

Disappearing thermal X-ray emission as a tell-tale signature of merging massive black hole binaries

Luke Major Krauth¹   Jordy Davelaar^{1,2,3,4}  Zoltán Haiman,^{1,2} John Ryan Westernacher-Schneider,⁵ Jonathan Zrake^{1,6}  and Andrew MacFadyen⁷ 

¹Department of Physics, Columbia University, New York, NY 10027, USA

²Department of Astronomy, Columbia University, New York, NY 10027, USA

³Astrophysics Laboratory, Columbia University, 550 W 120th Str, New York, NY 10027, USA

⁴Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA

⁵Leiden Observatory, Leiden University, PO Box 9513, NL-2300 RA Leiden, the Netherlands

⁶Department of Physics and Astronomy, Clemson University, Clemson, SC 29634, USA

⁷Center for Cosmology and Particle Physics, Physics Department, New York University, New York, NY 10003, USA

Accepted 2023 October 2. Received 2023 September 29; in original form 2023 April 5

ABSTRACT

The upcoming Laser Interferometer Space Antenna (LISA) is expected to detect gravitational waves (GWs) from massive black hole binaries (MBHB). Finding the electromagnetic (EM) counterparts for these GW events will be crucial for understanding how and where MBHBs merge, measuring their redshifts, constraining the Hubble constant and the graviton mass, and for other novel science applications. However, due to poor GW sky localization, multiwavelength, time-dependent EM models are needed to identify the right host galaxy. We studied merging MBHBs embedded in a circumbinary disc (CBD) using high-resolution two-dimensional simulations, with a Γ -law equation of state, incorporating viscous heating, shock heating, and radiative cooling. We simulate the binary from large separation until after merger, allowing us to model the decoupling of the binary from the CBD. We compute the EM signatures and identify distinct features before, during, and after the merger. Our main result is a multiband EM signature: we find that the MBHB produces strong thermal X-ray emission until 1–2 d prior to the merger. However, as the binary decouples from the CBD, the X-ray-bright minidisks rapidly shrink in size, become disrupted, and the accretion rate drops precipitously. As a result, the thermal X-ray luminosity drops by orders of magnitude, and the source remains X-ray dark for several days, regardless of any post-merger effects such as GW recoil or mass-loss. Looking for the abrupt spectral change where the thermal X-ray disappears is a tell-tale EM signature of LISA mergers that does not require extensive pre-merger monitoring.

Key words: accretion, accretion discs – black hole physics – hydrodynamics.

1 INTRODUCTION

Over the past decade, significant progress has been made in the detection of gravitational waves (GWs) by the Laser Interferometer Gravitational-Wave Observatory (Collaboration 2015), Virgo (Acernese et al. 2014) and the Kamioka Gravitational Wave Detector (Akutsu et al. 2021). These detectors are sensitive to GWs with frequencies between \sim 10 Hz and \sim 1 kHz, and have detected nearly 100 mergers involving stellar-mass black holes (BHs) and neutron stars. The Laser Interferometer Space Antenna (LISA; Amaro-Seoane et al. 2017) will access new frequencies between \sim 0.1 mHz and \sim 1 Hz, opening a new window to detect binaries with much larger masses and wider orbital separations.

While detecting these new sources and characterizing their GW signatures will be impactful by itself, it is recognised that the combined detection of GW and electromagnetic (EM) signatures will

enable a range of additional science. This so-called multimessenger combination can provide novel probes of alternative theories of gravity (such as those requiring extra dimensions, see de Rham, Melville & Tolley 2018), the mass of the graviton (such as in recent ghost-free massive gravity models, see Hassan & Rosen 2012), and the distance-redshift relation (including measurements of the Hubble constant H_0 , see Schutz 1986). On the astrophysics side, multimessenger detections will yield novel constraints on accretion physics and the co-evolution of galaxies and their central massive BHs (see e.g. Baker et al. 2019 for a brief review and references).

Possible EM signatures from merging MBHBs can be divided into two broad categories: pre- and post-merger (see e.g. Bogdanovic, Miller & Blecha 2022, for a comprehensive review). Before their merger, MBHBs are expected to be typically surrounded by a circumbinary disc (CBD), delivered to the nucleus by the preceding galaxy merger (Begelman, Blandford & Rees 1980). In recent years, a consensus has been reached by hydrodynamical simulations describing a MBHB + CBD system: while the binary creates a low-density central cavity, out to a few times its separation, the individual

* E-mail: LMK2202@columbia.edu

components are fed efficiently through narrow streams, fuelling individual ‘minidisks’. Pre-merger EM signatures can therefore originate from one of several locations: the CBD, the minidisks, or the streams, either colliding with each other or the wall of the cavity (see e.g. Westernacher-Schneider et al. 2022, and references therein).

Early work suggested that because the GW-driven inspiral time ($t_{\text{gw}} \propto a^4$) decreases more rapidly than the viscous inflow time (e.g. $t_{\text{visc}} \propto a^{7/5}$; Pringle 1991), the binary would outrun the circumbinary disc (Liu, Wu & Cao 2003; Milosavljević & Phinney 2005), starving and dimming the binary until the cavity re-fills in approximately a viscous time after the merger (which would take at least years or decades). However, long-term two-dimensional (2D) hydrodynamical simulations following the GW-driven inspiral did not reveal evidence for such a clear ‘decoupling’. Instead, they showed that the circumbinary gas could follow the inspiralling binary and feed the individual BHs all the way until the merger (Farris et al. 2015b; Tang, Haiman & MacFadyen 2018). Relativistic simulations of fewer orbits closer to the merger also suggest that efficient fuelling of the minidisks at these late states is possible, provided circumbinary gas still surrounds the binary (Noble et al. 2012; Bowen et al. 2017).

Pre-merger EM signatures include periodic brightness fluctuations on approximately the binary’s orbital timescale, driven by hydrodynamic modulations (see e.g. Artymowicz & Lubow 1996; Hayasaki, Mineshige & Sudou 2007; MacFadyen & Milosavljević 2008; Roedig et al. 2011; Shi et al. 2012; Noble et al. 2012; D’Orazio, Haiman & MacFadyen 2013; Farris et al. 2014, for early works) and/or kinematic Doppler (Bode et al. 2010; D’Orazio, Haiman & Schiminovich 2015; Haiman 2017) and lensing (D’Orazio & Di Stefano 2018; Davelaar & Haiman 2022a; Davelaar & Haiman 2022b) effects. The build-up of a non-axisymmetric density distribution near the edge of the cavity, known as a ‘lump’, modulates accretion on to the minidisks, causing periodicity on several times the orbital time-scale, as well as on beat periods between the binary and the lump (e.g. Shi et al. 2012; D’Orazio, Haiman & MacFadyen 2013; Farris et al. 2014; Shi & Krolik 2015; Noble et al. 2021; Mignon-Risse, Varniere & Casse 2023). Furthermore, mass exchange between the minidisks can introduce additional periodicity in the pre-merger phase (Bowen et al. 2017; Westernacher-Schneider et al. 2022).

Post-merger signatures are also expected to arise from the abrupt mass-loss and recoil velocity imparted to the remnant BH, caused by the burst of anisotropic GW emission at merger. These effects are nearly instantaneous, with the BH typically losing a few per cent of its mass (e.g. Tichy & Marronetti 2008) and recoiling at hundreds of km s^{-1} (e.g. Baker et al. 2008; Lousto & Zlochower 2013, and references therein), depending on the binary mass ratio, eccentricity, and BH spins. In the idealized case of a near-Keplerian, geometrically thin circumbinary disc whose mass is negligible compared with the BH remnant (a good approximation for MBHBs detectable by LISA, out to radii $\sim (10^2 - 10^4) \times r_g$, where $r_g = GM_{\text{bin}}/c^2$ is the gravitational radius according to the total BH mass M_{bin} , see e.g. Lippai, Frei & Haiman 2008), the change in the potential and the effective kick received by the gas cause spiral and/or concentric circular shock-waves to propagate outwards in the disc, as shown in analytic orbital calculations (e.g. Lippai, Frei & Haiman 2008; Schnittman & Krolik 2008; Shields & Bonning 2008; Penoyre & Haiman 2018), hydrodynamic simulations (O’Neill et al. 2009; Corrales, Haiman & MacFadyen 2010; Rossi et al. 2010; Megevand et al. 2009; Rosotti, Lodato & Price 2012), and also in relativistic simulations of non-Keplerian discs (Zanotti et al. 2010). The resulting ‘afterglows’ can emerge from weeks to months after the merger, and exhibit a characteristic time-dependent spectral energy distribution (SED).

In this paper, we revisit the above observational signatures accompanying the merger of a pair of massive BHs. Our motivation for this is two-fold.

First, to discover the pre-merger signatures, the ongoing merger, detected by LISA, would need to be sufficiently well localized to allow a prompt observational campaign by large field-of-view (FOV) instruments (Kocsis, Haiman & Menou 2008). According to recent forecasts, sky localization to within $\sim 10 \text{ deg}^2$ (the FOV of the Vera Rubin Observatory, see Vera C. Rubin Collaboration et al. 2020) will typically only be available 1–2 d before merger (Mangiagli et al. 2020). This motivates developing a better understanding of signals immediately preceding the merger. In particular, the aforementioned hydrodynamical simulations of pre-merger circumbinary disc dynamics lacked adequate spatial and temporal resolution near the merger, relied on simplified thermodynamics, and/or did not track the binary inspiral for long enough from the time of the decoupling to the merger.

Secondly, existing models of post-merger afterglows (e.g. Rossi et al. 2010; Corrales, Haiman & MacFadyen 2010) have employed many simplifying assumptions, such as a smooth, laminar, inviscid, near-Keplerian initial disc in a Newtonian potential, and either an isothermal or an adiabatic (i.e. non-radiative) equation of state. Since these works, it has been understood that the vicinity of the binary at merger is not devoid of gas; rather there is a distorted cavity that is filled with streams of gas (Farris et al. 2015b; Tang, Haiman & MacFadyen 2018). At present, there are no predictions for the EM signatures arising from the mass-loss and the recoil immediately following the merger from this inner region of the disturbed circumbinary disc.

In this paper, we address the above limitations by employing high-resolution hydrodynamical simulations with the fixed-mesh GPU-enabled code SAILFISH¹ to compute the EM signatures immediately preceding and following the merger. Our work extends previous findings by incorporating physical viscosity and a more realistic treatment of thermodynamics, directly solving the energy equation, using a Γ -law equation of state for the gas, and a physically motivated cooling prescription. Additionally, we simulate the binary inspiral for several hundred orbits, which allows us to follow the system from the fiducial decoupling epoch all the way past merger. We perform a suite of simulations varying the disc Mach number and viscosity, the binary mass, and grid resolution to test the robustness of our results.

Our main result is a novel multiband EM signature: we find that the fiducial $10^6 M_{\odot}$ MBHB model produces strong thermal X-ray emission until 1–2 d prior to merger (in agreement with previous works). However, at that stage, the minidisks are abruptly disrupted and the shocked streams disappear from the inner cavity. As a result, the X-ray luminosity drops by orders of magnitude, and the source remains X-ray dark for several days afterwards, regardless of any post-merger effects such as recoil or mass-loss. At the same time, the optical and infrared emission, dominated by circumbinary gas farther out, remains nearly steady. These distinct signatures could help identify LISA counterparts without the need for localizing the source on the sky more than 1–2 d prior to merger.

The remainder of this paper is organized as follows. In Section 2, we discuss our hydrodynamics code, numerical schemes, and the initial set-up for our list of runs. In Section 3.1, we present our main results on the gas dynamics and EM signatures before and after the merger. In Section 3.2, we discuss how our results depend on varying the most important parameters in our set-up. In Section 4,

we discuss observational issues and strategies. Finally, in Section 5, we summarize our main conclusions and the implications of this work.

2 HYDRODYNAMICAL SET-UP, POST-PROCESSING, AND MODELS

2.1 Hydrodynamical set-up

All simulations were run using the publicly available 2D GPU-enabled hydrodynamics code SAILFISH. In this section, we give a brief summary of the technical aspects of SAILFISH (for full details, see Westernacher-Schneider et al. 2022).

We solve the vertically integrated Newtonian fluid equations, keeping the lowest non-trivial order in powers of z/r under the conditions of a thin disc ($h/r \ll 1$) and mirror symmetry about $z = 0$. These equations read

$$\partial_t \Sigma + \nabla_j (\Sigma v^j) = S_\Sigma, \quad (1)$$

$$\partial_t (\Sigma v_i) + \nabla_j (\Sigma v^j v_i + \delta_i^j \mathcal{P}) = g_i + \nabla_j \tau_i^j + S_{p,i}, \quad (2)$$

$$\begin{aligned} \partial_t E + \nabla_j [(E + \mathcal{P}) v^j] &= v^j g_j + \nabla_j (v^i \tau_i^j) \\ &\quad - \dot{Q} + S_E, \end{aligned} \quad (3)$$

where Σ is the surface density, \mathcal{P} is the vertically integrated pressure, v^i is the mid-plane horizontal fluid velocity, $E = \Sigma \epsilon + (1/2) \Sigma v^2$ is the vertically integrated energy density, ϵ is the specific internal energy density at the mid-plane of the disc, g_i is the vertically integrated gravitational force density, and $\tau_i^j = \Sigma v \left(\nabla_i v^j + \nabla^j v_i - (2/3) \delta_i^j \nabla_k v^k \right)$ is the viscous stress tensor [in a form that is trace-free in a three-dimensional (3D) sense] with zero bulk viscosity, v is the kinematic shear viscosity, S_Σ , $S_{p,i}$, and S_E are mass, momentum, and energy sinks, and \dot{Q} is the local blackbody cooling prescription, assuming hydrogen dominates the gas density (see e.g Frank, King & Raine 2002), given as

$$\dot{Q} = \frac{8}{3} \frac{\sigma}{\kappa \Sigma} \left(\frac{m_p \mathcal{P}}{k_B \Sigma} \right)^4, \quad (4)$$

where σ is the Stefan–Boltzmann constant, $\kappa = 0.4 \text{ cm}^2 \text{ g}^{-1}$ is the opacity due to electron scattering, m_p is the proton mass, and k_B is the Boltzmann constant. Thermal conductivity is neglected.

We use a torque-free sink prescription (Dempsey, Muñoz & Lithwick 2020; Dittmann & Ryan 2021) to model the removal of gas by each point mass. We choose the sink radius equal to r_s (where $r_s \equiv 2GM_{\text{bh}}/c^2$ is the Schwarzschild radius of a single BH). This allows us to achieve sufficient resolution near the BHs (sink diameter = 8 cells).

We initialize the disc to the conditions described in Goodman (2003):

$$\Sigma = \Sigma_0 \left(\frac{\sqrt{r^2 + r_{\text{soft}}^2}}{a_0} \right)^{-3/5} \quad (5)$$

$$\mathcal{P} = \mathcal{P}_0 \left(\frac{\sqrt{r^2 + r_{\text{soft}}^2}}{a_0} \right)^{-3/2} \quad (6)$$

$$\vec{v} = \left(\frac{GM_{\text{bin}}}{\sqrt{r^2 + r_{\text{soft}}^2}} \right)^{1/2} \hat{\phi}, \quad (7)$$

where r_{soft} is the gravitational softening length, which mimics the vertically integrated component of the gravitational force in the plane of the disc, and is set equal to the sink radius, and a_0 is the initial binary separation.

Additionally, we initialize a cavity at $r = 2a_0$ by multiplying both Σ and \mathcal{P} by a window function $f(r)$, given by

$$f(r) = 10^{-4} + (1 - 10^{-4}) \exp \left(-2a_0 / (r^2 + r_{\text{soft}}^2)^{1/2} \right)^{30}. \quad (8)$$

We use a Γ -law equation of state $\mathcal{P} = \Sigma \epsilon (\Gamma - 1)$, where $\Gamma = 5/3$. While the gas is more likely to be radiation-dominated ($\Gamma = 4/3$), our choice allows us to compare with previous results most closely related to our study (Corrales, Haiman & MacFadyen 2010; Farris et al. 2015b; Tang, Haiman & MacFadyen 2018; Westernacher-Schneider et al. 2022, which all used $\Gamma = 5/3$), and to avoid the viscous and thermal instabilities associated with such $\Gamma = 4/3$ models (Lightman & Eardley 1974; Shakura & Sunyaev 1976, respectively). We use a Shakura–Sunyaev viscosity prescription $\nu = \alpha c_s h$ (Shakura & Sunyaev 1973) with a fiducial $\alpha = 0.1$. The sound speed is determined by $c_s^2 = \Gamma \mathcal{P} / \Sigma$. The disc's half thickness, h , is determined by $h = \sqrt{\mathcal{P}/\Sigma} / \tilde{\Omega}$ with $\tilde{\Omega} = \sqrt{GM_1/r_1^3 + GM_2/r_2^3}$, where r_1 and r_2 are the distances from the particle and the respective BHs. The disc mass is small compared to the mass of the binary, so self-gravity can be ignored.

Although hydrodynamic radiation pressure is typically important in the innermost regions of MBH discs, including it in simulations is technically challenging. Thus, we omit radiation pressure, and instead initialize our disc with a characteristic disc aspect ratio h/r that is similar to that expected when radiation pressure is accounted for. In practice, the characteristic disc aspect ratio is adjusted via choices of Σ_0 and \mathcal{P}_0 , and when interpreted literally, can imply an extremely super-Eddington accretion rate (D’Orazio, Haiman & MacFadyen 2013). The effective temperature of the disc, T_{eff} , is related to mid-plane temperature T by

$$T_{\text{eff}}^4 = \frac{4}{3} \frac{T^4}{\kappa \Sigma}, \quad (9)$$

In Section 2.2, we describe how in post-processing we scale the effective temperature down into the range expected when radiation pressure is accounted for. Although the disc aspect ratio develops self-consistently from a balance of heating and cooling in the simulation runs, we find its initialized characteristic value is representative of the fully developed system. We relate the accretion rate to the surface density and kinematic shear viscosity via $\dot{M} = 3\pi \Sigma v$ (Frank, King & Raine 2002).

The combined mass of the binary was chosen to be $M_{\text{bin}} = 10^6 M_\odot$, roughly matching the value where LISA is most sensitive (Amaro-Seoane et al. 2017). Simulations suggest that gas accretion drives binaries towards equal mass (see e.g. Farris et al. 2014; Duffell et al. 2020), so we simulate binaries with mass ratios $q \equiv m_2/m_1 = 1$. An initial separation of $a_0 = 50 r_g$ (where $r_g \equiv GM_{\text{bin}}/c^2$ is the gravitational radius for the total binary mass) was chosen to be wide enough for the binary to traverse the fiducial decoupling radius. More specifically, the initial viscous time, $t_v = 2/3 r^2/v$, in our fiducial run at $r = a_0$ is 16.54 d, which is shorter than the inspiral time of 27.84 d. The binary orbits are circular and the BHs are assumed to have zero spin.

The gravitational field of each individual BH is modelled by a Plummer potential,

$$\Phi_n = - \frac{GM_n}{\sqrt{r_n^2 + r_{\text{soft}}^2}}, \quad (10)$$

where M_n is the mass of the n th BH, r_n is the distance from a field point to the n th BH, and r_{soft} is the gravitational softening length scale introduced above. A pseudo-Newtonian potential, such as Paczyński-Wiita (Paczyński & Wiita 1980), would likely benefit

Table 1. Our suite of simulations. The bold denotes which parameter has been modified from the fiducial model in the first row. All models listed are run with no kick, kicks in four different directions with respect to the orientation of the cavity, and with and without mass loss in every case, *except Res_8k, which was only run in the most luminous kick direction (described in Section 3.1.2). When applied, the recoil velocity is 530 km s⁻¹ and the mass-loss is 3 per cent of the binary’s total initial mass for all simulations.

Model	Description	α	Mass (M _⊙)	Mach	Res (Δx)
Fiducial	Fiducial Model	0.1	10 ⁶	10	0.01
Alpha0.3	Higher viscosity	0.3	10 ⁶	10	0.01
Mass10 ⁸	Higher BH mass	0.1	10⁸	10	0.01
Mach30	Higher Mach number	0.1	10 ⁶	30	0.01
Res_1k	Lowest resolution	0.1	10 ⁶	10	0.04
Res_2k	Lower resolution	0.1	10 ⁶	10	0.02
Res_8k*	Highest resolution	0.1	10 ⁶	10	0.005

our study. However, when attempting to employ this potential, we noticed post-merger numerical effects in which the sinks did not appear to be removing material properly. As such, we instead opted to use a Plummer potential. We intend to resolve these numerical issues in future work, and perhaps explore more advanced pseudo-Newtonian potentials, in addition to Paczyński-Wiita.

The realistic state of the disc, including the eccentric cavity and lump in the cavity wall, takes time to develop from our idealized axisymmetric initial conditions. Thus, before initiating inspiral, we allow our system to evolve for 820 orbits on a fixed circular orbit, corresponding to ~ 6.3 viscous times at $r = a_0$. By comparison, the time to merger is roughly 25 per cent of this duration (approximately 220 initial orbital times).

The inspiral is implemented with the quadrupole approximation (Peters 1964), i.e. the binary separation follows

$$a(t) = a_0(1 - t/\tau)^{1/4}, \quad (11)$$

where $\tau = a_0^4/4\beta$ is the total GW inspiral time and $\beta \equiv (64/5)(G^3/c^5)m_1m_2(m_1 + m_2)$. For the binary separation of 50 r_g , this implies $\tau \sim 28$ d ~ 220 initial orbits ~ 350 total orbits. This initial separation ensures that we follow the binary through the decoupling limit. Once the BHs are within two Schwarzschild radii of each other, the BHs are instantly merged, with a new position at the origin and a sink radius of 2 r_s , double the radius of either initial BH.

Recoil is expected to occur nearly instantaneously at merger. Thus, in our kicked simulations, immediately after the merger an in-plane velocity of 530 km s⁻¹ is applied to our BH. This is a typical value, and also what was chosen in Corrales, Haiman & MacFadyen (2010), allowing our results to be compared to that work. Similarly, as mass-loss is also expected to occur approximately instantaneously at merger, we treat it as such, by reducing the total mass of the resulting BH by 3 per cent for our mass-loss simulations, in line with expected values. We examine kicks in four different directions with respect to the orientation of the cavity, as described below, as well as a reference model without any kick. All of these setups are examined with and without mass-loss. These variations, along with the additional model runs discussed in Section 2.3, are listed in Table 1. We use a uniform Cartesian grid with a square domain of side length 40 a_0 and 4000² cells, giving a resolution of $\Delta x = \Delta y = 0.01 a_0$. We use a buffer source term at the outer boundary off the grid, the purpose of which is to drive the solution to the initial conditions for the disc, preventing the outer boundary from propagating artefacts inwards. Additional tests include: changing the sink prescription inside the sink to immediately set fluid velocities to zero and the pressure and density to their respective floor values; altering the density and

pressure floor values; and running with different PLM (piece-wise linear method) values.² All of these tests reveal negligible effects on our scientific conclusions. We also perform a resolution test in our fiducial model to ensure convergence of our results (Appendix A).

2.2 Post-processing

In post-processing, the effective temperature is obtained from equations (4) and (9), leading to

$$\dot{Q} = 2\sigma T_{\text{eff}}^4, \quad (12)$$

where the factor of 2 comes from the fact that the disc cools through two faces (top and bottom). Since the effective temperature is related to the accretion rate via

$$T_{\text{eff}}^4 \propto \dot{M}, \quad (13)$$

the artificially high accretion rates required to achieve realistic characteristic disc aspect ratios in the gas-dominated models result in artificially high effective temperatures, which affect the EM luminosity and spectrum. Therefore, we correct for this in post-processing by uniformly re-scaling the effective temperature back down to our target system via the map $T_{\text{eff}}^4 \rightarrow T_{\text{eff}}^4/M_{\text{boost}}$, where M_{boost} is the dimensionless ratio of the accretion rate of the gas pressure model to that of the ‘target’ model where radiation pressure is accounted for.

We assume black-body emission from each cell of our domain, which allows us to compute the luminosity in different bands. We neglect Doppler effects, so our light curves are valid for observers who are viewing the disc face-on. We discuss how our results are modified when the light curves are modulated by relativistic Doppler boosting in Appendix B. We calculate the luminosity of an area element dA in the frequency band between ν_1 and ν_2 via

$$dL = \pi dA \int_{\nu_1}^{\nu_2} \frac{2h\nu^3/c^2}{\exp\left(\frac{h\nu}{kT_{\text{eff}}}\right) - 1} d\nu, \quad (14)$$

To obtain the total luminosity of the disc, we sum over the spatial domain, excluding the outer buffer region. We examine the bolometric luminosity, as well as the luminosities in three fixed bands $E_{\text{opt}} : 1.8\text{--}3.1$ eV, $E_{\text{UV}} : 3.1\text{--}124.0$ eV, $E_{\text{X-ray}} : 124.0$ eV – 124.0 keV.

²This controls the aggressiveness of slope-limiting in the hydrodynamic solver.

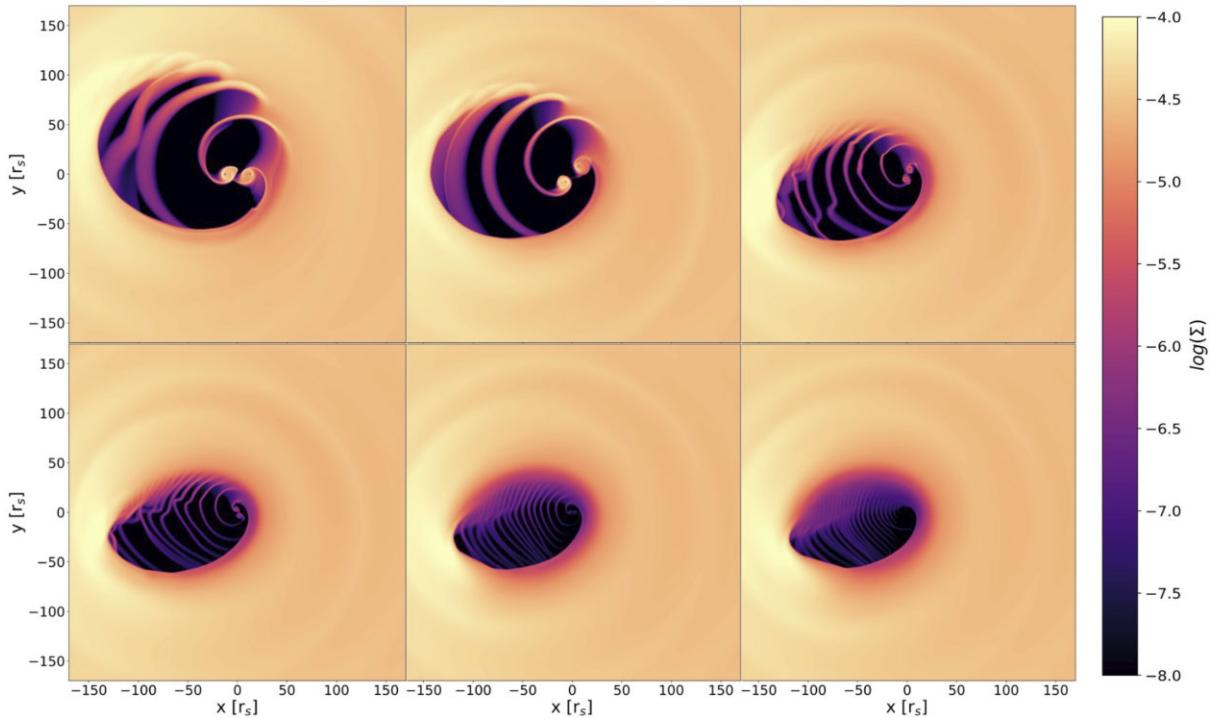


Figure 1. Snapshots of the logarithmic surface density from the beginning of inspiral to merger, shown (reading left to right, top to bottom) at 27.8 d, 13.9 d, 1 d, 5 h, and 1 h before merger, and at merger.

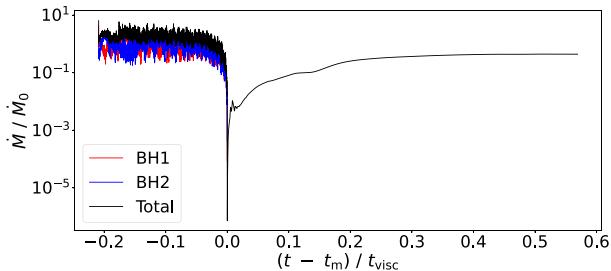


Figure 2. BH accretion rates before and after the time of merger, t_m , for the case with no post-merger recoil or mass-loss. Accretion rates are normalized by the steady Shakura–Sunyaev value around a single BH.

2.3 Suite of models

We also explore variants of our model with an increased viscosity of $\alpha = 0.3$, an increased total binary mass of $M_{\text{bin}} = 10^8 M_{\odot}$, and an increased Mach number of $\mathcal{M} = 30$. In each of these models, we again examine the cases with no kick, kicks in four directions, and with and without mass-loss in every case. Our full suite of models is shown in Table 1.

3 RESULTS

In this section, we describe the results of our simulation library, first our fiducial model, and then how our main results depend on model parameters.

3.1 Fiducial model

We present in this subsection the main results from the fiducial merger. We summarize the evolution up to the merger and how

this affects the EM emission, then highlight post-merger effects and comment on the nature of the periodicities we find in EM light curves.

3.1.1 Binary evolution up to merger

After running the simulation for about six viscous times, the CBD is in a settled state. Fig. 1 shows the logarithmic surface density for the \sim inner third of our domain. The top left panel, which shows the system just before inspiral is initiated, shows a clear eccentric cavity, represented by the underdense region which surrounds the binary. As the BHs orbit they pick up mass from the inner edge of the cavity, creating streams within the cavity. Parts of these streams are accreted on to two smaller accretion discs around each BH, the minidisks, while other parts are expelled towards the opposite side of the cavity. The minidisks are clearly visible as overdensities around each BH. All of these features have been found in previous simulations of accreting equal-mass binaries on fixed circular orbits.

The other five panels in Fig. 1 show the later stages of the binary inspiral, ending with the merger in the bottom right-hand panel. As binaries orbit, the minidisk tidal truncation radius tracks, on average, $\sim 1/3$ of the instantaneous separation (Paczynski 1977). As the binary shrinks, the minidisk tidal truncation radius also shrinks, generally leading to a decrease in their total mass. The frequency of stream creation naturally increases with the binary frequency, producing an increasing number of ejected streams visible in the cavity at later times. At merger, the minidisks are disrupted and disappear.

The mass accretion rate of both binary components is shown as a function of time in Fig. 2. As our BH would otherwise move out of the simulation domain in this time period, the model shown has no kick (nor mass loss). A significant drop in the accretion rate occurs at the time of merger. This indicates that as the minidisks surrounding the BHs disperse, the accretion also abates, at least temporarily.

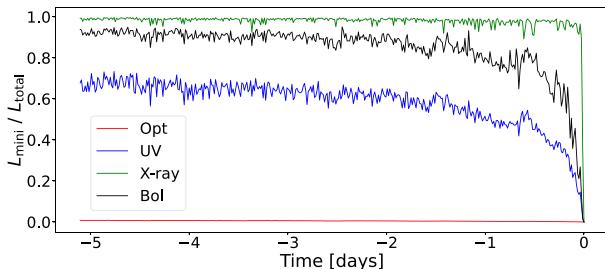


Figure 3. The fraction of the total luminosity contributed by the minidisks in different bands, as labelled. The minidisks produce nearly all of the X-rays but almost none of the optical emission.

Accretion begins to restore soon after merger, and appears to settle over nominal viscous time-scales.

Since mass accretion rates are not directly observable, it is interesting to investigate in which EM bands this minidisk depletion could leave an imprint. We first investigated in which EM frequency band the minidisks are brightest. Fig. 3 shows the fraction of optical, UV, X-ray, and bolometric luminosities that originate from the minidisks. Luminosity is considered to have originated from the minidisks if it is emitted from a circular region centred on the binary barycentre, whose radius is chosen to lie just outside the outermost part of either minidisk (calculated from a combination of the instantaneous binary separation and truncation radius). The bulk of the bolometric emission (black curve) comes from the minidisks. Practically all of the X-ray emission (green curve) comes from the minidisks, whereas practically all of the optical emission (red curve) comes from the CBD. At the moment of merger, a clear drop in the minidisks' fraction of X-ray emission is evident. Since the minidisks survive up until hours before the merger, and they dominate the X-ray emission, their destruction leads to a corresponding drop in X-ray emission. This can be seen in Fig. 4, which shows light curves before and after merger for a case with a post-merger kick. The inset shows a zoomed-in view of the 12 h before and after merger. The red and green shaded regions indicate 5-h windows before and after the merger. A drop of five orders of magnitude is seen in the X-ray luminosity during this period. In the post-merger phase, partial re-brightening of X-rays can arise from shocks formed in the CBD due to BH recoil. We consider such post-merger effects in detail next.

3.1.2 Post merger effects

We investigate post-merger effects by imparting mass-loss as well as kicks in various directions. To inform our expectations for the effect of kicks, we first consider a toy model, following a test particle on a Keplerian elliptical orbit, and kick it at different points along its orbit, corresponding to periapse, apoapse, and two points in between. At each point, we examine kicks in four directions, to obtain a total of 16 kicked orbits. These orbits do not capture hydrodynamical effects, but serve to guide expectations of what happens in the simulations. The left four panels of Fig. 5 show the corresponding trajectories. Note that Galilean invariance implies, for example, that a BH kick upwards is equivalent to kicking the test particle downwards (upper left panel). Particle trajectories are counterclockwise, with kicked trajectories shown as dashed lines with varied colours.

Taking the BH kick Up direction as an example, we see that the particles kicked at the Top, Apoapsis, and Bottom points have only slight perturbations to their trajectories. However, the particle kicked

at the Periapsis point, which receives a relative kick opposing its velocity, has a more substantial change to its orbit. In particular, the orbit becomes more circular, with a significant reduction in the apoapsis. We might expect that a fluid element with this fate would be diverted away from the cavity wall, towards the inner region of the low-density cavity. Thus, we might also expect this to be the least luminous post-merger kick direction, as there are decreased chances of interactions between the gas streams. Conversely, looking at the lower right panel, BH Kick Down, the particle at Periapsis now receives a boost parallel to its velocity, which significantly increases its apoapsis. We expect that in the simulation, this corresponds to fluid elements travelling deeper into the opposing wall of the circumbinary cavity. These additional collisions might in turn be expected to cause a greater overall luminosity. Finally, the luminosity resulting from the Left and Right kicks might be expected to be in between the Up and Down cases, and, because of their overall symmetry, may have similar values.

These expectations are borne out by our hydrodynamic simulations. The upper right panel shows light curves for the same four kick directions (determined with the respect to the cavity orientation). Mass-loss leads to overall lower post-merger luminosity, as seen in the lower right panel, in which we show the no-kick and kick-down results with and without mass-loss.

We note that Fig. 4 shows the case where the BH kick direction opposes the orbit of the gas at periapsis (i.e. fluid elements near periapsis receiving boosts to their velocities in the frame of the kicked BH) and without mass-loss. According to the above, this case exhibits maximal re-brightening in the post-merger phase. Yet, even in this optimal scenario, the drop in the X-ray band is evident.

3.1.3 Periodicities

Fig. 6 shows the Lomb–Scargle periodogram of the X-ray light curve (green) and the GW chirp (red) for the last 15 d before merger. The GW chirp is computed from the Peters (1964) quadrupole formula for a circular equal-mass binary at $z = 1$. Time is re-scaled by the instantaneous binary orbital period, resulting in a periodogram where the frequency is re-scaled by the instantaneous binary orbital frequency. This is done to assess the periodicities of our signal in relation to the binary's shrinking orbital period. We note three distinct peaks in the X-ray periodogram: a peak near $\omega/\omega_{\text{bin}} \approx 0.1$, a peak at or slightly above the binary frequency $\omega/\omega_{\text{bin}} = 1$, and a small peak in line with the GW chirp at twice the binary orbital frequency $\omega/\omega_{\text{bin}} = 2$. We can investigate these frequencies by studying different segments of the light curve.

Fig. 7 shows the X-ray light curve (green) and the GW chirp (red) for the last 15 d of inspiral in the top panel, zoomed-in views of the first and last day of this interval in the middle two panels, and Lomb–Scargle periodograms during these two 1-d windows (with frequency measured in units of the instantaneous binary orbital frequency and the low-frequency spike omitted) in the bottom two panels. In the top ‘panoramic’ panel, we see two distinct periodicities: a longer one on the order of a day, and a shorter one on the order of hours to minutes, both of which shorten over time.

The longer period is produced by a well-known overdensity in the cavity wall, called a lump, which propagates a wave pattern along the cavity wall in the prograde direction. The lump forms from streams ejected from the binary, building an asymmetric gas distribution on the opposing cavity wall. The associated wave pattern propagating along the cavity wall modulates the distance between the binary and the nearest side of the cavity wall, thereby modulating the rate of CBD feeding to the binary. When CBD feeding to the binary is in a trough

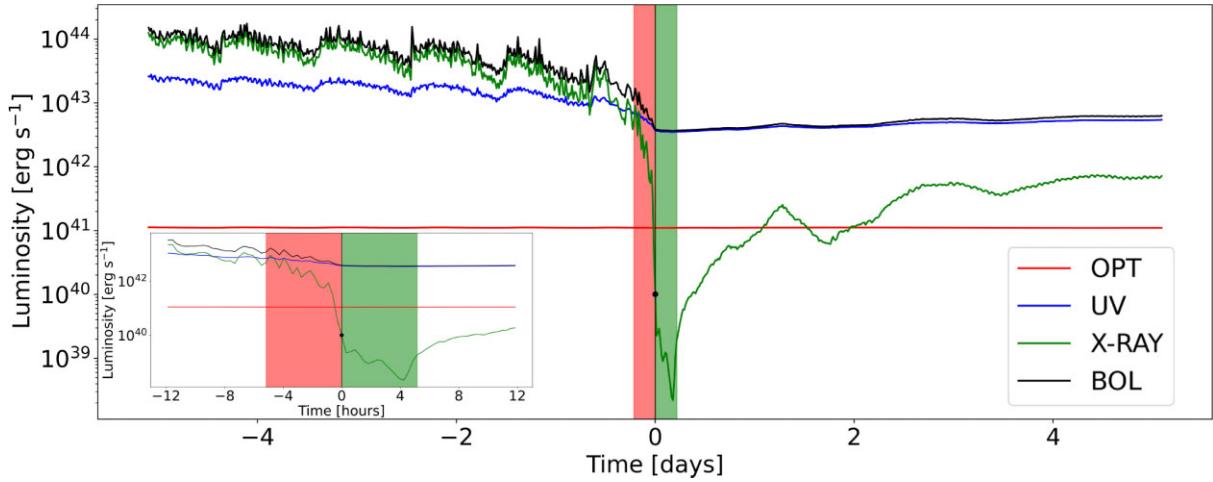


Figure 4. Optical, UV, X-ray, and bolometric light curves from 5 d before merger to 5 d after merger. The black vertical line is centred at the merger ($t = 0$), with a black dot marking the X-ray luminosity at the moment of the merger. The red/green zones indicate five hours before/after the merger. The inset shows a zoom-in version around the merger. This run includes a recoil kick (pointed downward in the bottom right panel of Fig. 1), contributing to the post-merger recovery of the luminosity, but does not include mass-loss, which inhibits recovery.

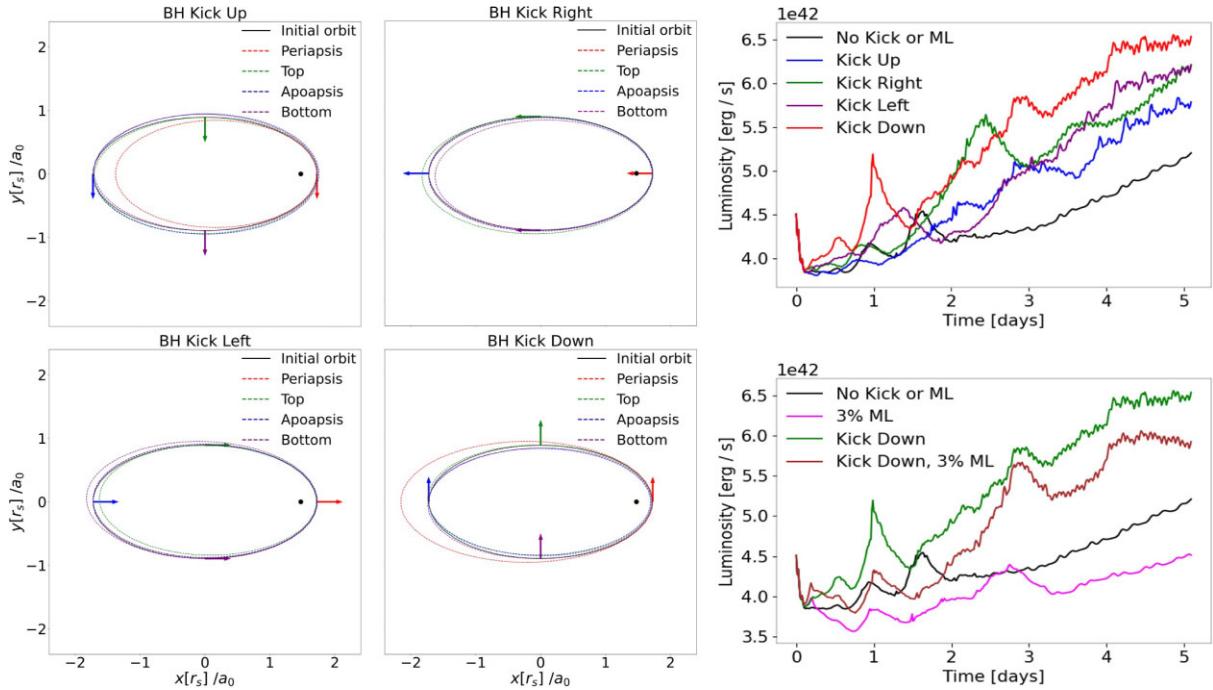


Figure 5. The left four panels correspond to four different kick directions for the BH at merger (Up, Right, Left, Down). In each panel, different colours show the perturbed post-merger orbital paths of test particles which suffered a recoil kick at four different positions along their initial orbit (at Periapsis, Top, Apoapsis, and Bottom). The top right panel shows the simulated post-merger bolometric light curves in our fiducial model with the same four different kick directions as well as no kick (or mass-loss) for comparison. This panel demonstrates that the highest (red) and lowest (blue) post-merger luminosities are produced by kicks in the Down and Up directions. As the left panel shows, these are the directions which re-direct test particles near their periapsis towards and away from the cavity wall, respectively. The bottom right panel shows the bolometric light curves for no kick or mass-loss, 3 per cent mass-loss, kick in the downwards (most luminous) direction, and both a kick in the downwards direction and a 3 per cent mass-loss. This panel demonstrates that mass loss tends to reduce the post-merger luminosity, both in the absence and the presence of a kick.

or a crest, the luminosity is minimal or maximal, respectively. This periodic process accounts for the low-frequency spike at $\omega/\omega_{\text{bin}} = 0.1$ in the initial 15-d periodogram.

The faster X-ray periodicities in Fig. 7 reflect mass trading activity between minidisks, which cause EM flares. Examining the middle left panel we see that the X-ray light curve has a

dominant frequency, which is evidently a near-orbital frequency $\omega/\omega_{\text{bin}} \approx 1$ in the corresponding periodogram. In the middle right panel, the dominant frequency is even faster, and becomes more rapid towards merger. The corresponding periodogram in the lower right panel implicates $\omega/\omega_{\text{bin}} = 2$, agreeing with the GW chirp frequency.

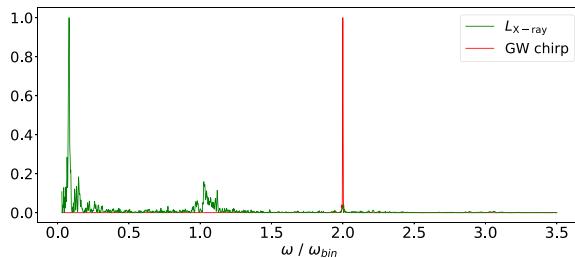


Figure 6. Lomb–Scargle periodograms, with maximum peaks normalized to unity, for the X-ray light curve (green) and GW chirp (red) during the 15 d before merger. We adopt a scaled frequency for the light curves, i.e. frequency is measured in units of the instantaneous binary orbital frequency.

Visual inspection of the gas evolution in the minidisks suggests that pulses of mass-trading between minidisks correspond to both the near-orbital and twice-per-orbit flares. At large separation, near-orbital cadence was observed with Γ -law gas models in recent work (Westernacher-Schneider et al. 2022, 2023), and was evident in accretion rates even earlier (Farris et al. 2015a). On the other hand, two flares per orbit in the late stages of the inspiral was reported previously with 2D simulations (e.g. Tang, Haiman & MacFadyen 2018), and a minidisk mass-exchange mechanism with twice-per-orbit cadence was found in 3D simulations of inspiral in relativistic potentials (Bowen et al. 2017). Given these past findings, it is natural to expect a transition between these two cadences during inspiral. This transition is also observed in some of our sensitivity tests, but we leave a more detailed investigation of this phenomenon to future work.

3.2 Parameter dependence

To assess the robustness of our results we ran simulations with different values of several parameters, including: increasing viscosity to $\alpha = 0.3$, increasing the total binary mass to $M_{\text{bin}} = 10^8 M_{\odot}$ and increasing the Mach number to $\mathcal{M} = 30$ (see Table 1). We also performed a resolution study, and tested if the disc remains geometrically thin during the full simulations, a necessary requirement for our vertically integrated 2D approach. We present these numerical validation studies in the Appendix.

Fig. 8 shows the fraction of the total luminosity emanating from the minidisks for each model. Fig. 9 shows the light curve in each band as labelled, and Fig. 10 shows the accretion rates (normalized by their respective accretion rates for Shakura–Sunyaev discs around a single BH).

In the high- α model, the minidisk luminosity fractions are comparable to the fiducial model. The presence of large intermittent downward spikes is caused by hot material leaving the minidisk luminosity zone, which is somewhat arbitrarily defined. Again, we see a several-order-of-magnitude drop in the accretion rate (Fig. 10) and the X-ray light curve (Fig. 9) at the time of the merger, as in our fiducial model. The higher viscosity refills the post-merger cavity more rapidly, leading to faster recovery of the luminosity compared to our fiducial model. The lump in this model is more pronounced, leading to a larger amplitude lump periodicity in the light curve. The lump period is also shorter, likely due a smaller cavity caused by shorter viscous times. Additionally, the drop in UV luminosity is delayed; in the fiducial model, the decline is significant around 0.7 d prior to merger, whereas the decline is delayed until ~ 0.3 d before merger in the $\alpha = 0.3$ model. Again, this is likely because the shorter viscous time allows the cavity wall to follow the binary

longer, keeping the minidisks more full until their disruption just before merger.

In the high BH-mass model, not only is the X-ray emission almost exclusively from the minidisks, as in the fiducial model, but the UV emission has increased to a much higher minidisk-fraction as well. This corresponds to a larger drop in the UV and bolometric luminosities near the time of merger. Again, we see a significant drop in the X-ray luminosity at merger, too, by a staggering 15 orders of magnitude. This is probably because there is more high-energy emission in a higher mass system, and it is biased towards the hottest regions.

In the higher Mach-number model, the minidisks starve sooner during the inspiral. This can likely be attributed to the increased cavity size, specifically on the near-side of the cavity wall, making it even easier for the shrinking binary to outrun the cavity wall (e.g. Ragusa, Lodato & Price 2016), and longer viscous time associated with the higher Mach number ($t_v \sim \mathcal{M}^2$ if α and binary parameters are held fixed, as we did). This is our only simulation that begins inspiral past the nominal decoupling epoch, inspiralling over only 0.2 viscous times compared to the other models in which the inspiral lasts 1.7 (fiducial and $M_{\text{bin}} = 10^8 M_{\odot}$ cases), and 5.1 ($\alpha = 0.3$ case) viscous times. This results in smaller minidisk fractions (Fig. 8) and causes the drop in the X-ray flux to begin already several days before merger (Fig. 9). However, the imprint of the longer lump period is more pronounced, which could help in the EM identification of these more massive binaries if a longer pre-merger temporal baseline is available. However, because the X-ray luminosity drops earlier, when LISA’s sky localization is still expected to be too poor to trigger EM monitoring (Mangiagli et al. 2020), one may need to rely on archival data to look for this flux decrease over several days prior to merger. Additionally, this is the only model in which the most luminous kick direction was less clear-cut; we show the ‘BH Kick Left’ case, which may be the most luminous in the post-merger phase, but each kick direction had a similar luminosity at the end of our 5-d post-merger window. The UV and bolometric luminosities are indistinguishable, since the UV band dominates, and they are remarkably steady across the merger. In fact, the small bump ~ 0.75 d after the merger only exists because of the kick imparted to the BH, otherwise, we find that these light curves remain flat. Together with our results from other models, this suggests that if an order of magnitude drop is observed in the UV near merger, then a much larger X-ray drop should accompany it, but if the UV is constant near merger, then the X-ray may already be below detectable levels.

Although it is too expensive with our current computational methods to evolve each of these models through the full recovery of the post-merger luminosity, in Fig. 10 we show the mass accretion rates over the same time windows as in Fig. 9, but with time measured in units of the initial orbital time. The accretion rates are shown for no kick nor mass-loss, as in Fig. 2. In all but the Mach 30 case, we see a several-order-of-magnitude drop in the accretion rate at the merger, whereas in the Mach 30 model it decreases sooner, undulating downwards until the moment of merger, similar to the luminosity evolution in Fig. 9.

Finally, in Fig. 11, we show each model’s X-ray light curve, individually normalizing the curves to their average value over one period about halfway through inspiral. As the binary shrinks, the tidal truncation radii of the minidisks are reduced, scaling linearly with the orbital separation. As a lowest order picture, one may expect each minidisk’s surface area to follow the binary separation r as $\propto r^2$. We show a curve representing this scaling, denoted in the legend as ‘ $\propto r^2$ ’. We see all but the Mach 30 light curves follow this $\propto r^2$ line, indicating that the rapid loss in luminosity can be attributed

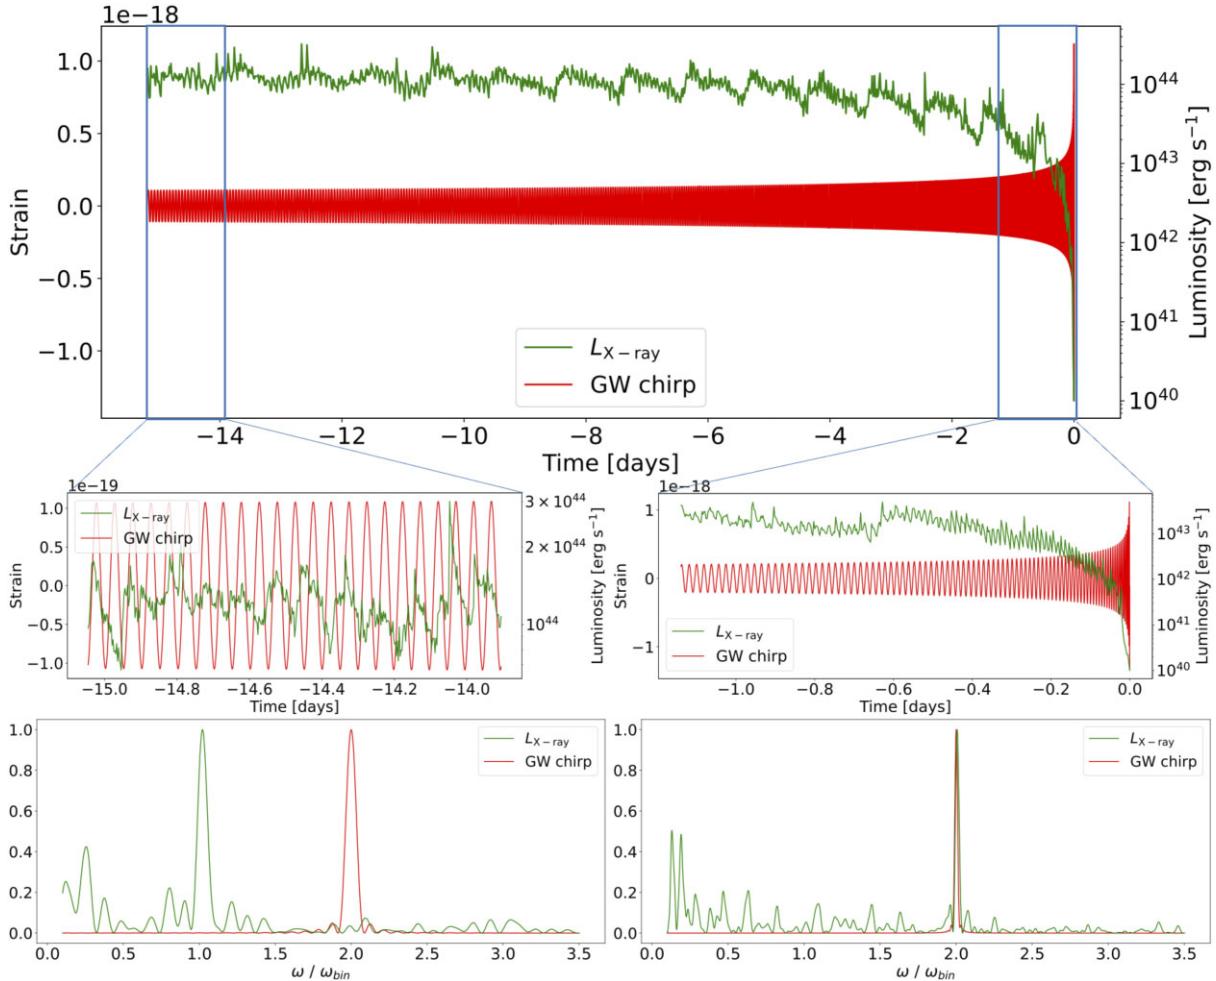


Figure 7. X-ray luminosity (green) and GW strain (red) 15 to 0 days (top), 15 to 14 d (middle left) and 1 to 0 days (middle right) before merger. Lomb–Scargle periodograms are shown on the lower left and right panels, with maximum peaks normalized to unity, for the X-ray light curve (green) and the GW chirp (red) during the same two time windows as the row above. Frequency is displayed in units of the instantaneous binary orbital frequency.

to the decreased surface area of the minidisks. One expects that an enhanced drop below the r^2 line will occur to the extent that the rate of CBD feeding is not sufficient to keep the minidisks saturated with material. This picture is consistent with the fact that the high-viscosity run ($\alpha = 0.3$) stays above the r^2 line the most, and the fact that the high-Mach run ($\mathcal{M} = 30$) drops below the r^2 line the most (the accretion rate is lowest in this run).

4 OBSERVATIONAL PROSPECTS TO IDENTIFY EM COUNTERPARTS OF LISA SOURCES

Identifying the host galaxy of merging MBHBs with networks of GW interferometers is difficult since they yield large sky localization uncertainties. In the case of LISA, weeks before merger, the approximate position of the MBHB can be determined from the relative amplitudes and phases of the two polarization components, from the periodic Doppler shift in the signal originating from the detector's heliocentric motion, and from the additional modulation in the signal originating from the detector's time-varying orientation (Cutler 1998). However, it is not until near merger that the sky position can be significantly narrowed via additional effects – namely spin-precession, LISA's pattern response (which becomes frequency-dependent and assists with localization; Rubbo, Cornish & Poujade

2004; Marsat, Baker & Canton 2021), and higher order harmonics (beyond the quadrupole; Baibhav, Berti & Cardoso 2020; Marsat, Baker & Canton 2021).

Mangiagli et al. (2020) have recently studied sky localization accuracy as a function of time prior to merger, using the most up-to-date LISA configuration and sensitivity curve, for a suite of MBHBs with varying total masses, mass ratios, spins, sky-positions and inclination angles, as well as different detector orientations. They find, in all cases, that the median error of sky-localization decreases with time, but the dispersion around it increases. In a case near our fiducial model, they find that even as little as 10 h before the merger, the typical uncertainties in the sky localization are still near 10 deg^2 , comparable to the field of view of large-scale optical/infrared survey telescopes, such as the Vera Rubin Observatory.

Lops et al. (2023) have analysed the implication of these findings for a triggered campaign to find EM counterparts before merger. They considered multiple apparatus parameter configurations for both the Advanced Telescope for High-energy Astrophysics (Athena; Nandra et al. 2013) and for the proposed Lynx mission (The Lynx Team 2018), and examined binaries with varying total masses and redshifts. They assessed which binaries would be observable and identifiable 10 h before, 1 h before, and at merger via X-ray detection. For a binary near the mass of our fiducial model, at $z < 1$, Lynx, with a

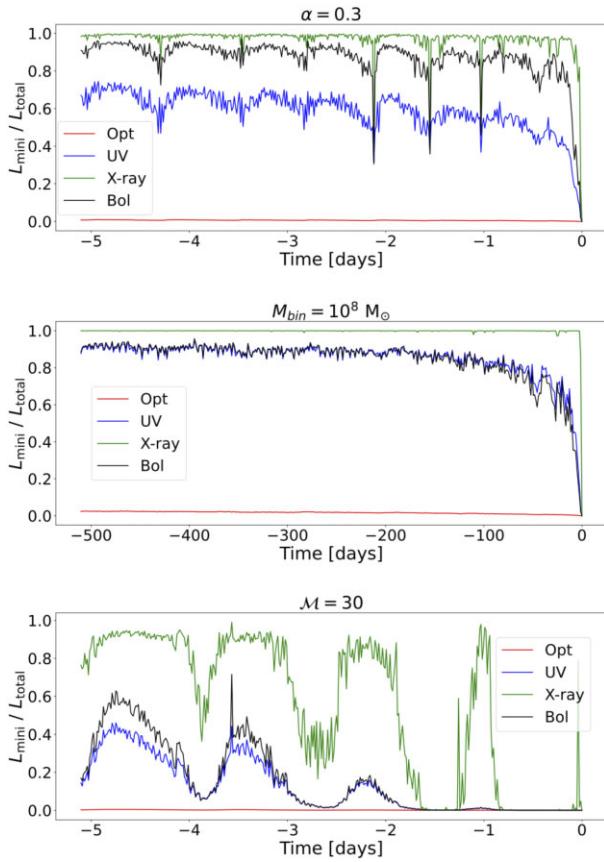


Figure 8. The fraction of each EM band originating from the minidisks for the models with a higher viscosity ($\alpha = 0.3$, top), higher BH mass ($M_{\text{bin}} = 10^8 M_{\odot}$, middle) and higher Mach number ($\mathcal{M} = 30$, bottom).

FOV of 0.1 deg^2 and a spatial angular resolution of 0.5 arcsec , can search the entire $\sim 10 \text{ deg}^2$ LISA error-box found by Mangiagli et al. (2020) with ~ 80 pointings, within 10 h prior to merger. Although Athena actually has a larger FOV of 0.4 deg^2 , it has a poorer angular resolution of $5\text{--}10 \text{ arcsec}$ and lower sensitivity, leading to higher exposure times required to detect the binaries. For our fiducial model of a $10^6 M_{\odot}$ binary, assuming it shines near Eddington luminosity with $10\text{--}40$ per cent of its bolometric flux in the X-ray band, within 10 h prior to merger Athena can tile the LISA error-box with ~ 5 pointings out to $z = 0.5$ for hard X-rays in the $2\text{--}10 \text{ keV}$ band, or with ~ 20 pointings out to $z = 1$ for soft X-rays in the $0.5\text{--}2 \text{ keV}$ band. When looking 10 h before the merger, our results suggest emission may peak near 2 keV (see Fig. B1) and decrease towards merger.

Dal Canton et al. (2019) also explored tiling the LISA error box in the days leading to merger, with the proposed wide-field X-ray telescope Transient Astrophysics Probe (TAP), which has a FOV of 1 deg^2 and angular resolution of 5 arcsec . While the focus of this paper was centred on detectability via the quasiperiodic Doppler modulation due to the orbital motion, their findings suggest a similar range of detectability, out to $z = 1$. The probability of detection for our fiducial model decreases significantly beyond this redshift. More encouragingly, however, they find that a fraction of their simulated sources are not detected because of the number of photons collected is insufficient to measure the quasiperiodic modulations. As our method does not rely on this modulation, it should allow the detection of fainter sources to higher redshifts.

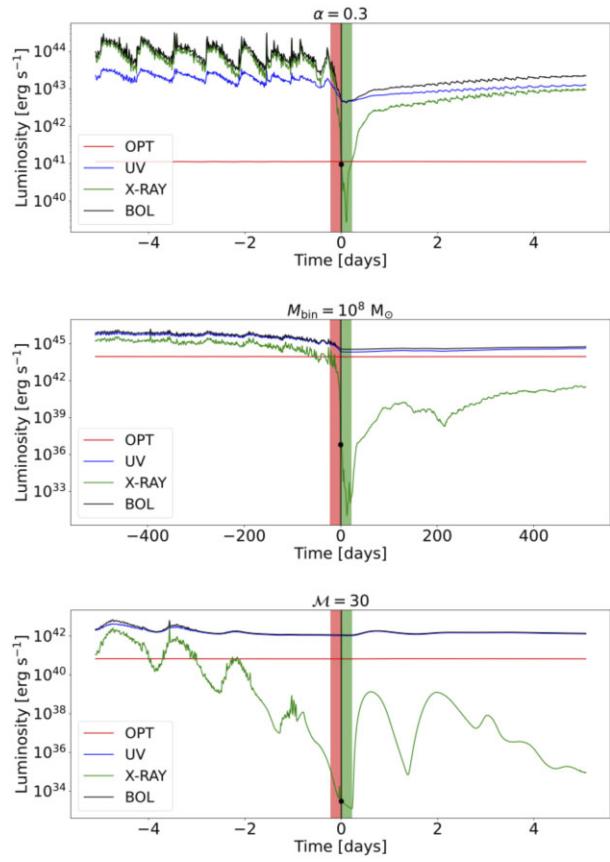


Figure 9. Luminosity in various bands for the models with a higher viscosity ($\alpha = 0.3$, top), higher BH mass ($M_{\text{bin}} = 10^8 M_{\odot}$, middle) and higher Mach number ($\mathcal{M} = 30$, bottom). Top and bottom panels show 5 d before and after the merger, and the red/green zones indicate 5 h before/after the merger. Due to the different time scales for the high-mass model, the middle panel shows 500 d before and after the merger, and the red/green zones indicate 500 h before/after the merger. In all panels, a black vertical line is centred at $t = 0$, with a black dot indicating the X-ray luminosity at the moment of merger.

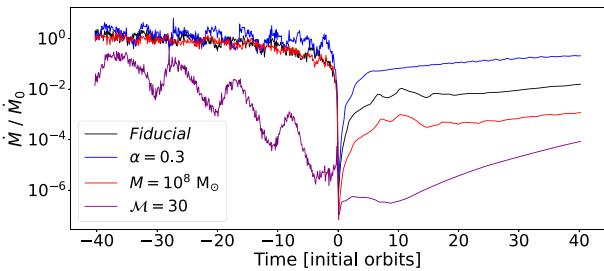


Figure 10. Accretion rates as a function of time (with time measured in units of the initial orbital time) for each model with no recoil or mass-loss, normalized by their respective accretion rates for Shakura–Sunyaev discs around a single BH.

Previous proposals for identifying the host galaxy before merger require monitoring the pre-merger EM source for a prolonged time, in order to identify periodicity in the correct EM counterpart. By contrast, in principle, looking for the drop in thermal X-ray luminosity requires just two data points – one just before or at merger, and one a few hours earlier – because the counterpart will undergo a spectral change where the thermal X-rays disappear but

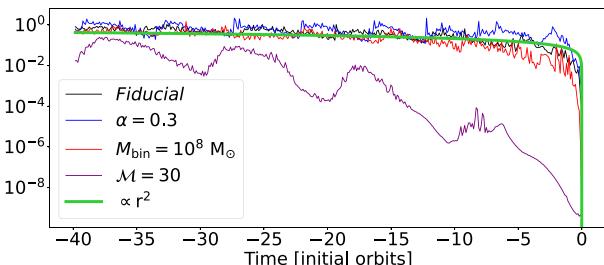


Figure 11. Normalized X-ray light curves in each model, and a (green) line proportional to the square of the binary separation r^2 , representing the expected evolution of the surface area of tidally truncated minidisks as a function of time in initial orbits.

stays relatively bright in the optical/UV. There is no reason for any other source among the many candidates in the LISA error volume to display this behaviour. This can significantly improve the chance of identifying the EM counterpart before merger occurs.

5 CONCLUSIONS

We performed 2D hydrodynamical simulations of inspiralling, accreting binaries, from the pre-decoupling epoch through the post-merger GW kick and mass-loss, and examined the corresponding EM signatures. Compared to previous works, our study includes a longer inspiral period, a more physical viscosity, and a more realistic treatment of thermodynamics, directly solving the energy equation with a Γ -law equation of state for the gas, and with a physically motivated cooling prescription. Additionally, compared to previous works (e.g. Farris et al. 2015b; Tang, Haiman & MacFadyen 2018), we use a much higher resolution which is important for understanding the late-stage minidisk destruction. From a suite of runs, we can draw the following conclusions:

(i) In all models prior to merger, the minidisks surrounding the BHs emit essentially no luminosity in the optical band, but account for a sizeable fraction of the UV band, and nearly all of the X-ray band. Most of the models in our suite, including the fiducial $M_{\text{bin}} = 10^6 M_{\odot}$ binary, evolve across the so-called nominal decoupling time, i.e. when the GW inspiral outpaces the viscous time-scale. However, we see the minidisks persist well past this nominal decoupling, often nearly until merger.

(ii) On the other hand, the minidisks are eventually disrupted, which leads to a significant drop in both the accretion rates and the X-ray luminosity (and a smaller drop in the UV luminosity). In all models, this happens very near the time of merger, except in the Mach 30 case. Only in the Mach 30 model did we see any appreciable drop beginning earlier, with the accretion and thermal X-ray luminosity gradually decreasing until the merger, undulating along the way.

(iii) Although BH recoil (or mass-loss) can cause a non-negligible increase in post-merger luminosity through additional shocks, the plummet of the thermal X-ray luminosity still does not recover in the next several days post-merger. Mass-loss tends to exacerbate the decrease in post-merger luminosity, extending the time needed to recover.

(iv) Sufficiently close to merger (within ≈ 10 h) and for certain binary masses and redshifts, Athena and Lynx could perform a full search of the LISA error box. Instead of requiring extensive pre-merger monitoring to extract a periodicity, our results suggest that as few as two data points are needed to identify the

source via its disappearing thermal X-ray emission just before merger.

(v) Accretion rates and light curves exhibit two periodicities, a multi-orbit one originating from the lump-modulation of accretion on to the minidisks, and another on shorter time-scales, arising from mass trading between minidisks. Well before merger this latter period is close to the binary's orbital time, while at late stages, it is reduced to half of the orbital period.

Overall, our most important finding is that the minidisks dominate the X-ray emission. The minidisks last until moments before merger, but are ultimately disrupted. The corresponding sudden drop in the X-ray flux can provide a tell-tale EM signature, which can aid in the identification of the unique counterpart of the LISA GW source. Importantly, unlike attempts to identify pre-merger EM periodicity through extended monitoring, identifying an X-ray dropout is feasible with as few as two data points.

Our hydrodynamical models omit some physical ingredients that could alter the results we presented in this work. The most important ones that merit further investigation are three dimensional effects, magnetic fields, general relativity (GR), radiative feedback, and emission and feedback from jets and/or coronae. Both 3D (e.g. Moody, Shi & Stone 2019) and general-relativistic magnetohydrodynamic (e.g. Combi et al. 2021; Noble et al. 2021; Avara et al. 2023) simulations indicate an increase in the variability of the light curve, although the fluctuation amplitudes (at most an order of magnitude) remain much smaller than the magnitude of the X-ray drop we predict here. GR would introduce several dynamical effects. A shallower post-Newtonian potential may also cause added variability and introduce an innermost stable circular orbit (ISCO), which can change the timing and magnitude of the X-ray drop in a way that also depends on the BH spins (e.g. Paschalidis et al. 2021; Combi et al. 2022). Ray-tracing, because of lensing near the BHs, results in weighting the portions of the minidisk that are closer to the BHs more heavily (e.g. Davelaar & Haiman 2022a), again increasing variability, but unlikely to be able to hide a several-order-of-magnitude drop in flux. Some 3D GRMHD findings (e.g. Gutiérrez et al. 2022) suggest that the thermal emission of the minidisks may peak at a lower frequency than we find here, perhaps moving the large drop to the UV band. Radiative feedback (e.g. del Valle & Volonteri 2018) could dismantle the minidisks sooner, again hastening when the drop occurs and obscuring its coincidence with the merger. Additionally, as the minidisks are disrupted, the BH magnetosphere generates a split-monopole that will reconnect any remaining magnetic flux (e.g. Bransgrove, Ripperda & Philippov 2021). Depending on the reconnection time-scales, this could generate a non-thermal X-ray component (e.g. Hakobyan, Ripperda & Philippov 2023), that could obscure the drop we observe in thermal X-ray emission. Interactions between misaligned jets may also lead to non-thermal emission at photon energies near (though perhaps not identical to) that of the minidisks (e.g. Gutiérrez et al. 2023). It is unclear if either dissipating jets or coronal contributions could obscure the drop in flux. More follow-up work will be needed to address these caveats.

As we were finalizing this manuscript for submission, we became aware of a closely related preprint posted on the arXiv (Dittmann, Ryan & Miller 2023), addressing the gas dynamics of inspiralling binaries, similar to the present study. There are several notable differences in the methodology and emphasis. In particular, Dittmann, Ryan & Miller (2023) use a different grid code (ATHENA++), and perform a more extensive parameter study, including the dependence on viscosity, and show that decoupling can occur inside the LISA band for sufficiently high viscosities. On the other hand, in addition to

examining mass accretion rates, our study also includes predictions for the thermal emission, and includes the post-merger evolution. Overall, the two studies find a consistent main result, namely that for comparable parameter choices (BH masses and viscosities), the accretion rate drops by several orders of magnitude just prior to merger, providing a unique new EM signature of LISA mergers.

ACKNOWLEDGEMENTS

We acknowledge support by National Science Foundation grant AST-2006176 (ZH) and National Aeronautics and Space Administration grant 80NSSC22K082 (AM and ZH). JD acknowledges support by a Joint Columbia/Flatiron Postdoctoral Fellowship. Research at the Flatiron Institute is supported by the Simons Foundation. This research was supported in part by the National Science Foundation under grant no. NSF PHY-1748958. This research has made use of National Aeronautics and Space Administration's Astrophysics Data System. *Software: PYTHON* (Oliphant 2007; Millman & Aivazis 2011), *SCIPY* (Jones et al. 2001), *NUMPY* (van der Walt, Colbert & Varoquaux 2011), and *MATPLOTLIB* (Hunter 2007).

DATA AVAILABILITY

The data underlying this article will be shared on reasonable request to the corresponding author.

REFERENCES

Acernese F. et al., 2014, *Class. Quant. Grav.*, 32, 024001
 Akutsu T. et al., 2021, *Prog. Theor. Exp. Phys.*, 2021, 05A101
 Amaro-Seoane P. et al., 2017, Laser Interferometer Space Antenna. preprint (arXiv:1702.00786)
 Artynowicz P., Lubow S. H., 1996, *ApJ*, 467, L77
 Avara M. J., Krolik J. H., Campanelli M., Noble S. C., Bowen D., Ryu T., 2023, preprint (arXiv:2305.18538)
 Baibhav V., Berti E., Cardoso V., 2020, *Phys. Rev. D*, 101, 084053
 Baker J. G., Boggs W. D., Centrella J., Kelly B. J., McWilliams S. T., Miller M. C., van Meter J. R., 2008, *ApJ*, 682, L29
 Baker J. et al., 2019, preprint (arXiv:1903.04417)
 Begelman M. C., Blandford R. D., Rees M. J., 1980, *Nature*, 287, 307
 Bode T., Haas R., Bogdanović T., Laguna P., Shoemaker D., 2010, *ApJ*, 715, 1117
 Bogdanović T., Miller M. C., Blecha L., 2022, *Living Rev. Relativ.*, 25, 3
 Bowen D. B., Campanelli M., Krolik J. H., Mewes V., Noble S. C., 2017, *ApJ*, 838, 42
 Bransgrove A., Ripperda B., Philippov A., 2021, *Phys. Rev. Lett.*, 127, 055101
 Collaboration L. S., 2015, *Class. Quant. Grav.*, 32, 074001
 Combi L., Armengol F. G. L., Campanelli M., Ireland B., Noble S. C., Nakano H., Bowen D. i., 2021, *Phys. Rev. D*, 104, 044041
 Combi L., Armengol F. G. L., Campanelli M., Noble S. C., Avara M., Krolik J. H., Bowen D. i., 2022, *ApJ*, 928, 187
 Corrales L. R., Haiman Z., MacFadyen A., 2010, *MNRAS*, 404, 947
 Cutler C., 1998, *Phys. Rev. D*, 57, 7089
 D’Orazio D. J., Di Stefano R., 2018, *MNRAS*, 474, 2975
 D’Orazio D. J., Haiman Z., MacFadyen A., 2013, *MNRAS*, 436, 2997
 D’Orazio D. J., Haiman Z., Schiminovich D., 2015, *Nature*, 525, 351
 Dal Canton T., Mangiagli A., Noble S. C., Schnittman J., Ptak A., Klein A., Sesana A., Camp J., 2019, *ApJ*, 886, 146
 Davelaar J., Haiman Z., 2022a, *Phys. Rev. D*, 105, 103010
 Davelaar J., Haiman Z., 2022b, *Phys. Rev. Lett.*, 128, 191101
 de Rham C., Melville S., Tolley A. J., 2018, *J. High Energy Phys.*, 2018, 83
 del Valle L., Volonteri M., 2018, *MNRAS*, 480, 439
 Dempsey A. M., Muñoz D., Lithwick Y., 2020, *ApJ*, 892, L29
 Dittmann A. J., Ryan G., 2021, *ApJ*, 921, 71
 Dittmann A. J., Ryan G., Miller M. C., 2023, *ApJ*, 949, L30
 Duffell P. C., D’Orazio D., Derdzinski A., Haiman Z., MacFadyen A., Rosen A. L., Zrake J., 2020, *ApJ*, 901, 25
 Farris B. D., Duffell P., MacFadyen A. I., Haiman Z., 2014, *ApJ*, 783, 134
 Farris B. D., Duffell P., MacFadyen A. I., Haiman Z., 2015a, *MNRAS*, 446, L36
 Farris B. D., Duffell P., MacFadyen A. I., Haiman Z., 2015b, *MNRAS*, 447, L80
 Frank J., King A., Raine D. J., 2002, *Accretion Power in Astrophysics*, 3rd edn. Cambridge Univ. Press, Cambridge
 Goodman J., 2003, *MNRAS*, 339, 937
 Gutiérrez E. M., Combi L., Noble S. C., Campanelli M., Krolik J. H., Armengol F. L., García F., 2022, *ApJ*, 928, 137
 Gutiérrez E. M., Combi L., Romero G. E., Campanelli M., 2023, preprint (arXiv:2301.04280)
 Haiman Z., 2017, *Phys. Rev. D*, 96, 023004
 Hakobyan H., Ripperda B., Philippov A. A., 2023, *ApJ*, 943, L29
 Hassan S. F., Rosen R. A., 2012, *Phys. Rev. Lett.*, 108, 041101
 Hayasaki K., Mineshige S., Sudou H., 2007, *PASJ*, 59, 427
 Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90
 Jones E., Oliphant T., Peterson P. et al., 2001, SciPy: Open source scientific tools for Python. <http://www.scipy.org/>
 Kocsis B., Haiman Z., Menou K., 2008, *ApJ*, 684, 870
 Lightman A. P., Eardley D. M., 1974, *ApJ*, 187, L1
 Lippai Z., Frei Z., Haiman Z., 2008, *ApJ*, 676, L5
 Liu F. K., Wu X.-B., Cao S. L., 2003, *MNRAS*, 340, 411
 Lops G., Izquierdo-Villalba D., Colpi M., Bonoli S., Sesana A., Mangiagli A., 2023, *MNRAS*, 519, 5962
 Lousto C. O., Zlochower Y., 2013, *Phys. Rev. D*, 87, 084027
 MacFadyen A. I., Milosavljević M., 2008, *ApJ*, 672, 83
 Mangiagli A. et al., 2020, *Phys. Rev. D*, 102, 084056
 Marsat S., Baker J. G., Canton T. D., 2021, *Phys. Rev. D*, 103, 083011
 Megevand M., Anderson M., Frank J., Hirschmann E. W., Lehner L., Liebling S. L., Motl P. M., Neilsen D., 2009, *Phys. Rev. D*, 80, 024012
 Mignon-Risse R., Varniere P., Casse F., 2023, *MNRAS*, 519, 2848
 Millman K. J., Aivazis M., 2011, *Comput. Sci. Eng.*, 13, 9
 Milosavljević M., Phinney E. S., 2005, *ApJ*, 622, L93
 Moody M. S. L., Shi J.-M., Stone J. M., 2019, *ApJ*, 875, 66
 Nakamura T. K., 2009, *EPL (Europhysics Letters)*, 88, 20004
 Nandra K. et al., 2013, preprint (arXiv:1306.2307)
 Noble S. C., Mundim B. C., Nakano H., Krolik J. H., Campanelli M., Zlochower Y., Yunes N., 2012, *ApJ*, 755, 51
 Noble S. C., Krolik J. H., Campanelli M., Zlochower Y., Mundim B. C., Nakano H., Zilhão M., 2021, *ApJ*, 922, 175
 O’Neill S. M., Miller M. C., Bogdanović T., Reynolds C. S., Schnittman J. D., 2009, *ApJ*, 700, 859
 Oliphant T. E., 2007, *Comput. Sci. Eng.*, 9, 10
 Paczynski B., 1977, *ApJ*, 216, 822
 Paczynsky B., Wiita P. J., 1980, *A&A*, 88, 23
 Paschalidis V., Bright J., Ruiz M., Gold R., 2021, *ApJ*, 910, L26
 Penoyre Z., Haiman Z., 2018, *MNRAS*, 473, 498
 Peters P. C., 1964, *Phys. Rev.*, 136, B1224
 Pringle J. E., 1991, *MNRAS*, 248, 754
 Ragusa E., Lodato G., Price D. J., 2016, *MNRAS*, 460, 1243
 Roedig C., Dotti M., Sesana A., Cuadra J., Colpi M., 2011, *MNRAS*, 415, 3033
 Rosotti G. P., Lodato G., Price D. J., 2012, *MNRAS*, 425, 1958
 Rossi E. M., Lodato G., Armitage P. J., Pringle J. E., King A. R., 2010, *MNRAS*, 401, 2021

Rubbo L. J., Cornish N. J., Poujade O., 2004, *Phys. Rev. D*, 69, 082003
 Schnittman J. D., Krolik J. H., 2008, *ApJ*, 684, 835
 Schutz B. F., 1986, *Nature*, 323, 310
 Shakura N. I., Sunyaev R. A., 1973, *A&A*, 24, 337
 Shakura N. I., Sunyaev R. A., 1976, *MNRAS*, 175, 613
 Shields G. A., Bonning E. W., 2008, *ApJ*, 682, 758
 Shi J.-M., Krolik J. H., 2015, *ApJ*, 807, 131
 Shi J.-M., Krolik J. H., Lubow S. H., Hawley J. F., 2012, *ApJ*, 749, 118
 Tang Y., Haiman Z., MacFadyen A., 2018, *MNRAS*, 476, 2249
 The Lynx Team, 2018, preprint (arXiv.1809.09642)
 Tichy W., Marronetti P., 2008, *Phys. Rev. D*, 78, 081501
 van der Walt S., Colbert S. C., Varoquaux G., 2011, *Comput. Sci. Eng.*, 13, 22
 Vera C. Rubin Observatory LSST Solar System Science Collaboration et al., 2021, *Bulletin of the American Astronomical Society*, 53, 236
 Westernacher-Schneider J. R., Zrake J., MacFadyen A., Haiman Z., 2022, *Phys. Rev. D*, 106, 103010
 Westernacher-Schneider J. R., Zrake J., MacFadyen A., Haiman Z., 2023, preprint (arXiv:2307.01154)
 Zanotti O., Rezzolla L., Del Zanna L., Palenzuela C., 2010, *A&A*, 523, A8

APPENDIX A: NUMERICAL VALIDATION

To test our numerical convergence, we perform three variants of our fiducial run (4000^2 grid cells): two models with lower resolution (1000^2 and 2000^2 grid cells), and one model with higher resolution (8000^2 grid cells). Fig. A1 shows X-ray light curves, where 4k (black) indicates the fiducial model. The fiducial model appears robust near the window of merger we are most interested in. Note that verifying a converged result in this figure is complicated by the fact that the lump phase at merger is not equivalent across each run. Thus, the deviation between the 4k and 8k runs beginning 2 d after merger does not imply a lack of convergence at those times. None the less, low resolution clearly overpredicts the post-merger luminosity significantly.

Since our 2D column-integrated approach requires that the discs are geometrically thin, we checked the local aspect ratio $h/r \sim 1/\mathcal{M}$, shown in Fig. A2. The overall h/r is consistent with our intended set-up, and the local disc aspect ratio remains valid. In the shock-heated regions, we see a rise in h/r , consistent with previous findings (e.g. Tang, Haiman & MacFadyen 2018); nevertheless, the disc remains thin, e.g. with at most $h/r \approx 0.13$. We also checked the $\mathcal{M} = 30$ simulation, shown in Fig. A3. We again see that the shock-heated regions show a rise in h/r , but the overall h/r stays consistent with the thin disc assumption.

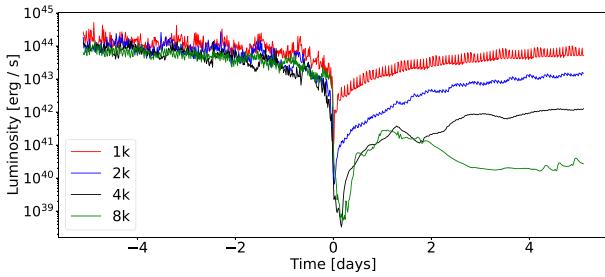


Figure A1. Resolution test of the X-ray luminosity for our fiducial model. The post-merger luminosity results from the downward BH with no mass-loss.

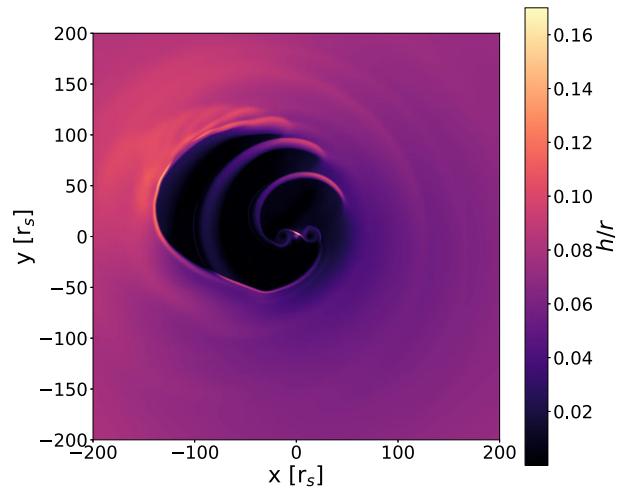


Figure A2. Map of the local disc thickness h/r in the fiducial model, just before inspiral is initialized.

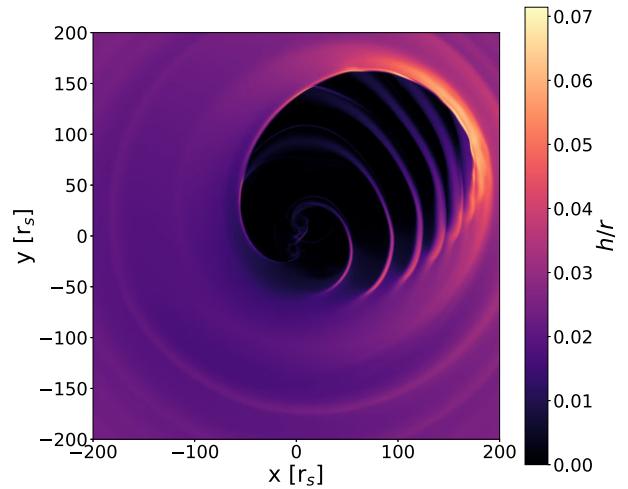


Figure A3. Map of the local disc thickness h/r for the $\mathcal{M} = 30$ run, just before inspiral is initialized.

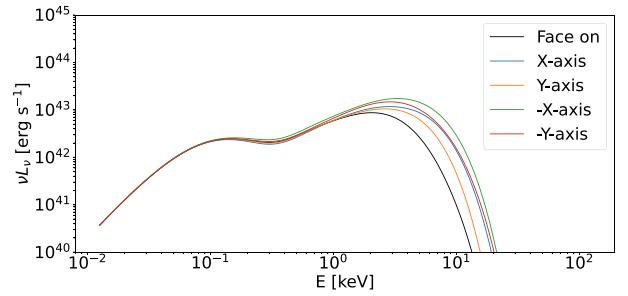


Figure B1. Spectra ~ 10 h before merger from four different viewing angles – the X , Y , $-X$, and $-Y$ axes –, incorporating relativistic Doppler effects. Independent of viewing angle, we see a brightening at higher frequencies, causing an overall hardening of the spectrum.

APPENDIX B: DOPPLER EFFECTS

Since this study was performed face-on, relativistic Doppler boosting was omitted, but given that $v \ll c$, this should be a minor effect. Here,

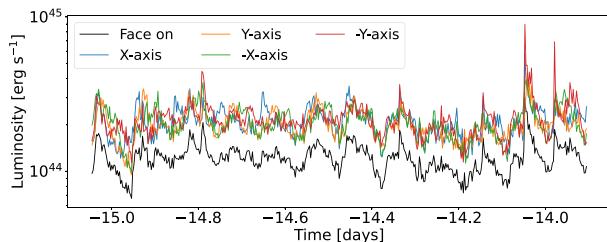


Figure B2. X-ray luminosity 15 to 14 d before merger, for viewing angles of face-on and along the X , Y , $-X$, and $-Y$ axes.

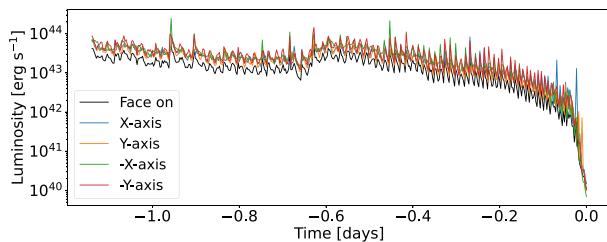


Figure B3. X-ray luminosity 1 to 0 d before merger, for viewing angles of face-on and along X , Y , $-X$, and $-Y$ axes.

we consider results for inclined systems, for which Doppler effects are not negligible.

First, in Fig. B1, we examine the spectra ~ 10 h before merger for four different azimuthal viewing angles lying in the equatorial plane (an observer along the X , Y , $-X$, and $-Y$ axes) to ascertain the upper

limit of these effects. The spectra are computed in the same way as before, except that we include a relativistic Doppler boost from each patch of the disc. Practically, this is performed by modifying the effective temperature, following equation 3 of Nakamura (2009). The figure shows that lower frequencies are unaffected, but the higher end of the spectrum tends to shift upwards in energy, independent of viewing angle. This is a result of the curvature in the Wien tail of the black-body spectrum, producing a larger brightening of the blue-shifted parts of the disc than the dimming of the patches red-shifted by a similar velocity (see also Tang, Haiman & MacFadyen 2018).

Next, in Figs B2 and B3, we examine the X-ray light curves from the same viewing angles, during the same time windows as in middle panels of Fig. 7. Corroborating what we see with the spectra, we see a boost to the overall luminosity in the X-ray band, independent of viewing angle. While there may be some increased variability, particularly at earlier times, we see a comparable signal and periodicity, and most importantly, we still see a significant drop in the X-ray luminosity at the time a merger.

Generically, the increased pre-merger luminosity due to the Doppler boost will help decrease the required exposure time for an X-ray triggered campaign, thereby increasing the mass and redshift at which the entire LISA error-box can be tiled, as discussed in Section 4. Specifically, in the case of Athena, however, because of a decreased sensitivity to harder X-rays, whether this would result in an overall increase or decrease to the exposure time will be system dependent.

This paper has been typeset from a TeX/LaTeX file prepared by the author.