



The late inspiral of supermassive black hole binaries with circumbinary gas discs in the LISA band

Yike Tang,¹★ Zoltán Haiman²★ and Andrew MacFadyen¹★

¹Center for Cosmology and Particle Physics, Physics Department, New York University, New York, NY 10003, USA

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

Accepted 2018 February 14. Received 2018 February 13; in original form 2018 January 5

ABSTRACT

We present the results of 2D, moving-mesh, viscous hydrodynamical simulations of an accretion disc around a merging supermassive black hole binary (SMBHB). The simulation is pseudo-Newtonian, with the BHs modelled as point masses with a Paczynski–Wiita potential, and includes viscous heating, shock heating, and radiative cooling. We follow the gravitational inspiral of an equal-mass binary with a component mass $M_{\text{bh}} = 10^6 M_{\odot}$ from an initial separation of $60r_g$ (where $r_g \equiv GM_{\text{bh}}/c^2$ is the gravitational radius) to the merger. We find that a central, low-density cavity forms around the binary, as in previous work, but that the BHs capture gas from the circumbinary disc and accrete efficiently via their own minidisks, well after their inspiral outpaces the viscous evolution of the disc. The system remains luminous, displaying strong periodicity at twice the binary orbital frequency throughout the entire inspiral process, all the way to the merger. In the soft X-ray band, the thermal emission is dominated by the inner edge of the circumbinary disc with especially clear periodicity in the early inspiral. By comparison, harder X-ray emission is dominated by the minidisks, and the light curve is initially more noisy but develops a clear periodicity in the late inspiral stage. This variability pattern should help identify the electromagnetic counterparts of SMBHBs detected by the space-based gravitational-wave detector LISA.

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

1 INTRODUCTION

Supermassive black holes (SMBH) are currently believed to reside in most, if not all galactic nuclei, and SMBH binaries (SMBHBs) are expected to be produced frequently in mergers of galaxies (e.g. Kormendy & Ho 2013). The resulting compact SMBHBs are likely embedded in a gaseous environment. Electromagnetic (EM) signatures of such compact SMBHBs can arise from their interaction with gas.

The space-based gravitational-wave (GW) detector LISA will be sensitive to SMBHBs with masses in the range 10^4 – $10^7 M_{\odot}$ (Amaro-Seoane et al. 2017). Identifying a GW source in the EM bands has considerable importance for astrophysics and cosmology (e.g. Phinney 2009). It has been proposed that periodic emission in the EM bands may track the orbital motion of SMBHBs throughout the late inspiral, and that this would then allow a unique identification of a LISA source (Kocsis, Haiman & Menou 2008). Comparing the EM and GW chirp signals would then also help probe the difference in the propagation speed of photons and gravitons

(Kocsis et al. 2008; Haiman 2017), complementing measurements of the graviton speed from phasing of the GWs alone (see e.g. the review by Will 2006 and references therein). A similar constraint has recently been derived from the time delay between the GWs and the accompanying gamma-ray burst arriving from the neutron-star merger GW170817 (Abbott et al. 2017). An SMBHB offers an improved and more robust measurement of this time delay; this is especially the case if Doppler modulations of the EM chirp are detected from the orbital motion of the BHs themselves, fixing the relative phase of the EM and GW signals (Haiman 2017).

In order to predict the EM signatures of SMBHB mergers in the LISA band, the first step is to resolve the gas dynamics during this late inspiral stage. The morphology of the gas, even on scales as large as that of the circumbinary disc, is sensitive to the behaviour of gas streams in the vicinity of the individual BHs (Tang, MacFadyen & Haiman 2017; Bowen et al. 2017a). One necessary requirement for reliably simulating the LISA stage is to resolve the accretion flows near the individual BHs, on scales down to that of the innermost stable circular orbit (ISCO; located at $6GM_{\text{bh}}/c^2 \equiv 6r_g$ for a non-spinning BH). Furthermore, the gas dynamics cannot be in a steady state, since the binary orbit is evolving rapidly. Depending on the component masses, for a circular, comparable-mass binary, the LISA band corresponds to the inspiral from an initial separation of

* E-mail: yt611@nyu.edu (YT); zoltan@astro.columbia.edu (ZH); macfadyen@nyu.edu (AMF)

$(50-200)GM/c^2$ all the way to the merger, and covers the last ≈ 100 to a few thousand orbits of the binary. The orbital evolution of the binary is likely overwhelmingly dominated by its gravitational radiation at this late stage (Haiman, Kocsis & Menou 2009; Dotti, Sesana & Decarli 2012; Tang et al. 2017).

Because of the dual requirements of high spatial resolution and large number of orbits, simulating the LISA stage in its entirety has remained beyond the capability of fully three-dimensional (GR)MHD simulations. They have generally focused instead on the last ≈ 10 orbits preceding the merger (Bode et al. 2010; Farris et al. 2012; Giacomazzo et al. 2012; Noble et al. 2012; Shi et al. 2012; Bowen et al. 2017b). *However, following the binary throughout the entire LISA band is feasible in 2D hydrodynamics.* Previous high-resolution Newtonian simulations have indeed followed binaries for thousands of orbits. However, these works have focused on large separations, well before the LISA stage (Cuadra et al. 2009; Krolik 2010; Roedig et al. 2012; Farris et al. 2015a). As such, they did not directly resolve accretion inside the ISCO (typically either ignoring accretion, or implementing accretion via a sink prescription).

In our earlier work (Farris et al. 2015b), we performed a 2D simulation of an inspiraling and merging BH binary, embedded in a viscous circumbinary disc. However, this earlier study was purely Newtonian, and employed an artificial sink that depleted the gas according to the fluid viscous time-scale. Moreover, the focus was on the overall accretion rate, and we did not present an analysis of the gas dynamics and did not compute the emerging EM emission and light curves.

In this paper, we focus on modelling the entire inspiral stage observable with LISA, extending our earlier work. We consider an equal-mass, non-spinning BH binary with a component mass of $M_{\text{bh}} = 10^6 M_{\odot}$ and initial separation $60r_{\text{S}}$. This corresponds to the last ≈ 100 orbits of the merger, or the last $\approx (1+z)/2$ week observed for a binary at redshift z . The binary is assumed to be on a circular orbit, and we allow the orbit to shrink due to GW emission all the way to merger. The binary is embedded in a viscous, optical thick α -disc with radiative cooling. We employ a pseudo-Newtonian potential for the black holes. Our spatial resolution is sufficient to resolve the individual BH's ISCO, allowing us to remove gas inside the event horizon physically. We analyse the dynamics and the time-dependent spectral signatures of the emerging thermal emission.

This paper is organized as follows. In Section 2, we summarize our numerical methods. In Section 3, we present our main findings, including disc spectra and light curves. We also examine the periodicities of the light curves separately in the soft and hard X-ray bands. In Section 4, we discuss the implications of our results, along with some caveats. Finally, in Section 5, we summarize our conclusions and indicate topics for future work.

2 NUMERICAL SET-UP

The model and numerical set-up of this work closely resembles that in our previous work (Farris et al. 2015b). We briefly outline the key aspects here and highlight the modifications we made for this study. We refer the reader to the above paper for further details.

All simulations are performed using the publicly available moving-mesh grid code **DISCO** in 2D (Duffell & MacFadyen 2012; Duffell 2016). We simulate an equal-mass SMBHB with $M_{\text{bin}} = 2M_{\text{bh}} = 2 \times 10^6 M_{\odot}$ and initial separation $a_0 = 60GM_{\text{bh}}/c^2$. At this separation, a typical $10^6 M_{\odot}$ binary at $z < 1$ can be localized by LISA on the sky to a few square degrees, giving the chance to a wide-field telescope to locate and identify the source (Kocsis et al.

2008; Lang & Hughes 2008). We assume the BHs are non-spinning so that $a_0 = 10r_{\text{isco}} = 30r_{\text{S}}$, where r_{S} is the Schwarzschild radius of an individual BH. The BHs are placed on the grid, and the vertically integrated fluid equations are solved assuming an α -viscosity with $\alpha = 0.1$.

Inside the cavity that forms around the binary (with a size roughly twice the initial binary separation), the grid spacing is roughly linear in radius and r_{isco} is resolved with ≈ 8 radial cells; there are ≈ 100 cells in total inside the ISCO. Outside the cavity, the radial grid spacing increases approximately logarithmically up to $30a_0$, allowing us to concentrate computational resources on the flow in the inner regions of the disc. The total number of radial cells is 512. At all radii, we chose our azimuthal grid spacing such that the aspect ratio of each cell is kept approximately unity, unless this requires a number of azimuthal grid cells $N_{\phi} > 512$, in which case we cap the number of ϕ -zones and set $N_{\phi} = 512$.

To mimic the effect of black hole accretion, gas inside the event horizon (roughly the innermost 2 or 3 radial cells around each BH) is depleted effectively instantaneously at each time-step using a sink prescription similar to Tang et al. (2017). We use a sink radius $r_{\text{sink}} = r_{\text{S}}$ and a sink time-scale $\approx 10^{-4}$ orbital period. In this paper, we assume the mass and orbit of the binary black hole are not affected by the accretion on the time-scale of our simulations.

Our simulation is pseudo-Newtonian, i.e. we model the black holes as point masses with the modified potential (Paczynski & Wiita 1980)

$$\psi(r) = -GM/(r - r_{\text{S}}). \quad (1)$$

Since the Paczynski–Wiita potential diverges at $r = r_{\text{S}}$, we employ an upper limit for the gravitational force: $|f(r)| = |f(1.5r_{\text{S}})|$ for $r < 1.5r_{\text{S}}$. We choose $1.5r_{\text{S}}$ because our cell size is $\approx 0.5r_{\text{S}}$.

We use a γ -law equation of state for the gas in the form $P = (\gamma - 1)e$, where P and e are the vertically integrated pressure and internal energy, respectively, and $\gamma = 5/3$ is the adiabatic index. Radiative cooling is incorporated naturally through the energy equation, as described in Farris et al. (2015a). We assume the disc is optically thick and geometrically thin (the optical depth due to electron scattering at $r \sim a$ near the time of decoupling is $\sim 10^5$; Haiman et al. 2009). These assumptions are consistent with our simulation results. The cooling rate is $q_{\text{cool}} = \frac{4\sigma}{3\tau_e} T^4$, where T is the mid-plane temperature and τ_e is the optical depth for electron scattering ($\tau_e = \Sigma \sigma_t / m_p$). For a Shakura–Sunyaev disc around a single BH, the cooling is in balance with the local viscous heating rate, $q_{\text{vis}} = 9/8\alpha P\Omega$. Following Farris et al. (2015a) and Farris et al. (2015b), we parametrize our simulations by specifying the Mach number at the fixed initial radius $r = a_0$, $\mathcal{M}_a = v_a / \sqrt{P(a)/\Sigma(a)}$, where $v_a \equiv (GM_{\text{bin}}/a)^{1/2}$ and Σ is the surface density. We then obtain the local cooling rate from the scaling

$$q_{\text{cool}} = \frac{9}{8} \alpha \Omega(a) P(a) \mathcal{M}_a^8 \left(\frac{P}{\Sigma v_a^2} \right)^4 \left(\frac{\Sigma}{\Sigma(a)} \right)^{-1}, \quad (2)$$

where Ω is the Keplerian angular frequency.¹ Changing \mathcal{M}_a is equivalent to changing the thickness h/r of the disc, and in this work we fix $\mathcal{M}_a = (h/r)^{-1} = 10$. A standard steady-state quasar disc (Shakura & Sunyaev 1973) with a near-Eddington accretion rate $\sim \dot{M}_{\text{Edd}}$ has a scale height of $h/r \sim 10^{-3}$, and is much thinner than the disc in our simulations ($h/r \sim 0.1$). Simulating an accretion

¹ We note there is a typographical error in equation (2) of Farris et al. (2015a); the correct exponent of the Mach number is \mathcal{M}_a^8 , as in equation (2) above.

disc with $h/r \sim 10^{-3}$ is prohibitively challenging, and our choice is typical of values employed in previous numerical works (e.g. D’Orazio, Haiman & MacFadyen 2013; Farris et al. 2014; Farris et al. 2015b; Ragusa, Lodato & Price 2016; Miranda, Muñoz & Lai 2017) have included slightly smaller h/r values ranging from 0.02 to 0.13) and allows for a direct comparison with those works. The optical thickness to electron scattering ($\tau_{\text{es}} > 10$) and free–free absorption ($\tau_{\text{ff}} > 10^4$) are both very high in our simulation domain (see e.g. fig. 10 in Corrales, Haiman & MacFadyen 2010). Therefore, throughout this paper, we compute the spectrum assuming multicolour thermal body radiation

$$L_v = \int \frac{2h\nu^3}{c^2(\exp\left[\frac{h\nu}{kT_{\text{eff}}(r, \phi)}\right] - 1)} dA, \quad (3)$$

where $T_{\text{eff}}(r, \phi)$ is the effective temperature obtained from the local cooling rate $q_{\text{cool}}(r, \phi)$ in the cell at (r, ϕ) via $T_{\text{eff}} = (q_{\text{cool}}/\sigma)^{0.25}$ with σ the Stefan–Boltzmann constant.

While computing the observed spectrum and light curve, we additionally consider the relativistic Doppler effect and the gravitational redshift. For simplicity, and to assess the maximal impact of the Doppler effect, we further assume the circumbinary disc is observed edge-on. Photons emitted by gas suffer a Doppler shift in frequency $D = [\Gamma(1 - \beta_{\parallel})]^{-1}$, where $\Gamma = (1 - \beta^2)^{-1/2}$ is the Lorentz factor, and β_{\parallel} the component of the velocity along the line of sight. The apparent flux F_v at a fixed observed frequency v is modified from the flux of a stationary source F_v^0 to $F_v = D^3 F_{v/D}^0$. For the gravitational redshift, we employ $D = (1 - r_s/R_e)^{-0.5}$, where r_s is the Schwarzschild radius and R_e is the distance between the source of emission and the BH. The above is an approximation to a fully relativistic ray-tracing in the binary’s dynamical metric (Schnittman et al. 2017). In particular, we ignore gravitational self-lensing (D’Orazio & Di Stefano 2017; Haiman 2017) and other relativistic effects, such as light-travel time modulations (i.e. Shapiro delay), which appear at order v/c .

Our disc has an initial surface density profile $\Sigma(r) = \Sigma_0(r/a)^{-0.5}$ and contains an artificial cavity with the density depleted inside the radius $2.5a_0$, as in Tang et al. (2017). Viscosity in the circumbinary disc transports angular momentum outwards, causing the cavity to refill. We first hold the binary on the initial circular orbit and evolve the system for several hundred orbits to reach a quasi-steady state. We then allow the binary separation to shrink due to quadrupole formula for GW emission (Peters 1964),

$$a(t) = a_0(1 - t/\tau)^{0.25}. \quad (4)$$

Here, $\tau = 7.22$ d is the GW inspiral time from the initial separation of $30GM_{\text{bin}}/c^2$. Finally, we employ a set of convenient code units, in which the initial binary period is $T(0) = 2\pi$ and $a_0 = 1$; as a result, $\tau = 385.17$ in code units.

3 RESULTS

As stated above, before allowing the black holes to inspiral and merge, we first hold the binary separation fixed, and evolve the system for several hundred orbits to reach a quasi-steady state. This quasi-steady state is similar to that in previous works (D’Orazio et al. 2013; Farris et al. 2014, 2015a; Tang et al. 2017) in that the cavity becomes lopsided, and two narrow streams periodically connect the minidisks around the individual BHs with the circumbinary disc. Similar to Farris et al. (2015b), we find that these features persist until right before the merger, with only a gradual decline in the size and surface density of the minidisks and the streams.

In Fig. 1, we show the 2D distributions of surface density and effective temperature at three different snapshots, corresponding to $t = 0$ (the beginning of the inspiral), 0.5τ (half-way to merger), and 0.99τ (just before the merger).

A spectrum computed at the beginning of the inspiral is also shown in Fig. 2. In this figure, we additionally show three distinct components of the spectrum, arising from outer disc ($r > 3a_0$), the shocked streams inside the cavity and shocked hot gas near the cavity wall ($a_0 < r < 3a_0$) and minidisks ($r < a_0$). At frequencies below $\sim 1\text{keV}$, the emission from the viscously heated outer disc dominates, while at higher frequencies, the emission mostly comes from the shock-heated gas in the minidisks, streams, and near the cavity wall. At frequencies between $\sim(1\text{--}20)\text{keV}$, the total spectrum shows a plateau with a small depression near $\sim 3\text{keV}$, resembling the ‘notch’ discussed in Roedig, Krolik & Miller (2014). The spectrum computed in Farris et al. (2015a) has a larger magnitude at the highest frequencies, and does not show a plateau. This is because in our simulation we resolve the event horizon, and gas near the black hole is efficiently cleared out, as it falls rapidly into the BH. In Farris et al. (2015a), a much slower sink prescription was implemented, causing hot, shock-heated gas to accumulate around each black hole, and brightening the high-frequency emission.

In Fig. 3, we show the total spectrum at three different snapshots, $t = 0$ (blue; same as in Fig. 2), 0.8τ (orange), and 0.99τ (green). The low-frequency part of the spectrum, which is dominated by the outer regions of the circumbinary disc, shows little change throughout the simulation. By comparison, the higher frequencies ($\gtrsim 1\text{keV}$), which arises from the minidisks or streams in the cavity, declines noticeably by 0.8τ ; the decline in the hard X-ray band becomes significant during the end of inspiral (0.99τ). Note that in this paper we chose $M_{\text{bin}} = 2 \times 10^6 M_{\odot}$, so that the GW signal is in the LISA band, and 0.99τ corresponds to the last ≈ 2 h of the inspiral.

The scaling relationships in our simulation can be calculated by equating the viscous heating $q_{\text{vis}} (\propto \alpha \Omega P)$ to the radiative cooling $q_{\text{cool}} (\propto \Sigma^{-1} T^4)$. Note that we fix $a_0 = 60GM_{\text{bh}}/c^2$, so that the orbital velocity is a constant, independent of M_{bin} , and the angular velocity $\Omega \propto M_{\text{bh}}^{-1}$. Since we have further fixed $\mathcal{M}_a = 10$, the mid-plane temperature T and ratio P/Σ are also fixed, i.e. $P \propto \Sigma$. Finally, with α fixed as well, we have $q_{\text{vis}} \propto \Sigma/M_{\text{bh}}$ and $q_{\text{cool}} \propto \Sigma^{-1}$. This then yields the scaling $\Sigma_0 \propto M_{\text{bh}}^{1/2}$ and the effective temperature $T_{\text{eff}} \propto \Sigma^{-1/4}$ scales as $M_{\text{bh}}^{-1/8}$. Note that this implies that the photon frequency scales with total binary mass as $\nu \propto T_{\text{eff}} \propto M_{\text{bin}}^{-1/8}$, as well. For a steady disc around a solitary BH, it would further scale with Mach number as $h\nu \propto \mathcal{M}_a^{-5/4}$ ($\Sigma \propto \mathcal{M}^{-3}$ and $T \propto \mathcal{M}^{-2}$). This suggests that for a thinner disc, the frequencies could be lower than shown for our fiducial $\mathcal{M}_a = 10$ in Figs 2 and 3. However, for gas near the cavity wall and minidisks, the variation of the disc thickness caused by its interaction with the binary BHs requires further study (e.g. Ragusa et al. 2016). In Fig. 4, we show h/r of our simulated disc. In the shock-heated region, h/r is increased to ~ 0.2 , while in the outer region h/r stays at ~ 0.1 .

In Fig. 5, we illustrate the effects of the relativistic Doppler shift and gravitational redshift on the observed spectra. As mentioned in Section 2, we assume the disc is observed edge-on, so that the Doppler effect is maximized. Note the spectra in Fig. 5 do not include the inclination angle. The observed spectrum will be dimmer by a factor of $\cos(\theta)$, where θ is the angle of inclination to the line of sight. We include four viewing angles in the plane of the disc, corresponding to the $\pm\hat{x}$ and $\pm\hat{y}$ directions. As the figure shows, at the low-frequency end ($h\nu \lesssim 1\text{keV}$), the spectrum is not modified by either the Doppler or gravitational shifts. At higher frequencies

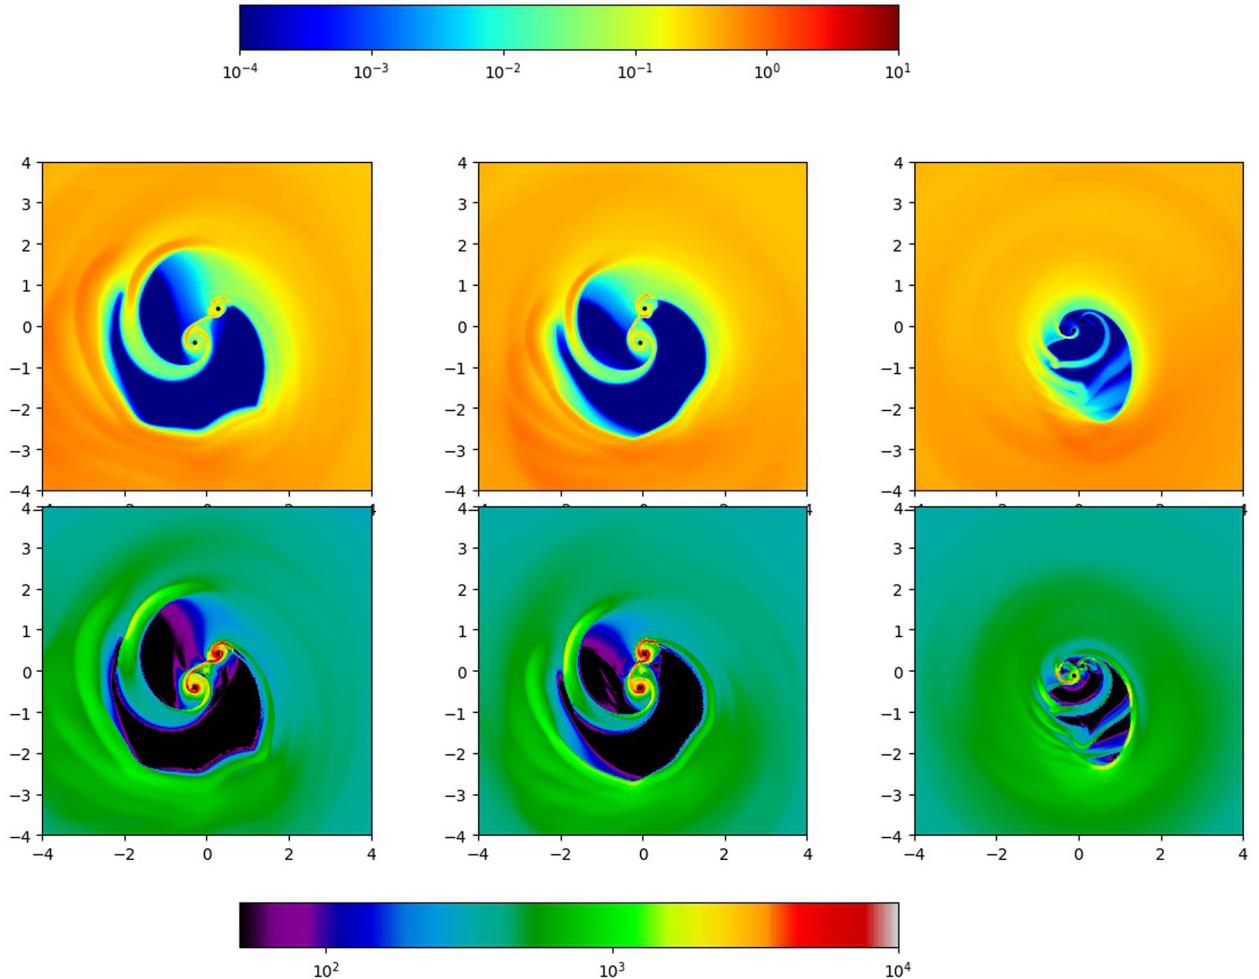


Figure 1. The top row shows snapshots of the surface density Σ/Σ_0 , where Σ_0 is the initial surface density at $r = a$, and the bottom row shows snapshots of the effective temperature T_{eff} in [eV]. The horizontal and vertical axes are in units of the initial binary separation a_0 . The snapshots are taken, from left to right, at the times $t = 0.0\tau$, 0.5τ , and 0.99τ , where $\tau = 7.22$ d corresponds to the merger. The cavity size decreases and the BHs accrete efficiently during the entire inspiral.

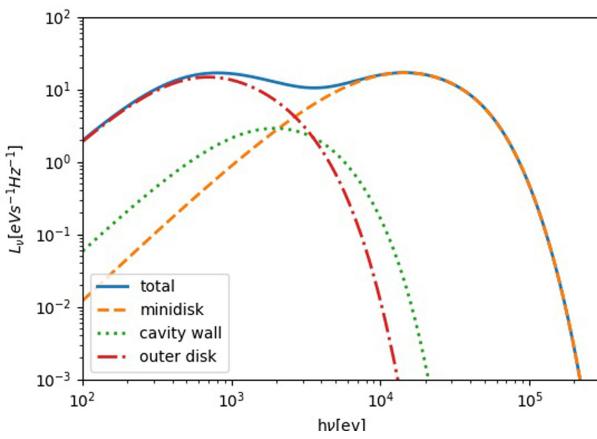


Figure 2. Thermal spectrum computed from a simulation snapshot at the beginning of the inspiral ($t = 0$). The full spectrum is shown by the blue curve. The distinct components arising from gas in the minidisks ($r < a$), streams and cavity wall ($a < r < 3a$) and from the outer regions of the circumbinary disc ($r > 3a$) are represented by the orange dashed line, green dotted line, and red dash-dotted line, respectively.

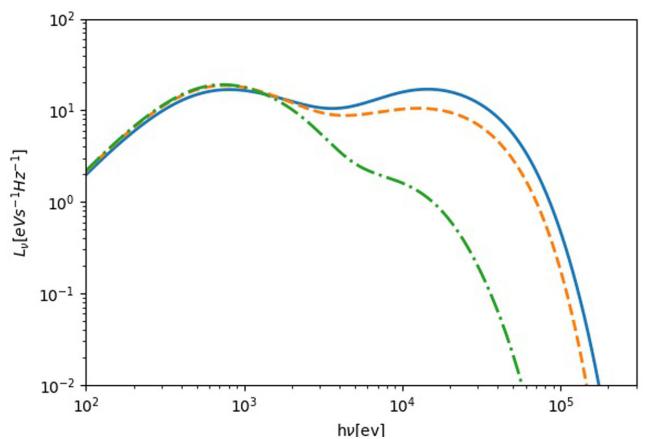


Figure 3. Total composite thermal spectra computed from simulation snapshots at $t = 0$ (blue curve), $t = 0.8\tau$ (orange dashed), and $t = 0.99\tau$ (green dot-dashed).

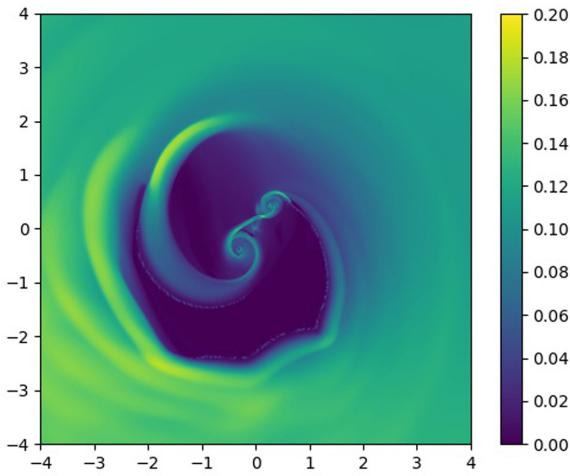


Figure 4. Local disc thickness h/r throughout the disc at $t = 0$.

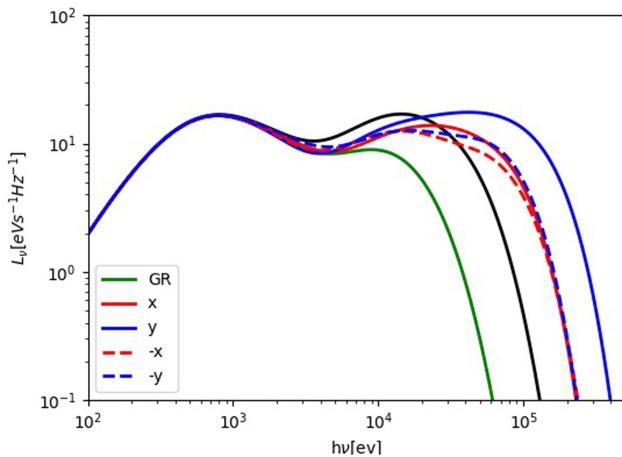


Figure 5. Spectra illustrating the impact of the special relativistic Doppler shift and the gravitational red/blue shift. For reference, the black curve shows the spectrum without considering either of these effects. The green curve shows the spectrum including only the gravitational redshift, while the other four spectra, shown in blue and red, include both relativistic effects (for the four different viewing directions $\pm\hat{x}$ and $\pm\hat{y}$, as labelled in the inset). All spectra are calculated from the snapshot at $t = 0$, when the orbital velocity of the BHs is $0.091c$.

($1\text{ keV} \lesssim h\nu \lesssim 20\text{ keV}$), the Doppler shift causes an overall dimming (i.e. red and blue curves), while at the highest frequencies ($h\nu \gtrsim 20\text{ keV}$), the spectrum is brightened significantly. Interestingly, we find that the dimming and the brightening retain their sign in all four different viewing directions; we offer a simple interpretation of this result in Section 4 below. For reference, the figure also shows a case with only the GR redshift (green curve), which is independent of viewing angle, and always reduces the observed flux.

In order to study the variability of the spectrum at different frequencies, we calculate the light curves in the soft (at 2 keV) and hard (at 10 keV) X-ray bands. In Fig. 6, we show the light curve at the beginning of the simulation and right before merger, with and without the Doppler effect. In the 2 keV band, we find clear periodicity in the light curve throughout the inspiral. In the 10 keV band, the inspiral light curve at the beginning is much noisier. We can still see some periodicity by eye, but it is harder to distinguish than in the 2 keV case. However, in the late stages of the inspiral,

the 10 keV band also shows clear periodicity at the orbital period. The light curves also show an overall see-saw modulation, with a period approximately eight times longer than t_{orb} . Furthermore, the light curves show a clear behaviour analogous to the GW chirp signal, in that the periodicity tracks the binary's increasing orbital frequency.

In order to identify the origin of the periodicity in the light curves, in Fig. 7 we decompose them into emission arising from $r < a$ (dominated by gas in/near the minidisks) and from $r > a$ (dominated by the circumbinary disc). This figure reveals that the emission from the circumbinary disc is clearly periodic from the beginning of the simulation all the way to the merger in both bands. This is because each BH periodically interacts with the cavity wall and creates a stream which hits the cavity wall, shock-heats the gas, and leads to enhanced X-ray emission. By comparison, the emission from the minidisks is noisy in the beginning, and develops clear periodicity only closer to the merger. One explanation for this phenomenon is that in the earlier stage the minidisks act as buffers of the accretion streams. Unlike steady-state α discs, the internal shocks in the minidisks create significant noise; their overall shape often differs significantly from a circular disc. In the late stage of the inspiral (in the last 60 h before the merger), as the truncation radius decreases, this buffer effect declines, and the narrow accretion streams eventually connect the cavity wall and event horizon almost directly (see the right-hand panel of Fig. 1). Therefore, the noise produced by the minidisks gradually disappears.

In the 2keV band, emission from the circumbinary disc dominates, and the light curve therefore also shows very clear periodicity, during the entire inspiral process. The opposite situation occurs in the 10keV band, in which the total emission is dominated by the minidisks. The light curve is noisy in the beginning, and shows much clearer periodicity in the end stage. From the bottom-most panel of Fig. 7, we see that the emission from the minidisks and circumbinary discs has the same period, but with different phases. Therefore, in the later stages, when the strength of the emission from these two regions is comparable, the total light curve sometimes shows two maximas in a single period.

Haiman (2017) and D’Orazio & Haiman (2017) proposed a toy model for the minidisks as two ‘light bulbs’, both of which are steady in their own rest frame. Therefore, unless observed perfectly face-on, the Doppler modulation will lead to periodicity at the orbital frequency. Note that the light curve would still be modulated, even in the case of an equal-mass binary, because the red and blueshifts of the thermal emission from the two BHs, with equal but opposite line-of-sight velocities, will not exactly cancel. However, in our simulation, we find that the internal structure of the minidisks is far from stationary, with shocks forming and streams of gas entering and leaving the fiducial Hill sphere. The emission from the minidisk regions is therefore quite noisy. As a result, the Doppler modulation may be difficult to extract from the observed light curve. We expect that this may still be feasible, given a GW template from LISA, with which the X-ray light curve can be cross-correlated, but this requires further analysis. Fortunately, we find that the shock-heated cavity wall provides a clear periodicity throughout the entire inspiral process. Moreover, since this periodicity does not originate from the Doppler effect (see Fig. 6), it should exist independently of the inclination angle.

In order to assess how well the periodicity in the light curves track the binary’s orbital period, we rescale the light-curve time series with the evolving binary orbital period. In Fig. 8, we plot the Lomb–Scargle periodograms with and without this time rescaling. The rescaled periodograms show peaks at $\approx 2 \times \Omega_{\text{orb}}$, both at the

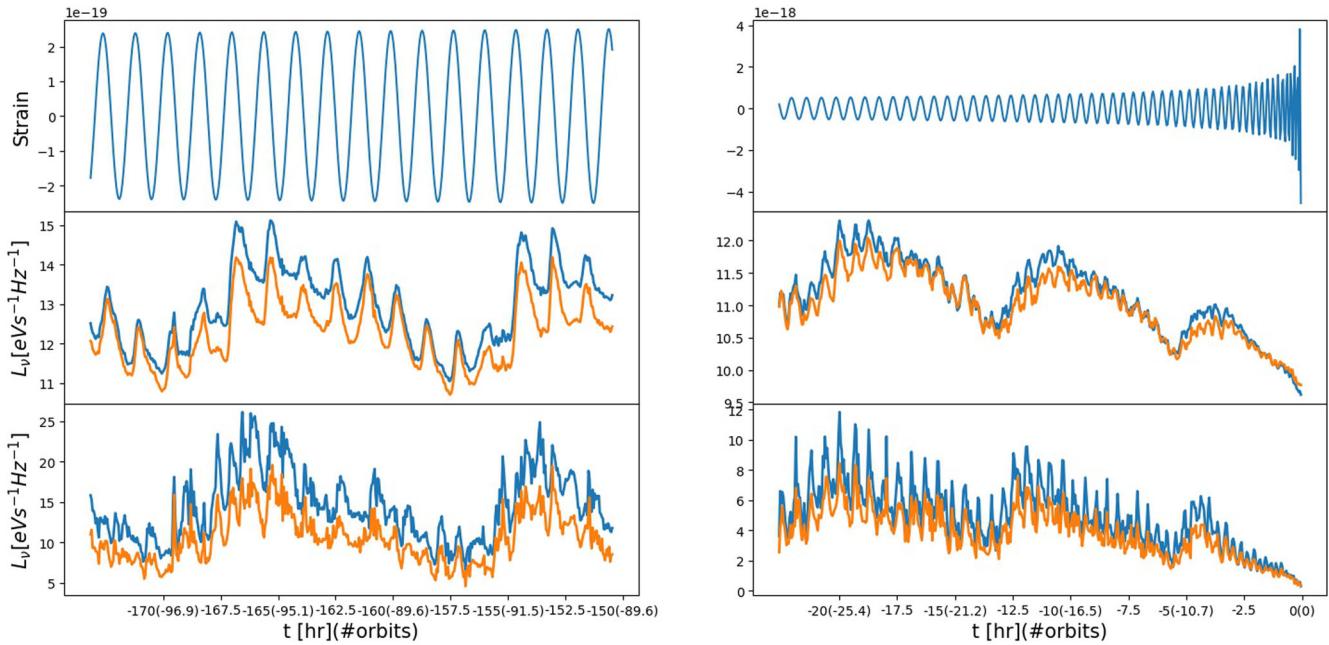


Figure 6. Gravitational wave strain (top panels) and X-ray light curves (middle and bottom) calculated at the beginning (left) and end (right) of the simulated inspiral in the LISA band. The top panel shows the GW strain observed edge-on from a binary at redshift $z = 1$. The middle panel corresponds to the 2 keV band and the bottom panel to 10 keV band. The blue curves show the light curves neglecting the relativistic Doppler shift or gravitational redshift. The orange curves show the light curves including these effects. While the Doppler effect imposes a small additional periodic modulation due to the binary's orbital motion, its dominant effect is a nearly uniform dimming of the light curve (this results from the asymmetry between the large blue and redshifts due to the internal motion of gas in the minidisks; see the text for discussion). The x -axis shows the look-back time in the source's rest frame, and the number of binary orbits (in parentheses).

beginning and the end of the inspiral, implying that this periodicity is able to follow the accelerating binary orbital frequency (i.e. the ‘chirp’) throughout the entire inspiral process. The first peak in the periodogram without rescaling arises from a dense lump in the lopsided cavity wall, as reported in a series of previous works (e.g. MacFadyen & Milosavljević 2008; Cuadra et al. 2009; Krolik 2010; Roedig et al. 2011; Noble et al. 2012; Shi et al. 2012; D’Orazio et al. 2013; Farris et al. 2014; Bowen et al. 2017a). This feature is smeared out and disappears after time rescaling, because it does not follow the binary’s frequency chirp. More specifically, the frequency of this low-frequency peak increases from $\approx 0.25\Omega_{\text{orb}}(0)$ to $\approx 0.4\Omega_{\text{orb}}(0)$, but this increase is slower than the evolving binary orbital frequency $\Omega_{\text{orb}}(t)$. Just prior to the merger, the period of this low-frequency modulation is around 15 orbits. Similarly, in Fig. 1, we see that the overall size of the cavity is shrinking, but cannot keep up with the binary separation, which decreases more rapidly. Nevertheless, this low-frequency modulation persists throughout the inspiral process, as we can clearly see in Fig. 6. In principle, the evolution of this low-frequency modulation could be measured in the X-ray light curves, and would probe the gas morphology near the central cavity.

4 DISCUSSION

One of the motivations of this work was to compute the modulation of the light curves due to relativistic Doppler effect. Naively, such a modulation inevitably arises from the same orbital motion that produces GWs, and therefore the GW and X-ray chirp signals should track one another with a known phase (Haiman 2017). This would be of particular interest for LISA, allowing a robust measurement of the speed difference of photons and gravitons, and more generally, a secure identification of the EM counterpart of a LISA binary.

Our results suggest that it may be difficult to extract this Doppler modulation from the X-ray light curve. We find efficient fuelling of both BHs through the inspiral process, with gas funnelled inside each BH’s Hill radius, forming tidally truncated ‘minidisks’. These minidisks persist nearly all the way to the merger, until the tidal truncation radius shrinks to a size comparable to that of the ISCOs (roughly 60 h prior to merger). However, the effective temperature, density, and velocity fields of the gas inside these minidisks are not stationary, and the net Doppler shift is dominated by these large ‘internal’ gas motions, rather than the orbital velocity of the BHs. Interestingly, as Figs 5 and 6 show, the most prominent effect of the Doppler shift in the 2–10 keV range is an overall dimming of the light curve, rather than a sinusoidal modulation. This is because the hot gas dominating this emission is located very close to each BH (near or inside its Hill radius), and the velocity of this gas is larger than the orbital speed of the BHs. Regardless of the viewing direction, there is always gas with both large blue and redshifts along the line of sight, and the total composite spectrum includes both of these shifts. The net result can be either an overall dimming or brightening, depending on the local slope and curvature of the spectrum. In other words, for the steeply curved blackbody spectral shape, the Doppler effect is not symmetric under switching the sign of the line-of-sight velocity. Fig. 5 shows that in the 2–10 keV range, the redshift and the corresponding dimming from receding gas dominate over the blueshift and brightening from approaching gas.

We emphasize that this result is sensitive to the spectral shape. One caveat is that, as mentioned above, due to numerical constraints our simulated disc is thicker than a typical active galactic nucleus (AGN) disc of a solitary BH. If we were to scale the simulation to correspond to a higher Mach number, this would shift the

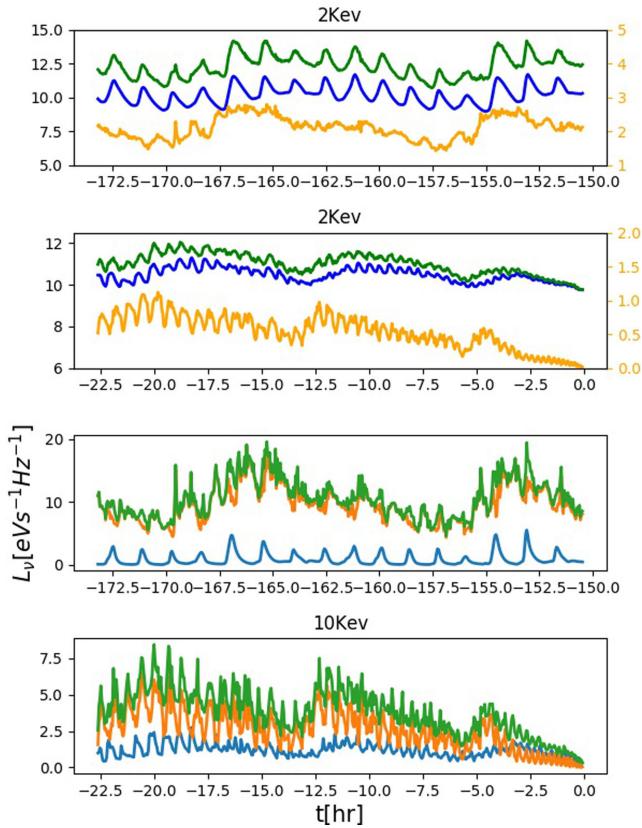


Figure 7. Decomposition of the light curves. The orange and blue curves correspond to emission arising from $r < a$ (dominated by gas in/near the minidisks) and from $r > a$ (dominated by the circumbinary disc). The green curves correspond to the total emission. The upper two panels show the light curves in the 2 keV band, and the bottom two panels in the 10 keV band. For clarity of display, the orange curves, which show sub-dominant emission from the minidisks in the soft band in the upper two panels, use the y-axis labels on the right side.

normalization of the photon energy. For example, the soft versus hard X-ray bands could become FUV and EUV bands for a disc with a Mach number of $\mathcal{M}_a \sim 1000$, instead of the $\mathcal{M}_a = 10$ we adopted (recall that $h\nu \propto \mathcal{M}_a^{-5/4}$ for single black hole accretion disc). On the other hand, the gas in the inner regions is being heated primarily by shocks, rather than viscosity, and we find that h/r in the inner regions of the flow remains ~ 0.2 throughout the simulation, justifying the scaling to lower Mach numbers (see Fig. 4). However, it requires further study to understand how this depends on the choice of \mathcal{M}_a . A higher \mathcal{M}_a could lead to a larger cavity and reduce the gas fuelling from the cavity wall (Ragusa et al. 2016). This may reduce the luminosity of the minidisks.

We also note that in equation (2), we assumed that the fluid is gas-pressure dominated; this assumption is likely violated in the shock-heated streams and minidisks.

Another caveat in our conclusions is that we focused on thermal emission. We have verified that the effective absorptive opacity, defined as $\tau_{\text{eff}} = [\tau_{\text{abs}}(\tau_{\text{abs}} + \tau_{\text{scat}})]^{1/2}$, where τ_{abs} is the optical depth to free-free absorption and τ_{scat} is the optical depth to electron scattering, remains $\tau_{\text{eff}} \sim 10^4 \gg 1$ throughout our simulation. This justifies considering a thermalized spectrum. However, empirically, the X-ray emission from an AGN consists of a combination of thermal disc emission and a hot corona, and in this paper, we have

not included the latter component. If the emission from a corona turns out to be more stable (i.e. if the corona lacks large internal bulk motions), the Doppler modulation in the X-ray band may be dominated by the BH's orbital motion, and much more prominent. Moreover, in this paper, we have focused on an equal-mass binary. We expect that the Doppler modulation would be more prominent in the case of low-mass ratio binaries, because the circum-primary and -secondary minidisks would be more stable, and there would be less cancellation between their emission.

The overall shape of the light curves we found show a clear 'X-ray chirp', caused by the interaction between the binary and the gas discs. The periodicity of the chirp tracks the shrinking binary orbital period t_{orb} , in tandem with the GW signal. Since the chirp is produced by hydrodynamical modulation, its absolute phase is not tied directly to the GW emission. Nevertheless, the speed difference between photons and gravitons should still be measurable, in principle, given a precise measurement of the phase-evolution of both the GW and the X-ray chirp signals. In the future, it would be interesting to compute the accuracy to which the phase difference can be measured in practice, if the absolute phase is unknown ab-initio, given realistic observational S/N from LISA, and from an X-ray instrument such as Athena,² Lynx,³ ULTRASAT,⁴ or EXTP (Zhang et al. 2016). In this paper, we neglect spin of the BHs, but the spin would change the GW signal significantly, and would further complicate a direct comparison of the X-ray light curve with the LISA signal. The ISCO for a spinning black hole is smaller than a non-spinning one, which may enhance the luminosity at high frequencies.

We have made several simplifying assumptions throughout this paper, which we intend to relax in our future work. We have focused here on an equal-mass binary, but we intend to study the dependence of the spectra and light curves on the binary mass ratio. Similarly, we intend to determine the dependence of these features on the disc thickness. We also used an α -prescription for the viscosity. We intend to perform magnetohydrodynamic (MHD) simulations in the future (e.g. Krolik 2010). In this work, we have also assumed that the binary BHs are on circular orbits, with zero eccentricity. On the other hand, eccentricity is likely to develop as a result of strong gas torques during the earlier stages of the inspiral (e.g. Cuadra et al. 2009; Roedig et al. 2011; Sesana, Gualandris & Dotti 2011; Roedig et al. 2012; Miranda, Muñoz & Lai 2017). In the future, we intend to study live binaries, and follow the development of eccentricity, and to assess the impact of eccentricity on the accretion rates and the predicted spectra and light curves.

In the present simulation, we followed a pseudo-Newtonian approach; a full relativistic simulation of the gas in the vicinity of the event horizon could modify our results. Finally, we started our current simulations somewhat past the nominal 'decoupling' stage. In the cartoon picture of the binary inspiral (e.g. Milosavljevic & Phinney 2005), once the binary is sufficiently compact that the GW inspiral time-scale is shorter than the viscous time-scale in the nearby disc, the BHs outpace the disc and leave the disc gas behind. The decoupling separation a_{dec} can be calculated with $t_{\text{gw}} = t_{\text{vis}}$, where $t_{\text{gw}} = \frac{5c^5 a^4}{64G^3 m_1 m_2 (m_1 + m_2)}$ is the GW inspiral time-scale (Peters 1964) and $t_{\text{vis}} = \frac{2(2a)^2}{3\nu}$ is the viscous time in an unperturbed disc at the radius $r = 2a$.

² See www.cosmos.esa.int/web/athena.

³ See www.astro.msfc.nasa.gov/lynx.

⁴ See www.weizmann.ac.il/ultrasat.

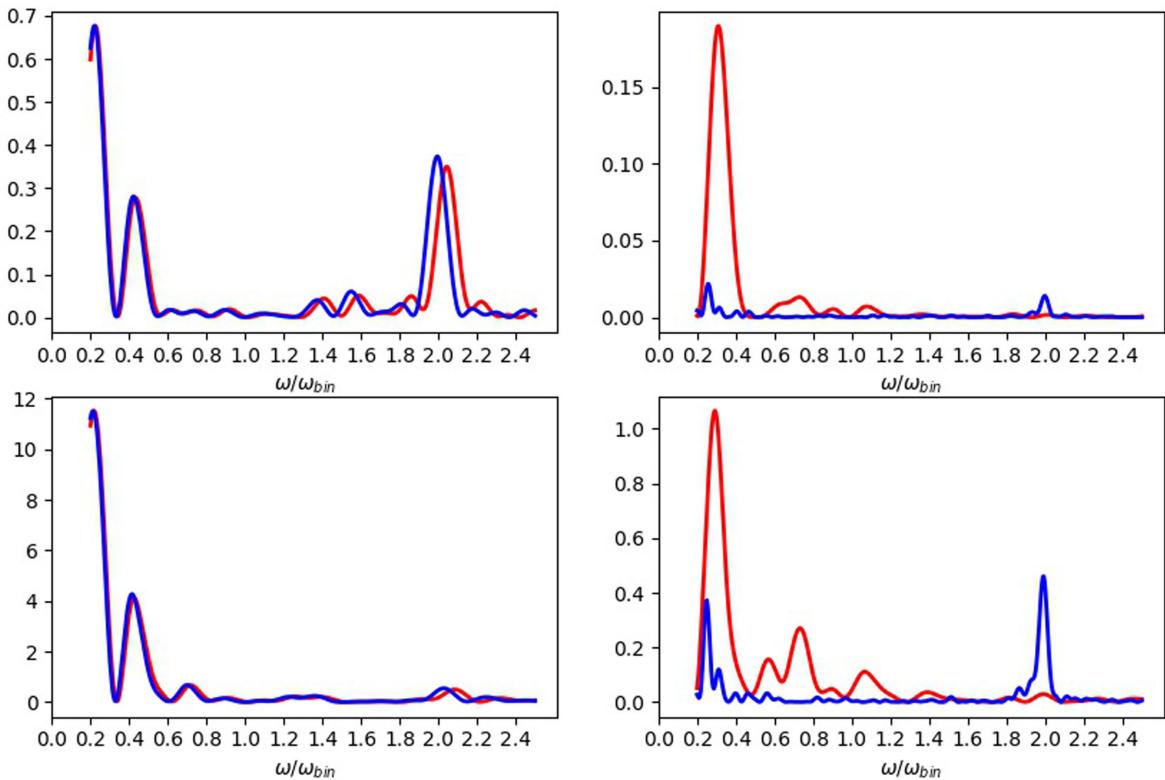


Figure 8. Lomb-Scargle periodograms of the light curves. The four panels correspond to those in Fig. 6, namely to the 2 and 10 keV bands (top and bottom, respectively) and to the beginning and end of the simulated inspiral (left and right, respectively). In each panel, the red curve corresponds to the periodogram of the original light curves shown in Fig. 6. The blue curves correspond to scaled light curves, in which time is measured in units of the instantaneous binary orbital period.

With the parameters ($\alpha = 0.1$; $\mathcal{M}_a = 10$) adopted in our simulation, we find the decoupling separation $a_{\text{dec}} = 128GM_{\text{bh}}/c^2$, which is somewhat larger than our initial separation of $60GM_{\text{bh}}/c^2$.

In this picture, our procedure of holding the binary's orbit initially fixed, and allowing it to reach a steady state would not be justified. However, we first note that even in a simple 1D toy model, the decoupling radius can only be taken as an order-of-magnitude estimate. For example, torques modify the density profile outside the binary's orbit, causing a pile-up, increasing the pressure, steepening the radial gradients, shortening the viscous time, and decreasing a_{dec} (e.g. Milosavljevic & Phinney 2005; Kocsis, Haiman & Loeb 2012). More importantly, from the results we find here, as well as from our earlier work (Farris et al. 2015b), we see that the cartoon picture is inaccurate, and the cavity can, in fact, shrink and follow the inspiraling binary, even well past the nominal decoupling. This is because the gas dynamics near the black holes is dominated by gravity and shocks, and not by viscosity, and the gas configuration is neither steady nor axisymmetric. More specifically, in Farris et al. (2015b), a simulation was started at the separation corresponding to $t_{\text{gw}} = 5t_{\text{vis}}$, i.e. safely before the nominal decoupling. Nevertheless, we found that the gas flows inwards, following the binary, and the accretion rate is only halved, even well past the decoupling, $0.1t_{\text{vis}}$ before the merger. Similarly, in our simulation here, the spectrum is only slightly modified even at 0.8τ (see Fig. 3). We therefore expect our way of setting up the initial condition to be reasonably accurate. However, one caveat to note here is that the nominal decoupling separation is larger for a disc with smaller aspect ratios h/r (i.e. large Mach numbers). This means that the reduction of accretion rate and luminosity in more realistic, thinner discs may

be more significant than the discs in our simulations. We intend to study this in future work, by beginning simulation near or prior to the decoupling stage and with larger Mach numbers.

5 CONCLUSIONS

We have performed 2D simulations of accretion on to an inspiraling and merging binary BH system, and studied the corresponding EM signatures. We have computed the evolving multicolour blackbody emission from snapshots of our simulations, including the effects of the relativistic Doppler shift and gravitational redshift. The most important conclusions of this work can be summarized as follows:

(1) *The light curves show clear periodicity.* Accretion on to the individual BHs has long been known to be periodic (Artymowicz & Lubow 1994), and it is not surprising that any luminosity produced by gas near the individual BHs is periodic. However, we have found that the emission arising from gas farther out, near the cavity wall, which dominates the softer X-ray bands, shows a similar (and even clearer) periodicity. This is because streams of gas, flung outwards by the BHs, periodically hit and shock the cavity wall. The frequency is twice the binary's instantaneous orbital frequency.

(2) *Distinct behaviour in different bands.* In the beginning of the simulation, the emission from the minidisks is noisy, but develops clear periodicity during the later stages. In the soft X-ray band, the emission from the circumbinary disc dominates, and the light curve shows a clear periodicity from the beginning of the simulation ($a = 60GM/c^2$). By comparison, in the hard X-ray band, the emission in the beginning is dominated by the minidisks, and periodicity

in the early-inspiral light curve is obscured by noise, with clearer periodicity developing only in the late stages.

(3) *Doppler modulation is sub-dominant.* We find that the Doppler modulation of the light curve does not follow a simple toy model of two moving light bulbs in orbit. This is because the effective temperature, density, and velocity field of the gas in the minidisks are not stationary, and the line-of-sight velocity of this gas is dominated by the large internal gas motions inside the minidisks, rather than the clean orbital motion of the BHs. We find that the most conspicuous effect of the Doppler shift is an overall dimming ($1\text{keV} \lesssim h\nu \lesssim 20\text{keV}$) or brightening ($h\nu \gtrsim 20\text{keV}$) of the light curves. We caution that these results are sensitive to the spectral shape and emission mechanism. In particular the Doppler effect could more closely track the binary's orbital motion for X-ray emission from more stable coronae around the individual BHs which do not have large internal motions.

(4) *The EM chirp follows the GW inspiral.* We have found that the minidisks persist until the very end of the inspiral process. The discs are truncated at the gradually shrinking size of the Hill radius. The thermal emission from the minidisks declines gradually, but remains overall significant all the way to the merger. The overall periodicities in both the soft and hard X-ray bands track the shrinking orbital period of the binary.

(5) *No ‘second decoupling’.* Fontecilla, Chen & Cuadra (2017) suggested that the inner (circum-primary) disc will become geometrically thick ($h/r \sim 1$) in the late stages of the inspiral, as a result of tidal heating. They find, using 1D analytical models and numerical simulation, that the viscous time in the inner disc decreases to $t_{\text{vis}} \sim t_{\text{GW}}$, and it effectively decouples from the binary. By comparison, in our 2D simulation, the size and surface density of the two minidisks are continuously adjusted, and eventually, in the late stages, the minidisks disappear, and narrow accretion streams instead fall inside the event horizon directly. The inspiral process does thicken the minidisks, but the aspect ratios remain below $h/r \sim 0.2$.

The simulations presented here are based on many simplifications. Nevertheless, our results suggest that an EM chirp signal may accompany the GW emission, with the EM period tracking that of the GW signal right up to the merger of an SMBH binary. A detection of this EM chirp could provide a potential ‘smoking gun’ evidence to uniquely identify the EM counterpart of an SMBH binary detected by LISA.

ACKNOWLEDGEMENTS

We thank Daniel D’Orazio, Paul Duffell, Julian Krolik, and Geoffrey Ryan for useful discussions. Financial support was provided by NASA through ATP grant NNX15AB19G and ADAP grant NNX17AL82G and by NSF grants 1715661 and 1715356. ZH also gratefully acknowledges support from a Simons Fellowship in Theoretical Physics and hospitality by NYU during his sabbatical leave when this work began. This work was performed in part at Aspen Center for Physics, which is supported by National Science Foundation grant PHY-1607611.

REFERENCES

Abbott B. P. et al., 2017, *ApJ*, 848, L13
 Amaro-Seoane P. et al., 2017, preprint ([arXiv:1702.00786](https://arxiv.org/abs/1702.00786))
 Artymowicz P., Lubow S. H., 1994, *ApJ*, 421, 651
 Bode T., Haas R., Bogdanović T., Laguna P., Shoemaker D., 2010, *ApJ*, 715, 1117
 Bowen D. B., Mewes V., Campanelli M., Noble S. C., Krolik J. H., Zilhão M., 2017a, *ApJ*, 853, L17
 Bowen D. B., Campanelli M., Krolik J. H., Mewes V., Noble S. C., 2017b, *ApJ*, 838, 42
 Corrales L. R., Haiman Z., MacFadyen A., 2010, *MNRAS*, 404, 947
 Cuadra J., Armitage P. J., Alexander R. D., Begelman M. C., 2009, *MNRAS*, 393, 1423
 D’Orazio D. J., Di Stefano R., 2017, *MNRAS*, 474, 2975
 D’Orazio D. J., Haiman Z., 2017, *MNRAS*, 470, 1198
 D’Orazio D. J., Haiman Z., MacFadyen A., 2013, *MNRAS*, 436, 2997
 Dotti M., Sesana A., Decarli R., 2012, *Adv. Astron.*, 2012
 Duffell P. C., 2016, *ApJS*, 226, 2
 Duffell P. C., MacFadyen A. I., 2012, *ApJ*, 755, 7
 Farris B. D., Gold R., Paschalidis V., Etienne Z. B., Shapiro S. L., 2012, *Phys. Rev. Lett.*, 109, 221102
 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
 Fontecilla C., Chen X., Cuadra J., 2017, *MNRAS*, 468, L50
 Giacomazzo B., Baker J. G., Miller M. C., Reynolds C. S., van Meter J. R., 2012, *ApJ*, 752, L15
 Haiman Z., 2017, *Phys. Rev. D*, 96, 023004
 Haiman Z., Kocsis B., Menou K., 2009, *ApJ*, 700, 1952
 Kocsis B., Haiman Z., Menou K., 2008, *ApJ*, 684, 870
 Kocsis B., Haiman Z., Loeb A., 2012, *MNRAS*, 427, 2680
 Kormendy J., Ho L. C., 2013, *ARA&A*, 51, 511
 Krolik J. H., 2010, *ApJ*, 709, 774
 Lang R. N., Hughes S. A., 2008, *ApJ*, 677, 1184
 MacFadyen A. I., Milosavljević M., 2008, *ApJ*, 672, 83
 Milosavljević M., Phinney E. S., 2005, *ApJ*, 622, L93
 Miranda R., Muñoz D. J., Lai D., 2017, *MNRAS*, 466, 1170
 Noble S. C., Mundim B. C., Nakano H., Krolik J. H., Campanelli M., Zlochower Y., Yunes N., 2012, *ApJ*, 755, 51
 Paczyński B., Wiita P. J., 1980, *A&A*, 88, 23
 Peters P. C., 1964, *Phys. Rev.*, 136, B1224
 Phinney E. S., 2009, Astro2010: The Astronomy and Astrophysics Decadal Survey, Science White Papers, National Research Council (2010) Washington, 235
 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
 Roedig C., Sesana A., Dotti M., Cuadra J., Amaro-Seoane P., Haardt F., 2012, *A&A*, 545, A127
 Roedig C., Krolik J. H., Miller M. C., 2014, *ApJ*, 785, 115
 Schnittman J. D., Dal Canton T., Camp J., Tsang D., Kelly B. J., 2017, *ApJ*, 853, 123
 Sesana A., Gualandris A., Dotti M., 2011, *MNRAS*, 415, L35
 Shakura N. I., Sunyaev R. A., 1973, *A&A*, 24, 337
 Shi J.-M., Krolik J. H., Lubow S. H., Hawley J. F., 2012, *ApJ*, 749, 118
 Tang Y., MacFadyen A., Haiman Z., 2017, *MNRAS*, 469, 4258
 Will C. M., 2006, *Living Rev. Relativ.*, 9, 3
 Zhang S. N. et al., 2016, in Jan-Willem A., den H., Tadayuki T., Marshall B., eds, Proc. SPIE Conf. Ser. Vol. 9905, Space Telescopes and Instrumentation 2016: Ultraviolet to Gamma Ray. SPIE, Bellingham, p. 99051Q

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