014; Fürst et al. 2016; Israel et al. 2017a,b) excluding black holes as the accreting objects in these three sources. It can now be said with some certainty that a large fraction of ULXs are accreting NSs (Kluźniak & Lasota 2015; King, Lasota & Kluźniak 2017; Wiktorowicz et al. 2017; Pintore et al. 2017). There have been a wide range of proposed values for the strength of the magnetic field from relatively low ($B < 10^9$ G) (Kluźniak & Lasota 2015), to moderate (10^{10} G $\lesssim B \lesssim 10^{13}$ G) (King et al. 2017; Walton et al. 2018), to high ($B > 10^{13}$ G), magnetar-like fields (Ekşi et al. 2015; Mushtukov et al. 2015b).

Accretion onto a NS is more complicated than accretion onto a black hole. NS have no event horizon. They have a surface layer and outer crust which can reach densities of up to 10^{11} g cm^{-3} at its base, where it is composed of fully ionized neutron rich nuclei. At larger radii, and lower densities ($\sim 10^7$ g cm^{-3}) the composition of the nuclei becomes less neutron rich. Below 10^4 g cm^{-3}, the nuclei are no longer fully ionized. Finally, near the surface, the outer crust is composed mainly of crystallized iron atoms reaching down to 10 g cm^{-3} (Chamel & Haensel 2008).

Gas accreting onto these outer layers is expected to slow down and release some of its kinetic energy (Syunyaev & Shakura 1986; Kluzniak & Wilson 1991; Narayan & Yi 1995; Narayan, Garcia & McClintock 1997; Inogamov & Sunyaev 1999; Sibgatullin & Sunyaev 2000; Popham & Sunyaev 2001; Mukhopadhyay & Fang 2002) as it spins up the star (Kluzniak & Wagoner 1985). This energy can be converted into radiation (normally X-rays) or transferred to the outflowing gas. Additionally, many NSs have strong magnetic fields, some of which can be strong enough to channel the accreting gas into dense accretion columns, depositing gas at the magnetic poles, forming hot spots at low accretion rates, $\lesssim 10^{17}$ g s^{-1}. Misalignment of the magnetic poles with the rotation axis

---

* E-mail: dabarca@camk.edu.pl

Published by Oxford University Press on behalf of the Royal Astronomical Society

causes the hot spots to rotate resulting observationally in X-ray (accretion powered) pulsars. At high accretion rates ($10^{17} \sim 10^{19}$ g s^{-1}), the gas is expected to experience a radiation shock and to form an accretion column above the NS surface through which gas sinks slowly through a dense radiation field to eventually settle on the NS surface (Basko & Sunyaev 1976; Mushtukov et al. 2015a,b; Revnivtsev & Mereghetti 2015). As the accretion rate increases, the accretion column starts to widen and spread over a surface roughly corresponding to the surface of the magnetosphere. Emission from the central regions and through the sides of the inner part of the accretion column can interact with the outer parts of the accretion column producing complicated pulse profiles (Mushtukov et al. 2018).

At even higher accretion rates ($\gtrsim 10^{19}$ g s^{-1}), in the context of ULXs, the accretion column spreads into an accretion curtain, a geometrically extended surface corresponding to an optically thick layer completely surrounding the pulsar magnetosphere which reprocesses all of the radiation generated near the NS surface strongly smoothing the pulse profile (Mushtukov et al. 2017). Observationally, in pulsating and non-pulsating sources, this manifests as a double blackbody with hot (>1 keV) and cold (<0.7 keV) components corresponding to the thermal emission from the accretion enveloping the magnetosphere, and thermal emission from the accretion disc truncated at the magnetosphere, respectively (Koliopanos et al. 2017).

Here, using numerical simulations, we try to see if non-pulsating ULXs can be explained by super-Eddington accretion onto NS. We ignore the effects of a stellar magnetic field, (for an accretion rate of 10 times the Eddington limit and a magnetic moment of $\mu < 10^{27}$ G cm^3 a simple calculation of the Alfvén radius shows the effects of the magnetic field are confined close to the NS surface), and consider only the effects of a hard surface. We use a general relativistic radiation magnetohydrodynamic (GRRMHD) code KORAL to capture the most relevant physical processes.

1.1 NS related accretion simulations

In this section, we mention some simulations that are related to NS accretion. The X-ray spectra from spherical accretion onto high- and low-mass NSs was computed from coupled hydrodynamic radiation transfer calculations, which were shown to yield results which differ strongly from a blackbody (Alme & Wilson 1973). Dhang, Sharma & Mukhopadhyay (2016) performed hydrodynamic simulations of spherical accretion onto a hard surface in one and two dimensions. The hard surface was modelled in two ways, with a reflective, and a 'leaky' boundary condition, the latter being where mass is allowed to cross the inner boundary at a fixed subsonic speed to model efficient cooling. This is important because not all works include a hard surface. More complicated Bondi–Hoyle (El Mellah & Casse 2015) and magnetic Bondi–Hoyle (Toropina, Romanova & Lovelace 2012) simulations have been performed, but they are without a hard inner boundary.

A 1.5D coupled radiative transfer and hydrodynamics calculation was performed by Kluzniak & Wilson (1991) which simulated the boundary layer between the NS and the accretion disc. The boundary layer was simulated by introducing an optically thin stream of plasma inside the innermost stable circular orbit (ISCO), where the infall velocity quickly becomes supersonic. The plasma decelerates in the upper layers of the boundary layer on the NS surface resulting in the creation of hard X-rays. Because velocities in the accretion gap are supersonic, such a calculation is valid without considering the contribution of the accretion disc.

A further work of interest is Kawashima et al. (2016), who performed radiation hydrodynamic simulations of the accretion column of a super-Eddington accretion NS using flux-limited diffusion, where the radiative flux follows the gradient of radiative energy density. They found sub-Eddington luminosities along the optically thick accretion column but super-Eddington luminosity when viewed from the sides, in agreement with Basko & Sunyaev (1976).

Romanova et al. (2012) have performed global MHD simulations of MRI-driven accretion onto magnetized stars. A number of interesting results are presented on the interaction between the stellar magnetic field and the accretion disc, however, the lack of strong gravity or radiation hydrodynamics means that the results do not accurately describe accretion onto NSs at large accretion rates.

Parfrey & Tchekhovskoy (2017) used an innovative method to run GRMHD simulations of accretion onto rotating magnetized NSs to model accreting millisecond pulsars. Their method interpolates between the normal GRMHD flow and the force-free magnetosphere. At the lowest magnetizations of the NS, they show that the magnetic field is crushed by the accretion flow, and accretion proceeds normally. Due to the lack of radiation, their simulations are scale-free. When scaling their system to the mildly super-Eddington accretion flow that we describe in this work, we find that indeed the magnetic field would be crushed at a magnetic moment of $\mu = 10^{26}$ G cm^3, even when rotating at millisecond periods.

Takahashi & Ohsuga (2017) have published the first global 2.5D GRRMHD simulations of accretion onto a NS and their work represents the state-of-the-art on the subject. They simulate super-Eddington accretion onto a magnetized NS with the radial flux and velocity set to zero at the inner boundary, as a means to model a NS ULX system. They report luminosities of about an order of magnitude above the Eddington limit with a significant amount of beaming which accurately describes a non-pulsating ULX source. In their simulation, the magnetic field is strong enough to truncate the disc leading to accretion along magnetic field lines. They observe some matter piling up at the inner boundary due to the inner boundary condition, but do not run the simulation for long duration, ($t_{\rm max} = 15\,000 t_g$), and so it is hard to say what the effect of the accumulation of gas has on the accretion disc.

Our work considers the context of a boundary layer (as opposed to an accretion column) with an accretion disc, with a sophisticated radiation treatment which deals with the optically thick and optically thin regimes. In our simulation, we will focus strictly on the effect of a reflective boundary. We will run our simulations for longer durations, $t \sim 160\,000\, t_g$, where $t_{\rm g} = GM/c^3$, to see what happens when a large amount of gas is accumulated in the vicinity of the NS.

2 NUMERICAL METHODS

We investigate accretion onto NSs using a sophisticated 3D general relativistic radiation magnetohydrodynamics solver, KORAL used extensively to study accretion onto black hole at high and low accretion rates, and other related phenomena. Details of the numerical implementation are given in Sądowski et al. (2013, 2015). Here, we describe the most relevant features.

2.1 Governing equations

The equations of GRRMHD, which can be written in their conservative form as

$$\nabla_\mu(\rho u^\mu) = 0, \qquad (1)$$

$$\nabla_\mu T^\mu{}_\nu = G_\nu, \tag{2}$$

$$\nabla_\mu R^\mu{}_\nu = -G_\nu, \tag{3}$$

$$\nabla_\mu (n u^\mu) = \dot{n}, \tag{4}$$

are solved in KORAL on a static 1D, 2D, or 3D mesh. Our mesh is a spherical 2.5D (2D axisymmetric) grid using a static metric, $g_{\mu\nu}$ with signature $(-\,+\,+\,+)$. Here, ρ is the gas rest-mass density in the comoving fluid frame, u^μ is the gas four-velocity, $T^\mu{}_\nu$ is the MHD stress-energy tensor given by

$$T^\mu{}_\nu = (\rho + u_{\rm int} + p + b^2) u^\mu u_\nu + \left(p + \frac{1}{2}b^2\right) \delta^\mu_\nu - b^\mu b_\nu. \tag{5}$$

Here, $u_{\rm int}$ is the internal energy of the gas, and $p = (\gamma - 1)u_{\rm int}$ is the gas pressure, calculated using the adiabatic index, $\gamma = 5/3$. The radiation stress-energy tensor is given by $R^\mu{}_\nu$ which is coupled to the gas stress-energy tensor by the radiation four-force, G_ν, making use of electron scattering and bremsstrahlung opacities as well as Comptonization (Sądowski & Narayan 2015), which evolves the photon number, n, by taking into account the creation and annihilation of photons by emission and absorption, $\dot{n}$, while conserving n for Compton scattering exchanges of energy. The radiation stress-energy tensor is completed using the M_1 closure scheme (Sądowski et al. 2013), which assumes there is a frame in which the radiation is isotropic. The M_1 scheme allows radiation to diffuse through gas at large optical depths, and to freely stream along geodesics at very low optical depths. The magnetic field four-vector, described by Gammie, McKinney & Tóth (2003), is given by b^μ, and it is evolved using the induction equation which, when written in the coordinate basis appears as

$$\partial_t(\sqrt{-g} B^i) = -\partial_j \left(\sqrt{-g} b^j u^i - b^i u^j\right). \tag{6}$$

Here, B^i is the normal magnetic field three-vector, which is related to the magnetic field four-vector by

$$b^t = B^i u^\mu g_{i\mu}, \tag{7}$$

$$b^i = \frac{B^i + b^t u^i}{u^t}, \tag{8}$$

for metric g_{ij} and metric determinant, g (Komissarov 1999).

2.2 Mean-field dynamo for 2.5D runs

One particularly useful tool implemented in KORAL is a mean-field magnetic dynamo which allows for axisymmetric 2D (2.5D) accretion disc simulations to be run for long durations without depleting the magnetic field due to turbulent dissipation, which normally occurs in axisymmetric simulations of MRI (Sądowski et al. 2015). The dynamo has been tested against 3D simulations and has been found to accurately approximate the disc's spatial properties, accretion rate, surface density, and angular momentum for example. The 2D dynamo disc does, however, have a tendency to overestimate the magnitude and variability of the radiative flux (Sądowski & Narayan 2016). Nevertheless, the advantages of being able to run a 2D simulation as opposed to 3D make the mean-field dynamo a valuable tool, allowing for almost a 100-fold speed up in runtime, and so we chose to implement it in this work in order to run long duration simulations. We also implement an adaptation to make the dynamo more suitable to simulations where we expect a significant amount of gas to accumulate at the inner simulation boundary. We include a smooth cut-off to deactivate the dynamo in cells with a specific angular momentum lower than 80 per cent of the Keplerian value.

3 NUMERICAL SET-UP

3.1 Initial conditions

We initialize our accretion disc in a typical way by starting with an equilibrium torus near the non-rotating black hole as given in Penna, Kulkarni & Narayan (2013). The torus is threaded with a weak magnetic field in loops of alternating polarity. The total pressure is distributed between gas and radiation assuming local thermal equilibrium. Once the simulation starts, the MRI quickly develops turbulence and accretion begins. We measure the rate of mass accretion in units of $L_{\rm Edd}/c^2$, where $L_{\rm Edd}$ is the Eddington luminosity and c is the speed of light. The initial torus is set up to give a constant accretion rate of about 200 $L_{\rm Edd}/c^2$ which would correspond to a luminosity of about 10 $L_{\rm Edd}$ for the efficiency of a Shakura–Sunyaev disc.

3.2 Boundary conditions

A common practice when simulating accretion onto stars is to ignore the effects of a hard surface and let gas flow through the inner boundary in a standard outflow boundary condition for the hydrodynamic quantities. This allows the simulation to approach a quasi-steady state and is useful for studying the interaction between the stellar magnetic field, and the accretion disc (Romanova et al. 2012; Čemeljić, Shang & Chiang 2013). This allows the star to behave somewhat like a black hole.[1] In order to study the the difference between an inflowing boundary condition at the inner edge, and a reflective boundary condition, as well as to have a natural, physical inner boundary as a baseline comparison, we run one simulation with a black hole as the inner boundary condition. This is achieved by choosing a horizon penetrating coordinate system (Kerr–Schild) and placing the inner boundary of the simulation sufficiently behind the event horizon. This simulation we call BHRUN.

In order to study the effects of the release of kinetic energy, which is expected to significantly impact the behaviour of the accretion disc and outflows, we implement a reflective boundary for the main simulation of our study. The reflective boundary at $r = r_{\rm in}$ is set up so that the reconstructed radial velocity of the gas at the inner boundary is opposite about zero ($u^r_l = -u^r_r$, where u^r_l, u^r_r are the left and right reconstructed radial velocities at the cell interface of the inner boundary), so that no gas is able to leave the domain. We also set the tangential velocities u^θ, u^ϕ in the ghost cells to zero. Note that this does not enforce the reconstructed tangential velocities to be exactly zero at the inner boundary, it does, however, effectively remove angular momentum and allows the gas to approach a non-rotating state at the inner boundary, and so we expect the formation of a boundary layer. Radiative quantities are treated in the same way. We expect these boundary conditions to more accurately reflect the behaviour of an accretion disc and a boundary layer around a NS than an inflow boundary condition, and so this simulation we call NSRUN.

---

[1] This is not entirely true, there are a variety of inner boundary conditions on the magnetic field quantities that have various effects on the absorption of the hydrodynamic quantities.

3.3 Grid

Both simulations are run in Kerr–Schild coordinates (although this is not strictly necessary for NSRUN) with a non-rotating, central mass of $M = 1.4\ \mathrm{M}_{\odot}$ (Schwarzschild space–time). We run two simulations on 2D, axisymmetric spherical grids with logarithmic spacing in radius, and increased resolution near the equatorial plane. The resolution for NSRUN is $N_r \times N_\theta \times N_\phi = 352 \times 240 \times 1$. The radial coordinate spans from $r_{\mathrm{in}} = 5r_{\mathrm{g}}$ to $r_{\mathrm{out}} = 5000r_g$, where $r_{\mathrm{g}} = GM/c^2$. The resolution was higher for BHRUN ($N_r \times N_\theta \times N_\phi = 384 \times 240 \times 1$) because we have to extend the inner boundary to below the event horizon at $r = 2r_{\mathrm{g}}$. The larger number of radial grid cells allows the two simulations to have comparable resolution at the same radii.

4 RESULTS

We have run two axisymmetric GRRMHD simulations, one with a reflective, non-rotating inner boundary at radius $r = 5r_{\mathrm{g}}$ NSRUN, and one with a black hole inner boundary, BHRUN. Snapshots and time averages of both simulations are shown in Figs 1 and 2. Each figure contains four images, each corresponding to a different time, $t_i \in \left\{7500\,t_g, 15\,000\,t_g, 30\,000\,t_g, 60\,000\,t_g\right\}$. Snapshots at these times are shown in the upper two quadrants of each image. We are also showing the time averaged structure of the simulations in the lower two quadrants of each image, all of which are averaged from $t = t_i - t_i/3$ to $t = t_i + t_i/3$. The right two quadrants of each image show rest mass density, and the left two quadrants show radiation energy density in the fluid frame.

The BHRUN simulation forms a typical, geometrically thick, super-Eddington accretion disc similar to what is seen in previous simulations of this type, see Yang et al. (2014), Jiang, Stone & Davis (2014) Sądowski et al. (2015), Sądowski & Narayan (2016), and Ogawa et al. (2017), for examples. As the disc reaches inflow equilibrium, one can see the formation of a funnel region along the polar direction. The funnel region is nearly devoid of gas but is filled with radiation and so one can expect the disc to appear very bright when viewed along the funnel.

The reflective inner boundary shows a much different scenario. Gas cannot pass through the inner boundary, so it accumulates a dense layer around the inner edge of the simulation. At late times, the density at the inner edge reaches nearly $10\ \mathrm{g\,cm^{-3}}$. The density at the inner edge increases gradually as more and more gas is compressed into the atmosphere. We can also see large amounts of radiation being accumulated at the inner boundary due to a combination of dissipation and advection. In general, hot flows of this type are optically thick to scattering, but optically thin to absorption. However, it is unclear a priori if the radiation can escape to infinity in any reasonable amount of time because gas is blown off the outer edges of the accumulating layer which forms a dense outflow. Photons produced near the disc must random walk through this thick outflow, some of them may even be scattered back through the disc. We are particularly interested in whether or not pulsations may be visible. With such a thick scattering dominated atmosphere, it is likely that pulsations would lose their coherence. We measure the scattering optical depth in the next section.

4.1 Transition layer and accretion disc structure

The inner boundary of the NS simulation causes the angular velocity of the gas to approach zero. This is seen in Fig. 3. The top panels of Fig. 3 show the density-weighted, θ-averaged profiles of the specific angular momentum of time averaged time intervals used in Figs 1 and 2 for NSRUN. Because NSRUN is computationally cheaper compared to BHRUN and because the behaviour at late times is interesting we run NSRUN for an additional $80\,000\ t_{\mathrm{g}}$ so that we can show an additional time interval in Fig. 3, $t_5 \sim 120\,000\,t_g$. The left-hand three panels show the mid-plane quantities, averaged or integrated over a range of θ from $\pi/2 - \pi/20$ to $\pi/2 + \pi/20$. The right-hand three panels are integrated in the polar region, over θ from $\pi/8 - \pi/20$ to $\pi/8 + \pi/20$. As the simulation evolves, we can see the formation of a transitional region which expands outwards as more gas is accumulated. The rate of expansion, measured by calculating the radius at which ℓ reaches 90 per cent of its Keplerian value, follows a $t^{0.85}$ power law. For accretion onto NSs at lower accretion rates, this transitional layer usually lies in a small belt around the star where the flow properties transition to match the stellar surface. This is the classical picture of a boundary layer. We can see from Fig. 2 that we have a much more extended, atmospheric layer, although the flow velocity does transition to match the surface of the NS, so we will refer to this region as the transition layer.

The process by which angular momentum is transferred from the accretion disc to a NS is complicated. Various viscous and magnetic processes are involved. We try to reproduce the effect of driving the tangential velocity to zero at the inner edge of the simulation with our inner boundary condition. The normal magnetoturbulent processes that transport angular momentum through the disc do not operate on the numerical level of the cell interface. Instead angular momentum is transported between cells according to the flux computed by the HLL Riemann solver. The angular momentum flux is made up of a hydrodynamic and a magnetic component, $T_{r\phi} = T_{r\phi}^{(\mathrm{HD})} + T_{r\phi}^{(\mathrm{mag})}$ where $T_{r\phi}^{(\mathrm{HD})} = (\rho + u_{\mathrm{int}} + p)u_r u_\phi$ and $T_{r\phi}^{(\mathrm{mag})} = b^2 u_r u_\phi - b_r b_\phi$. By far the largest contribution at the inner edge is the $\rho u_r u_\phi$ term. The inner boundary condition leads to a flux that is approximately, $T_{r\phi} \approx \rho u_r u_\phi/2$. Reflecting the angular velocity, for example, would lead to $T_{r\phi} \approx \rho u_r u_\phi$. We can see that the source of the torque at the inner edge is numerical in nature. A more detailed study of the effect of different boundary conditions on the transport of angular momentum through the transition layer is left for further studies. For now, we are satisfied the angular momentum transitions towards zero at the inner boundary through the mid-plane and along the disc axis.

The middle two frames of Fig. 3 show angular velocity averaged in the same way as the top panels of Fig. 3. We can see the expansion of the transition layer evolving into a quasi-flat region in angular velocity in the mid-plane indicating that some large scale coupling is causing the inner disc to rotate like a rigid body. The innermost region is driven towards zero angular velocity as is expected by the boundary condition.

The polar region shows different behaviour. The angular momentum in the polar regions increases with radius at early times, and evolves to be quasi-flat at late times in the inner region. Meanwhile, the angular velocity starts quasi-flat, and evolves to decrease with radius at late times, so the transition layer does not display rigid body behaviour near the poles.

As is seen in Fig. 2, the transition layer reaches very high densities, and so, even though it has very low angular velocity, it is still able to contain a significant amount of angular momentum, as shown in the bottom panels of Fig. 3. Here, we are plotting the angular momentum integrated at a particular radius, r, $\mathcal{J}(r) = 2\pi \int \rho u_\phi \sqrt{-g}\mathrm{d}\theta$. The total angular momentum in the tran-

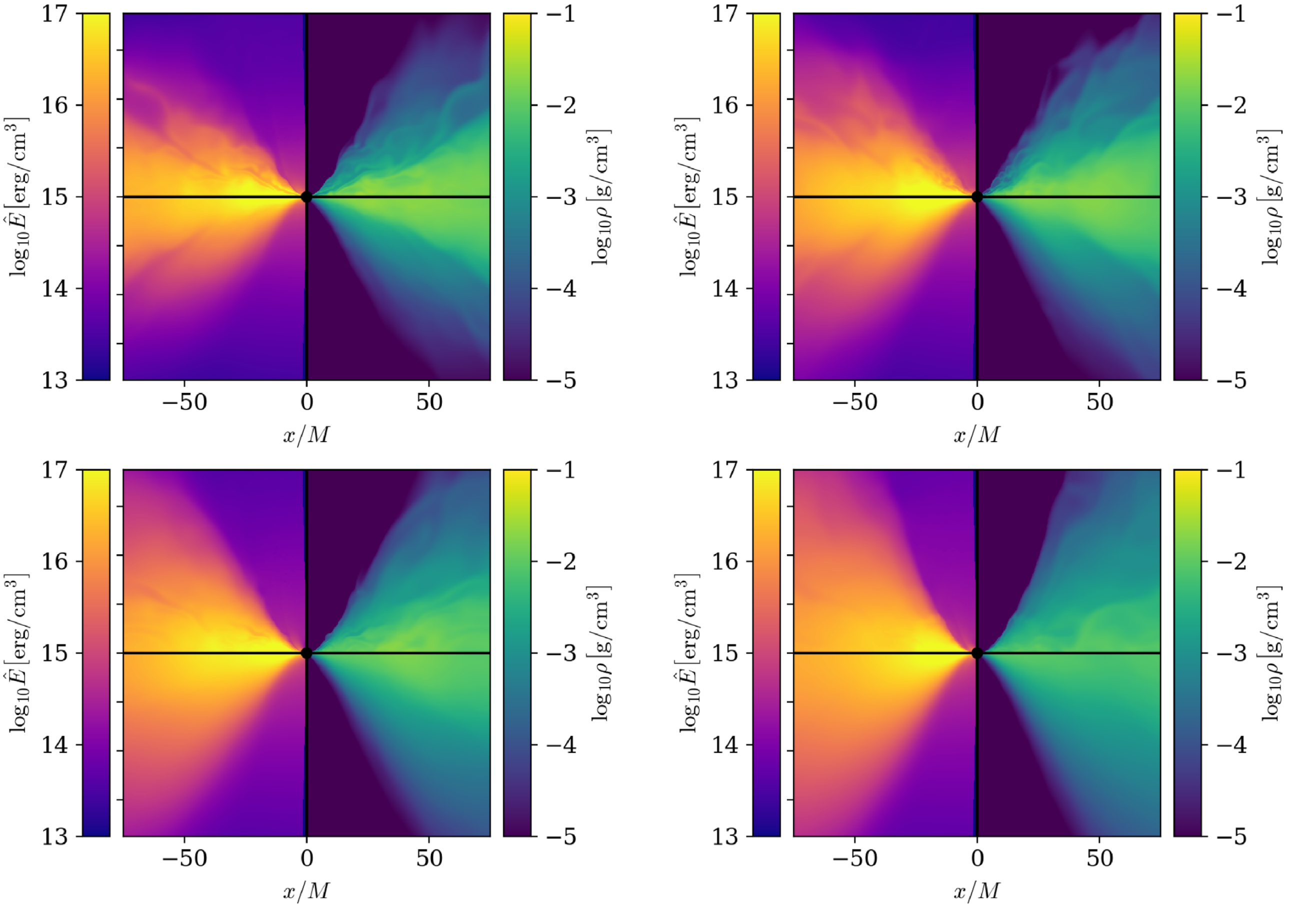

Figure 1. Here, we show snapshots of the 2.5D axisymmetric simulations of super-Eddington accretion from BHRUN. The left two quadrants of each image show radiative energy density in the fluid frame, and the right quadrants shows the rest mass density in the fluid frame. The upper quadrants of each image show the instantaneous state of the simulation at time t_i, and the bottom quadrants are time averaged from time $t_i - t_i/3$ to $t_i + t_i/3$. From left to right and top to bottom the times correspond to $t_i = 7500\,t_g$, $15\,000\,t_g$, $30\,000\,t_g$, and $60\,000\,t_g$.

sition layer increases with time, especially near the inner edge. Angular momentum is transported through the transition layer where it accumulates at the inner edge. It is likely that a stronger numerical torque at the inner edge would allow more angular momentum to flow through the inner boundary, thus smoothing $\mathcal{J}$ near the NSsurface. Dynamically, what is important are ℓ and Ω, both of which are largely below their corresponding Keplerian values, and so we do not expect a stronger torque at the inner boundary to have a significant effect on the evolution of the simulation.

The transition layer does not display much structure in the spatial distribution of gas density and radiation energy density, except for a strong radial gradient, even when the colour scales are suitably adjusted. The flow structure is much more informative.

In Fig. 4, we show the time-averaged spatial distribution of the poloidal velocities in the two upper panels for the time interval, t_4, from $40\,000\,t_g$ to $80\,000\,t_g$. All of the remaining figures are constructed with time averaged data over the same period. The poloidal velocity and accretion rate are defined as follows, $u^{\rm pol} = \sqrt{u^r u_r + u^\theta u_\theta}$, $\dot{M}_{\rm pol} = 2\pi r^2 \sin\theta\, \rho u^{\rm pol}$. For BHRUN, the flow is primarily directed inward in the disc and at the inner boundary. This is a boundary condition imposed by the space–time. In the polar region, one can see the transition from inflow to outflow at about radius of $r = 10r_g$. This is the stagnation radius, above which is a radiation driven outflow. Velocities in the polar region are relativistic, while velocities in the disc are around a few hundredths of the speed of light except in the very inner regions.

Again, NSRUN shows a much different inner structure. Gas flows through the transition layer in the equatorial plane where it meets the reflective boundary and is directed tangentially along the surface of the inner boundary until the polar regions where it is again redirected. The whole process forms two large eddies which seem to recycle the gas into the inner edge of the accretion flow, which could be responsible for the large-scale coupling seen in Fig. 3. The eddies seem to be connected to a conical outflow. We can see two streams of gas being launched from the two regions where the eddy circles back to the disc. It is important to note that this is the time average structure, the non-averaged flow being much more turbulent. The eddies are indicative of convective cells. Indeed, a calculation of the Schwarzschild stability parameter shows that the transition layer should be unstable to convection. Whether convection is driving the eddy motion is more difficult to say due to the extra source of momentum from the accretion flow.

The poloidal accretion rate is shown in the middle two panels of Fig. 4 over a large range in r and z. The accretion rate in the transition layer is nearly an order of magnitude higher than in the disc, indicating that the gas is recycled many times in the inner flow. For BHRUN, we see the typical picture of a nearly empty funnel region, indicating that most of the outflowing gas is ejected in a wind at larger radii.

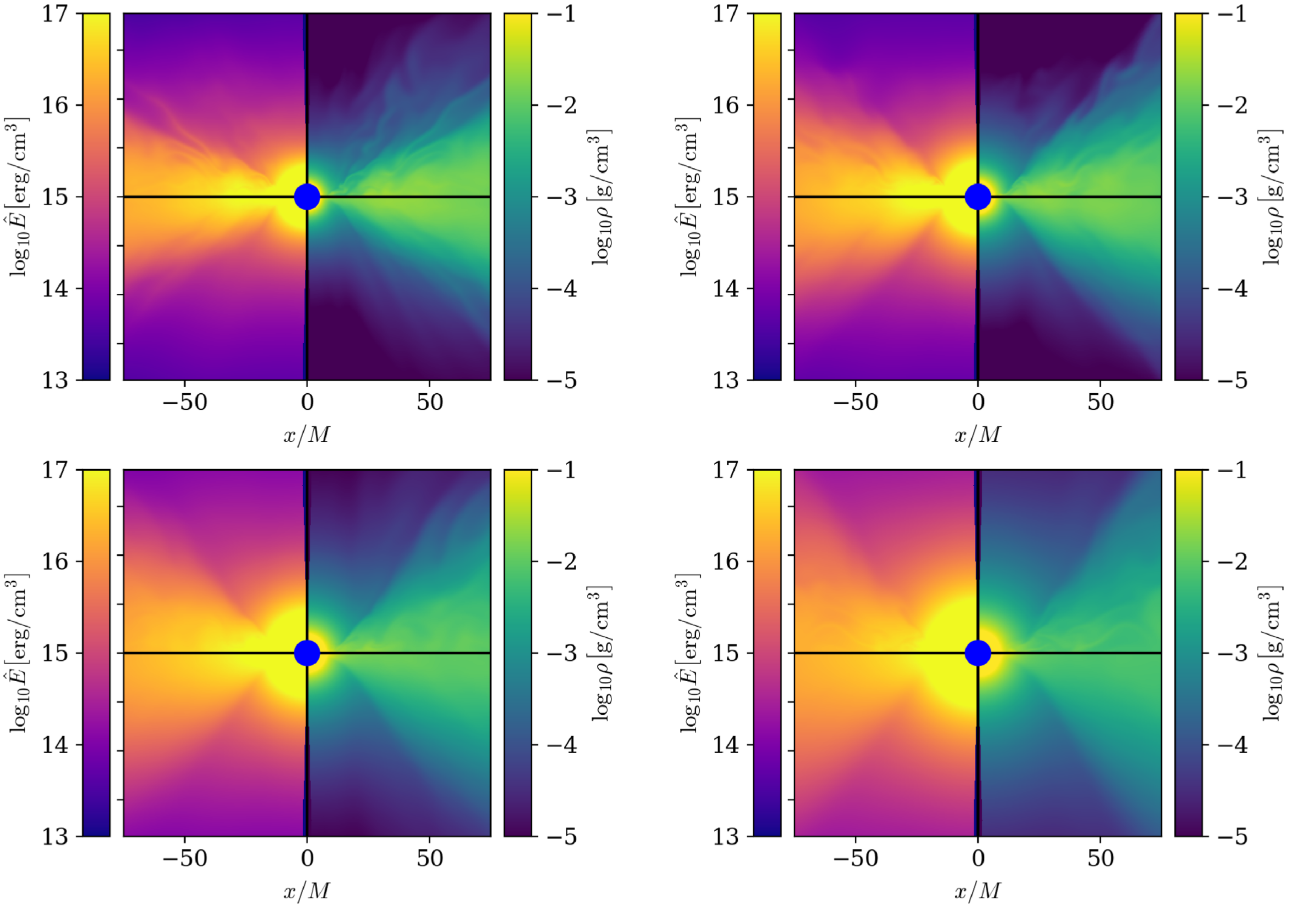

Figure 2. Here, using the same scheme as in Fig. 1, we show four snapshots of NSRUN, the simulation with the reflective boundary condition.

We expect a much different picture in the presence of strong magnetic field. For $\mu \gtrsim 10^{30}$ G cm^3, we expect the accretion flow to be directed along magnetic field lines and deposited at the poles for a dipolar magnetic field as was seen in Takahashi & Ohsuga (2017). The magnetic field would also arrest any outflows from the innermost regions, so it is likely that an optically thin funnel would also be able to form in this case, leading to much larger observed luminosities.

We show the poloidal accretion rate over an even larger range of r and z in the lower two panels of Fig. 4. We can no longer see the structure of the conical outflows in NSRUN which was decollimated at larger radii. We also show a contour of the relativistic Bernoulli number in green, Be = 0, which we define and discuss in more detail in Section 4.2. We can clearly see that more material is launched in NSRUN than in BHRUN.

The region behind the initial torus in our simulation is unphysical. In a real physical system, this region would be occupied by the gas flowing from the companion star. We exclude this region in a somewhat arbitrary way. We pick the last streamline which originates from the surface defined by Be = 0 (shown in white) which does not turn back towards the disc mid-plane at larger radii. The gas inside this region is expected to flow back towards the disc or to the companion star, and so we exclude this region from all of our calculations.

In Fig. 5, we show the time averaged (again, over t_4) spatial structure of both discs. The colour map of density has been expanded to cover a wider range of densities with the trade off being a lack of detail. We show two zoom levels. Red, white, and green contours of the scattering optical depth are also shown. From green to white to red, the contours correspond to $\tau_{\rm sca} = 0.01, 0.1, 1, 10, 100$. Here, $\tau_{\rm sca}$ is calculated in a simple way,

$$\tau_{\rm sca}(r) = \int_r^{r_{\rm out}} \rho\kappa_{\rm es}\sqrt{g_{rr}}\,{\rm d}r, \qquad (9)$$

where $\kappa_{\rm es} = 0.34$ g cm^{-2}. For reference, the root-mean-squared polar scale height of the disc defined by $h = \sqrt{\int(\theta-\pi/2)^2\rho{\rm d}\theta/\int\rho{\rm d}\theta}$. For BHRUN (left), super-Eddington accretion leads to a thick disc with strong outflows and an optically thin funnel region reaching all the way down to the inner boundary. Again, we see a different picture for NSRUN (right). First, as is also seen in Fig. 2 a large amount of gas is deposited in a transition layer around the inner edge. In addition, a large amount of gas is ejected and the entire domain is filled with a thick outflow. Measuring the optical depth shows that this outflow is extremely optically thick. The last contour of the scattering optical depth visible in the figure is given by the red line which corresponds to $\tau_{\rm sca}$ of 100. A zoom of the entire simulation domain is shown in the bottom two panels. We can see the photosphere for NSRUN lies on average at a couple thousand $r_{\rm g}$, which is indicated by the white line.

In Fig. 6, we show the radial profiles of both simulations. BHRUN and NSRUN are shown in blue and orange, respectively, for the time interval t_4. The green lines shows the later evolution of NSRUN over

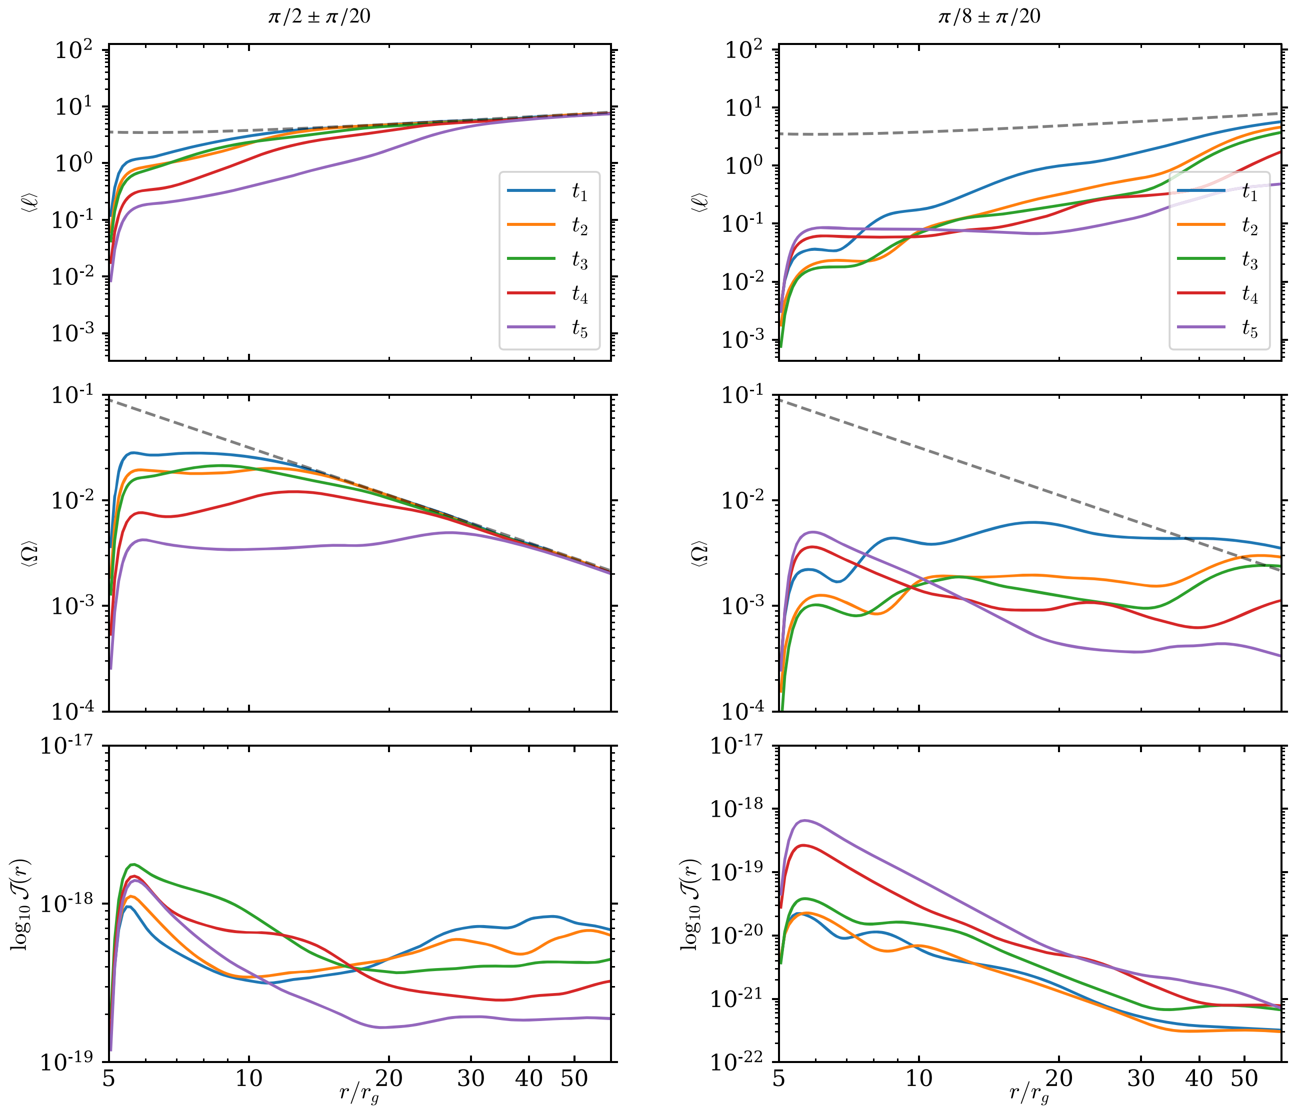

Figure 3. We plot angular momentum related quantities in the disc mid-plane ($\theta = \pi/2$, left) and in the polar region of the simulation ($\theta = \pi/8$, right). The plots were averaged/integrated over θ over a range of angles corresponding to $\pi/10$ radians. Top: Here we show a time series of the specific angular momentum of the gas, $\ell = u_\phi$ for NSRUN. The plotted quantities correspond to a ρ-weighted θ-average of the time averaged simulation data. The durations for the averaging are the same as for Figs 1 and 2, $t_1 = 7\,500\,t_g$, $t_2 = 15\,000\,t_g$, $t_3 = 30\,000\,t_g$, $t_4 = 60\,000\,t_g$ with an additional time interval centred at, $t_5 = 120\,000\,t_g$. The Keplerian angular momentum is shown by the dashed grey line. Middle: Here is the same plot as above except for angular velocity, $\Omega = u^\phi/u^t$. Bottom: Here, we are plotting the total angular momentum density integrated along the θ direction.

time interval t_5. The first quantity shown is the surface density,

$$\Sigma = \int_{\pi/2-h}^{\pi/2+h} \rho\, r \sin\theta \mathrm{d}\theta \tag{10}$$

integrated within one scale height. The three remaining plots correspond to density weighted, θ averages, which for a given quantity f, appears as

$$\langle f \rangle = \frac{\int_{\pi/2-h}^{\pi/2+h} \rho\, f \sqrt{-g}\mathrm{d}\theta}{\int_{\pi/2-h}^{\pi/2+h} \rho \sqrt{-g}\mathrm{d}\theta}. \tag{11}$$

The scale height captures a majority of the mass in BHRUN and can be a reasonable approximation for the boundary of the accretion disc. It becomes a less useful tool, however, in the inner edge of NSRUN, where the distribution of density becomes roughly spherical, and the distinction between disc and transition layer requires more information. Nevertheless, this only affects the weighting of the quantities, which, if they are also roughly spherical, should still give a good measure of the radial profile. We can see that the transition layer surface density follows a power law before settling into the accretion disc which matches that of the black hole at about radius $r = 30\,r_g$. The second panel shows u^r/c, the density weighted radial component of the four velocity. The inflow velocity for NSRUN remains non-relativistic. The local maximum at the inner edge is related to the circularization pattern seen in Fig. 4. The third panel shows density-weighted temperature, T. The temperature inside of the atmosphere approaches $10^{8.5}$ K, and also follows a power law at the inner edge. We also measured the growth rate of the temperature on the surface of the NS and found that it grows with a $T(t) \propto t^{1/4}$ dependence. The last panel shows the density weighted ratio of

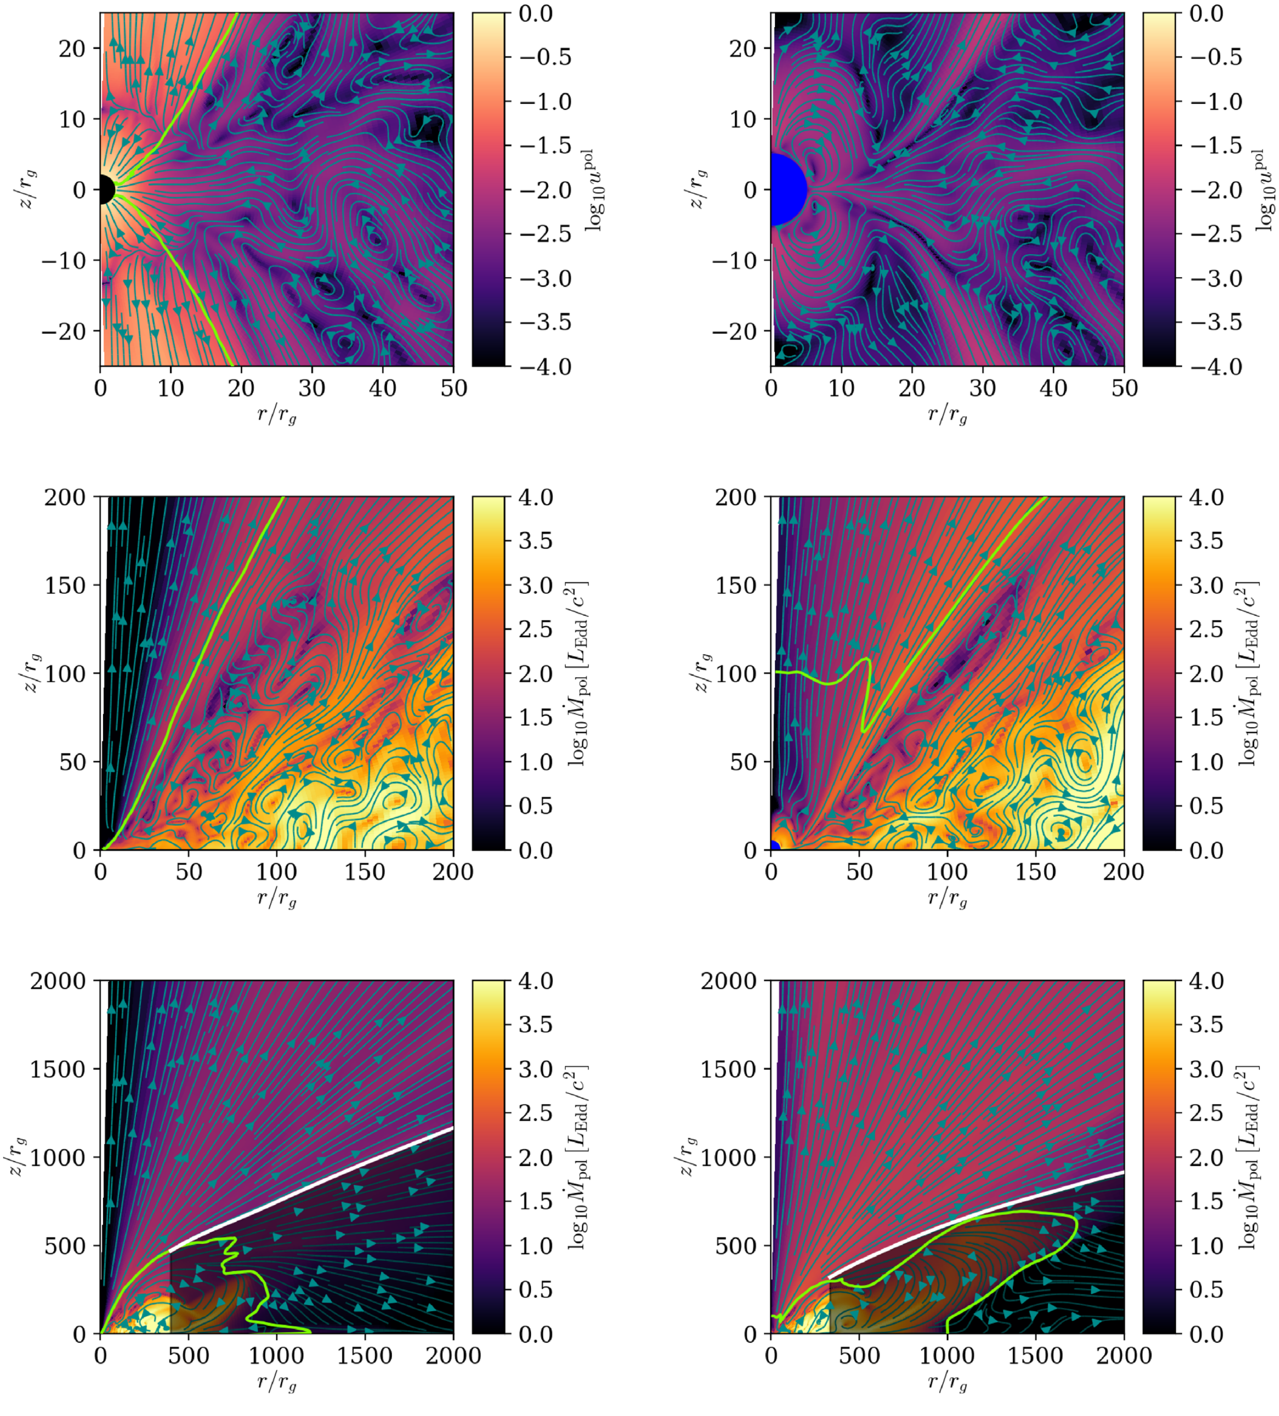

Figure 4. Here, we show the time-averaged poloidal velocity flow structure of the black hole-like simulation (upper left) and reflective simulation (upper right). The magnitude of the poloidal velocity ($u^{\rm pol} = \sqrt{u^r u_r + u^\theta u_\theta}$) is given by the colour map. The streamlines follow the poloidal velocity vectors for all plots in the figure. The middle two images show the mass accretion rate in the poloidal plane ($\dot{M}_{\rm pol} = 2\pi r^2 \sin\theta\, \rho u^{\rm pol}$) for BHRUN simulation (middle left) and NSRUN (middle right). The two bottom images are the same as for the middle images, except for a larger range of r and z. We show a contour of Be = 0 in green and in white we show the last streamline which originates from the Be = 0 surface. The shaded region below, we exclude from our analysis.

angular velocity to the Keplerian angular velocity, $\Omega/\Omega_{\rm k}$. We can also see the angular velocity transition towards zero as observed in Fig. 3. The accretion disc is mildly sub-Keplerian.

4.2 Outflow and luminosities

Accurate measurements of the outflow from black hole and NS accretion discs are important for estimating feedback on the surrounding media of these systems. We measure the radiative and kinetic luminosities of our accretion flows as well as the mass outflow rates. We purposely chose a very large simulation domain so that we can measure these quantities out to larger radii, however, these measurements are meaningless without indication that the simulation has reached some sort of equilibrium. Typically, this is measured by computing the viscous time of time averaged data over a particular interval, and seeing at which radius the viscous time

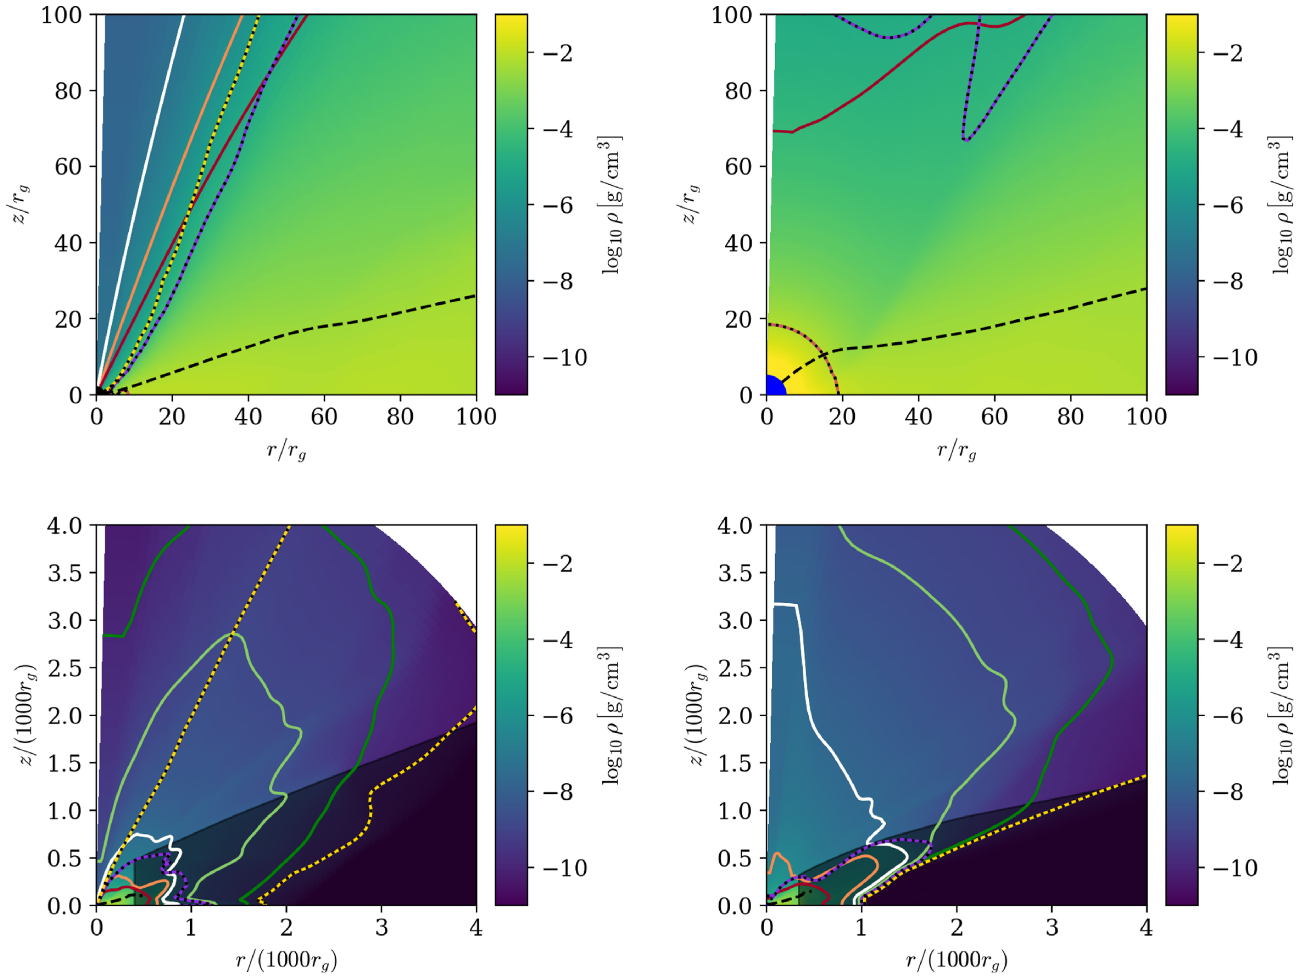

Figure 5. Here, we show the time-averaged structure of the BHRUN (top left) and the NSRUN (top right). Expanded views of the simulations are shown in the bottom panels. The colour map shows density. The coloured contours correspond to optical depth to scattering. From green to red, the values correspond to $\tau_{\rm sca} = 0.01, 0.1, 1, 10, 100$. The while line shows $\tau_{\rm sca} = 1$. The striped red, purple, and yellow lines show contours of Be $= -0.05$, 0, 0.05 respectively. The dashed black line shows the rms scale height, h. We have also darkened the region contained by the last streamline to originate from the Be $= 0$ surface.

is equal the length of that time interval (Narayan et al. 2012). This roughly corresponds to about $r = 80r_{\rm g}$ for our simulations averaged over time interval t_4, evidence of which can be seen in Fig. 6, where the averaged radial velocity changes sign at about 100 $r_{\rm g}$, an indication that inflow equilibrium has not been reached passed this radius. This means that we can trust the results of our simulation inside radius $r = 80r_{\rm g}$.

Outside $80r_{\rm g}$, however, we can still trust the results of the outflow, as long as it is causally connected to the converged region inside radius $r = 80r_{\rm g}$. Because the velocity of the outflow is much higher than that of the inflow, this corresponds to much larger radii where we can reliably measure the outflow. To quantify this, we measure the density weighted average velocity, u^r as a function of radius. We take the average over the θ coordinate only for cells with Be > 0 to reasonably track outflowing gas. Here, Be is the relativistic Bernoulli number,

$$\mathrm{Be} = -\frac{T^{t}{}_{t} + R^{t}{}_{t} + \rho u^{t}}{\rho u^{t}}. \qquad (12)$$

In steady state, Be is conserved along streamlines and gas with positive Be is energetic enough to escape to infinity, and so it is a reasonable parameter to define the outflow. Separating the outflow from the disc flow is not precise, so one should not put too much influence on the Bernoulli parameter. For example, we can see from Fig. 4 that the outflow from near the NS transition layer appears to change Bernoulli parameter. We have investigated this further and found that when looking at the non-time-averaged data in this particular region, the flow does not appear to be steady, and the Be $= 0$ surface is highly variable. For steady flows (such as the outflow at larger radii), the Bernoulli parameter should still be appropriate for calculating the causally connected radius, $r_{\rm cc}$.

We use time-averaged data for the time interval, t_4, from 40 000 $t_{\rm g}$ to 80 000 $t_{\rm g}$, and multiply the time and θ-averaged velocity by the duration of the time interval, $\Delta t = 40\,000\,t_{\rm g}$ to find the causally connected radius, $r_{\rm cc}$, for an outflow of that velocity. This assumes a constant velocity along the outflowing trajectory, however, since velocity tends to decrease with radius, the causally connected radius we find is a lower limit. Then, as long as the causally connected radius is larger than the radius where we are measuring the velocity, we can believe the results of the outflow. A plot of $r_{\rm cc}$ is shown in Fig. 7. We can see that the outflow of BHRUN is reliable throughout the entire outflow region excluding the region behind the initial torus shown in Fig. 4. For NSRUN, it is reliable to about $r \approx 4000r_{\rm g}$. It is important to note that this analysis does not take into account changes in the nature of the outflow with time, but only that the outflow is causally connected to the central region of the simulation.

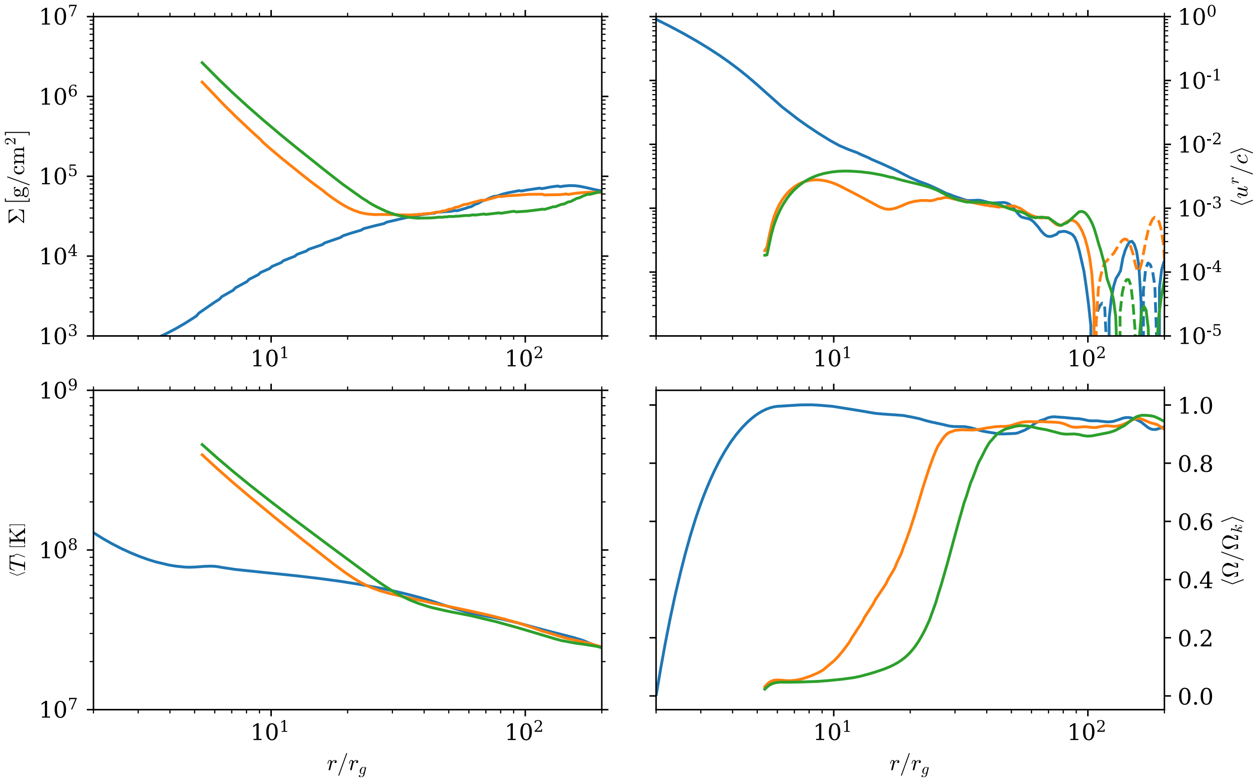

Figure 6. Here, we show radial profiles of surface density, Σ, density-weighted, θ-averaged inflow velocity, the density-weighted, θ-averaged temperature, and the density-weighted, θ-averaged ratio of angular momentum to the Keplerian angular momentum. The data is time averaged over the range, t_4 from $40\,000\,t_{\rm g}$ to $80\,000\,t_{\rm g}$ and θ-averaged over one scale height, h. Plots are shown for the both BHRUN (blue) and NSRUN (orange). We also plot in green the same quantities from NSRUN time averaged over the interval t_5 from $80\,000\,t_{\rm g}$ to $160\,000\,t_{\rm g}$.

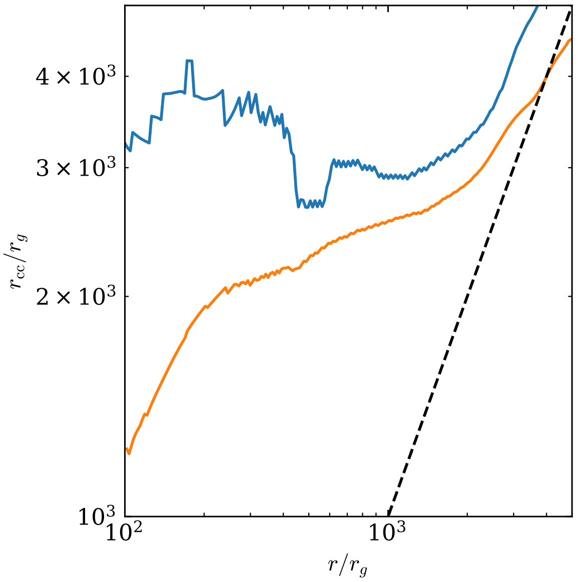

Figure 7. Here, we demonstrate the convergence of the outflow by plotting the radius causally connected to the inner region of the disc, $r_{\rm cc}$, as a function of the radius at which measure the velocity, r. For $r_{\rm cc} > r$, we can reliably measure the outflow of the simulation. For reference, $r = r$ is shown by the black dashed line. As with Fig. 6, blue corresponds to BHRUN and orange to NSRUN. If $r_{\rm cc} > r$, then the gas had enough time to reach radius r.

An additional issue which must be addressed is contamination of the results by the initial condition. Our initial condition is an equilibrium torus where the total pressure has been distributed between gas and radiation. The torus is constructed assuming radiation pressure domination, which implies an adiabatic index, $\gamma = 4/3$. However, the effective adiabatic index depends on the particular mix of gas and radiation at a particular position, and so does not everywhere equal 4/3, especially in cooler parts of the torus. For this reason, the torus is not in perfect hydrostatic equilibrium. There is also a certain amount of gas that is blown off of the outer edges of the torus due to radiation pressure. However, the whole torus was constructed with Be < 0, and so due to the lack of dissipation or viscosity in the regions in the outer parts of the torus where the MRI has yet to develop, there is a reasonable expectation that Be remains less than zero. Therefore, by measuring the outflowing quantities over regions with Be > 0, we are reasonably screening the contamination by the initial condition. Additionally, this region overlaps significantly with the area excluded by the last streamline from the Be $= 0$ surface. We can see which region this corresponds to by examining the purple dashed line in Fig. 5. This region is relatively small compared to the computational domain of the simulation. A more robust diagnostic would be for example, evolving a tracer along with the simulation to track the evolution of gas. Then, it would be possible to check whether the gas originated from the inner, converged regions of the simulations or not.

In Fig. 8, we show the angular distribution of radiative flux, $F^r = R^r{}_t$, as a function of θ. The units of the flux are given in the inferred isotropic luminosity by multiplying by $4\pi r^2$. We plot the flux at four different radii, $50r_{\rm g}$, $500r_{\rm g}$, $2000r_{\rm g}$, and $4000r_{\rm g}$. We

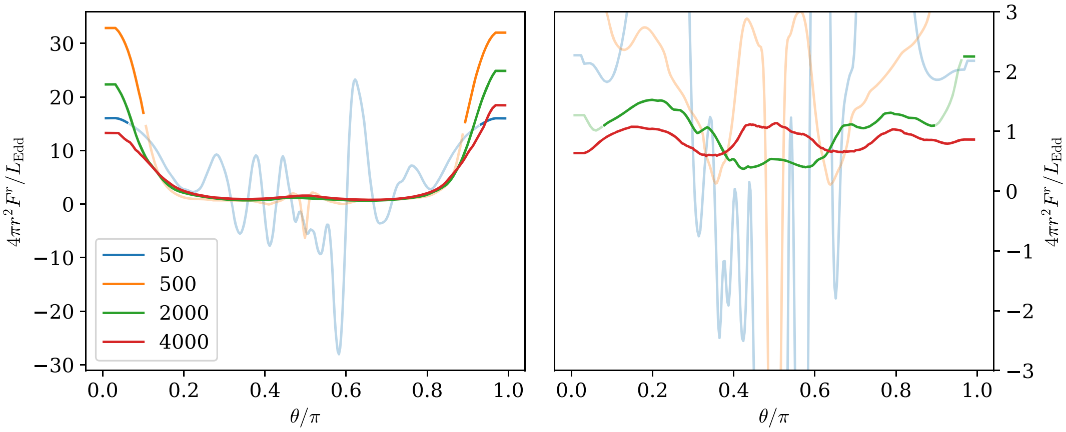

Figure 8. Here, we show the radial distribution of the flux in units of the apparent isotropic luminosity, $4\pi r^2 F^r/L_{\rm Edd}$, as a function of the line of sight for BHRUN (left) and NSRUN (right). Each line corresponds to the observable flux at a particular radius shown in the legend in units of $r_{\rm g}$. The faded regions correspond to areas where the scattering optical depth is greater than one.

have emphasized the regions where the gas is optically thin, as the radiation in the optically thick regions is not expected to reach the observer. The optically thick regions should not be completely ignored, however, it is still possible for the outflowing matter to add or subtract from the radiation field, and so it is difficult to say what the distribution of radiation will look like at infinity. The BHRUN behaves as expected. The optically thin funnel region produces locally super-Eddington fluxes. We can see that the radiation is highly collimated, even at large radii. The intensity of radiation increases with radius at lower radii, and decreases at larger radii, as the disc and outflowing gas emit and absorb radiation. NSRUN is more complicated, however. Due to the thick outflows the photosphere is pushed nearly to the edge of the simulation domain. Only at very large radii can we measure a flux of locally about Eddington. Larger accretion rates may be required to reach higher luminosities, although it is hard to say whether or not the increase in outflowing gas will cancel any increase in luminosity. Indeed, in the case of black holes, it is well known that at super-Eddington accretion rates, an increase in the accretion rates corresponds to a decrease in the radiative efficiency so that the luminosity scales logarithmically, $L \propto \log(1 + \dot{M}/\dot{M}_{\rm Edd})$ (Shakura & Sunyaev 1973). It may even be interesting to try smaller accretion rates in case the outflow scales differently than the luminosity in the presence of a reflective boundary condition.

Besides observational properties, the implications of black hole feedback play an important role in the astrophysics of star formation, star clusters, and galaxies. In Fig. 9, we show the luminosity of kinetic energy, defined as,

$$L_{\rm KE} = 2 \times 2\pi \int_0^{\theta_{\rm out}} -(u_{\rm t} + \sqrt{-g_{\rm tt}})\rho u^r \sqrt{-g}\, {\rm d}\theta, \qquad (13)$$

where $\theta_{\rm out}$ is the angle at which Be = 0. By integrating over positive Be, we are choosing only the gas which is energetic enough to reach infinity. The extra factor of two reflects the fact that we include the contribution from both sides of the equatorial plane. In the non-relativistic limit, the integrand approaches $1/2\rho v^2 v^r$, where v^i is the three-velocity, v^2 is the square of the three-velocity, and v^r is the radial component of the three-velocity. NSRUN shows a large amount of material ejected into the surroundings of the NS star environment. A large amount of it remains bound and so only at around $r = 100 r_{\rm g}$ is $\theta_{\rm out} > 0$. NSRUN everywhere has a less energetic outflow than BHRUN, levelling off at $L_{\rm KE} \approx 0.4 L_{\rm Edd}$. For BHRUN, the gas is energetic enough to escape to infinity in the funnel region even at radii below the stagnation radius, however the velocity vector is directed inward and so the flux of kinetic energy is into the inner boundary, thus the transition to negative values in the kinetic luminosity.

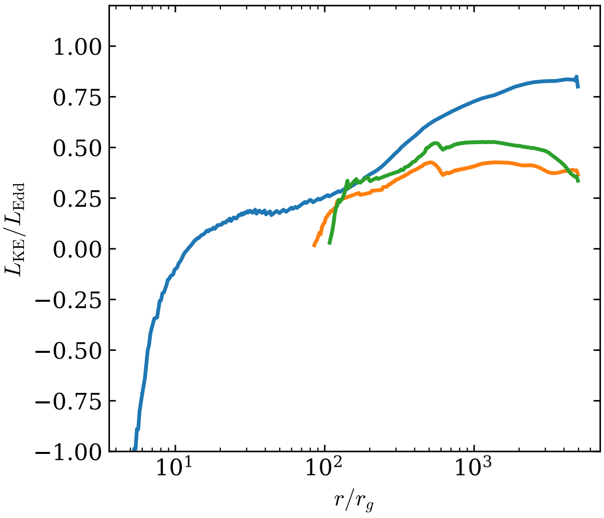

Figure 9. Here, we plot the outflowing luminosity of kinetic energy $L_{\rm KE}$, as a function of radius. The kinetic energy flux is integrated over the outflowing region defined by positive relativistic Bernoulli number, Be. BHRUN is shown in blue, and NSRUN is shown in orange averaged over time interval, t_4, and green averaged over t_5.

The total radiative output of the simulation is difficult to measure directly, especially for optically thick flows. Radiation and gas can still exchange energy even after the gas has left the simulation domain. We chose to run our simulations with a very large outer radius, so that the photosphere is still contained in the simulation. This allows us to measure the amount of radiative energy outflowing from optically thin and outflowing regions. Radiation flowing through optically thin regions is expected to escape to infinity, and so is an effective lower limit on the total radiative output. We define the optically thin region by the angle at a particular radius where the radial scattering optical depth is equal to 1. For each radius, we

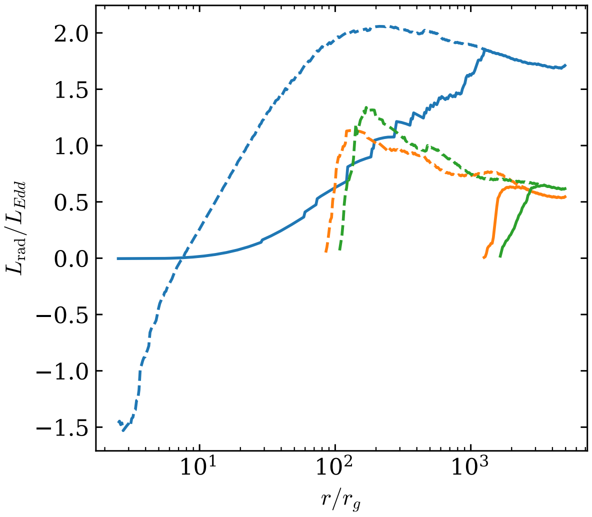

Figure 10. Radiative luminosities are shown for BHRUN (blue) and NSRUN (orange) averaged over the time period t_4. The last time period for NSRUN, t_5 is shown in green. The solid lines correspond to the luminosity integrated over angles that are optically thin, $L_{\rm thin}$, while the dashed lines show the luminosity integrated over the angles with positive Be, $L_{\rm out}$.

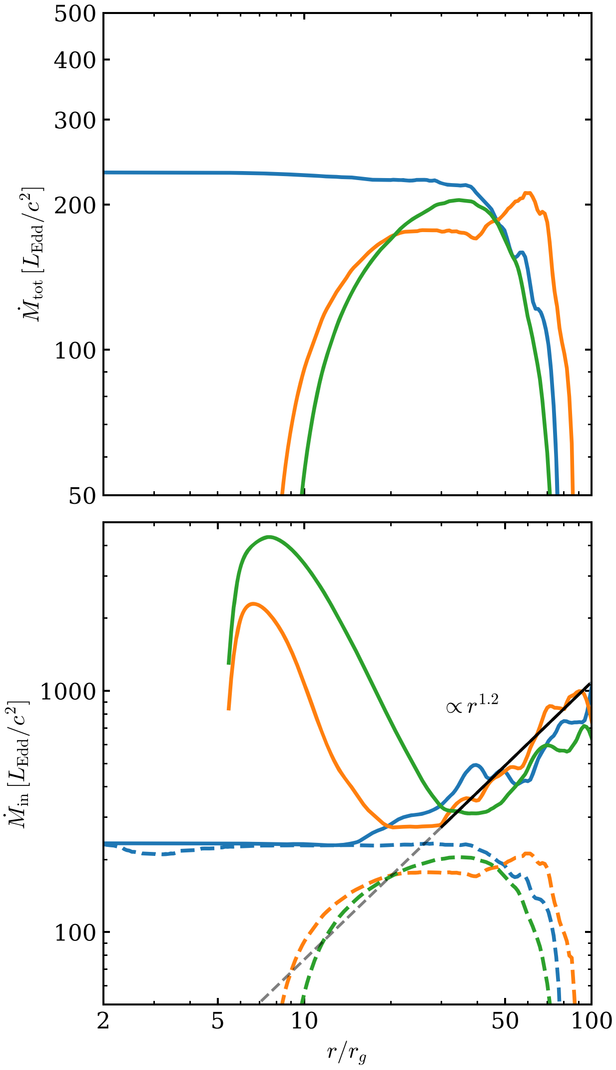

Figure 11. Top: Here, we show the total accretion rate, $\dot{M}_{\rm tot}$, of BHRUN (blue) and NSRUN (orange) computed from data averaged over the time interval, t_4. In green, we show data from NSRUN time averaged over the interval t_5. Bottom: Here, we plot the mass accretion rate of the inflow, $\dot{M}_{\rm in}$, again for BHRUN and NSRUN from t_4 and t_5. We include two definitions of the inflow. The first corresponds to the solid lines which are defined as the integral over all cells with negative u^r, $\dot{M}_{{\rm in},u^r}$. The dashed line shows the the integral over all cells with negative Be, $\dot{M}_{{\rm in},Be}$. A power law fit to $\dot{M}_{{\rm in},u^r}$ for NSRUN from t_4 is shown in black. The solid black line shows the radii over which the fit was performed, the dashed part is an extension of the power law for reference.

integrate the radial radiative flux over the optically thin region,

$$L_{\rm thin} = 2 \times 2\pi \int_0^{\theta_{\rm thin}} -R^r{}_t \sqrt{-g}\,{\rm d}\theta. \qquad (14)$$

A reasonable upper limit of the radiative luminosity, is the integral of F^r over the region of $\mathrm{Be} > 0$, $L_{\rm out}$, as in the calculation of $L_{\rm KE}$. This gas is energetic enough to reach infinity, and so it is possible for the outflowing gas to eventually release its trapped photons. Note that it is possible for this gas to produce additional photons, but this is not thought to contribute significantly to the radiative luminosity, so $L_{\rm out}$ is not a strict upper limit. The outflowing luminosity is then given as,

$$L_{\rm out} = 2 \times 2\pi \int_0^{\theta_{\rm out}} -R^r{}_t \sqrt{-g}\,{\rm d}\theta. \qquad (15)$$

The outflowing and optically thin radiative luminosities are shown in Fig. 10. In general, BHRUN shows about 2–3 times as high luminosity as NSRUN. As shown in Figs 1 and 2, NSRUN produces a much larger amount of radiation energy. However, the vast majority is trapped in the optically thick outflow and transition layer. We also observe a general decrease in luminosity at radii larger than about $300r_{\rm g}$. This is because of radiation flowing into the excluded regions behind the initial torus. The optical depth of this region is low and so radiation can escape from the outflow into this region. It is hard to say how the radiation would behave in this region in reality, it would depend on the extent and thickness of the accretion disc. We do not seem much change in the radiative luminosity over the t_5 time interval.

Super-Eddington accretion flows are known to have strong radiation driven outflows (Ohsuga et al. 2005; Hashizume et al. 2015; Fiacconi et al. 2017). We can measure this by examining the accretion rate in a few different ways. First in the top panel of Fig. 11, we measure the total (net) mass accretion rate as a function of radius,

$$\dot{M}_{\rm tot} = 2\pi \int_0^{\pi} -\rho u^r \sqrt{-g}\,{\rm d}\theta. \qquad (16)$$

BHRUN, which is shown in blue, gives a flat accretion rate out to about $r \lesssim 50r_{\rm g}$, which roughly corresponds to the region where the simulation has achieved steady state. NSRUN, shown in orange, is also nearly flat for a small region between $20r_{\rm g}$ and $50r_{\rm g}$. At low radii the simulation is accumulating mass, which corresponds to a non-flat slope of $\dot{M}_{\rm tot}$. In general, $\dot{M}_{\rm tot}$ is lower for NSRUN than for BHRUN. This is not due to a lower accretion rate in the disc, but due to the fact that the gas that is normally lost through the inner boundary of BHRUN either accumulates in the transition layer where it is recycled into the inner accretion flow, or is ejected into the outflow, the latter of which contributes to a lower value of $\dot{M}_{\rm tot}$. We also show $\dot{M}_{\rm tot}$ for NSRUN averaged over the time interval t_5. We can see that the flat region indicating an accretion disc in steady state is no longer present. This is due to two effects. The first is the transition layer increasing in radius. The second is due to the fact that the outflows are increasing at higher radii and so the total accretion rate over the whole sphere decreases.

The second panel of Fig. 11 shows the inflowing accretion rate as measured in two different ways. The first is a more naive measurement of the inflow accretion rate, (17)

$$\dot{M}_{\mathrm{in},u^r} = 2\pi \int_{u^r<0} -\rho u^r \sqrt{-g}\mathrm{d}\theta, \tag{17}$$

where we simply sum over individually inflowing cells. These are the solid curves in the second panel, and in general, BHRUN and NSRUN agree with each other outside of the NS transition layer. This indicates that they are accreting at the same accretion rate, the fundamental scale of accretion, and we can expect the same behaviour from their accretion discs. The inflowing accretion rate, $\dot{M}_{\mathrm{in},u^r}$, is a good measure to determine the mass accretion rate in the disc when most or all of the disc turbulence is averaged out, then all of the gas in the disc should have $u^r < 0$, and all of the gas in the outflow should have $u^r > 0$. This is not the case for the data we show except at small radii. For flows where the turbulence is not fully averaged out, $\dot{M}_{\mathrm{in},u^r}$ is more of an upper limit. A super-Eddington accretion flow is expected to have an accretion rate that increases linearly with r (Shakura & Sunyaev 1973), and we can see from the black line, that this dependence is nearly reproduced in NSRUN. We obtain $\dot{M}_{\mathrm{in},u^r} \propto r^{1.2}$.

We can still see a large amount of gas flowing through the NSRUN accretion disc for the t_5 interval with nearly the same slope as for t_4, although at slightly lower values. This is likely due to the structure of the initial torus, which over very long periods of time, does not supply the accretion disc with gas at the same rate, slightly dropping with time.

The dashed lines correspond to the integral of the total mass flux over the regions with negative Bernoulli number, (18)

$$\dot{M}_{\mathrm{in,Be}} = 2\pi \int_{\mathrm{Be}<0} -\rho u^r \sqrt{-g}\mathrm{d}\theta. \tag{18}$$

Most of the gas at lower radii has negative Bernoulli number, whether or not it is part of the disc or the outflow. From Fig. 4, we can see streamlines that are initially outflowing with negative Be which turn around and return to the disc mid-plane, but we can also see streamlines which seem to change from negative to positive Be, and which continue to flow outwards. It is important to note that Be is only conserved along streamlines for steady flows. Closer to the disc, the flow is very much non-steady, and so the streamlines do not always reproduce the gas trajectory. For this reason, one must be cautious when using Be as a diagnostic for outflow. It is necessary to know something about the flow, and so we put more weight on Be farther from the disc where the flow is more steady. In this way, Be is more useful as a diagnostic for outflows at larger radii. We still include plots of $\dot{M}_{\mathrm{in,Be}}$ for completeness, although they should mostly match $\dot{M}_{\mathrm{tot}}$ at lower radii. The departure of $\dot{M}_{\mathrm{in,Be}}$ from $\dot{M}_{\mathrm{in},u^r}$ for BHRUN is also evident from the top-left panel of Fig. 4, where there is gas which has positive Be but is directed into the black hole.

Fig. 12 shows the outflowing accretion rate, measured in the same way as the inflowing accretion rate, $\dot{M}_{\mathrm{in},u^r}$, but with opposite sign, $\dot{M}_{\mathrm{out},u^r}$. Again, we show BHRUN in blue and NSRUN in orange. We also plot in green the mass outflow rate for NSRUN over the last interval, t_5 from 80 000t_{g} to 160 000t_{g} which was not run for the BHRUN, because it is more computationally demanding. At lower radii, $\dot{M}_{\mathrm{out},u^r}$ more or less follows $\dot{M}_{\mathrm{in},u^r}$, this is because, in inflow equilibrium, the difference between them, $\dot{M}_{\mathrm{tot}}$, should be constant. As the accretion rate drops with a decrease in radius, the outflow rate must also drop so that no gas can accumulate at any radius, at least for BHRUN. We have also de-emphasized $\dot{M}_{\mathrm{out},u^r}$ at radii larger than $100r_g$, where the simulation has not reached inflow equilibrium.

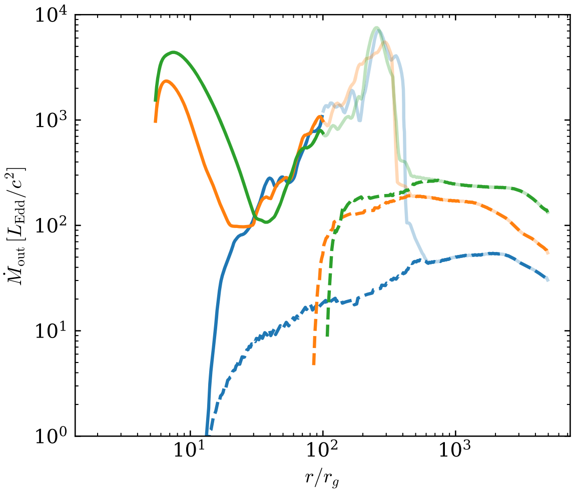

Figure 12. Here, we show the outflowing mass accretion rate, $\dot{M}_{\mathrm{out}}$. $\dot{M}_{\mathrm{out},u^r}$ is given by the solid lines and $\dot{M}_{\mathrm{out,Be}}$ is given by the dashed lines. Blue lines correspond to the BHRUN and orange to NSRUN both using time-averaged data from the interval t_4. The green lines shows data from NSRUN time-averaged over the interval, t_5.

For measuring the mass outflow that is expected to reach infinity, the relevant quantity is $\dot{M}_{\mathrm{out,Be}}$, the total mass outflow rate in the positive energy region, the region which is separate from the accretion disc. These are the dashed curves in Fig. 12. $\dot{M}_{\mathrm{out,Be}}$, is calculated in the same way as, $\dot{M}_{\mathrm{in,Be}}$ except for Be > 0. In general, NSRUN shows nearly an order of magnitude larger outflow rate than BHRUN at its peak when measured over the same time period. This is why the outflows are so optically thick. The maximum outflowing accretion rate for NSRUN nearly matches $\dot{M}_{\mathrm{in}}$ at the edge of the transition layer. If the outflow were to converge to this value ($\sim$200 L_{Edd}/c^2) at all radii $\gtrsim$1000r_g, then this would be a strong indication that all the inflowing matter is eventually ejected. We do not observe this however, and moreover we observe that mass is still accumulating at the inner edge. While the simulation was run for a long time, around (160 000t_{g}), the transition layer did not seem to reach a steady state. We can see the outflow at large radii become more flat for t_5, this also indicates that the outflow while apparently increasing with time, may slowly be approaching at least a quasi-steady state, and that outflows which are the same magnitude as the accretion rate in the disc can be expected for long periods of time.

5 DISCUSSION

We have run two long duration GRRMHD simulations, one of accretion onto a Schwarzschild black hole and one onto a NS (non-magnetized, non-spinning). Our black hole simulation, BHRUN, is used as a baseline to compare against the NS simulation for two reasons. One, it is very similar to the inflow boundary conditions used in many previous simulations. Secondly, it is a simple, well-studied system, without an artificial boundary condition, and is directly comparable to other previous simulations of its type and the model more closely matches the behaviour of the physical object it represents. It is immediately apparent that the simulations are quite different. The differences in the accretion disc structure are apparent inside radius $r = 30r_g$, where the flow in NSRUN begins to change into the transition layer. It is then inferred that the differences in the outflow are due to the various physical processes occurring in the inner region of the simulation. We observed from Fig. 3 that the transition layer increases with size as time goes on.

The simulation was run for 120 000 t_g, which only corresponds to 1 s of physical time, which is only one full pulsation for the pulsation periods observed in pulsating ULX systems. These observations were integrated over kilosecond time scales. The point being that even though our simulations are run for a long time (when compared to other GRRMHD simulations), they would have to be run much longer in order to capture the effects of the varying radiation source of a rotating magnetized star. Such long duration simulations are likely out of the question for 3D simulations, likely still difficult for our 2D simulations, and impossible for 2D simulations without any sort of dynamo. Nonetheless, we are free to speculate.

We observe that the radius of the outer boundary of the transition layer increases with a time dependence that is well described by a $t^{0.85}$ power law, and we have no reason not to expect this dependence to continue until either the mass supply is quenched, the radiation pressure becomes too strong, or the central density and temperature become large enough that some other cooling mechanism becomes effective, e.g., neutrino cooling or photon cooling through outflows. The central temperature grows with a $t^{1/4}$ power law, and so neutrino cooling, which may become relevant around 10^9 K may need to be considered (depending on the central density) after a few (4–8) times the simulation length. We expect the other processes to also take much longer than a few times the current simulation time, so it would be interesting to run an even longer duration simulation, in order to study the time evolution of the outflow and luminosity as the simulation evolves. In our case, the length of our simulation was chosen so that the outflow would be causally connected to the inflow equilibrium region of the inner simulation to radii past the photosphere. A non-physical limitation would be the transition layer growing faster than inflow equilibrium. A potential solution would be to run a black hole like simulation for long time and then to restart the simulation with a hard inner boundary.

Our main conclusion from our work is that the presence of a hard, non-rotating surface alone is likely not sufficient to produce enough radiation that can escape to infinity to explain the luminosities observed in ULX sources. The large amount of ejected gas not only obscures the innermost emitting regions, but also decollimates the radiation, further reducing the observed luminosity. Thus, it is vital that we consider the effects of magnetic fields. This was done by Takahashi & Ohsuga (2017), who found a bolometric luminosity to be about an order of magnitude larger than ours, the simulation is short, however, $t_{\rm stop} = 15\ 000 t_g$, so more work would be required to study the effect of a large amount of gas being accreted onto the magnetic field.

It would be interesting to study the case of accretion onto a non-magnetized NS with a different initial condition for the disc magnetic field. Our initial condition contains magnetic field loops of alternating polarity, and reconnection is allowed on a numerical level. Thus, no significant amount of magnetic flux accumulates on the inner edge of the simulation, indeed the magnetic pressure is several order of magnitude lower than the radiation pressure in the transition layer. Running a simulation with only one magnetic field loop, which normally leads to accretion in the magnetically arrested disc state in black hole simulations (McKinney, Dai & Avara 2015; Narayan, Sądowski & Soria 2017), would likely also lead to a build-up of magnetic flux in the transition layer which could be dynamically important.

One interesting application of our study could be an extension to supersoft sources (SSS) (Long, Helfand & Grabelsky 1981; van den Heuvel et al. 1992) and ultraluminous supersoft sources (ULS) (Urquhart & Soria 2016). ULSs are characterized by peak temperatures of around 10^6 K, with bolometric luminosities at a few times 10^{39} erg s^{-1} and photospheres at radii of around 10^4 km. SSSs are classified similarly with lower temperatures ($\sim 10^5$ K, $\sim 10^6$ K) and luminosities ($\sim 10^{36}$ erg s^{-1}, $\sim 10^{38}$ erg s^{-1}). The effective temperature of NSRUN can be computed by measuring the radiative energy density around the photosphere which is on average around a few times 10^6 K. The radius of the photosphere is around 5000 km. These features are indicative of ULS sources, although the bolometric luminosity is about an order of magnitude too low, closer to the higher end for SSSs. While SSSs can be well explained by nuclear burning on the surface of a white dwarf (van den Heuvel et al. 1992), ULSs, like ULXs have been harder to model. It would be interesting in future work to do post-processing of NSRUN to generate spectra to see if they match any of the more specific characteristics of ULSs or SSSs.

5.1 Comparision to Takahashi, Mineshige & Ohsuga (2018)

After this work was completed, we became aware of an important publication by Takahashi et al. (2018) (TMO18), who obtained similar results to ours by running a shorter simulation (7000 t_g), demonstrating the feasibility of super-Eddington accretion onto a non-magnetized NS. It is vital that we make a comparison to the work by TMO18, as the simulations they have run should be directly comparable to ours. TMO18 have performed two simulations, one onto a black hole, and one onto a reflective inner boundary as in our work. Both codes are GRRMHD solvers with M_1 closure. The code used is the same as used by the authors in their previous work. Both our simulations and those of TM018 show a power-law behaviour at the inner edge of the simulation. TMO18 report a significant increase in the rate of ejected mass over their black hole simulation. Indeed, we also infer a roughly 10 times increase in the mass outflow rate for NSRUN over BHRUN. The authors also report an increase in the radius of the photosphere over the black hole case. We find a much larger increase in the photosphere however, reaching out to a couple thousand r_g. This is likely due to the length of our simulations, which are run for more than ten times the length of the simulations in TMO18, and so we believe we are able to make a more accurate statement about the mass outflow to larger radii, and about the radiative luminosity, which TMO18 does not address.

One interesting difference is found in the strong change in the angular velocity in the transition layer, seen in our simulations. TMO18 show the centripetal force, which can be used as a proxy for Ω (indeed, in the Newtonian case it is proportional to Ω^2). TMO18 show that the centripetal force transitions from super-Keplerian to about half Keplerian. This contrasts against our simulations where Ω transitions to practically zero at the inner edge of the simulation, even at early times. This difference could be due to differences in the inner boundary condition for u^ϕ. When we average our data over the same time interval as is described in TMO18 (3000–5000 t_g) and examine Ω/Ω_K, we see that it transitions from unity at about 10 r_g to 0.3 at the stellar surface ($5r_g$). This corresponds to the same range of radii over which the centripetal force of TMO108 transitions to about 0.6 times the Keplerian value which would indicate an $\Omega/\Omega_K \approx 0.8$, more than twice the angular velocity ratio of our simulation. This indicates that we have a stronger torque at the inner edge of our simulation, and so it is probably easier for gas to accumulate in the transition layer. It will be beneficial in the future to study angular momentum transfer through the inner boundary for different boundary conditions.

5.2 Caveats

It is important to stress that our simulations are a first attempt at measuring the effects of a hard surface in the context of a global, radiative, MRI-driven, super-Eddington accretion flow. More physics is planned in future implementations. Indeed, there are a number of issues that must be addressed before we may consider these results truly robust.

First, our simulations are performed in 2D axisymmetry. Sądowski & Narayan (2016) have shown that non-axisymmetric effects can lead to a lower luminosity accretion flow around a black hole. Since the luminosities we measured were on the low end of what is expected for ULX objects, we cannot be sure if a 3D simulation would be bright enough. It is possible that this difference would not occur in the presence of hard surface, but a fully 3D simulation would be required to confirm.

Secondly, we have neglected the effects of rotation. Accretion onto a rapidly rotating star will release less energy than accretion onto a stationary star. The gas will also be less bound however. For a 1 s rotational period, the effects of rotation at the NS surface should be negligible; however, the acceleration of gravity scales as r^{-2} while centripetal acceleration increases with r. At larger radii, this may become relevant if the accreted atmosphere rotates nearly uniformly.

6 SUMMARY

We have found that a reflective, non-rotating boundary at the inner edge of an accretion disc simulation has a significant effect on the behaviour and structure of the inner disc as well as its emission and outflows. We observe large amounts of gas accumulating on the inner boundary in a transition layer, where the angular velocity transitions from its Keplerian value to near zero. We also found lower rates in the outflow of kinetic energy over the black hole case. However, we did measure larger mass outflow rates, affecting the release of radiation to the observer. In fact, the hard surface of a non-magnetized NS leads to lower radiative luminosities in super-Eddington flows relative to black holes, and with radiation decollimated to the point where they are not likely to explain even the lowest luminosity ULXs, although they may be applicable to ULSs or SSSs. This work is a first step in a larger plan to study accretion processes around neutron stars. Further work includes studying different accretion rates, including rotation, and eventually the addition of stellar magnetic fields.

ACKNOWLEDGEMENTS

The first author thanks Ramesh Narayan, Maciek Wielgus, and Jean-Pierre Lasota for advice and illuminating conversations. Research supported in part by the Polish National Science Centre grants 2013/08/A/ST9/00795 and 2017/27/N/ST9/00992. Computations in this work were performed on the Prometheus machine, part of the PLGrid infrastructure.

REFERENCES

Alme M. L., Wilson J. R., 1973, ApJ, 186, 1015
Bachetti M. et al., 2014, Nature, 514, 202
Balbus S. A., Hawley J. F., 1991, ApJ, 376, 214
Basko M. M., Sunyaev R. A., 1976, MNRAS, 175, 395
Bradshaw C. F., Fomalont E. B., Geldzahler B. J., 1999, ApJ, 512, L121
Čemeljić M., Shang H., Chiang T.-Y., 2013, ApJ, 768, 5
Chamel N., Haensel P., 2008, Living Rev. Relativ., 11, 10
Colbert E. J. M., Mushotzky R. F., 1999, ApJ, 519, 89
Dhang P., Sharma P., Mukhopadhyay B., 2016, MNRAS, 461, 2426
Ekşi K. Y., Andaç İ. C., Çıkıntoğlu S., Gençali A. A., Güngör C., Öztekin F., 2015, MNRAS, 448, L40
El Mellah I., Casse F., 2015, MNRAS, 454, 2657
Fiacconi D., Pinto C., Walton D. J., Fabian A. C., 2017, MNRAS, 469, L99
Fürst F. et al., 2016, ApJ, 831, L14
Gammie C. F., McKinney J. C., Tóth G., 2003, ApJ, 589, 444
Hashizume K., Ohsuga K., Kawashima T., Tanaka M., 2015, PASJ, 67, 58
Inogamov N. A., Sunyaev R. A., 1999, Astron. Lett., 25, 269
Israel G. L. et al., 2017a, Science, 355, 817
Israel G. L. et al., 2017b, MNRAS, 466, L48
Jiang Y.-F., Stone J. M., Davis S. W., 2014, ApJ, 796, 106
Kawashima T., Mineshige S., Ohsuga K., Ogawa T., 2016, PASJ, 68, 83
King A. R., Davies M. B., Ward M. J., Fabbiano G., Elvis M., 2001, ApJ, 552, L109
King A., Lasota J.-P., Kluźniak W., 2017, MNRAS, 468, L59
Kluźniak W., Lasota J.-P., 2015, MNRAS, 448, L43
Kluzniak W., Wagoner R. V., 1985, ApJ, 297, 548
Kluzniak W., Wilson J. R., 1991, ApJ, 372, L87
Koliopanos F., Vasilopoulos G., Godet O., Bachetti M., Webb N. A., Barret D., 2017, A&A, 608, A47
Komissarov S. S., 1999, MNRAS, 303, 343
Long K. S., Helfand D. J., Grabelsky D. A., 1981, ApJ, 248, 925
McKinney J. C., Dai L., Avara M. J., 2015, MNRAS, 454, L6
Mukhopadhyay B., Fang L. Z., 2002, Int. J. Mod. Phys. D, 11, 1305
Mushtukov A. A., Suleimanov V. F., Tsygankov S. S., Poutanen J., 2015a, MNRAS, 447, 1847
Mushtukov A. A., Suleimanov V. F., Tsygankov S. S., Poutanen J., 2015b, MNRAS, 454, 2539
Mushtukov A. A., Suleimanov V. F., Tsygankov S. S., Ingram A., 2017, MNRAS, 467, 1202
Mushtukov A. A., Verhagen P. A., Tsygankov S. S., van der Klis M., Lutovinov A. A., Larchenkova T. I., 2018, MNRAS, 474, 5425
Narayan R., Yi I., 1995, ApJ, 452, 710
Narayan R., Garcia M. R., McClintock J. E., 1997, ApJ, 478, L79
Narayan R., Sądowski A., Penna R. F., Kulkarni A. K., 2012, MNRAS, 426, 3241
Narayan R., Sądowski A., Soria R., 2017, MNRAS, 469, 2997
Ogawa T., Mineshige S., Kawashima T., Ohsuga K., Hashizume K., 2017, PASJ, 69, 33
Ohsuga K., Mori M., Nakamoto T., Mineshige S., 2005, ApJ, 628, 368
Parfrey K., Tchekhovskoy A., 2017, ApJ, 851, L34
Penna R. F., Kulkarni A., Narayan R., 2013, A&A, 559, A116
Pintore F., Zampieri L., Stella L., Wolter A., Mereghetti S., Israel G. L., 2017, ApJ, 836, 113
Popham R., Sunyaev R., 2001, ApJ, 547, 355
Revnivtsev M., Mereghetti S., 2015, Space Sci. Rev., 191, 293
Romanova M. M., Ustyugova G. V., Koldoba A. V., Lovelace R. V. E., 2012, MNRAS, 421, 63
Shakura N. I., Sunyaev R. A., 1973, A&A, 24, 337
Shklovsky I. S., 1967, ApJ, 148, L1
Sibgatullin N. R., Sunyaev R. A., 2000, Astron. Lett., 26, 772
Syunyaev R. A., Shakura N. I., 1986, Sov. Astron. Lett., 12, 117
Sądowski A., Narayan R., 2015, MNRAS, 454, 2372
Sądowski A., Narayan R., 2016, MNRAS, 456, 3929
Sądowski A., Narayan R., Tchekhovskoy A., Zhu Y., 2013, MNRAS, 429, 3533
Sądowski A., Narayan R., Tchekhovskoy A., Abarca D., Zhu Y., McKinney J. C., 2015, MNRAS, 447, 49
Takahashi H. R., Ohsuga K., 2017, ApJ, 845, L9
Takahashi H. R., Mineshige S., Ohsuga K., 2018, ApJ, 853, 45

Toropina O. D., Romanova M. M., Lovelace R. V. E., 2012, MNRAS, 420, 810

Urquhart R., Soria R., 2016, MNRAS, 456, 1859

van den Heuvel E. P. J., Bhattacharya D., Nomoto K., Rappaport S. A., 1992, A&A, 262, 97

Walton D. J. et al., 2018, MNRAS, 473, 4360

Wiktorowicz G., Sobolewska M., Lasota J.-P., Belczynski K., 2017, ApJ, 846, 17

Yang X.-H., Yuan F., Ohsuga K., Bu D.-F., 2014, ApJ, 780, 79

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

4 Paper 3: Beamed emission from a neutron-star ULX in a GRRMHD simulation

In this letter (Abarca et al. 2021), I run a 2D axisymmetric GRRMHD simulation of super-Eddington accretion onto a neutron star with a 2×10^{10} G dipolar magnetic field. In order to handle the large magnetizations present in the magnetosphere I implement the method from Parfrey and Tchekhovskoy (2017) and adapt it to work with GRRMHD simulations. I also use a boundary condition which is meant to model gas hitting the surface, becoming shocked and releasing a fraction (in this case 0.75) of its kinetic energy as outflowing radiation.

The accretion flow forms a disk with an accretion rate of $\sim 20\, M_\odot$ which is truncated at about $r \approx 10\, GMc^{-2}$ by the magnetic field. At this point the flow is driven along magnetic field lines forming accretion columns. As the gas hits the inner boundary it slows down due to the boundary condition and becomes shocked. A large amount of radiation is released at the base of the column and this radiation becomes collimated by the outflowing gas so that when it reaches the observer it appears to be originating form a source many times brighter. I measure peak isotropic luminosities about 140 L_{Edd} while our lower limit for the actual luminosity is around $2.5\ L_{\mathrm{Edd}}$, implying a maximum beaming factor of $b \approx 1/50$. This is quite a large value and due to the limitations of the simulation there are reasons to believe that the actually inferred isotropic luminosity maybe be a few times lower, but even taking that into account, I have shown that an accreting neutron star with a relatively weak field can emit beamed emission and should be a good candidate for a ULX.

The Astrophysical Journal Letters, 917:L31 (7pp), 2021 August 20 https://doi.org/10.3847/2041-8213/ac1859

Beamed Emission from a Neutron-star ULX in a GRRMHD Simulation

David Abarca[1], Kyle Parfrey[2], and Włodek Kluźniak[1]
[1] Nicolaus Copernicus Astronomical Center of the Polish Academy of Sciences, Bartycka 18, 00-716 Warsaw, Poland; dabarca@camk.edu.pl
[2] School of Mathematics, Trinity College Dublin, Dublin 2, Ireland

Received 2021 June 25; revised 2021 July 22; accepted 2021 July 28; published 2021 August 20

Abstract

We perform a global 2.5D general-relativistic radiation magnetohydrodynamic simulation of supercritical accretion onto a neutron star with a 2×10^{10} G dipolar magnetic field, as a model of a neutron-star-powered ultraluminous X-ray source (ULX). We compute a lower limit on the total luminosity of $\sim 2.5\, L_{\rm Edd}$, and find the radiation to be highly beamed by the accretion disk outflows. The apparent isotropic luminosity, which is a function of the viewing angle, reaches a maximum above $100\, L_{\rm Edd}$, consistent with the luminosities observed in ULXs.

Unified Astronomy Thesaurus concepts: Ultraluminous x-ray sources (2164); Magnetohydrodynamical simulations (1966); Neutron stars (1108)

1. Introduction

Ultraluminous X-ray sources (ULXs) are extragalactic, non-active-galactic-nucleus X-ray sources with luminosities exceeding 10^{39} erg s^{-1} (Kaaret et al. 2017). Observations, in a handful of ULXs, of coherent pulsations of ~ 1 s periodicities (Trudolyubov 2008; Bachetti et al. 2014; Motch et al. 2014; Fürst et al. 2016; Israel et al. 2017a, 2017b; Townsend et al. 2017; Tsygankov et al. 2017; Brightman et al. 2018; Carpano et al. 2018; Doroshenko et al. 2018; Fürst et al. 2018; Heida et al. 2019; Chandra et al. 2020) have shown that at least some of these sources are powered by slowly rotating neutron stars accreting above their critical limits, $\dot{M}_{\rm Edd} = \eta L_{\rm Edd}/c^2$, where $\eta \sim 0.2$ is the binding energy per unit mass at the surface of the neutron star (Syunyaev & Shakura 1986) and $L_{\rm Edd} = 4\pi G M m_p c/\sigma_T$ is the Eddington luminosity of an object with mass M.

For low accretion rates η is also the expected radiative efficiency, and the luminosity is proportional to the mass accretion rate, $L/L_{\rm Edd} = \dot{M}/\dot{M}_{\rm Edd}$. However, due to their extremely large optical depths, accretion disks with $\dot{M} \gtrsim \dot{M}_{\rm Edd}$ can no longer cool efficiently. The accretion flow traps photons and the advection of radiation becomes the primary mode of energy transport in the disk (Begelman 1978; Abramowicz et al. 1988; Sądowski & Narayan 2016; Czerny 2019). The large concentration of photons launches a radiation-pressure-driven outflow, which originates at the radius where the radiation flux through the surface of the disk becomes super-Eddington (Shakura & Sunyaev 1973). The outflow extends from this radius (referred to as the spherization radius) down to the inner edge of the accretion disk. Because of advection, the value of the spherization radius will differ somewhat from its classic thin-disk value (Shakura & Sunyaev 1973). In fact, a substantial fraction of the photons in the radiation-pressure-dominated inner disk will be advected to the vicinity of the stellar surface and released there. How it escapes to infinity is the major focus of this Letter.

King et al. (2001) suggested long before the first pulsating ULXs were observed that the outflows from a super-Eddington disk could collimate the emission released near the compact object, in a manner similar to the collimation predicted by the thick disk model of the Warsaw group (Abramowicz et al. 1978; Paczyński & Wiita 1980). The system would then appear to be very bright when viewed face-on, and the inferred isotropic luminosity $L_{\rm iso} = F/(4\pi d^2)$, where F is the radiation flux measured by the observer and d is the distance to the source, would be much larger than the true luminosity, L, i.e., the total emitted radiation power.

An additional interesting feature of pulsating ULXs is their unusually high spin-up rates. Kluźniak & Lasota (2015) inferred a dipole field of $\sim 10^9$ G from the spin-up rate of M82 X-2. As more pulsating ULXs were found, all with high spin-ups, a model was formed that incorporated the period, spin-up, and luminosity to predict the magnetic field strength and intensity of the beaming (King et al. 2017; King & Lasota 2019, 2020, hereafter referred to as the KLK model). Besides very small values of the beaming factor $b = L/L_{\rm iso}$—implying a high degree of beaming—the model also predicts dipole magnetic fields in the range of 10^9–10^{13} G, with most values falling between 10^{10} and 10^{11} G.

In order to model a neutron star accreting at super-Eddington accretion rates, it is necessary to run general-relativistic radiative magnetohydrodynamic (GRRMHD) simulations. As of writing this Letter, there is only one such global simulation (Takahashi & Ohsuga 2017) that includes a stellar magnetic field, and two that do not (Abarca et al. 2018; Takahashi et al. 2018).

The simulation discussed in Takahashi & Ohsuga (2017), while an important result, has some shortcomings. It is unclear how the highly magnetized regions are treated, or what effects the numerical density floor or background atmosphere have on the emerging radiation. The simulation is run for a rather short duration, not allowing adequate time for the outflows to reach a steady state. To overcome these issues, all of which could potentially contaminate measurements of the luminosity and flux distribution, we introduce a scheme that captures the highly magnetized regions of the simulation more realistically. Such a scheme was introduced by Parfrey & Tchekhovskoy (2017) and we have implemented it in the GRRMHD code `Koral`. We ensure the numerical floors do not affect the emerging radiation and run the simulation for a much longer duration.

We wish to investigate the degree to which the radiation produced by an accreting magnetized neutron star is beamed. Small values of b and lower neutron-star magnetic fields would support the KLK model, while values of b near unity would indicate that some other configuration must be responsible for such high observed luminosities. Even if we do find the radiation to be highly beamed, a direct comparison to the KLK model would not be very informative considering the requirement that the magnetospheric and spherization radii be quite close to each other.

A further caveat to consider is that we have aligned the magnetic dipole axis with the disk axis, so even if we included rotation, we would not expect to produce pulsations. However, the population of non-pulsed ULXs is much larger than the population of pulsed ULXs, and there is no reason to believe that some of these also may not be powered by neutron stars.

At the accretion rates we are considering, it has been predicted that the flow to the magnetosphere forms an optically thick "accretion curtain" shielding the stellar surface outside the polar regions (Mushtukov et al. 2017). For a rotating, misaligned dipole, this would produce smooth pulse profiles as the radiation is only able to escape through the thin funnel region along the magnetic poles and this may significantly lower the pulse fraction through multiple scatterings off the funnel wall, making a strong pulsed fraction inconsistent with strong beaming (Mushtukov et al. 2021). A fully 3D simulation with a misaligned dipole would be necessary to adequately address this possible inconsistency.

In Section 2 we describe the numerical methods and the simulation setup. In Section 3 we describe the results of the simulation. In Section 4 we discuss the effects of beaming and the expected observed luminosity of the system, and in Section 5 we summarize the results and our conclusions.

2. Numerical Methods

We use the code `Koral` (Sądowski et al. 2013, 2015), which solves the conservation equations of GRRMHD on a static grid in a fixed metric, $g_{\mu\nu}$. The evolution equations are given by

$$\nabla_\mu(\rho u^\mu) = 0, \tag{1}$$

$$\nabla_\mu T^{\mu}{}_{\nu} = G_\nu, \tag{2}$$

$$\nabla_\mu R^{\mu}{}_{\nu} = -G_\nu, \tag{3}$$

$$\nabla_\mu F^{*\mu\nu} = 0. \tag{4}$$

The equations correspond to conservation of mass, conservation of total energy-momentum with coupling of matter and radiation provided by the radiation four-force, G_ν, (Mihalas & Mihalas 1984), and the source-free Maxwell's equations. Conservation of mass depends on ρ, the baryon rest-mass density, and u^μ, the gas four-velocity. The stress-energy tensor for a magnetized gas is given by

$$T^{\mu}{}_{\nu} = (\rho + p + u_{\rm int} + b^2)u^\mu u_\nu + (p + b^2/2)\delta^\mu_\nu - b^\mu b_\nu, \tag{5}$$

which makes use of $u_{\rm int}$, the gas internal energy, $p = (\gamma - 1)u_{\rm int}$, the gas pressure (where $\gamma = 5/3$ is the adiabatic index), and the magnetic field four-vector, $b^\mu = \frac{1}{2}u_\nu F^{*\mu\nu}$.

Written in terms of the Hodge dual of the Faraday tensor, $F^{*\mu\nu}$, the source-free Maxwell equations correspond to the induction equation (spatial components), and the divergence-free condition of the magnetic field (t component), both of which are evolved using the flux constrained transport algorithm (Tóth 2000).

Equation (3) only evolves the $R^{t\mu}$ components of the radiation stress-energy tensor using G_ν, which includes opacities defined by electron scattering, Comptonization, and bremsstrahlung absorption. The spatial components, R^{ij}, are computed using the M_1 closure scheme (Mihalas & Mihalas 1984; Sądowski et al. 2013), which makes the approximation that there exists a frame with four-velocity $u^\mu_{\rm R}$ in which the radiation can be considered to be isotropic. In that frame, the radiation has energy density $\bar{E}$. We can then write the radiation tensor in the lab frame as

$$R^{\mu\nu} = \frac{4}{3}\bar{E}u^\mu_{\rm R}u^\nu_{\rm R} + \frac{1}{3}g^{\mu\nu}\bar{E}, \tag{6}$$

which can be inverted to give the radiation primitive variables $\bar{E}$ and $u^\mu_{\rm R}$ in terms of $R^{t\mu}$.

Since we study a nonrotating star[3] it is sufficient to use the Schwarzschild metric with a coordinate system that is logarithmic in radius and stretches from $r = 5\,r_g$ to $r = 1000\,r_g$, where $r_g = GM\,c^{-2}$ is the gravitational radius defined by the neutron-star mass M. For the purposes of this calculation we take the radius of the star to be $R = 5r_g$ and we specify[4] its mass at the canonical value $1.4\,M_\odot$. While modern studies with the NICER instrument indicate $R \approx 6.4\,r_g$ at $M \approx 1.3\,M_\odot$ (Riley et al. 2019), the adopted value of M or R does not affect our results qualitatively. In this work the geometry of accretion is controlled by the magnetosphere.[5] Unless otherwise specified, we adopt units where $G = c = 1$. Our simulation is run in 2D axisymmetry with resolution in r and θ corresponding to [512, 510] cells.

We initialize the simulation with an equilibrium torus (Penna et al. 2013) threaded by a single loop of magnetic field that feeds gas to the star at a rate of $\sim 20\dot{M}_{\rm Edd}$. The initial maximum $\beta = (p_{\rm gas} + p_{\rm rad})/p_{\rm mag}$ (the ratio of gas plus radiation to magnetic pressures) in the torus is equal to 10. We initialize a stellar dipole field with a maximum field strength on the stellar surface of 2×10^{10} G using the potential given in Wasserman & Shapiro (1983).

Outside the torus, the gas is initialized to a low-density background. This creates a large contrast in the magnetic and rest-mass energy densities. The ratio of these two quantities, the magnetization, $\sigma = b^2/(2\rho)$, provides an indication of where the numerical scheme should start to break down, with $\sigma \gg 1$ regions being especially prone to error/instabilities. Our simulation is initialized with a peak magnetization of $\sigma = 10^4$, and in order to evolve the system we implement the method described in Parfrey & Tchekhovskoy (2017). The basic idea of the scheme is to divide the gas into contributions from the real general-relativistic magnetohydrodynamic (GRRMHD) flow and from the numerical floor that keeps σ from becoming too large. When gas is dominated by the numerical floor, the density and internal energy are adjusted to their background levels and the velocity parallel to the magnetic field, as measured by the stationary observer, is reduced. We provide some additional adjustments that improve the scheme's robustness in the presence of radiation. We reduce the scattering and absorption opacities of the gas dominated by the numerical floor, and we balance energy gain/loss from round-off errors in the magnetic field by respectively subtracting/adding

---

[3] The observed ULX periods of several seconds make the spin of the neutron star negligible in our simulations, which typically run for about $\sim 10^5\,GM\,c^{-3}$, i.e., ~ 0.7 s of physical time.

[4] The results of radiative simulations do not scale with the stellar mass, which therefore must be specified.

[5] In general R/r_g varies with the stellar mass, and it is only for weakly magnetized stars (with a dipole moment about two orders of magnitude below what we assume here) that the character of flow depends on whether R is larger or smaller than the radius of the marginally stable orbit (Kluzniak & Wagoner 1985).

radiation energy during the conserved-to-primitive variable inversion.

We also introduce a new boundary condition that attempts to mimic the hard surface of the neutron-star crust. We treat the gas as in Parfrey & Tchekhovskoy (2017), allowing it to fall through the inner boundary unimpeded. Then, on a cell-by-cell basis, we measure the flux of kinetic, thermal, and radiative energy flowing through the boundary and return a fraction (albedo) of that energy as outflowing radiation. In the calculations reported in this Letter, that fraction corresponds to 75%, but a full study of this boundary condition for different values of the albedo is underway. An important consideration is the actual flux of radiation that crosses the inner boundary, which is controlled by the Riemann solver. The ghost cells are set to reflect 75% of the inflowing energy *in the ghost cells*. The HLL Riemann solver should pick a value that is roughly halfway between these two[6] fluxes and so in actuality we expect about 12.5% of the radiation flux to escape from the domain through the inner boundary.

We run the simulation for 80,000 t_g where $t_g = GMc^{-3}$. Normally, in 2D axisymmetry, the absence of a dynamo driven by turbulence from the magnetorotational instability (MRI) leads to decay of the magnetic field. This is remedied with the use of a mean-field dynamo that restores the magnetic field in the accretion disk as a model for how it would be regenerated in 3D (Sądowski et al. 2015).

3. Simulation Results

The field lines of the stellar dipole are deformed to wrap around the initial torus, and so far from the star they are out of equilibrium. As the simulation starts, the magnetic field quickly relaxes to a stable configuration enveloping the torus while the closed loops near the star are virtually unchanged.

The torus begins to evolve as the MRI builds up and the gas begins to accrete. When the gas reaches the stellar magnetic field, it forces it inward, raising the magnetic pressure until it balances the ram pressure, at which point the gas begins to slide along magnetic field lines, forming accretion columns. As the gas hits the inner boundary, it is shocked and a large amount of radiation begins to leave the base of the column perpendicularly through the column's sides. As the simulation progresses, the accretion disk converges at progressively larger radii to its steady-state solution, launching outflows that collimate the radiation released in the columns and inner parts of the disk by confining it to a funnel-like region about the polar axis.

A snapshot from the simulation at time $t = 32{,}000\,t_g$ is shown in the upper panel of Figure 1. The lower panel depicts a time average from $t = 40{,}000\,t_g$ to $t = 80{,}000\,t_g$. The left half of the panel shows $\widehat{E}$, the radiation energy density in the fluid frame, and the right half of the panel shows the gas rest-mass density.

The magnetic field of the torus is oriented to be opposite in direction to the dipole field when they meet, leading to reconnection that allows gas to flow smoothly into the accretion columns (Parfrey & Tchekhovskoy 2017). The snapshot shows the remnant of a loop from the torus that had just reconnected with the stellar dipole in the disk midplane, indicated by the bold contour in the upper panel.

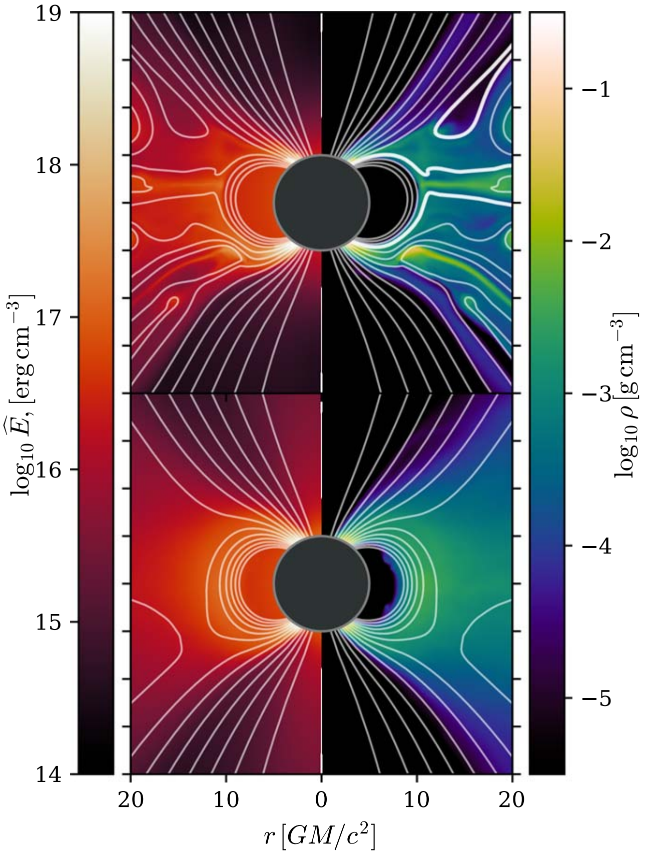

Figure 1. Snapshot (upper) and time-averaged (lower) plots of the radiation energy density (left), and the gas rest-mass density (right). We plot equispaced contours of A_ϕ, the ϕ component of the vector potential, which correspond to poloidal field lines due to the axisymmetric nature of the problem. A bold contour shows the remnant of a torus loop that had just reconnected with the neutron star's dipole.

The flow is quite turbulent. The snapshot captures a moment of lower luminosity before a high-density parcel of gas below the disk midplane enters the column and collides with the stellar surface, raising the luminosity significantly. The long-term effect of successive gas parcels hitting the surface and becoming shocked contributes to the steep radial gradient of radiation energy and gas density at the base of the column. This effect is also apparent from the difference in radiation energies in the polar region between the two panels. It is also evident that the polar region is largely devoid of gas. The gas is confined by the magnetic field to midlatitudes, strongly contrasting to what was observed in Abarca et al. (2018), where the absence of a stellar magnetic field allowed outflowing gas to fill the whole domain.

We can study the outflowing gas in more detail by considering the azimuthally integrated radial fluxes of gas and radiation, which are shown in Figure 2. The left half of the panel is the quantity $2\pi r\,\sin(\theta) F^r$ in units of $[L_{\rm Edd}\, r_g^{-1}]$, where $F^r = R^r{}_t$ is the radial component of the radiation flux (or momentum depending on the factor of c). One can then integrate by eye over $r\,{\rm d}\theta$ to estimate the luminosity. In a similar way, the gas momentum is integrated into the poloidal plane giving $2\pi r\,\sin(\theta)\rho u^r$ in units of $[L_{\rm Edd}\, r_g^{-1} c^{-2}]$.

Three contours are included. The solid green line shows the photosphere defined by $\tau_r = 1$, where τ_r is the scattering optical depth found by integrating from the outer boundary of the

---

[6] The left-biased flux is determined by the ghost cells and the right-biased flux by the domain.

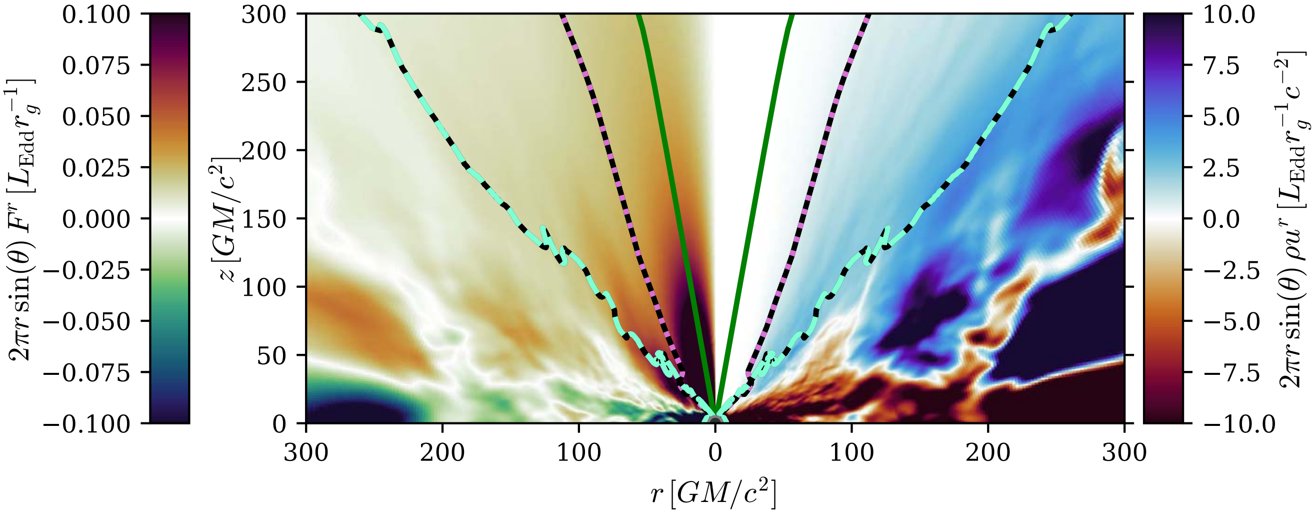

Figure 2. Radial radiative flux, $F^r = R^r_t$ (left), and gas momentum density, ρu^r (right), integrated into the poloidal plane. Overplotted in solid green is the photosphere at electron scattering depth unity, as measured radially, $\tau_r = 1$. The dashed pink and black line shows the photosphere as measured along the θ direction from the axis. The dashed teal and black line shows the zero-Bernoulli surface, Be = 0.

simulation to radius r. We can then assume that all of the outflowing radiation above this surface will reach the observer giving us a lower limit on the luminosity if we integrate the flux over the θ coordinate. By definition we can see very little gas above this surface, and all of the radiation that is between this line and the axis is expected to reach infinity. The polar region is completely dominated by a nearly radial flow of radiation escaping the inner region of the simulation. Surprisingly, a significant fraction of outgoing radiation is excluded by the $\tau_r = 1$ surface. There is very little momentum density in the corresponding region so one would expect the gas to be optically thin, and for this to be reflected in the $\tau_r = 1$ contour. However, the outflows do not flow exactly radially, nor in perfectly straight lines, so the gas at larger radii is obscuring the inner region. The flux at low radii, however, does not know about this gas and freely streams over a larger range of angles than indicated by the $\tau_r = 1$ surface. At some point (in our simulation between radius 300 and 400 r_g) the radiation scatters off the outflowing gas, becoming more confined. This is precisely the radiation-collimation effect that should lead to large apparent luminosities.

Another consideration is the small density gradient in the radial direction. This leads to the location of the photosphere being quite sensitive to small variations in the density and the precise value of the scattering cross section. The draconian approach of including 100% of the flux on one side of the contour and excluding 100% of the flux below the contour may not be appropriate for estimating the total luminosity, as the photons located immediately above and below the contour have almost the same probability of reaching the observer.

We can also measure the optical depth by integrating along θ from the axis. This surface, $\tau_\theta = 1$, is shown in Figure 2 by the densely dashed pink and black line. Because τ_θ is more useful for measuring the amount of radiation that leaves the accretion columns (since radiation escapes the column along the θ direction), it might provide a more accurate representation of the radiation that can reach the observer. The gradient of density along θ is much stronger than in the radial direction so there is much less uncertainty in the location of the photosphere in this direction. We can see that τ_θ is a good indicator for separating the very strong radiation flow near the axis from the less intense radiation flow in the gas outflows. Also, τ_θ is not affected by the geometry of the outflow at large radii; however, the question remains what happens to the radiation at large radii, or the radiation that scatters off of the side of the outflow near the accretion column?

The last curve, shown in loosely dashed teal and black represents the surface where the relativistic Bernoulli parameter,

$$Be = -\frac{T^t{}_t + R^t{}_t + \rho u^t}{\rho u^t}, \tag{7}$$

is equal to zero. This surface approximately splits the domain into energetically bound and unbound regions. Outflowing gas along this contour would be able to reach infinity with zero specific energy if it absorbed all of the radiative energy at its location. In reality, at some point the outflow should become diffuse enough that the radiation escapes. We can therefore use the zero-Bernoulli surface to define a region above which we can integrate the radiation flux to get an upper limit on the luminosity. If the gas rapidly becomes optically thin, the luminosity will be close to the integral of flux above this surface. It is more likely that the radiation deposits some of its momentum into the outflowing gas, lowering the luminosity. One could argue that it is also possible for the gas to cool contributing even more radiation to the total and exceeding this upper limit. However, we can verify that, at least in the simulation domain at radii larger than ∼70 r_g, the outflows almost exclusively absorb radiation.[7]

If we perform the integrals of radiative flux over spherical shells bounded by these three surfaces, we can plot the luminosity for each measure as a function of radius as shown in Figure 3. In addition to the three luminosities described above, for reference we also plot the total luminosity as integrated over the entire domain. Formally we can define the luminosities as

---

[7] Radiation transfers energy from the hot inner region to the cooler adiabatically expanding gas in the outer regions.

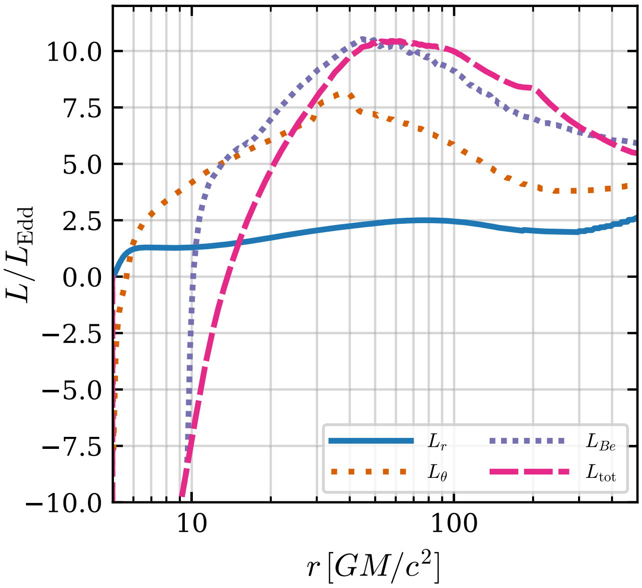

Figure 3. Four measures of the luminosity are plotted as a function of radius. In solid blue is L_r, loosely dotted orange shows L_θ, densely dotted purple is $L_{\rm Be}$, and dashed pink is $L_{\rm tot}$.

follows:

$$L_r = \int_{\tau_r<1} R^r{}_t\sqrt{-g}\ {\rm d}\theta, \tag{8}$$

$$L_\theta = \int_{\tau_\theta<1} R^r{}_t\sqrt{-g}\ {\rm d}\theta, \tag{9}$$

$$L_{\rm Be} = \int_{\rm Be>0} R^r{}_t\sqrt{-g}\ {\rm d}\theta, \tag{10}$$

$$L_{\rm tot} = \int_0^\pi R^r{}_t\sqrt{-g}\ {\rm d}\theta. \tag{11}$$

In steady state, the luminosity of a central radiation source would be constant (apart from redshift factors) with radius if radiative energy were conserved. However, the presence of gas, which can absorb and emit radiation, can change the shape of the luminosity curve even in steady state. Additionally, because optical depth is defined along coordinates, and not along the path of photons, it is possible to arbitrarily add or subtract to the luminosity curve if the average path of the photons is more complex, such as near the accretion columns.

We have run the simulation for a sufficiently long duration and taken a sufficiently long time average that the accretion disk should have reached a steady state out to radius 60–70 r_g and most of the turbulence should average away. The disk converges outward as the simulation runtime approaches the viscous time at a particular radius. The outflows, which have much larger velocity, converge much faster. A weak convergence condition for the outflow can be given by $r/v^r < t_{\rm avg}$, where $v^r = u^r/u^t$ is the coordinate velocity, and $t_{\rm avg}$ is the time period over which the simulation data were averaged. Our data are averaged over a long enough period of time that nearly the entirety of the gas outflow is able to satisfy the convergence condition. However, one must also take into account the origin of the outflow. A large portion of the outflowing gas originates from regions of the disk that have not yet converged, and as we explain later, this introduces uncertainty into some measures of the luminosity, especially at larger radii.

Despite all of this, L_r, shown as the solid blue line, is somewhat flat out to about radius $\sim 500\,r_g$ and so we believe that $L_r \approx 2.5\,L_{\rm Edd}$ is a good measure of the total radiative output of the simulation or at least a suitable lower limit. The steady rise up to radius $80\,r_g$ is probably due to radiation being emitted in and emerging from the outflows. Beyond, it is hard to determine whether the fluctuations are geometrical, or due to the unconverged nature of the simulation at large radii.

$L_{\rm Be}$, $L_{\rm tot}$, and L_θ, all show negative values near the star. This is related to the well-understood phenomenon of photon trapping in super-Eddington accretion disks (Ohsuga et al. 2002; Sądowski & Narayan 2016). Most of the radiation is advected inward by the optically thick gas before it can diffuse out of the disk. In our neutron-star case some of this energy is released at the surface. It would appear the photon-trapping effect is so strong that even with an albedo of 75%, inflowing radiation in the accretion columns dominates energy transport near the stellar surface. $L_{\rm tot}$ is also largely dominated by the advection of photons in the accretion columns and continues to decrease all the way to the surface.

Both L_r and L_θ rise steeply over the first r_g or so above the star due to the radiation shock. L_θ continues to rise as the accretion columns and accretion disk add to the luminosity. L_θ includes radiation released from the outflow below radius $\sim 70\,r_g$. At larger radii G_t switches sign, and the radiation contributing to L_θ passes through enough gas to deposit almost half of its momentum into the outflow lowering the luminosity to a local minimum of $\sim 3.75 L_{\rm Edd}$. The location of the $\tau_\theta = 1$ surface is unaffected by the outer boundary so the steady rise in L_θ above $r \sim 200\,r_g$ is likely due to the gas becoming steadily thinner, allowing more of the flux to contribute to the luminosity.

A similar effect is seen in $L_{\rm Be}$ and $L_{\rm tot}$. They rise sharply with increasing r up to $r \sim 60\,r_g$, as there is a significant amount of radiation advected with the outflow, and then drop as momentum starts to be transferred to the gas. $L_{\rm Be}$ is integrated over regions of the outflow that originate from parts of the disk that have not yet converged, especially beyond radius $r \sim 100\,r_g$, which increases the uncertainty in its value, especially at larger radii. $L_{\rm Be}$ largely follows $L_{\rm tot}$, although this appears to be largely a coincidence and is due to the equal amounts of radiation flowing inward and outward over the region where $Be < 0$.

4. Beamed Emission

The most important quantity, which is the signature of all ULXs, is a large apparent isotropic luminosity, that is, $L_{\rm iso} = 4\pi d^2 F$, measured from an observed flux, F, emitted by an object at a distance, d, away from the observer (neglecting cosmological effects). While it is impossible to predict $L_{\rm iso}$ reliably from the simulation data without sophisticated radiation post-processing, we can estimate it at a few locations in the simulation and see how it changes with radius. Plotted in the left panel of Figure 4 is $L_{\rm iso}$ as a function of viewing angle for different radii. The figure is presented in polar coordinates to emphasize the beaming pattern. One can immediately see that the vast majority of the flux is confined within 20° of the axis. The peak $L_{\rm iso}$ lies along the axis, and while it decays with radius (from the continuous green line through the dashed lines to the dotted one), it appears to be converging to a large value well above $100\,L_{\rm Edd}$ that is about 2×10^{40} erg s^{-1} for a canonical $1.4\,M_\odot$ neutron star.

The apparent luminosity is clearly bright enough for the neutron star to qualify as a ULX. To compare the measured

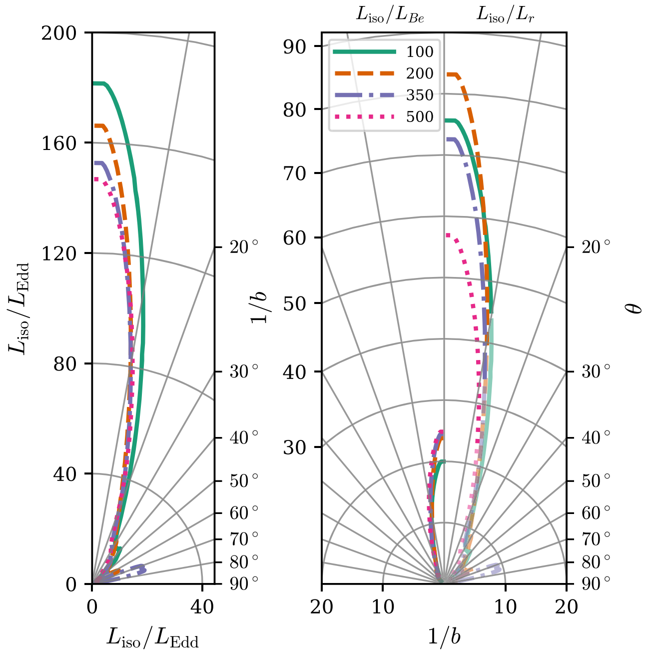

Figure 4. Inferred isotropic flux (left) and beaming pattern (right) as a function of viewing angle, measured from the pole. For the left panel, the radial coordinate corresponds to the isotropic luminosity, $L_{\rm iso} = 4\pi r^2 F^r$. For the right panel, the radial coordinate measures beaming, $1/b$, i.e., the inferred isotropic luminosity as a function of angle divided by the total luminosity. The left side of the right panel is a lower limit for $1/b$ obtained with the total luminosity, $L_{\rm Be}$. The right side of the right panel assumes the total luminosity is L_r, which provides an upper limit on $1/b$. All the quantities are measured at specific radii in the simulation, as indicated by the legend.

degree of beaming with the KLK model, we need to choose a quantity to function as the total luminosity. We consider L_r and $L_{\rm Be}$ as lower and upper limits, respectively (although we are confident that the true luminosity is much closer to L_r). These also correspond to lower and upper limits for the beaming factor $b = L/L_{\rm iso}$.

Near the axis, $L_{\rm iso}/L_{\rm Be} \sim 25$, and is more or less constant with radius. This already exceeds $1/b$ as computed from the KLK model (King et al. 2017; King & Lasota 2019, 2020) for all the sources they included. If we instead consider $L_{\rm iso}/L_r$, then $1/b$ shoots up to above 80. It falls at larger radii, to around 60.[8]

Note that the beaming factor is a function of the angle. In general, $1/b$ is proportional to flux, which is a function of θ, and lower viewing angles will tend to display more extreme beaming (corresponding to smaller values of b).

An important consideration is the effect that axisymmetry has on our simulation. Sądowski & Narayan (2016) compared a 2D axisymmetric accretion disk simulation onto a black hole to an identical simulation in 3D. The authors found that both of the simulations were remarkably similar, although the authors did notice that the degree of beaming was overestimated in the 2D simulation by a factor of about 2. We expect a similar degree of exaggerated beaming because the mechanism that collimates radiation is the same in both works. This could lower our estimated apparent isotropic flux to be around 10^{40} erg s^{-1}, which is still well in the realm of observed ULX luminosities.

Regardless of the measure used to compute the total luminosity, we have shown that a low-magnetic-field neutron star can produce emission that is sufficiently beamed to produce a ULX. The closest theoretical model to our simulation, the KLK model, also predicts large amounts of beaming, although not as extreme as we observe. The two models need not completely agree since they differ in several ways. The disk in our simulation has a very large spherization radius when compared to the magnetospheric radius. The KLK model requires that these two radii be much closer. This would substantially affect the outflows, as they are only launched below $r_{\rm sph}$ and above $r_{\rm M}$. Including the spherization radius in the domain is challenging, first, because it takes a long time for the simulation to converge to such large radii, and second, because beyond $r_{\rm sph}$, the disk should be similar to a thin disk, and thin disks are notoriously difficult to simulate.

One issue, although one that we already have plans to remedy, is the limitations of `Koral` when simulating the radiation field. `Koral` is a gray code that uses the M_1 closure scheme to transport radiation. M_1 works well for large extended sources, but when radiation originating from two or more distinct locations collide, the beams interact. For us, this is most problematic in the region directly outside of the accretion columns. The radiation flows toward the axis and is then directed upward due to the polar boundary condition, which is reflective. In reality, we expect the beams from the accretion columns to scatter off of the opposite wall formed by the gas outflows. After enough scatterings, the radiation should be collimated and largely moving along the axis. From Figure 1 we can see that the radiation already appears collimated as soon as it leaves the accretion column.

To get a more accurate representation of the radiation field, we need to go beyond the M_1 scheme. `HEROIC` (Narayan et al. 2016) is a postprocessing radiative transfer code that could be used to recompute the radiation field and spectrum for an observer at infinity as a function of viewing angle. While such calculations are out of the scope of this Letter, we plan to apply `HEROIC` postprocessing analysis to the rest of our simulations in a future publication.

5. Summary and Conclusions

We performed a 2D axisymmetric GRRMHD simulation of accretion onto a neutron star with a 2×10^{10} G dipolar magnetic field. The combination of the hard surface and confinement of the gas into accretion columns by the stellar magnetic field near the stellar surface allows the flow to release radiative energy at a rate of several times the Eddington limit. The fraction of this energy that is able to reach the observer, as opposed to being absorbed by the outflows, is difficult to calculate, but a lower limit of the observable luminosity should correspond to about $2.5\,L_{\rm Edd}$. The radiation easily escapes into the polar region, which is largely devoid of gas due to a combination of the magnetic field and rotation of the outflowing gas that collimates the radiation flow. While a more precise calculation of the radiation field is required due to the limitations of the simulation, our results show that this escaping radiation will be highly beamed. The apparent isotropic luminosity of the source observed pole-on should be on the order of $100\,L_{\rm Edd}$. This is encouraging if we wish to

---

[8] As r increases τ_r feels the effect of a finite outer boundary, so it is possible that the computed value of L_r may be slightly overestimated, leading to a slightly underestimated $1/b$ at large radii.

interpret the accreting system as a model of a neutron-star-powered ULX.

When compared to the KLK model (King et al. 2017; King & Lasota 2019, 2020), we find that the intensity of the beaming is larger, although we have reason to believe that postprocessing would show a less intensely beamed distribution of radiation at infinity. Furthermore, our simulation does not model the same system as considered by KLK. The distance between the Alfvén radius and spherization radius is large. We hope to produce simulations in future studies that can reproduce additional observable features of ULXs and that can provide more information about the nature of the magnetic field in pulsating and nonpulsating ULXs.

D.A. thanks Aleksander Sądowski and Andrew Chael for guidance on working with the `Koral` code; Maciek Wielgus, Miljenko Čemeljić, Jean-Pierre Lasota, and Alexander Tchekhovskoy for their advice, suggestions, and useful conversations; and Katarzyna Rusinek-Abarca for support, companionship, and patience. D.A. was supported in part by Polish National Science Center (NCN) Preludium grant 2017/27/N/ST9/00992. Research supported in part by NCN grant 2019/33/B/ST9/01564. Computations in this work were carried out on the Cyfronet Prometheus cluster, part of the PLGrid computing network.

Facility: Cyfronet Prometheus.

Software: Numpy (van der Walt et al. 2011; Harris et al. 2020), Scipy (Virtanen et al. 2020), matplotlib (Hunter 2007), pandas (McKinney 2010), Koral (Sądowski et al. 2013, 2015).

ORCID iDs

David Abarca https://orcid.org/0000-0002-9202-8734
Kyle Parfrey https://orcid.org/0000-0001-6173-0099
Włodek Kluźniak https://orcid.org/0000-0001-9043-8062

References

Abarca, D., Kluźniak, W., & Sądowski, A. 2018, MNRAS, 479, 3936
Abramowicz, M., Jaroszynski, M., & Sikora, M. 1978, A&A, 63, 221
Abramowicz, M. A., Czerny, B., Lasota, J. P., & Szuszkiewicz, E. 1988, ApJ, 332, 646
Bachetti, M., Harrison, F. A., Walton, D. J., et al. 2014, Natur, 514, 202
Begelman, M. C. 1978, MNRAS, 184, 53
Brightman, M., Harrison, F. A., Fürst, F., et al. 2018, NatAs, 2, 312
Carpano, S., Haberl, F., Maitra, C., & Vasilopoulos, G. 2018, MNRAS, 476, L45
Chandra, A. D., Roy, J., Agrawal, P. C., & Choudhury, M. 2020, MNRAS, 495, 2664
Czerny, B. 2019, Univ, 5, 131
Doroshenko, V., Tsygankov, S., & Santangelo, A. 2018, A&A, 613, A19
Fürst, F., Walton, D. J., Harrison, F. A., et al. 2016, ApJL, 831, L14
Fürst, F., Walton, D. J., Heida, M., et al. 2018, A&A, 616, A186
Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Natur, 585, 357
Heida, M., Lau, R. M., Davies, B., et al. 2019, ApJL, 883, L34
Hunter, J. D. 2007, CSE, 9, 90
Israel, G. L., Belfiore, A., Stella, L., et al. 2017a, Sci, 355, 817
Israel, G. L., Papitto, A., Esposito, P., et al. 2017b, MNRAS, 466, L48
Kaaret, P., Feng, H., & Roberts, T. P. 2017, ARA&A, 55, 303
King, A., & Lasota, J.-P. 2019, MNRAS, 485, 3588
King, A., & Lasota, J.-P. 2020, MNRAS, 494, 3611
King, A., Lasota, J.-P., & Kluźniak, W. 2017, MNRAS, 468, L59
King, A. R., Davies, M. B., Ward, M. J., Fabbiano, G., & Elvis, M. 2001, ApJL, 552, L109
Kluźniak, W., & Lasota, J.-P. 2015, MNRAS, 448, L43
Kluzniak, W., & Wagoner, R. V. 1985, ApJ, 297, 548
McKinney, W. 2010, Proc. SciPy, 9, 56
Mihalas, D., & Mihalas, B. W. 1984, Foundations of Radiation Hydrodynamics (New York: Oxford Univ. Press)
Motch, C., Pakull, M. W., Soria, R., Grisé, F., & Pietrzyński, G. 2014, Natur, 514, 198
Mushtukov, A. A., Portegies Zwart, S., Tsygankov, S. S., Nagirner, D. I., & Poutanen, J. 2021, MNRAS, 501, 2424
Mushtukov, A. A., Suleimanov, V. F., Tsygankov, S. S., & Ingram, A. 2017, MNRAS, 467, 1202
Narayan, R., Zhu, Y., Psaltis, D., & Sądowski, A. 2016, MNRAS, 457, 608
Ohsuga, K., Mineshige, S., Mori, M., & Umemura, M. 2002, ApJ, 574, 315
Paczyński, B., & Wiita, P. J. 1980, A&A, 500, 203
Parfrey, K., & Tchekhovskoy, A. 2017, ApJL, 851, L34
Penna, R. F., Kulkarni, A., & Narayan, R. 2013, A&A, 559, A116
Riley, T. E., Watts, A. L., Bogdanov, S., et al. 2019, ApJL, 887, L21
Sądowski, A., & Narayan, R. 2016, MNRAS, 456, 3929
Sądowski, A., Narayan, R., Tchekhovskoy, A., et al. 2015, MNRAS, 447, 49
Sądowski, A., Narayan, R., Tchekhovskoy, A., & Zhu, Y. 2013, MNRAS, 429, 3533
Shakura, N. I., & Sunyaev, R. A. 1973, A&A, 24, 337
Syunyaev, R. A., & Shakura, N. I. 1986, SvAL, 12, 117
Takahashi, H. R., Mineshige, S., & Ohsuga, K. 2018, ApJ, 853, 45
Takahashi, H. R., & Ohsuga, K. 2017, ApJL, 845, L9
Tóth, G. 2000, JCoPh, 161, 605
Townsend, L. J., Kennea, J. A., Coe, M. J., et al. 2017, MNRAS, 471, 3878
Trudolyubov, S. P. 2008, MNRAS, 387, L36
Tsygankov, S. S., Doroshenko, V., Lutovinov, A. A., Mushtukov, A. A., & Poutanen, J. 2017, A&A, 605, A39
van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, CSE, 13, 22
Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, NatMe, 17, 261
Wasserman, I., & Shapiro, S. L. 1983, ApJ, 265, 1036

Bibliography

David Abarca and Włodek Kluźniak. Radial oscillations of a radiation-supported levitating shell in Eddington luminosity neutron stars. MNRAS, 461(3):3233–3238, September 2016. doi: 10.1093/mnras/stw1432.

David Abarca, Włodek Kluźniak, and Aleksander Sądowski. Radiative GRMHD simulations of accretion and outflow in non-magnetized neutron stars and ultraluminous X-ray sources. MNRAS, 479(3):3936–3951, Sep 2018. doi: 10.1093/mnras/sty1602.

David Abarca, Kyle Parfrey, and Włodek Kluźniak. Beamed Emission from a Neutron-star ULX in a GRRMHD Simulation. ApJ, 917(2):L31, August 2021. doi: 10.3847/2041-8213/ac1859.

M. Abramowicz, M. Jaroszynski, and M. Sikora. Relativistic, accreting disks. A&A, 63:221–224, February 1978.

M. A. Abramowicz, B. Czerny, J. P. Lasota, and E. Szuszkiewicz. Slim accretion disks. ApJ, 332: 646–658, September 1988. doi: 10.1086/166683.

Marek A. Abramowicz, George F. R. Ellis, and Antonio Lanza. Relativistic Effects in Superluminal Jets and Neutron Star Winds. ApJ, 361:470, October 1990. doi: 10.1086/169211.

Angelo Marcello Anile. *Relativistic fluids and magneto-fluids : with applications in astrophysics and plasma physics*. 1989.

Jonathan Arons. Photon Bubbles: Overstability in a Magnetized Atmosphere. ApJ, 388:561, April 1992. doi: 10.1086/171174.

M. Bachetti, F. A. Harrison, D. J. Walton, B. W. Grefenstette, D. Chakrabarty, F. Fürst, D. Barret, A. Beloborodov, S. E. Boggs, F. E. Christensen, W. W. Craig, A. C. Fabian, C. J. Hailey, A. Hornschemeier, V. Kaspi, S. R. Kulkarni, T. Maccarone, J. M. Miller, V. Rana, D. Stern, S. P. Tendulkar, J. Tomsick, N. A. Webb, and W. W. Zhang. An ultraluminous X-ray source powered by an accreting neutron star. Nature, 514:202–204, October 2014. doi: 10.1038/nature13791.

Steven A. Balbus. Magnetohydrodynamics of Protostellar Disks. *arXiv e-prints*, art. arXiv:0906.0854, June 2009.

Steven A. Balbus and John F. Hawley. A Powerful Local Shear Instability in Weakly Magnetized Disks. I. Linear Analysis. ApJ, 376:214, July 1991. doi: 10.1086/170270.

Steven A. Balbus and John F. Hawley. Instability, turbulence, and enhanced transport in accretion disks. *Rev. Mod. Phys.*, 70:1–53, Jan 1998. doi: 10.1103/RevModPhys.70.1. URL `https://link.aps.org/doi/10.1103/RevModPhys.70.1`.

James M. Bardeen, William H. Press, and Saul A. Teukolsky. Rotating Black Holes: Locally Nonrotating Frames, Energy Extraction, and Scalar Synchrotron Radiation. ApJ, 178:347–370, December 1972. doi: 10.1086/151796.

M. M. Basko and R. A. Sunyaev. The limiting luminosity of accreting neutron stars with magnetic fields. MNRAS, 175:395–417, May 1976. doi: 10.1093/mnras/175.2.395.

Peter A. Becker and Michael T. Wolff. Thermal and Bulk Comptonization in Accretion-powered X-Ray Pulsars. ApJ, 654(1):435–457, January 2007. doi: 10.1086/509108.

V. S. Beskin, S. V. Chernov, C. R. Gwinn, and A. A. Tchekhovskoy. Radio Pulsars. Space Sci. Rev., 191 (1-4):207–237, October 2015. doi: 10.1007/s11214-015-0173-8.

Lars Bildsten. Theory and observations of Type I X-Ray bursts from neutron stars. In Stephen S. Holt and William W. Zhang, editors, *Cosmic Explosions: Tenth AstroPhysics Conference*, volume 522 of *American Institute of Physics Conference Series*, pages 359–369, June 2000. doi: 10.1063/1.1291736.

Donato Bini, Robert T. Jantzen, and Luigi Stella. The general relativistic Poynting Robertson effect. *Classical and Quantum Gravity*, 26(5):055009, March 2009. doi: 10.1088/0264-9381/26/5/055009.

D. A. Bollimpalli and W. Kluźniak. Radial modes of levitating atmospheres around Eddington luminosity neutron stars. MNRAS, 472(3):3298–3303, December 2017. doi: 10.1093/mnras/stx2140.

D. A. Bollimpalli, M. Wielgus, D. Abarca, and W. Kluźniak. Atmospheric oscillations provide simultaneous measurement of neutron star mass and radius. MNRAS, 487(4):5129–5142, August 2019. doi: 10.1093/mnras/stz1597.

M. Brightman, F. A. Harrison, F. Fürst, M. J. Middleton, D. J. Walton, D. Stern, A. C. Fabian, M. Heida, D. Barret, and M. Bachetti. Magnetic field strength of a neutron-star-powered ultraluminous X-ray source. *Nature Astronomy*, 2:312–316, February 2018. doi: 10.1038/s41550-018-0391-6.

I. Caballero and J. Wilms. X-ray pulsars: a review. Mem. Soc. Astron. Italiana, 83:230, January 2012.

S. Carpano, F. Haberl, C. Maitra, and G. Vasilopoulos. Discovery of pulsations from NGC 300 ULX1 and its fast period evolution. MNRAS, 476(1):L45–L49, May 2018. doi: 10.1093/mnrasl/sly030.

Andrew A. Chael, Ramesh Narayan, and Aleksander Sądowski. Evolving non-thermal electrons in simulations of black hole accretion. MNRAS, 470(2):2367–2386, September 2017. doi: 10.1093/mnras/stx1345.

Deepto Chakrabarty, Edward H. Morgan, Michael P. Muno, Duncan K. Galloway, Rudy Wijnands, Michiel van der Klis, and Craig B. Markwardt. Nuclear-powered millisecond pulsars and the maximum spin frequency of neutron stars. Nature, 424(6944):42–44, July 2003. doi: 10.1038/nature01732.

N. Chamel and P. Haensel. Physics of Neutron Star Crusts. *Living Reviews in Relativity*, 11:10, December 2008. doi: 10.12942/lrr-2008-10.

Amar Deo Chandra, Jayashree Roy, P. C. Agrawal, and Manojendu Choudhury. Study of recent outburst in the Be/X-ray binary RX J0209.6-7427 with AstroSat: a new ultraluminous X-ray pulsar in the Magellanic Bridge? MNRAS, 495(3):2664–2672, July 2020. doi: 10.1093/mnras/staa1041.

W. Coburn, W. A. Heindl, R. E. Rothschild, D. E. Gruber, I. Kreykenbohm, J. Wilms, P. Kretschmar, and R. Staubert. Magnetic Fields of Accreting X-Ray Pulsars with the Rossi X-Ray Timing Explorer. ApJ, 580(1):394–412, November 2002. doi: 10.1086/343033.

E. J. M. Colbert and R. F. Mushotzky. The Nature of Accreting Black Holes in Nearby Galaxy Nuclei. ApJ, 519:89–107, July 1999. doi: 10.1086/307356.

J. M. Corral-Santana, J. Casares, T. Muñoz-Darias, F. E. Bauer, I. G. Martínez-Pais, and D. M. Russell. BlackCAT: A catalogue of stellar-mass black holes in X-ray transients. A&A, 587:A61, March 2016. doi: 10.1051/0004-6361/201527130.

R. Courant, K. Friedrichs, and H. Lewy. Über die partiellen Differenzengleichungen der mathematischen Physik. *Mathematische Annalen*, 100:32–74, January 1928. doi: 10.1007/BF01448839.

W. G. Dixon. *Special relativity: the foundation of macroscopic physics.* 1978.

V. Doroshenko, S. Tsygankov, and A. Santangelo. Orbit and intrinsic spin-up of the newly discovered transient X-ray pulsar Swift J0243.6+6124. A&A, 613:A19, May 2018. doi: 10.1051/0004-6361/201732208.

K. Y. Ekşi, İ. C. Andaç, S. Çıkıntoğlu, A. A. Gençali, C. Güngör, and F. Öztekin. The ultraluminous X-ray source NuSTAR J095551+6940.8: a magnetar in a high-mass X-ray binary. MNRAS, 448: L40–L42, March 2015. doi: 10.1093/mnrasl/slu199.

R. F. Elsner and F. K. Lamb. Accretion by magnetic neutron stars. I. Magnetospheric structure and stability. ApJ, 215:897–913, August 1977. doi: 10.1086/155427.

Charles R. Evans and John F. Hawley. Simulation of Magnetohydrodynamic Flows: A Constrained Transport Model. ApJ, 332:659, September 1988. doi: 10.1086/166684.

L. G. Fishbone and V. Moncrief. Relativistic fluid disks in orbit around Kerr black holes. ApJ, 207: 962–976, August 1976. doi: 10.1086/154565.

J. Frank, A. King, and D. Raine. *Accretion power in astrophysics.*, volume 21. 1992.

F. Fürst, D. J. Walton, F. A. Harrison, D. Stern, D. Barret, M. Brightman, A. C. Fabian, B. Grefenstette, K. K. Madsen, M. J. Middleton, J. M. Miller, K. Pottschmidt, A. Ptak, V. Rana, and N. Webb. Discovery of Coherent Pulsations from the Ultraluminous X-Ray Source NGC 7793 P13. ApJ, 831: L14, November 2016. doi: 10.3847/2041-8205/831/2/L14.

F. Fürst, D. J. Walton, M. Heida, F. A. Harrison, D. Barret, M. Brightman, A. C. Fabian, M. J. Middleton, C. Pinto, V. Rana, F. Tramper, N. Webb, and P. Kretschmar. A tale of two periods: determination of the orbital ephemeris of the super-Eddington pulsar NGC 7793 P13. A&A, 616:A186, September 2018. doi: 10.1051/0004-6361/201833292.

C. F. Gammie, J. C. McKinney, and G. Tóth. HARM: A Numerical Scheme for General Relativistic Magnetohydrodynamics. ApJ, 589:444–457, May 2003. doi: 10.1086/374594.

Sergei K. Godunov and I. Bohachevsky. Finite difference method for numerical computation of discontinuous solutions of the equations of fluid dynamics. *Matematičeskij sbornik*, 47(89)(3):271–306, 1959. URL `https://hal.archives-ouvertes.fr/hal-01620642`.

Amiram Harten, Peter D. Lax, and Bram Van Leer. On upstream differencing and godunov-type schemes for hyperbolic conservation laws. *SIAM Review*, 25(1):35–61, 1983. ISSN 00361445. URL `http://www.jstor.org/stable/2030019`.

M. Heida, R. M. Lau, B. Davies, M. Brightman, F. Fürst, B. W. Grefenstette, J. A. Kennea, F. Tramper, D. J. Walton, and F. A. Harrison. Discovery of a Red Supergiant Donor Star in SN2010da/NGC 300 ULX-1. ApJ, 883(2):L34, October 2019. doi: 10.3847/2041-8213/ab4139.

A. Hewish, S. J. Bell, J. D. H. Pilkington, P. F. Scott, and R. A. Collins. Observation of a Rapidly Pulsating Radio Source. Nature, 217(5130):709–713, February 1968. doi: 10.1038/217709a0.

Shigenobu Hirose, Omer Blaes, and Julian H. Krolik. Turbulent Stresses in Local Simulations of Radiation-dominated Accretion Disks, and the Possibility of the Lightman-Eardley Instability. ApJ, 704(1):781–788, October 2009. doi: 10.1088/0004-637X/704/1/781.

Igor V. Igumenshchev, Marek A. Abramowicz, and Ramesh Narayan. Numerical Simulations of Convective Accretion Flows in Three Dimensions. ApJ, 537(1):L27–L30, July 2000. doi: 10.1086/312755.

Igor V. Igumenshchev, Ramesh Narayan, and Marek A. Abramowicz. Three-dimensional Magnetohydrodynamic Simulations of Radiatively Inefficient Accretion Flows. ApJ, 592(2):1042–1059, August 2003. doi: 10.1086/375769.

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

N. A. Inogamov and R. A. Sunyaev. Spread of matter over a neutron-star surface during disk accretion. *Astronomy Letters*, 25:269–293, May 1999.

G. L. Israel, A. Belfiore, L. Stella, P. Esposito, P. Casella, A. De Luca, M. Marelli, A. Papitto, M. Perri, S. Puccetti, G. A. R. Castillo, D. Salvetti, A. Tiengo, L. Zampieri, D. D'Agostino, J. Greiner, F. Haberl,

G. Novara, R. Salvaterra, R. Turolla, M. Watson, J. Wilms, and A. Wolter. An accreting pulsar with extreme properties drives an ultraluminous x-ray source in NGC 5907. *Science*, 355:817–819, February 2017a. doi: 10.1126/science.aai8635.

G. L. Israel, A. Papitto, P. Esposito, L. Stella, L. Zampieri, A. Belfiore, G. A. Rodríguez Castillo, A. De Luca, A. Tiengo, F. Haberl, J. Greiner, R. Salvaterra, S. Sandrelli, and G. Lisini. Discovery of a 0.42-s pulsar in the ultraluminous X-ray source NGC 7793 P13. MNRAS, 466:L48–L52, March 2017b. doi: 10.1093/mnrasl/slw218.

Y.-F. Jiang, J. M. Stone, and S. W. Davis. A Global Three-dimensional Radiation Magneto-hydrodynamic Simulation of Super-Eddington Accretion Disks. ApJ, 796:106, December 2014. doi: 10.1088/0004-637X/796/2/106.

Victoria M. Kaspi, Mallory S. E. Roberts, and Alice K. Harding. *Isolated neutron stars*, page 279–340. Cambridge Astrophysics. Cambridge University Press, 2006. doi: 10.1017/CBO9780511536281.008.

Shoji Kato, Jun Fukue, and Shin Mineshige. *Black-hole accretion disks*. January 1998.

T. Kawashima, S. Mineshige, K. Ohsuga, and T. Ogawa. A radiation-hydrodynamics model of accretion columns for ultra-luminous X-ray pulsars. PASJ, 68:83, October 2016. doi: 10.1093/pasj/psw075.

Tomohisa Kawashima and Ken Ohsuga. Super-critical column accretion on to strongly magnetized neutron stars in ULX pulsars. PASJ, 72(1):15, February 2020. doi: 10.1093/pasj/psz136.

A. King, J.-P. Lasota, and W. Kluźniak. Pulsing ULXs: tip of the iceberg? MNRAS, 468:L59–L62, June 2017. doi: 10.1093/mnrasl/slx020.

A. R. King, M. B. Davies, M. J. Ward, G. Fabbiano, and M. Elvis. Ultraluminous X-Ray Sources in External Galaxies. ApJ, 552:L109–L112, May 2001. doi: 10.1086/320343.

Andrew King and Jean-Pierre Lasota. No magnetars in ULXs. MNRAS, 485(3):3588–3594, May 2019. doi: 10.1093/mnras/stz720.

Andrew King and Jean-Pierre Lasota. Pulsing and non-pulsing ULXs: the iceberg emerges. MNRAS, 494(3):3611–3615, April 2020. doi: 10.1093/mnras/staa930.

Richard I. Klein, Jonathan Arons, Garrett Jernigan, and Juliana J. L. Hsu. Photon Bubble Oscillations in Accretion-powered Pulsars. ApJ, 457:L85, February 1996. doi: 10.1086/309897.

W. Kluźniak and J.-P. Lasota. An ultraluminous nascent millisecond pulsar. MNRAS, 448:L43–L47, March 2015. doi: 10.1093/mnrasl/slu200.

W. Kluzniak and R. V. Wagoner. Evolution of the innermost stable orbits around accreting neutron stars. ApJ, 297:548–554, October 1985. doi: 10.1086/163550.

W. Kluzniak and J. R. Wilson. Hard X-ray spectra from gap accretion onto neutron stars. ApJ, 372: L87–L90, May 1991. doi: 10.1086/186030.

S. S. Komissarov. A Godunov-type scheme for relativistic magnetohydrodynamics. MNRAS, 303: 343–366, February 1999. doi: 10.1046/j.1365-8711.1999.02244.x.

P. Kosec, C. Pinto, D. J. Walton, A. C. Fabian, M. Bachetti, M. Brightman, F. Fürst, and B. W. Grefenstette. Evidence for a variable Ultrafast Outflow in the newly discovered Ultraluminous Pulsar NGC 300 ULX-1. MNRAS, 479(3):3978–3986, September 2018. doi: 10.1093/mnras/sty1626.

Debora Lančová, David Abarca, Włodek Kluźniak, Maciek Wielgus, Aleksander Sądowski, Ramesh Narayan, Jan Schee, Gabriel Török, and Marek Abramowicz. Puffy Accretion Disks: Sub-Eddington, Optically Thick, and Stable. ApJ, 884(2):L37, October 2019. doi: 10.3847/2041-8213/ab48f5.

Jean-Pierre Lasota. *Black Hole Accretion Discs*, volume 440, page 1. 2016. doi: 10.1007/978-3-662-52859-4_1.

James M. Lattimer. The nuclear equation of state and neutron star masses. *Annual Review of Nuclear and Particle Science*, 62(1):485–515, 2012. doi: 10.1146/annurev-nucl-102711-095018. URL `https://doi.org/10.1146/annurev-nucl-102711-095018`.

Randall J. LeVeque. *Finite Volume Methods for Hyperbolic Problems*. Cambridge Texts in Applied Mathematics. Cambridge University Press, 2002. doi: 10.1017/CBO9780511791253.

C. D. Levermore. Relating Eddington factors to flux limiters. J. Quant. Spec. Radiat. Transf., 31(2): 149–160, February 1984. doi: 10.1016/0022-4073(84)90112-2.

W. H. G. Lewin, W. D. Vacca, and E. M. Basinska. Precursors to X-ray bursts : the result of expansion and subsequent contraction of the neutron star's photosphere. ApJ, 277:L57–L60, February 1984. doi: 10.1086/184202.

Walter H. G. Lewin, Jan van Paradijs, and Ronald E. Taam. X-Ray Bursts. Space Sci. Rev., 62(3-4): 223–389, September 1993. doi: 10.1007/BF00196124.

Q. Z. Liu, J. van Paradijs, and E. P. J. van den Heuvel. Catalogue of high-mass X-ray binaries in the Galaxy (4th edition). A&A, 455(3):1165–1168, September 2006. doi: 10.1051/0004-6361:20064987.

Q. Z. Liu, J. van Paradijs, and E. P. J. van den Heuvel. A catalogue of low-mass X-ray binaries in the Galaxy, LMC, and SMC (Fourth edition). A&A, 469(2):807–810, July 2007. doi: 10.1051/0004-6361:20077303.

D. Lynden-Bell and J. E. Pringle. The evolution of viscous discs and the origin of the nebular variables. MNRAS, 168:603–637, September 1974. doi: 10.1093/mnras/168.3.603.

J. F. Mahlmann, M. A. Aloy, V. Mewes, and P. Cerdá-Durán. Computational general relativistic force-free electrodynamics. I. Multi-coordinate implementation and testing. A&A, 647:A57, March 2021. doi: 10.1051/0004-6361/202038907.

Jonathan C. McKinney. General relativistic force-free electrodynamics: a new code and applications to black hole magnetospheres. MNRAS, 367(4):1797–1807, April 2006. doi: 10.1111/j.1365-2966.2006.10087.x.

M. J. Middleton, M. Brightman, F. Pintore, M. Bachetti, A. C. Fabian, F. Fürst, and D. J. Walton. On the magnetic field in M51 ULX-8. MNRAS, 486(1):2–9, June 2019. doi: 10.1093/mnras/stz436.

D. Mihalas and B. W. Mihalas. *Foundations of radiation hydrodynamics*. 1984.

B. Mishra, P. C. Fragile, L. C. Johnson, and W. Kluźniak. Three-dimensional, global, radiative GRMHD simulations of a thermally unstable disc. MNRAS, 463(4):3437–3448, December 2016. doi: 10.1093/mnras/stw2245.

Charles W. Misner, Kip S. Thorne, and John A. Wheeler. *Gravitation*. 1973.

Takahiro Miyoshi and Kanya Kusano. A multi-state HLL approximate Riemann solver for ideal magnetohydrodynamics. *Journal of Computational Physics*, 208(1):315–344, September 2005. doi: 10.1016/j.jcp.2005.02.017.

C. Motch, M. W. Pakull, R. Soria, F. Grisé, and G. Pietrzyński. A mass of less than 15 solar masses for the black hole in an ultraluminous X-ray source. Nature, 514(7521):198–201, October 2014. doi: 10.1038/nature13730.

B. Mukhopadhyay and L. Z. Fang. Hydrodynamical Study of Advective Accretion Flow Around Neutron Stars. *International Journal of Modern Physics D*, 11:1305–1319, 2002. doi: 10.1142/S0218271802002244.

A. A. Mushtukov, V. F. Suleimanov, S. S. Tsygankov, and J. Poutanen. The critical accretion luminosity for magnetized neutron stars. MNRAS, 447:1847–1856, February 2015a. doi: 10.1093/mnras/stu2484.

A. A. Mushtukov, V. F. Suleimanov, S. S. Tsygankov, and J. Poutanen. On the maximum accretion luminosity of magnetized neutron stars: connecting X-ray pulsars and ultraluminous X-ray sources. MNRAS, 454:2539–2548, December 2015b. doi: 10.1093/mnras/stv2087.

Alexander A. Mushtukov, Valery F. Suleimanov, Sergey S. Tsygankov, and Adam Ingram. Optically thick envelopes around ulxs powered by accreating neutron stars. *Monthly Notices of the Royal Astronomical Society*, 467(1):1202–1208, 2017. doi: 10.1093/mnras/stx141. URL `+http://dx.doi.org/10.1093/mnras/stx141`.

Alexander A. Mushtukov, Adam Ingram, Matthew Middleton, Dmitrij I. Nagirner, and Michiel van der Klis. Timing properties of ULX pulsars: optically thick envelopes and outflows. MNRAS, 484 (1):687–697, March 2019. doi: 10.1093/mnras/sty3525.

Alexander A. Mushtukov, Simon Portegies Zwart, Sergey S. Tsygankov, Dmitrij I. Nagirner, and Juri Poutanen. Pulsating ULXs: large pulsed fraction excludes strong beaming. MNRAS, 501(2): 2424–2429, February 2021. doi: 10.1093/mnras/staa3809.

R. Narayan and I. Yi. Advection-dominated Accretion: Underfed Black Holes and Neutron Stars. ApJ, 452:710, October 1995. doi: 10.1086/176343.

R. Narayan, M. R. Garcia, and J. E. McClintock. Advection-dominated Accretion and Black Hole Event Horizons. ApJ, 478:L79–L82, April 1997. doi: 10.1086/310554.

R. Narayan, A. Sądowski, R. F. Penna, and A. K. Kulkarni. GRMHD simulations of magnetized advection-dominated accretion on a non-spinning black hole: role of outflows. MNRAS, 426: 3241–3259, November 2012. doi: 10.1111/j.1365-2966.2012.22002.x.

Ramesh Narayan and Insu Yi. Advection-dominated Accretion: A Self-similar Solution. ApJ, 428:L13, June 1994. doi: 10.1086/187381.

Scott C. Noble, Charles F. Gammie, Jonathan C. McKinney, and Luca Del Zanna. Primitive Variable Solvers for Conservative General Relativistic Magnetohydrodynamics. ApJ, 641(1):626–637, April 2006. doi: 10.1086/500349.

I. D. Novikov and K. S. Thorne. Astrophysics of black holes. In *Black Holes (Les Astres Occlus)*, pages 343–450, January 1973.

Ken Ohsuga, Shin Mineshige, Masao Mori, and Masayuki Umemura. Does the Slim-Disk Model Correctly Consider Photon-trapping Effects? ApJ, 574(1):315–324, July 2002. doi: 10.1086/340798.

Ken Ohsuga, Masao Mori, Taishi Nakamoto, and Shin Mineshige. Supercritical Accretion Flows around Black Holes: Two-dimensional, Radiation Pressure-dominated Disks with Photon Trapping. ApJ, 628:368–381, July 2005. doi: 10.1086/430728.

B. Paczyński and P. J. Wiita. Thick accretion disks and supercritical luminosities. A&A, 500:203–211, August 1980.

K. Parfrey and A. Tchekhovskoy. General-relativistic Simulations of Four States of Accretion onto Millisecond Pulsars. ApJ, 851:L34, December 2017. doi: 10.3847/2041-8213/aa9c85.

Alessandro Patruno and Anna L. Watts. Accreting Millisecond X-ray Pulsars. In Tomaso M. Belloni, Mariano Méndez, and Chengmin Zhang, editors, *Astrophysics and Space Science Library*, volume 461 of *Astrophysics and Space Science Library*, pages 143–208, January 2021. doi: 10.1007/978-3-662-62110-3_4.

R. F. Penna, A. Kulkarni, and R. Narayan. A new equilibrium torus solution and GRMHD initial conditions. A&A, 559:A116, November 2013a. doi: 10.1051/0004-6361/201219666.

Robert F. Penna, Aleksander Sądowski, Akshay K. Kulkarni, and Ramesh Narayan. The Shakura-Sunyaev viscosity prescription with variable α (r). MNRAS, 428(3):2255–2274, January 2013b. doi: 10.1093/mnras/sts185.

Ciro Pinto, Matthew J. Middleton, and Andrew C. Fabian. Resolved atomic lines reveal outflows in two ultraluminous X-ray sources. Nature, 533(7601):64–67, May 2016. doi: 10.1038/nature17417.

F. Pintore, L. Zampieri, L. Stella, A. Wolter, S. Mereghetti, and G. L. Israel. Pulsator-like Spectra from Ultraluminous X-Ray Sources and the Search for More Ultraluminous Pulsars. ApJ, 836:113, February 2017. doi: 10.3847/1538-4357/836/1/113.

José A. Pons and Daniele Viganò. Magnetic, thermal and rotational evolution of isolated neutron stars. *Living Reviews in Computational Astrophysics*, 5(1):3, December 2019. doi: 10.1007/s41115-019-0006-7.

R. Popham and R. Sunyaev. Accretion Disk Boundary Layers around Neutron Stars: X-Ray Production in Low-Mass X-Ray Binaries. ApJ, 547:355–383, January 2001. doi: 10.1086/318336.

J. E. Pringle. Accretion discs in astrophysics. ARA&A, 19:137–162, January 1981. doi: 10.1146/annurev.aa.19.090181.001033.

Pablo Reig. Be/X-ray binaries. Ap&SS, 332(1):1–29, March 2011. doi: 10.1007/s10509-010-0575-8.

M. Revnivtsev and S. Mereghetti. Magnetic Fields of Neutron Stars in X-Ray Binaries. Space Sci. Rev., 191:293–314, October 2015. doi: 10.1007/s11214-014-0123-x.

Luciano Rezzolla and Olindo Zanotti. *Relativistic Hydrodynamics*. 2013.

George B. Rybicki and Alan P. Lightman. *Radiative Processes in Astrophysics*. 1986.

A. Sadowski. Slim accretion disks around black holes. *arXiv e-prints*, art. arXiv:1108.0396, August 2011.

A. Sądowski and R. Narayan. Photon-conserving Comptonization in simulations of accretion discs around black holes. MNRAS, 454:2372–2380, December 2015. doi: 10.1093/mnras/stv2022.

A. Sądowski and R. Narayan. Three-dimensional simulations of supercritical black hole accretion discs - luminosities, photon trapping and variability. MNRAS, 456:3929–3947, March 2016. doi: 10.1093/mnras/stv2941.

A. Sądowski, R. Narayan, A. Tchekhovskoy, and Y. Zhu. Semi-implicit scheme for treating radiation under M1 closure in general relativistic conservative fluid dynamics codes. MNRAS, 429:3533–3550, March 2013. doi: 10.1093/mnras/sts632.

A. Sądowski, R. Narayan, A. Tchekhovskoy, D. Abarca, Y. Zhu, and J. C. McKinney. Global simulations of axisymmetric radiative black hole accretion discs in general relativity with a mean-field magnetic dynamo. MNRAS, 447:49–71, February 2015. doi: 10.1093/mnras/stu2387.

N. I. Shakura and R. A. Sunyaev. Black holes in binary systems. Observational appearance. A&A, 24: 337–355, 1973.

I. S. Shklovsky. On the Nature of the Source of X-Ray Emission of Sco XR-1. ApJ, 148:L1, April 1967. doi: 10.1086/180001.

N. R. Sibgatullin and R. A. Sunyaev. Energy Release on the Surface of a Rapidly Rotating Neutron Star during Disk Accretion: A Thermodynamic Approach. *Astronomy Letters*, 26:772–778, December 2000. doi: 10.1134/1.1331157.

Aleksander Sądowski. Thin accretion discs are stabilized by a strong magnetic field. MNRAS, 459(4): 4397–4407, July 2016. doi: 10.1093/mnras/stw913.

Gary A Sod. A survey of several finite difference methods for systems of nonlinear hyperbolic conservation laws. *Journal of Computational Physics*, 27(1):1–31, 1978. ISSN 0021-9991. doi: https://doi.org/10.1016/0021-9991(78)90023-2. URL `https://www.sciencedirect.com/science/article/pii/0021999178900232`.

Jae Sok Oh, Hongsu Kim, and Hyung Mok Lee. Finite size effects on the Poynting-Robertson effect: A fully general relativistic treatment. New A, 16(3):183–186, April 2011. doi: 10.1016/j.newast.2010.07.008.

Volker Springel. E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh. *Monthly Notices of the Royal Astronomical Society*, 401(2):791–851, 01 2010. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2009.15715.x. URL `https://doi.org/10.1111/j.1365-2966.2009.15715.x`.

A. Stahl, M. Wielgus, M. Abramowicz, W. Kluźniak, and W. Yu. Eddington capture sphere around luminous stars. A&A, 546:A54, October 2012. doi: 10.1051/0004-6361/201220187.

Tod Strohmayer and Lars Bildsten. *New views of thermonuclear bursts*, volume 39, pages 113–156. 2006.

Tod E. Strohmayer and Craig B. Markwardt. Evidence for a Millisecond Pulsar in 4U 1636-53 during a Superburst. ApJ, 577(1):337–345, September 2002. doi: 10.1086/342152.

Tod E. Strohmayer, William Zhang, Jean H. Swank, Alan Smale, Lev Titarchuk, Charles Day, and Umin Lee. Millisecond X-Ray Variability from an Accreting Neutron Star System. ApJ, 469:L9, September 1996. doi: 10.1086/310261.

R. A. Syunyaev and N. I. Shakura. Disk Accretion onto a Weak Field Neutron Star - Boundary Layer Disk Luminosity Ratio. *Soviet Astronomy Letters*, 12:117–120, April 1986.

Ewa Szuszkiewicz. Slim accretion discs with different viscosity prescriptions. MNRAS, 244:377–383, May 1990.

Aleksander Sądowski, Maciek Wielgus, Ramesh Narayan, David Abarca, Jonathan C. McKinney, and Andrew Chael. Radiative, two-temperature simulations of low-luminosity black hole accretion flows in general relativity. *Monthly Notices of the Royal Astronomical Society*, 466(1):705–725, 12 2016. ISSN 0035-8711. doi: 10.1093/mnras/stw3116. URL `https://doi.org/10.1093/mnras/stw3116`.

H. R. Takahashi and K. Ohsuga. General Relativistic Radiation MHD Simulations of Supercritical Accretion onto a Magnetized Neutron Star: Modeling of Ultraluminous X-Ray Pulsars. ApJ, 845:L9, August 2017. doi: 10.3847/2041-8213/aa8222.

H. R. Takahashi, S. Mineshige, and K. Ohsuga. Supercritical Accretion onto a Non-magnetized Neutron Star: Why is it Feasible? ApJ, 853:45, January 2018. doi: 10.3847/1538-4357/aaa082.

Y. Tawara, T. Kii, S. Hayakawa, H. Kunieda, K. Masai, F. Nagase, H. Inoue, K. Koyama, F. Makino, K. Makishima, M. Matsuoka, T. Murakami, M. Oda, Y. Ogawara, T. Ohashi, N. Shibazaki, Y. Tanaka, S. Miyamoto, H. Tsunemi, K. Yamashita, and I. Kondo. A very long X-ray burst with a precursor from XB 1715-321. ApJ, 276:L41–L44, January 1984. doi: 10.1086/184184.

E.F. Toro. *Riemann Solvers and Numerical Methods for Fluid Dynamics*. 2009.

Gábor Tóth. The $\nabla\cdot$ B=0 Constraint in Shock-Capturing Magnetohydrodynamics Codes. *Journal of Computational Physics*, 161(2):605–652, July 2000. doi: 10.1006/jcph.2000.6519.

L. J. Townsend, J. A. Kennea, M. J. Coe, V. A. McBride, D. A. H. Buckley, P. A. Evans, and A. Udalski. The 2016 super-Eddington outburst of SMC X-3: X-ray and optical properties and system parameters. MNRAS, 471(4):3878–3887, November 2017. doi: 10.1093/mnras/stx1865.

S. P. Trudolyubov. XMM-Newton discovery of transient X-ray pulsar in NGC 1313. MNRAS, 387(1): L36–L40, June 2008. doi: 10.1111/j.1745-3933.2008.00478.x.

S. S. Tsygankov, V. Doroshenko, A. A. Lutovinov, A. A. Mushtukov, and J. Poutanen. SMC X-3: the closest ultraluminous X-ray source powered by a neutron star with non-dipole magnetic field. A&A, 605:A39, September 2017. doi: 10.1051/0004-6361/201730553.

B. van Leer. Towards the Ultimate Conservative Difference Scheme. V. A Second-Order Sequel to Godunov's Method. *Journal of Computational Physics*, 32(1):101–136, July 1979. doi: 10.1016/0021-9991(79)90145-1.

Bram van Leer. Towards the Ultimate Conservation Difference Scheme. II. Monotonicity and Conservation Combined in a Second-Order Scheme. *Journal of Computational Physics*, 14(4):361–370, March 1974. doi: 10.1016/0021-9991(74)90019-9.

Bram van Leer. Towards the Ultimate Conservative Difference Scheme. IV. A New Approach to Numerical Convection. *Journal of Computational Physics*, 23:276, March 1977. doi: 10.1016/0021-9991(77)90095-X.

D. J. Walton, M. Bachetti, F. Fürst, D. Barret, M. Brightman, A. C. Fabian, B. W. Grefenstette, F. A. Harrison, M. Heida, J. Kennea, P. Kosec, R. M. Lau, K. K. Madsen, M. J. Middleton, C. Pinto, J. F. Steiner, and N. Webb. A Potential Cyclotron Resonant Scattering Feature in the Ultraluminous X-Ray Source Pulsar NGC 300 ULX1 Seen by NuSTAR and XMM-Newton. ApJ, 857(1):L3, April 2018. doi: 10.3847/2041-8213/aabadc.

I. Wasserman and S. L. Shapiro. Masses, radii, and magnetic fields of pulsating X-ray sources : is the "standard" model self-consistent ? ApJ, 265:1036–1046, Feb 1983. doi: 10.1086/160745.

Anna L. Watts. Thermonuclear Burst Oscillations. ARA&A, 50:609–640, September 2012. doi: 10.1146/annurev-astro-040312-132617.

Rainer Weinberger, Volker Springel, and Rüdiger Pakmor. The AREPO public code release. *The Astrophysical Journal Supplement Series*, 248(2):32, jun 2020. doi: 10.3847/1538-4365/ab908c. URL `https://doi.org/10.3847/1538-4365/ab908c`.

M. Wielgus, A. Stahl, M. Abramowicz, and W. Kluźniak. Oscillations of the Eddington capture sphere. A&A, 545:A123, September 2012. doi: 10.1051/0004-6361/201220228.

M. Wielgus, W. Kluźniak, A. Sądowski, R. Narayan, and M. Abramowicz. Stable, levitating, optically thin atmospheres of Eddington-luminosity neutron stars. MNRAS, 454(4):3766–3770, December 2015. doi: 10.1093/mnras/stv2191.

Maciek Wielgus, Aleksander Sądowski, Włodek Kluźniak, Marek Abramowicz, and Ramesh Narayan. Levitating atmospheres of Eddington-luminosity neutron stars. MNRAS, 458(4):3420–3428, June 2016. doi: 10.1093/mnras/stw548.

G. Wiktorowicz, M. Sobolewska, J.-P. Lasota, and K. Belczynski. The Origin of the Ultraluminous X-Ray Sources. ApJ, 846:17, September 2017. doi: 10.3847/1538-4357/aa821d.

Lizhong Zhang, Omer Blaes, and Yan-Fei Jiang. Radiative MHD simulations of photon bubbles in radiation-supported magnetized atmospheres of neutron stars with isotropic Thomson scattering. MNRAS, 508(1):617–636, November 2021. doi: 10.1093/mnras/stab2510.