



Electromagnetic Emission from Supermassive Binary Black Holes Approaching Merger

Stéphane d'Ascoli^{1,2}, Scott C. Noble^{3,4}, Dennis B. Bowen¹, Manuela Campanelli¹, Julian H. Krolik⁵, and Vassilios Mewes¹

¹ Center for Computational Relativity and Gravitation, Rochester Institute of Technology, Rochester, NY 14623, USA

² École Normale Supérieure, 24 rue Lhomond, F-75005 Paris, France

³ Department of Physics and Engineering Physics, The University of Tulsa, Tulsa, OK 74104, USA; scott.c.noble@nasa.gov

⁴ Gravitational Astrophysics Laboratory, NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

⁵ Department of Physics and Astronomy, Johns Hopkins University, Baltimore, MD 21218, USA

Received 2018 June 5; revised 2018 August 6; accepted 2018 August 7; published 2018 October 2

Abstract

We present the first relativistic prediction of the electromagnetic emission from the surrounding gas of a supermassive binary black hole system approaching merger. Using a ray-tracing code to post-process data from a general relativistic 3D magnetohydrodynamic simulation, we generate images and spectra, and analyze the viewing angle dependence of the light emitted. When the accretion rate is relatively high, the circumbinary disk, accretion streams, and mini-disks combine to emit light in the UV/extreme-UV bands. We posit a thermal Compton hard X-ray spectrum for coronal emission; at high accretion rates, it is almost entirely produced in the mini-disks, but at lower accretion rates it is the primary radiation mechanism in the mini-disks and accretion streams as well. Due to relativistic beaming and gravitational lensing, the angular distribution of the power radiated is strongly anisotropic, especially near the equatorial plane.

Key words: accretion, accretion disks – black hole physics – galaxies: nuclei – magnetohydrodynamics (MHD) – radiative transfer

1. Introduction

1.1. Context

Electromagnetic (EM) observations of supermassive binary black holes (SMBBHs) and their environments have the potential to provide critical new information about both galaxy evolution and strong-field gravity. Unlike merging stellar-mass binary black holes (BBHs) recently discovered by the LIGO–Virgo Collaboration (Abbott et al. 2016a, 2016b, 2016c, 2017), SMBBHs may often merge in gas-rich environments (Dotti et al. 2009; Pfister et al. 2017; Goulding et al. 2018) and can therefore be EM-bright during all stages of the coalescence process.

Direct detection of SMBBHs through gravitational wave (GW) emission may be accomplished by orbiting GW observatories, but not any time soon (Amaro-Seoane et al. 2017). Pulsar Timing Array observations could detect GW radiation from SMBBHs, but the GW frequencies to which they are sensitive correspond only to the weakly relativistic regime for the most massive SMBBHs ($\gtrsim 10^9 M_\odot$; Shannon et al. 2015). Identification of photons from SMBBHs by some of the many EM telescopes now operating could jump-start this field, sharply refining our estimates of the population and evolution of SMBBHs, as well as guiding planning and development of space-based GW observatories.

Observational efforts to identify SMBBHs to date have been defined by qualitative guesses about what observable properties might be distinctive. One approach to finding true SMBBHs is focused on high-resolution imaging, possible only via radio-frequency Very Long Baseline Interferometry (Tremblay et al. 2016). The recent discovery of possible orbital motion in the radio galaxy 0402+379 presents an exciting new prospect of probing a SMBBH's kinematics (Bansal et al. 2017). Another approach rests upon the hope that some aspects of their light may exhibit periodic variability (Graham et al. 2015; Liu et al. 2015, 2016). The latter

approach is made difficult by the fact that the monitoring programs rarely cover more than a few cycles of the candidate periods identified, thus providing only weak evidence for periodicity. The results presented here are a first step toward establishing more physically grounded predictions of distinctive spectral and timing properties of these intriguing systems.

1.2. Prior Work

The structure of a circumbinary disk when the binary mass ratio $q \equiv M_2/M_1 \gtrsim 0.02$ is well-established: if the binary semimajor axis is a and its eccentricity is e , a gap forms within a radius from the center of mass $\approx 2a(1 + e)$ because closed orbits enclosing the binary do not exist at smaller radii, while an ordinary accretion disk occupies radii $\gg 2a$ (Pringle 1991; MacFadyen & Milosavljević 2008; Roedig et al. 2011; Shi et al. 2012). Although early work treating a 1D model of such a system argued that torques exerted by the binary would prevent any accretion through the gap (Pringle 1991), detailed 2D and 3D simulations have shown that, although matter can pile up near $r \approx 2a$, ultimately inflow equilibrium is achieved so that the mass accretion rate across the gap matches that in the outer parts of the circumbinary disk (MacFadyen & Milosavljević 2008; Noble et al. 2012; Roedig et al. 2012; Shi et al. 2012; D'Orazio et al. 2013; Farris et al. 2014; Shi & Krolik 2015). Matter crosses the gap in narrow streams, whose ultimate destinations depend upon the matter's specific angular momentum. Gas with specific angular momentum close to the circular orbit value at the circumbinary disk's inner edge spends enough time in the gap that binary torques propel it back to $r \approx 2a$, where it shocks against the disk; gas with significantly lower angular momentum is created by deflection in these shocks and plunges into the zone of the binary (Shi & Krolik 2015). Once the streams find their way close to the binary, they join one of the “mini-disks,” individual accretion disks each centered on one of the partners in the binary

(Bowen et al. 2017, 2018; Ryan & MacFadyen 2017; Tang et al. 2017). Although it was initially expected that accretion from the circumbinary disk to the mini-disks would be cut off when the timescale on which the binary orbit shrinks due to GW radiation becomes shorter than the accretion timescale in the inner region of the circumbinary disk (Milosavljević & Phinney 2005), simulations have shown that this cut-off does not actually occur (Noble et al. 2012; Farris et al. 2015a) because the very shortness of the binary lifetime means that only material very close to the edge of the circumbinary disk needs to be drawn upon to feed the flow.

Radiation can arise in any of these regions: the circumbinary disk, the streams, and the mini-disks. Its energy source is dissipation of kinetic energy and magnetic field energy which, in turn, is drawn from mass moving into deeper portions of the gravitational potential. Several different mechanisms can contribute to this dissipation. In ordinary accretion disks, the majority of the heat is generated by dissipation at the short-lengthscale end of the inertial cascade associated with magnetohydrodynamic (MHD) turbulence stirred by the magnetorotational instability (MRI; Balbus & Hawley 1998). A smaller portion can be generated by magnetic reconnection and related effects in the atmospheres of disks, their “corona” (Noble & Krolik 2009). These two processes are responsible for most of the dissipation in the circumbinary disk and possibly in the mini-disks. In addition, however, in the context of binary accretion, shocks can contribute in several ways. There are the shocks already mentioned, where outward-moving streams strike the inner edge of the circumbinary disk; their luminosity has been previously discussed (Noble et al. 2012; Tang et al. 2018). There can also be shocks where inward-moving streams strike the outer edge of a mini-disk (Roedig et al. 2014; Farris et al. 2015b). If the mini-disks are sufficiently hot (sound speed c_s at least $\sim 0.1 v_{\text{orb}}$, for v_{orb} the speed of a circular orbit at the relevant location), tidal interactions and stream impacts can generate spiral shocks of substantial amplitude within the mini-disks (Ju et al. 2016; Bowen et al. 2017; Ryan & MacFadyen 2017).

Radiation can also be created in the course of the merger proper, but most calculations of it so far have been conducted only at the level of “proof of principle” (Bode et al. 2010, 2012; Palenzuela et al. 2010; Farris et al. 2011, 2012, 2014; Giacomazzo et al. 2012; Gold et al. 2014a, 2014b; Kelly et al. 2017), with Farris et al. (2012) being the first fully relativistic general relativistic MHD (GRMHD) simulation of accreting BBHs. These numerical relativity calculations did not develop mini-disks, possibly because they did not integrate long enough or they considered binaries at too close a separation, even though Nature would have had ample time to do so before the binary became so close.

More has been accomplished concerning the epoch of approach to merger. Using analytic estimates for disk structure, Roedig et al. (2014) argued that when the binary separation is at least several tens of gravitational radii ($r_g \equiv GM/c^2 = M$ for $G = c = 1$; here M is the total binary mass), but close enough that the circumbinary disk is able to radiate a significant luminosity (i.e., $a \lesssim 300M$), there should be a “notch” in the thermal spectrum due to the weakness of radiation from the accretion streams crossing the gap. This notch might appear anywhere from the near-IR to the near-UV, depending on parameters. They further predicted that in this phase of SMBBH evolution there should be a substantial hard X-ray

component due to Compton cooling of the gas shocked when an accretion stream strikes the outer edge of a mini-disk.

A number of papers adopting 2D hydrodynamics, and assuming both accretion stress and dissipation are described by a phenomenological “ α ” viscosity, have explored accretion onto mini-disks over long enough periods of time for the circumbinary disk to reach equilibrium out to large multiples of a (MacFadyen & Milosavljević 2008; Farris et al. 2014, 2015a; Ryan & MacFadyen 2017; Tang et al. 2018). In all but one of the papers of this group, the fluid dynamics were Newtonian, and took place in a Paczynski–Wiita potential; in Ryan & MacFadyen (2017), a single mini-disk was studied with 2D GR hydrodynamics in a Schwarzschild spacetime perturbed by Newtonian tidal gravity to approximate the influence of the other black hole (BH).

These simulations have yielded predictions of the emitted radiation, doing so by describing the cooling rate in terms of their α parameter combined with a disk dynamical temperature, defined by ignoring radiation pressure and scaled by an assumed disk Mach number (again, Ryan & MacFadyen 2017 is the exception: in this paper the surface brightness in the fluid rest-frame is defined to be the thermal rate at a temperature defined by the ratio of vertically integrated gas pressure to surface density, but divided by the vertical optical depth). A principal result of this series of papers is the prediction of a spectrum comprising three quasi-Planckian thermal peaks, one at $\simeq 1$ keV arising principally from the circumbinary disk, another (somewhat weaker) component at $\simeq 3$ keV emitted by the streams, and a third at $\simeq 20$ keV radiated by the mini-disks, but fading over time (Tang et al. 2018).

Recently, Bowen et al. (2017, 2018) reported the first simulations of mini-disk dynamics when the binary separation is small enough (a few tens of M) that the orbit evolves due to GW emission. Using a fully relativistic spacetime for a binary comprising an equal-mass pair of non-spinning BHs, they encountered several surprises. Because the relativistic gravitational potential between the two BHs becomes shallower than in the Newtonian regime, the mini-disks stretch toward the L1 point and the amount of gas passing—or “sloshing”—back and forth between them increases sharply when the separation is $\lesssim 30M$. The sloshing is quasi-periodically modulated at a frequency $\simeq 2\text{--}2.75 \Omega_{\text{bin}}$, where Ω_{bin} is the binary’s orbital frequency. Although tidal effects in Newtonian binaries are known to induce $m = 2$ spiral waves in mini-disks, the leading-order post-Newtonian (PN) terms induce strong $m = 1$ features. Perhaps most surprisingly, when the separation is as small as $20M$, the inflow time in the mini-disks is so short that their mass responds strongly to modulation of their supply rate on the binary orbital timescale.

1.3. Our Work

Here we will make use of the data reported in Bowen et al. (2018), produced using the HARM3D code (Noble et al. 2009), to make detailed predictions of both the spectrum and the time dependence of the light emitted. Because the HARM3D code is both intrinsically conservative and uses a local cooling function to radiate nearly all the heat produced, whether generated by turbulent dissipation or shocks, the luminosity we predict is automatically consistent with the energy budget of the flow. These predictions are, however, dependent upon two assumptions about the fluid-frame spectrum: where the gas is optically thick, we assume it radiates a local blackbody spectrum;

and where the gas is optically thin, we assume it radiates hard X-rays in a manner similar to active galactic nuclei (AGNs), emitting a thermal Compton spectrum with temperature $kT = 0.2m_e c^2 \simeq 100$ keV.

The remainder of this paper is organized as follows. In Section 2, we specify the means by which we determine the EM emission from the simulation data and transport it through the binary's spacetime to a simulated observer. Then in Section 3, we describe the results of our ray-tracing calculations. In Section 4 we discuss the implications of our findings, and summarize them in Section 5.

2. Methodology

Calculating the radiation observed at infinity produced by gas in the state determined by the simulation requires a number of steps, and the use of different codes and techniques. In Section 2.1 we briefly describe the numerical details and assumptions behind the simulation; further details are given in a separate paper focused on its analysis (Bowen et al. 2018). We then explain our model for the thermodynamics assumed in the simulation in Section 2.2 and provide the specifics of the radiative transfer solution in Section 2.3.

Space and time coordinates are reported in units of the total BBH mass, M , assuming geometrized units $G = c = 1$. Stated times of snapshots are elapsed times from the start of the simulation, generally quoted in terms of the initial binary orbital period $t_{\text{bin}} \simeq 600M$.

2.1. GRMHD Simulation Details

Our calculation assumes the existence of an accretion flow around an inspiraling equal-mass binary with an initial separation $a_0 = 20M$. In the immediate vicinity of each BH, the spacetime is well modeled as a boosted and perturbed BH spacetime (Detweiler 2005; Poisson 2005) in horizon-penetrating coordinates (Johnson-McDaniel et al. 2009). Elsewhere in the domain, the spacetime may be described using PN theory (Blanchet 2014). The PN solution is combined with BH perturbation theory via asymptotic matching in regions of common validity to produce a global analytic spacetime, which has been shown to reasonably satisfy the Einstein field equations and is described in full detail in Mundim et al. (2014) and Ireland et al. (2016). The entire spacetime is described in terms of PN harmonic (PNH) coordinates, which are also the coordinates used in our ray-tracing calculations. The PN description is sufficiently high-order to self-consistently contain both gravitational radiation and the consequent orbital evolution.

The MHD simulation was performed using the HARM3D code (Noble et al. 2009), which evolves the magnetized matter on a background spacetime (which can be, as it is here, time-dependent) through conservation of baryon number density, conservation of stress-energy, and the Maxwell induction equation (see Noble et al. 2009 for more details). In the initial state, a circumbinary disk (whose properties are taken from the quasi-steady state $t = 50,000M$ snapshot of the RunSS simulation described in Noble et al. 2012) occupies the region outside $\approx 2a$ from the center of mass, while identical mini-disks fill the BHs' Roche lobes, each threaded by a weak poloidal magnetic field (that is, poloidal with respect to the central BH and a polar axis parallel to the orbital axis). If there were no magnetic field, both disks would be close to dynamical

equilibrium. The magnetic field eventually destabilizes the disks through magnetic winding and the development of the MRI; the resulting magnetic stress helps the disks accrete and evolve. Details regarding the mini-disk construction can be found in Bowen et al. (2017, 2018).

2.2. Thermodynamic Model

Because the simulation's flux-conservative numerical methods ensure that all dissipated energy is captured and turned into heat, we are able to self-consistently predict the amount of light emitted by our simulation. If there were no losses, the heat retained by the fluid would continually add vertical pressure support in the disk and geometrically thicken it. The thermal energy would then be accreted into the BH or be carried out from the disk by a wind. By adding a loss term to the energy equation, we can create a more realistic structural model for the accretion flow (i.e., disks with constant aspect ratio supported by gas pressure) while also evaluating the amount of energy available for radiation.

This loss term mimics a bolometric (frequency-integrated) cooling rate and appears as a sink term in the gas's equation of motion: $\nabla_\lambda T^\lambda_\mu = -\mathcal{L}_c u_\mu$. Its recorded value serves as a bolometric source to our radiation transfer solution, which is performed as a post-processing step. This procedure contrasts with other post-processing methods in which the gas's emissivity is determined by the temperature found in the simulation without consideration of radiative losses (e.g., Bode et al. 2010; Farris et al. 2011; Kelly et al. 2017).

As in Noble et al. (2012), the cooling function is designed to cool the gas toward the initial entropy ($S_0 = 0.01$) at a rate per unit volume

$$\mathcal{L}_c = \frac{\rho\epsilon}{t_{\text{cool}}} \left(\frac{\Delta S}{S_0} + \left| \frac{\Delta S}{S_0} \right| \right)^{1/2}, \quad (1)$$

where ρ is the rest-mass density, ϵ is the specific internal energy, $\Delta S \equiv S - S_0$, and t_{cool} is the cooling timescale. Choosing the target entropy as its initial value allows us to measure how much total heat is produced by internal dissipation from the simulation's onset to its completion.

We set four distinct regions of cooling timescales following Bowen et al. (2017). In the outermost region, $r \geq 1.5a$, the gas is cooled over a time equal to the local Keplerian orbital period: $t_{\text{cool}} = 2\pi(r + M)^{3/2}/\sqrt{M}$. Here, r is the PNH radial coordinate distance from the center of mass (Blanchet 2014). In the vicinity of an individual BH of mass m_i , where $r_i \leq 0.45a$, the cooling time is set using the local Boyer-Lindquist coordinates: $t_{\text{cool}} = 2\pi r_{\text{BL}}^{3/2}/\sqrt{m_i}$. For the coordinate mappings relating the PNH coordinate system to the local Boyer-Lindquist coordinates, see Bowen et al. (2017). In the remaining portion of the simulation domain, between the mini-disks and the circumbinary disk, the cooling time is set to the value found at the inner edge of the outer region, $t_{\text{cool}}(r = 1.5a)$.

Sample density and cooling function distributions in the equatorial plane as well as in a poloidal slice through one of the BHs are shown in Figures 1 and 2, respectively. To avoid contamination from excessively bright regions caused by artificially high-entropy zones in floor states, our ray-tracing calculations neglect emission from fluid cells with $\rho < 10^{-4}$ in code units. This density scale was consistently identified throughout the simulation's duration to be where gas

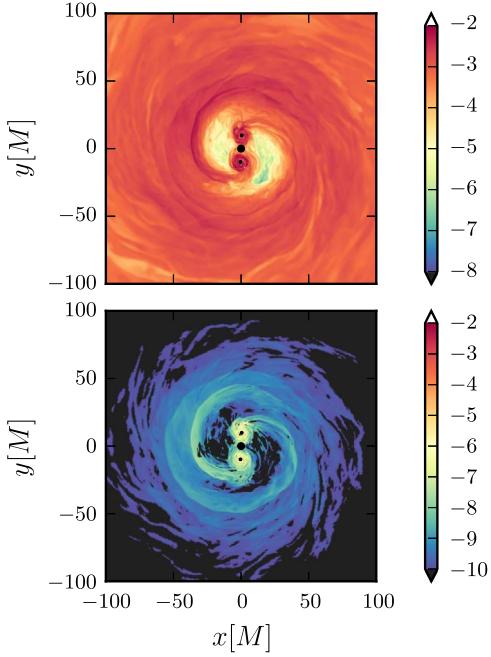


Figure 1. Snapshot at $t = 1030M$ of the rest-mass density ρ (top) and cooling function \mathcal{L}_c (bottom) in the equatorial plane of the HARM3D simulation, using a logarithmic color scale for each. The BHs' horizons are denoted as black circles displaced from the origin, while the black circle in the center represents the coordinate cutout at the origin. The horizontal and vertical coordinates are in the PNH Cartesian coordinates. We have set $\mathcal{L}_c = 0$ (black) where it is ignored in the radiative transfer calculation.

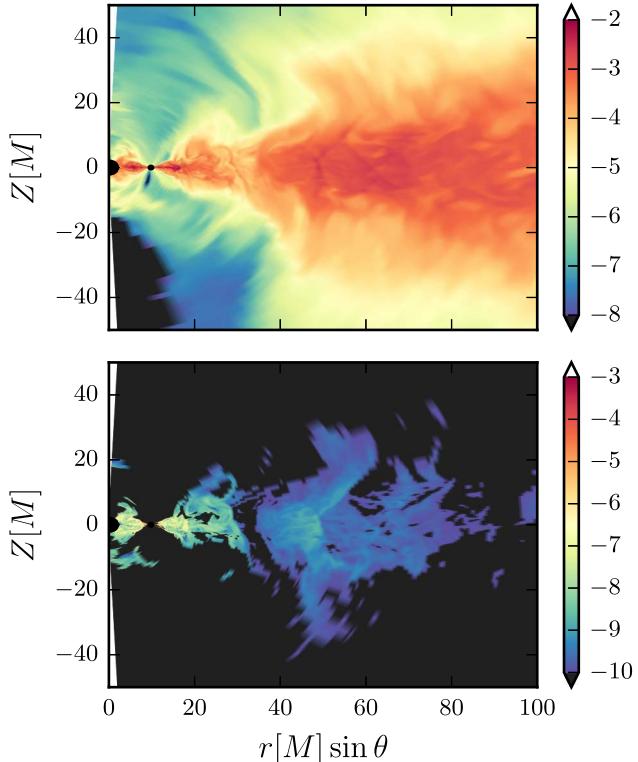


Figure 2. Snapshot at $t = 1030M$ of the rest-mass density ρ (top) and cooling function \mathcal{L}_c (bottom) in a poloidal plane through one of the BHs; the color scale is logarithmic for both quantities. The BHs' horizons are denoted as black circles displaced from the origin, while the black circle in the center represents the coordinate cutout at the origin. The horizontal and vertical coordinates are in the PNH Cartesian coordinates. We have set $\mathcal{L}_c = 0$ (black) where it is ignored in the radiative transfer calculation.

transitioned from a dense, turbulent state characteristic of a disk to a more laminar, tenuous flow characteristic of a corona. The “cutout” region covering the polar-coordinate origin at the center of mass appears in the center in black in the plots of Figures 1 and 2 (see Section 2.3.2 for further details).

2.3. Ray-tracing Method

After the simulation data have been generated, they are post-processed using a general relativistic ray-tracing code called BOTHROS (Noble et al. 2007). This code has been used to calculate the EM emission from a variety of single BH accretion simulations (Noble & Krolik 2009; Noble et al. 2009, 2011) and is used here for the first time in a dynamical spacetime. We provide a brief summary of the code before continuing with a description of the new aspects necessary for the work presented here.

BOTHROS allows a user to produce time- and frequency-dependent images of gas and is specifically tailored to systems including single and binary BHs. It approximates radiation as freely moving light rays—or null-like geodesics—within the system’s curved spacetime. The code uses an observer-to-source approach, shooting photons from a distant pinhole camera in various directions through the source volume. For a fixed camera location, tracing photons backward is advantageous computationally, as only the light rays received by the observer are calculated. Each ray that is launched ultimately contributes a spectrum, I_ν , to each pixel in the simulated camera. The camera can be positioned at any point in space, often specified in spherical coordinates $\{r_{\text{cam}}, \theta_{\text{cam}}, \phi_{\text{cam}}\}$. The integral of I_ν over the pixels produces the locally imaged flux spectrum $F_\nu(r_{\text{cam}}, \theta_{\text{cam}}, \phi_{\text{cam}})$.

If the integrated optical depth along a ray never reaches unity, the geodesic is terminated when it either exits the simulation domain or reaches a distance $r_i < 1.001 r_{\text{hor}i}$ from the i th BH with horizon radius $r_{\text{hor}i}$. Otherwise, the geodesic is terminated at the photosphere. From either kind of termination point, the radiation transfer equation is integrated along the geodesic in the opposite direction back to the camera. Its efficiency is large enough that, using 16 cores, BOTHROS is capable of producing images with resolution matching the simulation resolution in only a few minutes per frequency per simulation snapshot. The ability to process 3D time-dependent GRMHD data efficiently gives the code an advantage over others that rely on analytic models or angle-averaged data (e.g., Zhu et al. 2012) and that use more time-consuming (though more versatile) Monte Carlo approaches (Dolence et al. 2009; Schnittman & Krolik 2013; Schnittman et al. 2016).

2.3.1. Geodesic Calculation

Because of the time-dependent spacetime, a unique geodesic must be calculated for each pixel and snapshot. The field of view and the number of pixels determines the angular resolution of the resultant images and the initial conditions of the geodesics. The quantity r_{max} sets the extent of the field of view, so the angular field of view is therefore $\approx r_{\text{max}}/r_{\text{cam}}$. The simulation domain is a sphere of radius $260M$; depending on how much of the domain we want to ray-trace, we set r_{max} somewhere between $0M$ and $260M$.

To mimic an observer at infinity, the camera must be sufficiently distant for the spacetime to be nearly flat and for the rays shot out to be nearly parallel in the regions of interest. To

achieve these conditions, we place the camera at $r_{\text{cam}} = 1000M$ from the center of mass. In order to confirm that this is an appropriate choice, we have checked that $F_{\nu}(r_{\text{cam}}, \theta_{\text{cam}}, \phi_{\text{cam}})$ changes by less than 1% when moving r_{cam} from 10^3M to 10^5M . We chose the former setting for our image generation because it is computationally less demanding—a serious concern when processing $\mathcal{O}(10^4)$ snapshots, each with $\mathcal{O}(10^6)$ pixels.

We can adjust freely the polar (or inclination) angle θ_{cam} and the azimuthal angle of the camera ϕ_{cam} . The inclination angle is of crucial importance: viewing the SMBBH face-on ($\theta_{\text{cam}} = 0^\circ$) is qualitatively different from viewing it edge-on ($\theta_{\text{cam}} = 90^\circ$). As $\theta_{\text{cam}} \rightarrow 90^\circ$, images become more dependent on ϕ_{cam} (or the phase of the orbit) because relativistic beaming and double-lens effects introduce strong azimuthal dependence.

We use a Lagrangian approach to integrate the geodesic equation:

$$\frac{d^2x^\mu}{d\lambda^2} + \Gamma_{\alpha\beta}^\mu \frac{dx^\alpha}{d\lambda} \frac{dx^\beta}{d\lambda} = 0, \quad (2)$$

where $\Gamma_{\alpha\beta}^\mu$ are the Christoffel symbols, and λ is the affine parameter. In practice, Equation (2) is cast in first-order form, providing us with a set of eight ordinary differential equations to solve at each spacetime point for each ray, four equations each for the ray's position x^μ and the 4-velocity N^μ . However, one equation is eliminated because we find the time-component N^t from the null-like condition the ray must satisfy: $N^\mu N_\mu = 0$, choosing the positive N^t solution so that the ray points forward in time. The initial direction of the spatial components, N^i , is chosen so that the ray points toward the camera's focus at the camera's position.

In this work, we use BOTHROS in the fast-light approximation, which considers the simulation data frozen in time as the photon travels through it. The assumption simplifies the simulation data-processing, as it allows us to ray-trace one time-slice of data at a time. However, we still include the time dependence of the metric by including its time derivatives in Equation (2). The validity of this approach will be discussed in Section 4.4.

Because the camera is located at a distance $r_{\text{cam}} \gg a$ (where a is the binary separation), photons must travel large distances of empty (low emissivity and absorptance) space of little physical interest. For computational efficiency and without loss of accuracy, we may take larger steps through such regions than, for example, near the BHs by exploiting an adaptive stepsize control mechanism. The fifth-order Cash–Karp algorithm (Press et al. 1992) is used because of its high-order accuracy, stepsize adaptivity, and ability to handle stiff conditions. We permit a maximum relative error of 10^{-6} , which we have demonstrated—through a convergence study—yields $<1\%$ relative error in any quantities reported herein. Figure 3 illustrates the results of our geodesic calculations in the equatorial plane of the binary spacetime and gives some visual intuition of the metric's influence on the rays.

2.3.2. Processing the Simulation Data

The cells in our simulation are distributed nonuniformly in space using a time-dependent transformation between a uniformly discretized numerical coordinate system (x'^μ) —analogous to a 3D block of memory—and the actual spatial coordinates of the cells expressed in a “physical” spherical

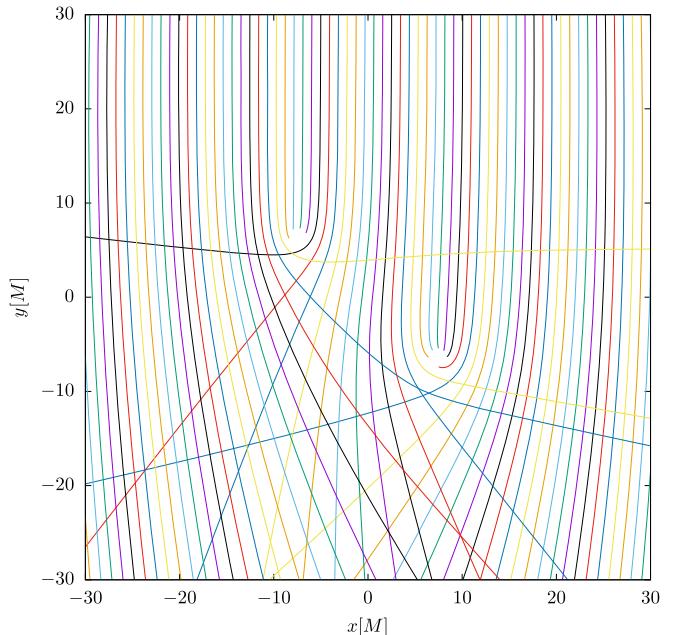


Figure 3. Representation of 100 geodesics in the equatorial plane, distinguished using color. The x and y coordinates shown are in the PNH Cartesian coordinate system.

coordinate system. The nonuniform system allows us to efficiently resolve small features near the BHs as well as the larger scale dynamics in the circumbinary disk. We call this system “warped” or dual-fisheye coordinates (Zilhão & Noble 2014). The simulation data are stored on disks as a set of files, one file for each time slice or snapshot. Each contains all the necessary 3D grid functions $(\rho, \mathcal{L}_c, v^i)$, where v^i are the 3-velocity components of the gas in the simulation's coordinates.

Because the geodesic integration may be done independently from any particular simulation's coordinate system and there are no symmetries to exploit, we chose to generate the geodesics in the coordinates used by our metric, namely the Cartesian PNH coordinate system (x^μ) . The radiation transfer equation, however, depends on both the geodesic information and the simulation data, so processing HARM3D simulation data requires us to transform between x^μ and x'^ν , a problem requiring the solution of a nonlinear algebraic set of equations. Grid function data are interpolated to points along the ray, and vector quantities are transformed to the PNH Cartesian basis in which the geodesics are expressed.

For each point along a ray, the interpolation proceeds by first converting the ray's coordinates into the simulation's coordinate system, so that we may efficiently look up the subset of data needed. Once the set of eight cells surrounding a geodesic point are found, the grid function values $(\rho, \mathcal{L}_c, v^i)$ from these cells are linearly interpolated to this geodesic point. To match the simulation's resolution, we add a point (if necessary) on the geodesic at each HARM3D cell encountered and interpolate in the same way. Sometimes the Cash–Karp algorithm used to integrate the radiation transfer equation requires values at intermediate points along the ray; we compute these by interpolating the stored values using the fourth-order Lagrangian method.

The simulation's grid excludes “cutout” or “excised” regions around the origin and z -axis in order to avoid coordinate singularities. Even though geodesics are free to move through

these cutout regions, no data are stored there. This means that they do not contribute to the ray's resultant flux or optical depth, so they are treated as vacua for the sake of the transfer integration. We note that rays traveling through the cutout may also travel through regions with gas and record a significant integrated opacity.

Once the GRMHD simulation data are read and interpolated, the 3-velocity (v^i) of the gas is used to calculate its 4-velocity, u^μ , in global PNH inertial coordinates. The fluid's 4-velocity allows us to calculate the fluid-frame frequency of the photon:

$$\nu = -k^\mu u_\mu, \quad (3)$$

where k^μ is the photon's 4-momentum. Because the spacetime and the fluid's velocity field are inhomogeneous and dynamic, a photon's locally measured frequency can vary significantly along its path.

2.3.3. Assigning Units

Lengths and times in the simulation are defined in units of M (with $G = c = 1$), but because we performed the GRMHD simulation in the Cowling approximation (neglecting the fluid's self-gravity), the unit of gas mass is undefined. This fact allows us to set the total mass M of the binary and the physical density (or mass) scale of the gas independently when converting simulation data from code units to physical units.

To define a physical lengthscale appropriate to a SMBH, we set $M = 10^6 M_\odot$. Instead of setting the gas density scale directly, we derive it from a more intuitive quantity, the accretion rate \dot{M} . To scale the accretion rate, we first calculate the accretion rate in code units:

$$\dot{M}(r, t) = -\int \rho u^r \sqrt{-g} d\theta d\phi, \quad (4)$$

where u^r is the radial component of the 4-velocity, and the integration is performed on spherical surfaces of fixed radius. In the circumbinary region, these spheres are centered on the center of mass and the radial component is that of the global inertial PNH basis. Once $\dot{M}(r, t)$ in code units is found, the density scale can be set by converting \dot{M} to cgs units as described in Appendix B. Because the simulation did not achieve inflow equilibrium in the mini-disks, while the inner portion of the circumbinary disk in RunSS did, we used the average \dot{M} over the radial range $2a < r < 4a$ in our initial condition for units definition. We scale this value to a fraction of the Eddington accretion rate in order to explore a range of optical thicknesses using the parameter $\dot{m} \equiv \dot{M}/\dot{M}_{\text{Edd}}$, where $\dot{M}_{\text{Edd}} = L_{\text{Edd}}/(\eta c^2)$ with nominal radiative efficiency $\eta = 0.1$ and $L_{\text{Edd}} = 1.2 \times 10^{38} M/M_\odot \text{ erg s}^{-1}$. We will start by studying high accretion rate flows ($\dot{m} = 0.5$) in Section 3.2, then move to low accretion rate flows ($\dot{m} = 8 \times 10^{-4}$) in Section 3.3.

The absence of inflow equilibrium in the mini-disks affects some of our results; its implications will be discussed in Section 4.4. We completely ignore emission from $r > 150M$, where the circumbinary disk is also out of inflow equilibrium, but the neglected luminosity is much smaller.

In order to faithfully recover the flux of a simulated pointing, we need to ensure that all parts of the simulation data are adequately sampled by the rays cast through it. This means that a given snapshot's flux, integrated over all the pixels, must be converged

with respect to the number of pixels used for a fixed field of view. We have found that a resolution of ~ 7 pixels/ M for the $r < 60M$ region and a coarser resolution of ~ 2 pixels/ M for the $60M < r < 150M$ region is necessary to compute the flux (at a given time and frequency) to $\sim 1\%$ accuracy.

2.3.4. Radiation Transfer Solution

The specific intensity measured at the camera is found by integrating the Lorentz-invariant form of the transfer equation along a geodesic:

$$\frac{\partial I}{\partial \lambda} = j - \alpha I, \quad (5)$$

where λ is the affine parameter, and I, j , and α are, respectively, the Lorentz invariant intensity, emissivity, and absorption coefficient. See Appendix A for details on the choice of affine parameter and a derivation of this equation.

The light coming from an optically thick medium is effectively radiated from its photosphere (the surface at which its optical depth passes through unity). If a geodesic encounters a photosphere, we integrate the transfer equation from there to the camera with initial condition $I_0 = I_{\text{photosphere}}$; otherwise the integration starts from the end of the geodesic with $I_0 = 0$. In other words, we have to integrate the equation only in optically thin regions.

In order to create spectra, we vary the frequency in cgs units at the camera, ν_∞ . At each point X along the geodesic, we calculate the corresponding Doppler and gravitationally shifted frequency ν_X via

$$\nu|_X = -F(\nu_\infty)(k^\mu u_\mu)|_X, \quad (6)$$

where $F(\nu_\infty)$ is the conversion factor from numerical units to cgs units (see Appendix B), k^μ is the 4-velocity of the photon, and u^μ is the 4-velocity of the gas in global inertial coordinates. Given a fluid-frame model for j_ν and α_ν , we can then find $j(X)$ and $\alpha(X)$ from Equation (15) to discretely integrate Equation (5). The integration is performed using the same fifth-order Cash–Karp algorithm used to integrate the geodesic equations.

2.3.5. Opacity Model

The radiative model chosen in a region strongly depends on whether the gas is thermalized. Gas inside a disk can generally be considered thermalized if the vertically integrated *effective* optical depth of the disk, $\tau_{\text{eff}} \sim \sqrt{\tau_a(\tau_a + \tau_s)}$, is much larger than unity, where τ_a and τ_s are the optical depths from absorptive and scattering processes, respectively. At the densities and temperatures of interest here, the dominant source of opacity is electron scattering, so we neglect absorptive processes. Even though we do not calculate τ_a and, rigorously speaking, $\tau = \tau_a + \tau_s \gg 1$ does not ensure thermalization, we make the reasonable assumption that there is still enough absorption in the disks for them to be thermalized.

We therefore assume a gray (frequency-independent) Thomson opacity for electron scattering:

$$\alpha_\nu = \kappa_T \rho, \quad (7)$$

where ρ is the gas density and $\kappa_T = \sigma_T/m_H$ is the Thomson opacity (obtained by dividing the Thomson cross section by the

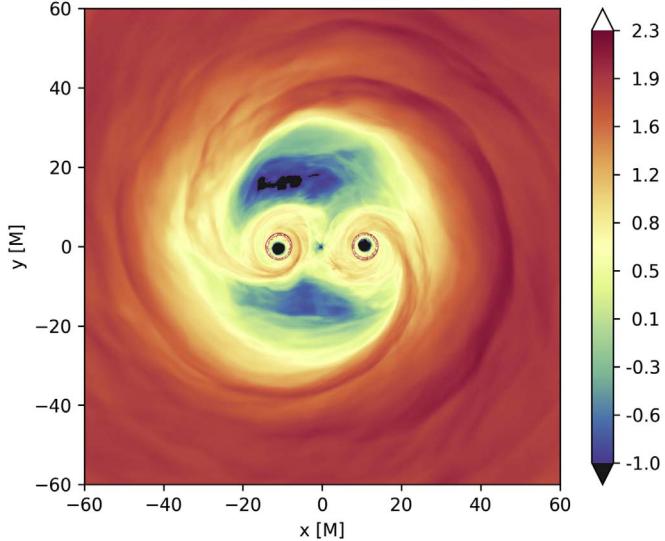


Figure 4. \log_{10} of the optical depth at $\dot{m} = 0.5$, $r_{\max} = 60M$, $\theta_{\text{cam}} = 0^\circ$, and $t = 1180M \simeq 2t_{\text{bin}}$. The rings of large optical depth circling the horizons correspond to the photon spheres, where the geodesics wrap around the BHs multiple times, accumulating extra optical depth.

mass of a hydrogen atom). Each segment along the ray contributes a Lorentz-invariant optical depth differential equal to

$$d\tau = \alpha_\nu \, ds = \alpha \, d\lambda. \quad (8)$$

Here we are interested in the optical depth between the observer and the material we simulated, so we integrate Equation (8) from the camera to the source along the geodesic.

In Figure 4, we show the calculated optical depth in the high accretion rate case at inclination $\theta_{\text{cam}} = 0^\circ$, which effectively corresponds to a vertical integration. We see that the image bifurcates neatly into two zones: that whose geodesics encounter so much gas (in mini-disks, accretion streams, to the circumbinary disk) that $\tau \gg 1$, and that whose geodesics traverse only cavities, so that $\tau \ll 1$ even after integrating to the far end of the geodesic.

2.3.6. Emissivity Model

Below the photosphere (the $\tau = 1$ surface), we assume the disk's gas is in thermal equilibrium. We therefore initialize the specific intensity at the photosphere with a blackbody spectrum,

$$I_\nu = B_\nu(\nu, T_{\text{eff}}) = \frac{2h\nu^3}{c^2} \frac{1}{e^{\frac{h\nu}{kT_{\text{eff}}}} - 1}. \quad (9)$$

The effective temperature, T_{eff} , is the temperature associated with the local radiative cooling flux (\mathcal{F}) at the photosphere and can be found using the Stefan–Boltzmann law:

$$T_{\text{eff}} = (\mathcal{F}/\sigma)^{1/4}, \quad (10)$$

where σ is the Stefan–Boltzmann constant. The flux is found by integrating the cooling function vertically inside the photosphere:

$$\mathcal{F} = \frac{1}{2} \int_{\tau>1} \mathcal{L}_c \, ds = \frac{1}{2} \int_{\tau>1} \mathcal{L}_c \, \nu \, d\lambda, \quad (11)$$

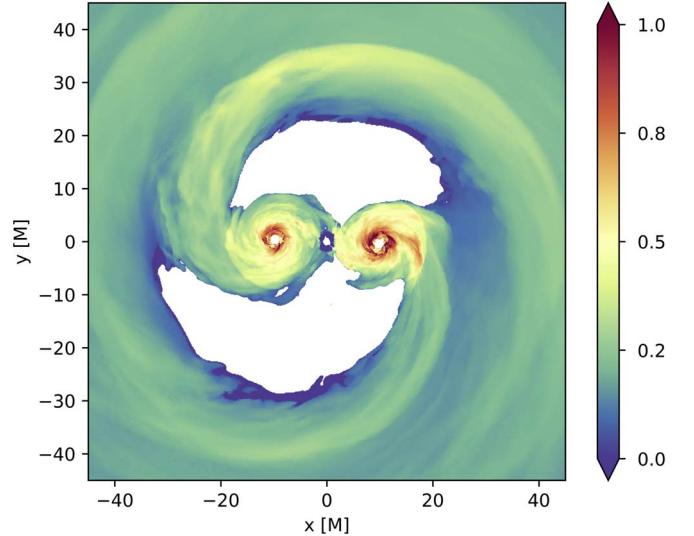


Figure 5. Effective temperature on the photosphere, projected into the binary's orbital plane, at $\dot{m} = 0.5$, $t = 1030M$. $\log_{10}(T_{\text{eff}}/T_0)$ in the fluid frame is shown, where $T_0 = 5 \times 10^5$ K. The effective temperature at infinity is altered by gravitational redshift and Doppler-boosting; the former dominates for face-on views, so the observed effective temperature seen at $\theta_{\text{cam}} \simeq 0^\circ$ would be rather lower near the BHs than shown here. Uncolored (white) areas within the cavity region lack the opacity necessary to surpass the photosphere criterion.

where the factor of $1/2$ comes from the fact that the disk has two photospheric surfaces from which heat can escape.⁶ This integral is approximately vertical through the disk for views with $\theta_{\text{cam}} \simeq 0^\circ$; we restrict our exploration of optically thick models to this viewing angle because applying our ray-tracing method at other viewing angles would locate the photosphere at an artificially high altitude from the disk midplane. This is because the disk is stratified vertically and extends indefinitely radially outward. Rays traveling horizontally pass through more material before reaching the actual photosphere, which is the surface at which the gas becomes transparent in any direction—not necessarily the same as the ray's but one that is almost always vertically outward due to the gradient of the disk's stratification. Figure 5 gives a representation of the photosphere and the effective temperature at its surface as viewed face-on.

In evaluating this image, it should be recognized that in those regions where the Thomson optical depth is $\sim 1\text{--}10$ —most of the mini-disks' surface area and part of the accretion streams—a substantial part of the dissipation may take place in regions that actually lie outside the thermalization photosphere. Consequently, our approximation overestimates the thermal luminosity and underestimates the luminosity arising from non-thermalized regions.

Outside thermalized regions, the predominant radiative process is inverse Compton scattering. In such regions, the dimensionless temperature $\Theta \equiv kT/m_e c^2$ rarely exceeds $\simeq 0.2$ because the Compton scattering energy-loss rate increases as electrons become relativistic and further plasma cooling can be accomplished by pair production (see the pedagogical review in

⁶ Even though we locate the $\tau = 1$ surface as the photosphere, we ignore it if the total optical depth along the ray, τ_{tot} , is $\tau_{\text{tot}} < 2$. This condition ensures that the “top” photospheric surface (i.e., the one found by integrating through the disk from above) lies above the “bottom” photospheric surface (i.e., the one found by integrating through the disk from below).

Krolik 1999). When the optical depth is small, the energy spectrum of photons is exponentially cut off above Θ and follows a power law below Θ . In observed AGNs, the power-law (for intensity) varies from $\simeq -0.3$ to $\simeq -1.3$ (Trakhtenbrot et al. 2017); for simplicity, we set it to -0.5 . We then take the emissivity to be

$$j_\nu \propto \mathcal{W}_\nu = \left(\frac{h\nu}{kT}\right)^{-1/2} e^{-\frac{h\nu}{kT}} \quad (12)$$

with $kT = 100$ keV, or $\Theta \simeq 0.2$. We normalize the spectrum in such a way that the bolometric emissivity matches the cooling function at every point: $\int j_\nu d\nu d\Omega = \mathcal{L}_c$. This gives:

$$j_\nu = \frac{\mathcal{L}_c}{4\pi A} \mathcal{W}_\nu, \quad (13)$$

$$A = \int d\nu \mathcal{W}_\nu = \frac{kT}{h} \sqrt{\pi}. \quad (14)$$

In these optically thin regions, the transfer equation is integrated starting with I_ν of the disk at the photosphere (Equation (9)) or zero if this geodesic does not encounter a photosphere. For j_ν and α_ν , it uses the emissivity (Equation (13)) and the scattering opacity (Equation (7)).

3. Results

As mentioned in Section 2.3.3, we chose a total system mass $M = 10^6 M_\odot$ to define the length scale and made two choices of gas density scale through two choices of accretion rate in Eddington units: $\dot{m} = 0.5$ (high, Section 3.2) and $\dot{m} = 8 \times 10^{-4}$ (low, Section 3.3).

3.1. Optical Depth Images

Thomson optical depth maps for $\dot{m} = 0.5$ at a variety of polar angles and times illustrate the basic geometry of the system (Figure 6); these images can be readily scaled to other accretion rates because they are linearly proportional to \dot{m} . The 16 panels show snapshots seen from four different polar angles at four equally spaced times spanning $150M$, a little more than a quarter of a binary orbit.

The face-on view (0° inclination) provides an approximate view of the surface density of the gas. The circumbinary disk is generally quite optically thick ($\tau \gtrsim 50$), especially in the overdensity or “lump” feature near its inner edge (Noble et al. 2012; Shi et al. 2012). As is usual for disks around binaries with order-unity mass-ratios, the region within $\simeq 2a$ of the center of mass has very low density except in a pair of spiral streams and in a pair of mini-disks, one surrounding each member of the binary. However, there is also a high optical depth ring around each BH. Rays reaching us from this close to a BH wrap around it several times before heading off to infinity, acquiring additional optical depth by traversing extra path length. The characteristic magnitude of the optical depth increases with inclination, reaching $\simeq 600$ – 900 at $\theta_{\text{cam}} = 90^\circ$, as the path through the BBH system is longer by a factor of $\sin(\theta_{\text{cam}})^{-1}$. Note, however, that the optical depth we measure at edge-on views is not meaningful because real rays would traverse parts of the disk well outside our simulation domain.

At intermediate viewing angles (e.g., the 39° inclination shown in the second row of Figure 6), the optical depth images still show the circumbinary disk geometry clearly. However, at large viewing angles ($\gtrsim 71^\circ$ inclination, bottom two rows of

Figure 6), gravitational effects distort the image very strongly. There is a region of high optical depth below the BHs, where the photons travel through the circumbinary disk twice: starting above the circumbinary disk on the far side of the BHs, they pass downward behind the BHs through the disk, curve through the cavity underneath, and are finally gravitationally pulled upward (by the BHs) and traverse a second time through the circumbinary disk toward the camera. The low optical depth region above the BHs arises from those photons that travel over the BHs and the circumbinary disk and then curve upward through the cavity, avoiding the dense gas found in the disks altogether.

Subtler relativistic features also appear at large viewing angles. A mass moving along the line of sight to an observer creates a gravitationally lensed image of a source on its far side that is smaller than its true size if the BH approaches the observer (Heyrovský 2005). More exotic but perhaps less apparent, the appearance of a secondary image of one BH on the opposite side of the other BH can be noticed; this image is due to the extreme light deflection close to the horizons. A good example can be seen in the third frames of the bottom two rows of Figure 6, where a small oval feature forms on the left side of the BH on the left. Regrettably, it is highly unlikely that any of these imaging features will be spatially resolvable in the foreseeable future.

3.2. High Accretion Rate

Our high accretion rate case, $\dot{m} = 0.5$, is designed to demonstrate the interplay between optically thick and optically thin regions. In the former, dissipated heat emerges in a thermalized spectrum; in the latter, it is radiated by inverse Compton scattering from a very hot electron population.

Source-integrated features such as spectra may be observable soon. To compute spectra from our data, we take two time-averages of the flux from face-on viewing angles, one over the second binary orbit, the other over the third. The averaging suppresses statistical fluctuations. We prevent blurring of the different components of the disk in the images by rotating the camera orientation ϕ_{cam} at the binary orbital frequency. We choose a face-on view at which our criterion for distinguishing thermalized from coronal regions, namely whether the ray optical depth is greater or less than unity, is well justified. At higher inclinations it becomes increasingly suspect for two reasons. First, the optical depth unity point on a geodesic is found above the actual photosphere because the path-length is $\propto \sec \theta_{\text{cam}}$. Second, actual systems viewed edge-on may be obscured by material at distances not included in our simulation.

Figure 7 shows the spectral luminosity density in this face-on view. We define it as $L_\nu = 4\pi r_{\text{cam}}^2 \int I_\nu \cos(\psi) d\Omega$, where ψ is the angle between the geodesic's direction at the camera and the line of sight to the center of mass. Images of the system surface brightness at various frequencies in the face-on view are shown in Figure 8.

Like classical AGN spectra, this spectrum can also be described in terms of two components: a thermal UV/soft X-ray portion and a coronal hard X-ray spectrum. The thermal UV originates from the photospheres of the system; the hard X-rays are emitted in optically thin regions, predominantly on the top and bottom surfaces of the disks. All three locales—the circumbinary disk, the accretion streams, and the mini-disks—contribute to both the thermal and coronal spectral components.

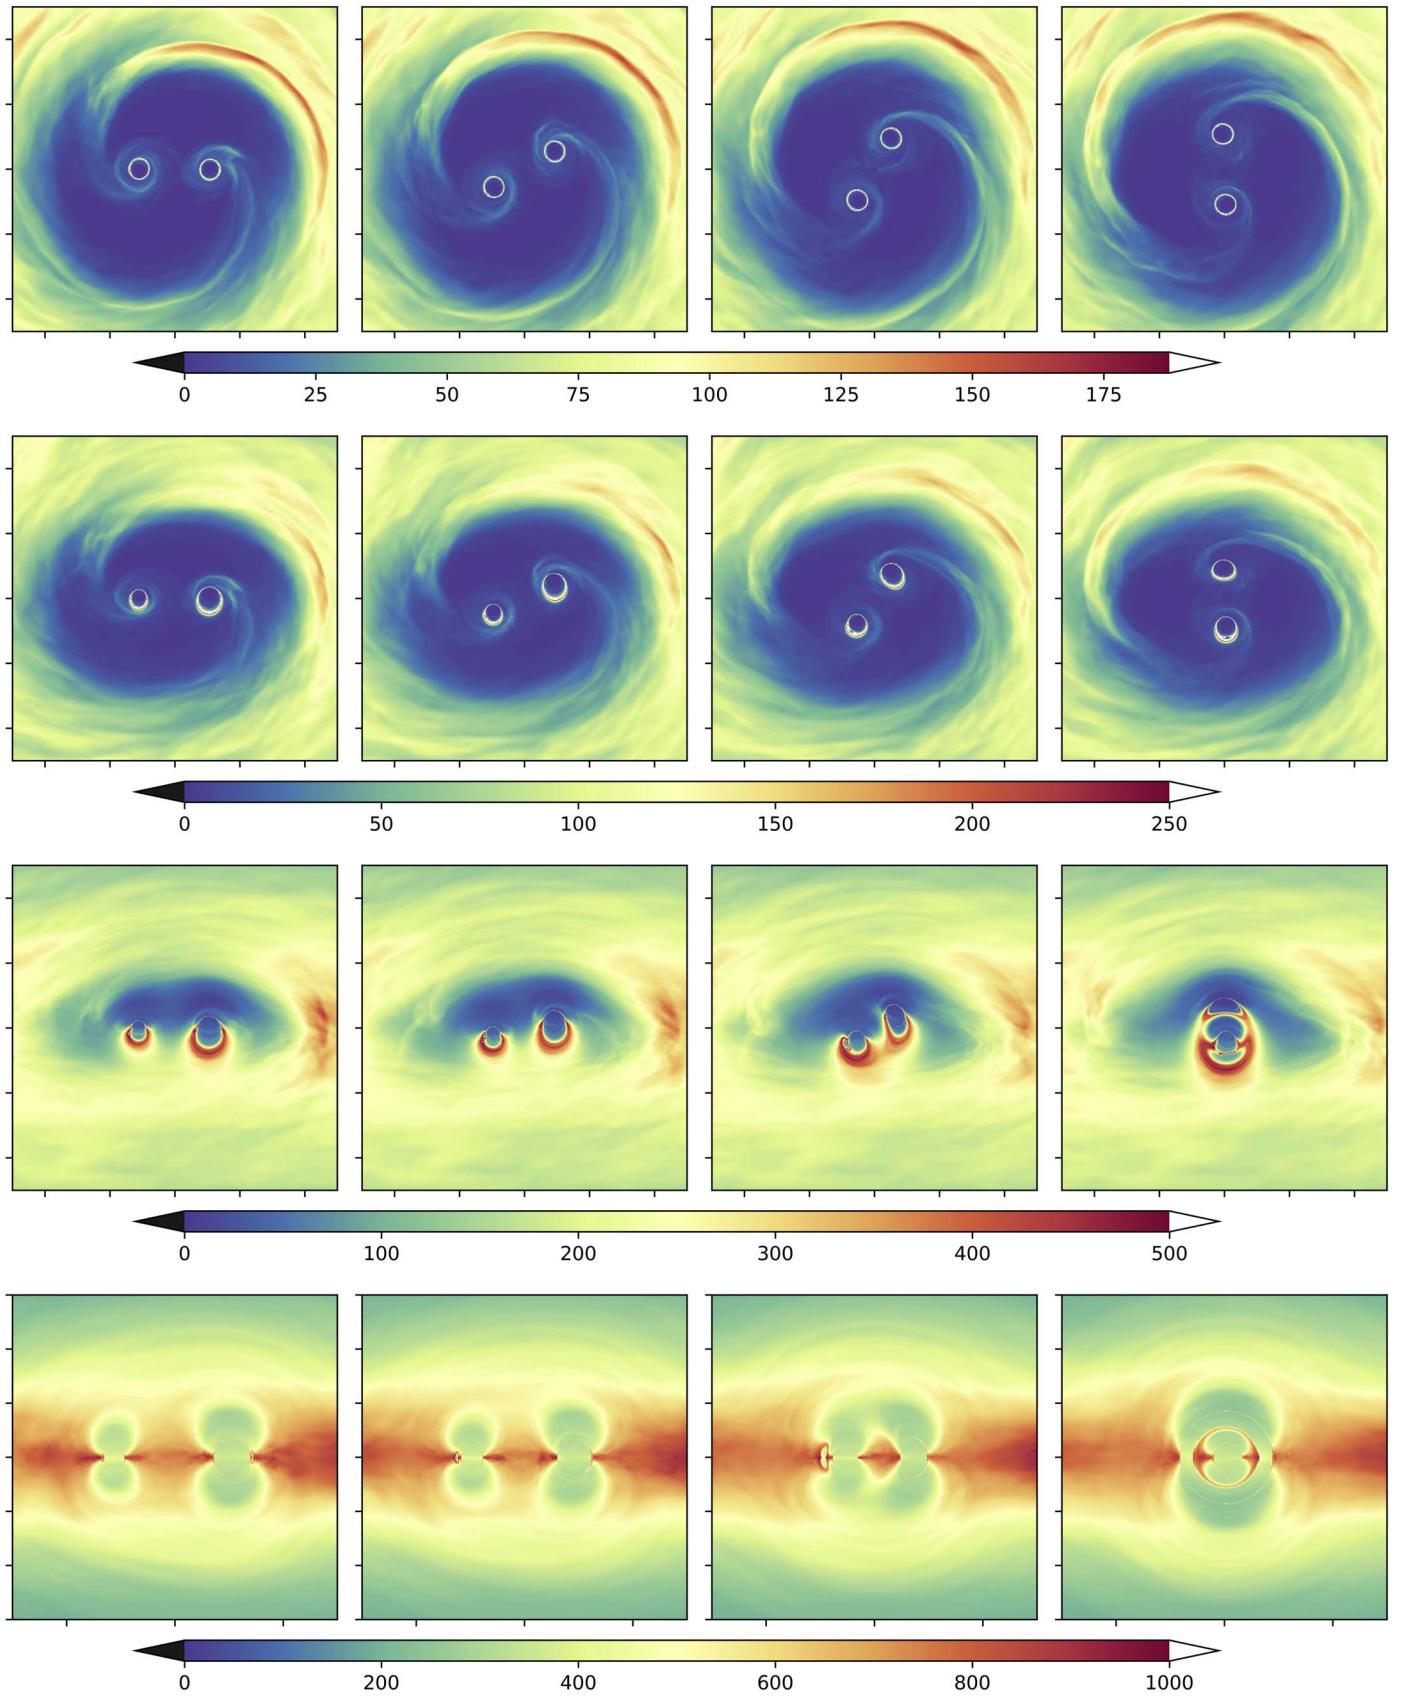


Figure 6. Optical depth integrated along the line of sight at $\theta_{\text{cam}} = 0^\circ, 39^\circ, 71^\circ$, and 90° from top to bottom, with $\dot{m} = 0.5$. Snapshot times are $1030M, 1080M, 1130M$, and $1180M$ from left to right. The outermost radius shown is $50M$, except at $\theta_{\text{cam}} = 90^\circ$, where it is set to $30M$ to focus on the complex relativistic effects at play. The space between tickmarks on all the axes is $20M$.

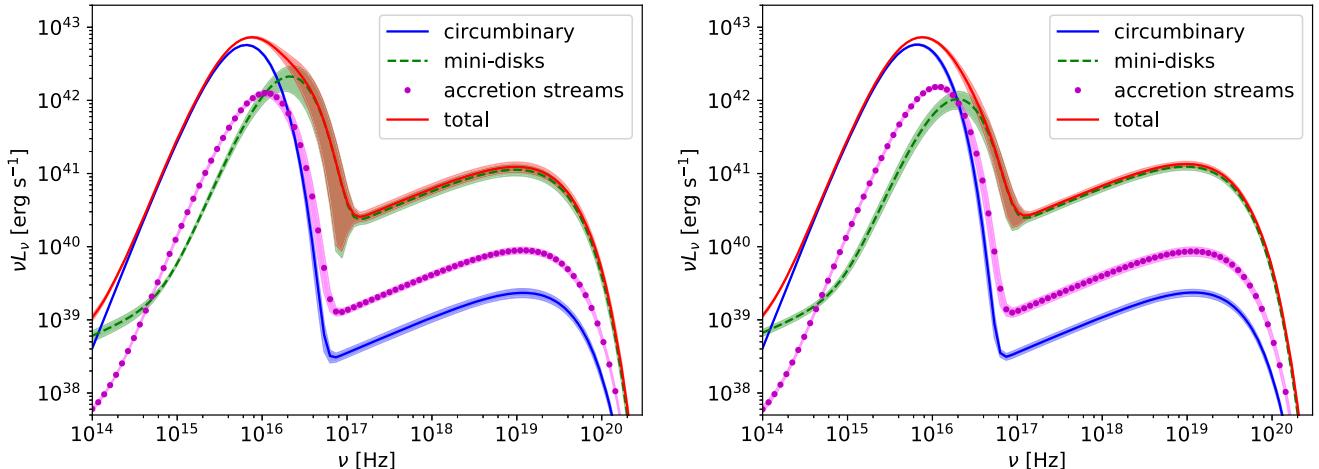


Figure 7. Time-averaged luminosity (νL_ν) spectrum obtained at $\theta_{\text{cam}} = 0^\circ$ and $r_{\text{cam}} = 1000M$ with $\dot{m} = 0.5$ using simulation data from the second orbit (left) and the third orbit (right). We have separated contributions from the mini-disk regions ($r < a$), the accretion streams ($a < r < 2a$) and the circumbinary region ($r > 2a$). The shaded region around each curve represents the temporal variability of each component, one standard deviation above and below, using a cadence of $10M$ for each orbit (60 samples). The cusps on the lower side of the mini-disks' shaded region in the left-hand panel represent points lying off the scale because the mini-disks' thermal emission fluctuates by an order of magnitude in the second orbit.

However, as remarked in Section 2.3.6, our assignment of all dissipation below the unit optical depth surface tends to transfer power from the Comptonized hard X-ray component to the thermal component, and this effect is particularly strong in the mini-disks and the accretion streams.

The emitted power is dominated by the thermal UV, with only $\sim 1\%$ radiated in hard X-rays; this ratio may, however, be exaggerated by our simple emission model. More surprisingly, the single greatest contribution ($\approx 65\%$) comes from the circumbinary disk rather than the mini-disks. Because we simulate a binary whose separation is only $20M$, the binding energy of an orbit at $2a = 40M$ is 0.0125 in rest-mass terms; this is more than half the effective radiative efficiency found in our simulation. The mini-disks account for most of the remainder (25%) in the second orbit, but share the luminosity almost evenly with the accretion streams in the third orbit. The mini-disks are less luminous than would be expected for the several reasons enumerated in the previous paragraph.

It is unsurprising that the thermal peaks from the three regions should be found at frequencies that increase gradually from the circumbinary disk to the accretion streams to the mini-disks. In time-steady ordinary accretion disks, the effective temperature is $\propto r^{-3/4} R_R^{1/4}(r)$, where R_R is a correction factor accounting for the net angular momentum flux and relativistic corrections. This relation might be a reasonable approximation within both the mini-disks and the circumbinary disk if r is defined as the distance to the near BH in the mini-disks and the distance to the center of mass in the circumbinary disk. However, the “notch” separating the circumbinary disk and mini-disk thermal spectra predicted by Roedig et al. (2014) is not apparent. This can likely be attributed to the comparative faintness of the mini-disks in a system with binary separation as small as the one we have analyzed (see Section 4.1 for the arguments supporting this contention).

A more detailed analysis of where different frequencies are radiated is aided by the images of Figure 8. The UV surface brightness (first panel) in optically thick regions varies hardly at all from the circumbinary disk to the accretion streams to the mini-disks, but the larger area of the circumbinary disk makes it the primary contributor to the luminosity in this band.

However, because the mini-disks are warmer than the circumbinary disk, their thermal spectrum remains bright farther into the extreme-UV (second panel). This image further reveals that, especially in the third orbit, a sizable part of the dissipation occurring in the mini-disks takes place in spiral shocks. In the soft X-ray band (third panel), the principal contributor is the extreme Wien tail of the thermal emission from the mini-disks. Finally, in the hard X-ray band (fourth panel), the emission is dominated by the optically thin component in the corona, which is strongly concentrated in the innermost rings of the mini-disks. Again, we caution that the mini-disk radiation produced in our model may overestimate the thermal component at the expense of the Comptonized X-rays.

The first two panels in Figure 8 also show that nearly all the light attributed to the accretion streams in Figure 7 is associated with the shock that occurs when the accretion stream, having been strongly torqued by the binary's gravity, is flung outward and strikes the inner edge of the circumbinary disk.

A number of these comments are in agreement with the 2D “ α -viscosity” hydrodynamics simulations of Farris et al. (2015b). They, too, found enhanced emission due to shocks between the accretion streams and both the inner edge of the circumbinary disk and the outer edges of the mini-disks. In addition, because they assumed all radiation was thermal, they placed all this light in the UV/extreme-UV. However, they also found the unshocked streams had high surface brightness, a result attributable to the “ α -viscosity” creating dissipation even in laminar regions if they contained significant shear. In addition, their separation of “mini-disks” from “cavity” from “circumbinary disk” is different from ours, so the separate luminosity contributions cannot be directly compared. Although the images of surface brightness are qualitatively similar to those of Tang et al. (2018) at times well before merger in their simulations, and the mass of the system they simulate is only twice that of ours, the effective temperatures Tang et al. (2018) find (~ 1 keV in the circumbinary disk, tens of keV in the mini-disks) are much higher than ours. It is possible this contrast arises because their simulation treats a case with an accretion rate $\gtrsim 10^4 \times$ ours (in order to support

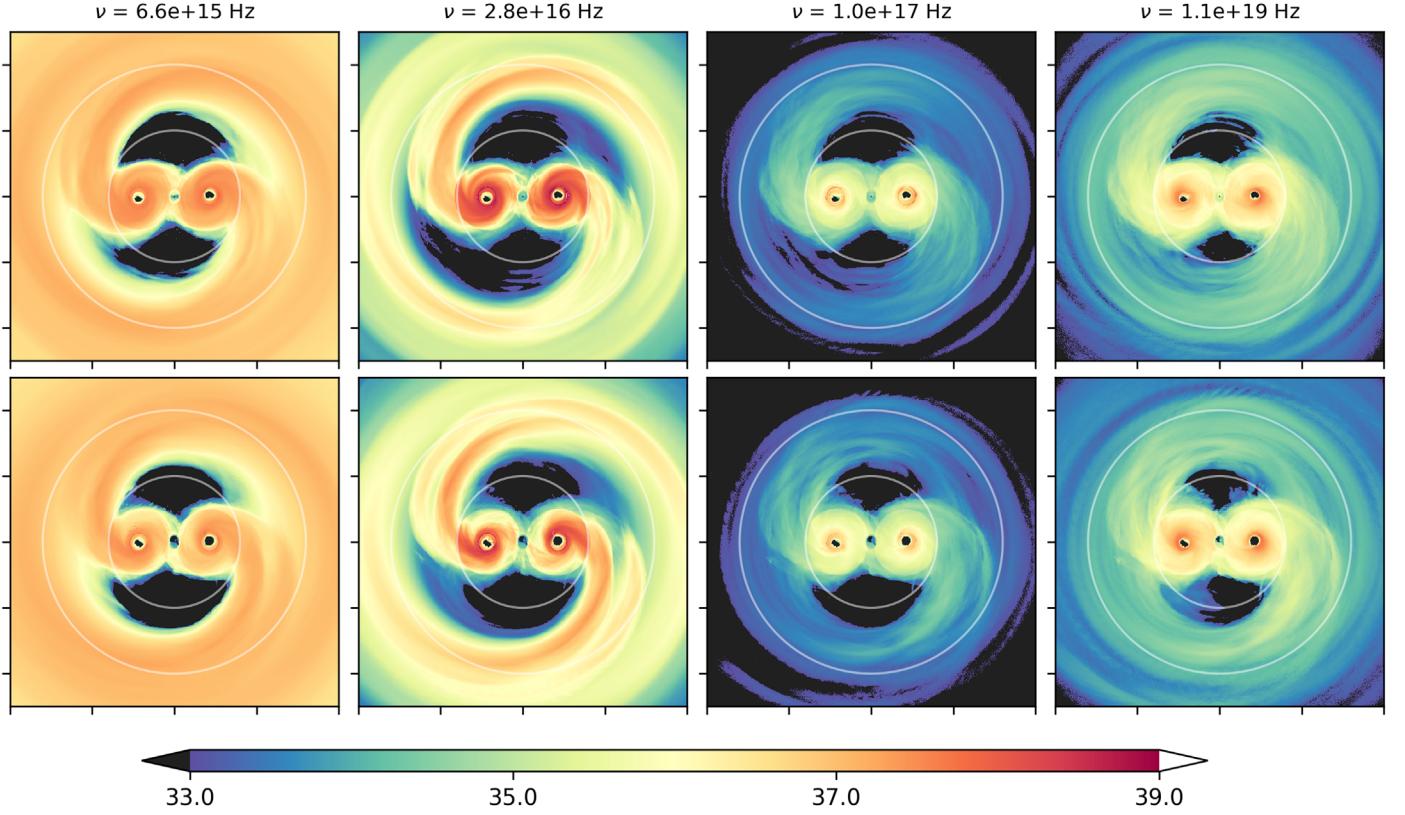


Figure 8. Images of time-averaged spectral power ($4\pi r_{\text{cam}}^2 \nu I_\nu$) [erg s $^{-1}$] at various frequencies with $\dot{m} = 0.5$, showing the transition from the blackbody-dominated regime to the inverse Compton scattering-dominated regime. The time averages were performed over the second orbit (top row) and the third orbit (bottom row) separately. From left to right, we encounter (i) the circumbinary dominated UV, (ii) mini-disk-dominated soft X-rays, (iii) X-rays near the boundary between thermal and corona dominance, and (iv) the mini-disk corona-dominated hard X-rays. The two white circles in each panel mark $r = a$ and $r = 2a$. The space between tickmarks on all the axes is $20M$.

temperatures $\gtrsim 10 \times$ higher on a similar thermally radiating surface area), but they explicitly state neither the accretion rate they found nor the magnitude of their initial gas surface density.

3.3. Low Accretion Rate

As explained earlier, we have complete freedom to choose the inclination of the camera only if the whole domain is optically thin. To enable further study of inclination effects, we also studied a low accretion rate, $\dot{m} = 8 \times 10^{-4}$. Because the density is $625 \times$ lower in this case than in the high accretion rate case, the system stays optically thin along all directions except those within $\approx 20^\circ$ of the orbital plane. For high inclination angles ($\theta_{\text{cam}} \gtrsim 45^\circ$), the camera's azimuthal angle also becomes important. Note that here the azimuthal angle is measured in the corotating frame of the binary, with $\phi_{\text{cam}} = 90^\circ$ when the camera is aligned with the two BHs. If, as is more likely in reality, the camera is stationary in an inertial frame, the camera's effective azimuthal angle varies periodically as the binary members orbit.

We first display images of very hard X-ray (10^{19} Hz, ≈ 400 keV) intensity at three different inclination angles, each viewed at four different times spread over a quarter-period (Figure 9). These images somewhat resemble those of optical depth presented in Section 3.1 because the cooling function, like the scattering coefficient, is proportional to the density. However, while the optical depth of the inner circumbinary disk is generally $\approx 5 \times$ the optical depth of the mini-disks, the

X-ray surface brightness of the mini-disks is $\approx 5 \times$ that of the circumbinary disk. This is, of course, because the rate of energy dissipation per unit mass is much higher, so close to the BHs.

The spectrum of the optically thin case looks very much like the high-energy ($\nu > 3 \times 10^{16}$ Hz) portion of the spectrum shown in Figure 7 because we assumed that the fluid-frame coronal spectrum is the same everywhere. Relativistic effects, however, lead to significant viewing-angle dependence.

Although the shape of the spectrum depends only very weakly on viewing angle, relativistic effects alter the angular distribution of its intensity strongly enough to be observationally interesting. Figures 10 and 11 display how the observed bolometric flux depends on both the polar and azimuthal angle of an observer. We show only bolometric luminosity because the spectral shape is much less sensitive to viewing angle than its overall level.

At most inclinations ($\theta_{\text{cam}} \lesssim 70^\circ$), the flux is almost independent of azimuthal angle, but increases with inclination due to relativistic beaming of light emitted by gas moving toward the observer. However, for nearly edge-on viewing angles ($\theta_{\text{cam}} \gtrsim 70^\circ$), the angular dependence becomes more complex due to three effects. First, lensing of the farther BH by the nearer one when $\phi_{\text{cam}} \approx 90^\circ$ or 270° significantly brightens the image; the peak flux can be a factor of 3 higher than at other azimuthal angles with the same inclination. Second, relativistic beaming increases sharply with greater polar angle. Third, even for \dot{m} as low as 8×10^{-4} , the circumbinary disk can intercept light along rays passing close to the orbital plane. Smaller values of \dot{m} would diminish the range of angles around the

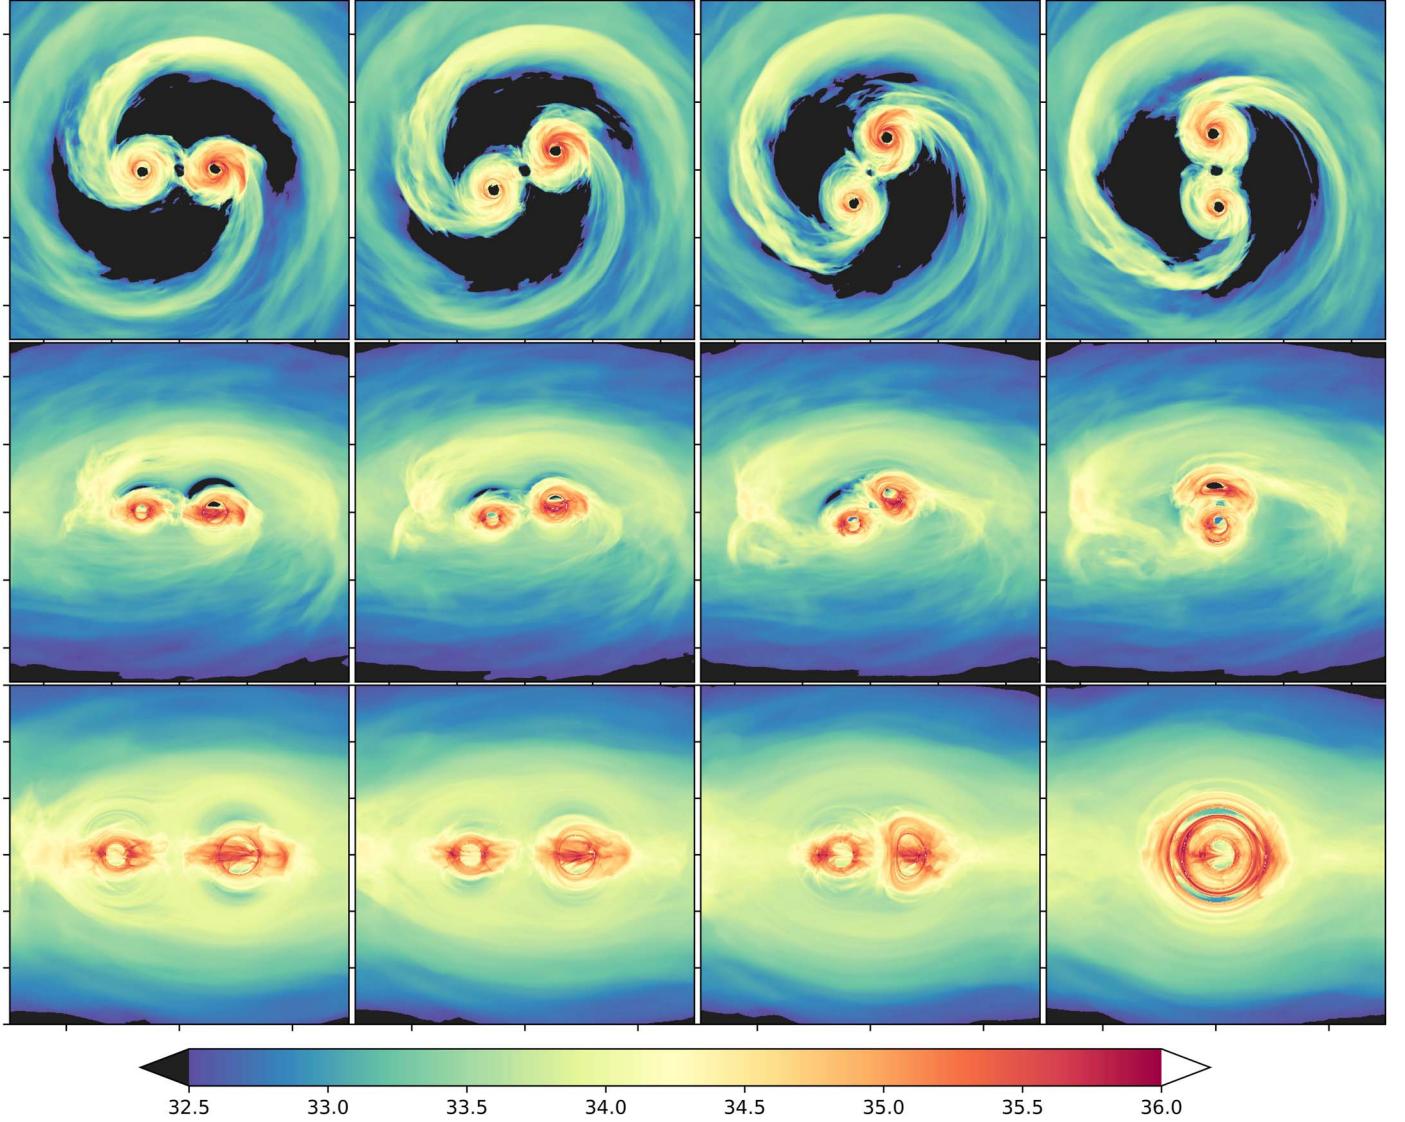


Figure 9. \log_{10} of spectral power ($4\pi r_{\text{cam}}^2 \nu I_\nu$) [erg s^{-1}] at $\nu = 10^{19} \text{ Hz}$, $\theta_{\text{cam}} = 0^\circ, 71^\circ, 90^\circ$ (from top to bottom) and $t = 1030M, 1080M, 1130M, 1180M$ (from left to right), with $\dot{m} = 8 \times 10^{-4}$. Again, the width of each image is $50M$, except at $\theta_{\text{cam}} = 90^\circ$ where it is set to $30M$ to give a better view of the gravitational lensing. The space between tickmarks on all the axes is $20M$.

plane affected by optical depth. Because the bright peak is due to lensing, it is not particularly affected by the optical depth through the disk. However, at other azimuthal angles, the disk cuts off the rise in flux due to Doppler beaming, running almost all the way around the orbital plane in which the flux at $\theta_{\text{cam}} = 90^\circ$ is almost a factor of 2 smaller than that at $\theta_{\text{cam}} = 70^\circ$.

4. Discussion

4.1. Spectral Features

There has been much discussion in the literature regarding the imprint on the thermal spectrum that may be created by the gap formed around a binary, with suggestions ranging from a sharp cut-off at the temperature of the circumbinary disk's inner edge (in early work assuming there is little accretion from the circumbinary disk to the binary: Gultekin & Miller 2012; Kocsis et al. 2012; Tanaka et al. 2012; Tanaka & Haiman 2013), to a deep notch between frequencies corresponding to

the temperature of the circumbinary disk's inner edge and those corresponding to the temperature of the mini-disks' outer edge (Roedig et al. 2014), to a gentle change of slope in this region (Farris et al. 2015b) or a distinct notch, but centered at several keV (Tang et al. 2018). The degree to which such a feature appears hinges on the contrast between the highest temperature achieved in the circumbinary disk and the lowest temperature found in the mini-disks; more precisely, the relevant contrast is between the temperatures of regions radiating thermally. Applying simple equilibrium Newtonian disk theory to this situation, as in Roedig et al. (2014), leads to a temperature ratio across the gap of $\simeq 3$.

In our high accretion rate example, this ratio is smaller, only $\simeq 1.25-2$ (see Figure 5), smoothing the spectrum so that the notch almost disappears (see Figure 8). The diminished temperature ratio results from a combination of effects, none of them present in simple disk theory.

One is a higher temperature strip along the inner edge of the circumbinary disk (also visible in Figure 5) due to the shock

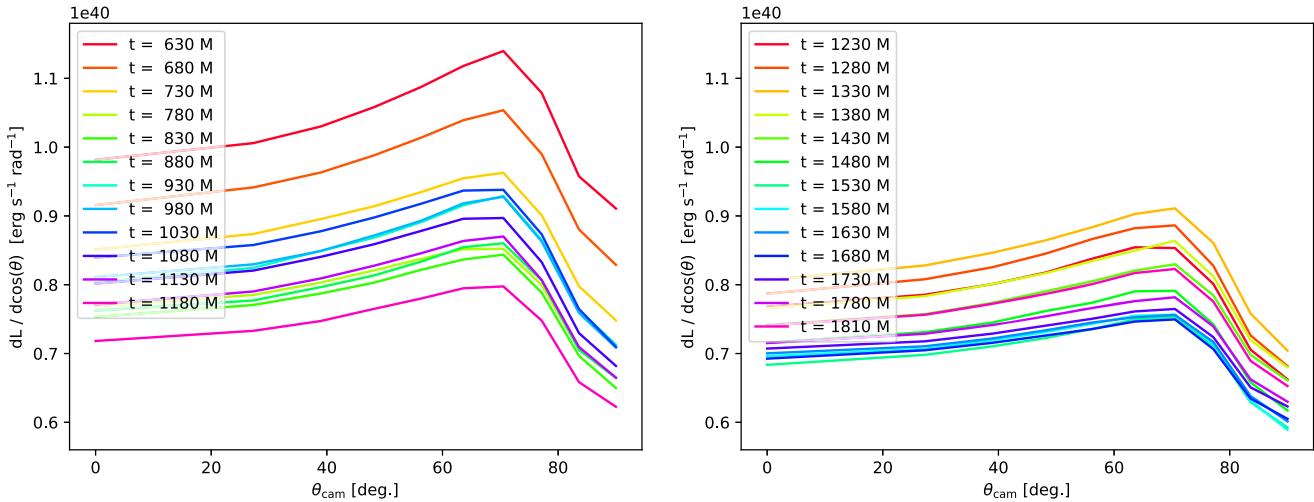


Figure 10. Dependence of the bolometric luminosity on the observer's angle of inclination (θ_{cam}) from the $+z$ -axis for $\dot{m} = 8 \times 10^{-4}$. The luminosity was averaged over azimuthal viewing angle (ϕ_{cam}), and time-averaged over the second orbit (left) and the third orbit (right). Curves at 12 equal intervals of time during each orbit are shown.

driven by accretion streams flung outward by the binary's torque. Also found in the work of Farris et al. (2015b), this strip can partially fill the notch if it radiates thermally.

The small binary separation we have studied ($a = 20M$) also leads to a weaker temperature contrast: the outer edges of these mini-disks are only a factor of 2 outside their innermost stable circular orbit (ISCO) radii. So close to the ISCO, the dissipation rate per unit area rises inward considerably more gradually than the classical r^{-3} scaling. Incorporating this correction into standard disk theory for spin-less BHs (as in our simulations) diminishes the predicted temperature ratio from $\simeq 3$ to $\simeq 2$, but if the BHs had near-maximal spin, there would be essentially no alteration to the temperature ratio because the ISCO angular momentum is rather smaller for rapidly spinning BHs than for non-spinning BHs, and this flattening of the dissipation profile is due to diminution of radial contrast in angular momentum. For essentially the same reason, at radii only a few times that of the ISCO, radial pressure gradients can accelerate inflow without dissipation (Beloborodov & Illarionov 2001; Krolik et al. 2005). Thus, the surface brightness not far outside the ISCO can be significantly depressed relative to the classical r^{-3} scaling. In addition, the short inflow time in the mini-disks when the binary separation is small leads to a state in which they are close to inflow equilibrium with respect to their instantaneous mass accretion rates, but this rate can be either larger or smaller than their mean share of the accretion rate through the circumbinary disk.

Thus, in the circumstances posited here, the notch is likely to be weak. However, it could partly re-emerge in binaries with separations larger by a factor of 2–3 or more or in binaries in which the BHs spin rapidly because in both instances the outer rim of the mini-disks is farther from the ISCO. It might be further re-excavated if more of what is considered “disk” in our analysis were, in fact, optically thin enough to make thermalization of its emission imperfect. Even at an accretion rate $\dot{m} = 0.5$, the Thomson optical depth through the outer rim of the mini-disks, where the accretion streams strike, is only $\simeq 5$ –10. Particularly at the very high post-shock temperature associated with these shocks, the absorption opacity is likely too low to thermalize the radiation spectrum. As a result, energy would be shifted from radiation at the lowest temperature of the mini-disks to much harder photons, while

the higher-temperature thermally radiating regions of the mini-disks would be unaffected.

Hard X-ray production is generic to accreting BHs, and is usually attributed to inverse Compton scattering by electrons heated by magnetic dissipation events in low-density material above the accretion disk (Schnittman et al. 2013). At low accretion rates, the optical depth of the mini-disks, the accretion streams, and the inner regions of the circumbinary disk would be so small that essentially all the radiated power should be in this band. Post-processing of simulation data relevant to single BHs with high accretion rates indicates that $\sim 10\%$ of the heating takes place in regions optically thin to Thomson scattering, and is therefore radiated as Compton-scattered hard X-rays (Noble & Krolik 2009; Schnittman et al. 2013). In our high accretion rate case, the hard X-ray luminosity is a similar fraction of the mini-disk luminosity. As we have already remarked elsewhere, more realistic emissivity models may lead to an augmentation of the X-ray luminosity. In addition, because this simulation had a sizable cut-out at the system center of mass, we were unable to see the “sloshing” motion described in Bowen et al. (2017); the shocks in this part of the flow may also produce X-rays.

4.2. Radiative Efficiency

Integrating over the face-on spectrum, we find that the emitted power in the high accretion case is $\simeq 0.1L_{\text{Edd}}$ rather than the $0.5L_{\text{Edd}}$ that might be expected from $\dot{m} = 0.5$. Several factors contribute to this discrepancy (which applies equally well to the low accretion rate case). Roughly half can be attributed to our definition of \dot{m} , which assumes 10% radiative efficiency, whereas these are Schwarzschild BHs, for which the canonical radiative efficiency, the binding energy at the ISCO, is only 5.7% (the actual radiative efficiency may reach $\gtrsim 6\%$ according to Noble et al. 2011). A smaller decrease in the luminosity observed in the face-on view is caused by relativistic effects: beaming into the equatorial plane and gravitational redshift (Noble et al. 2011 found the polar suppression in a Schwarzschild spacetime to be $\simeq 10\%$).

In addition, as we have already emphasized, each mini-disk's outer edge is only twice as far from its BH as its ISCO. With

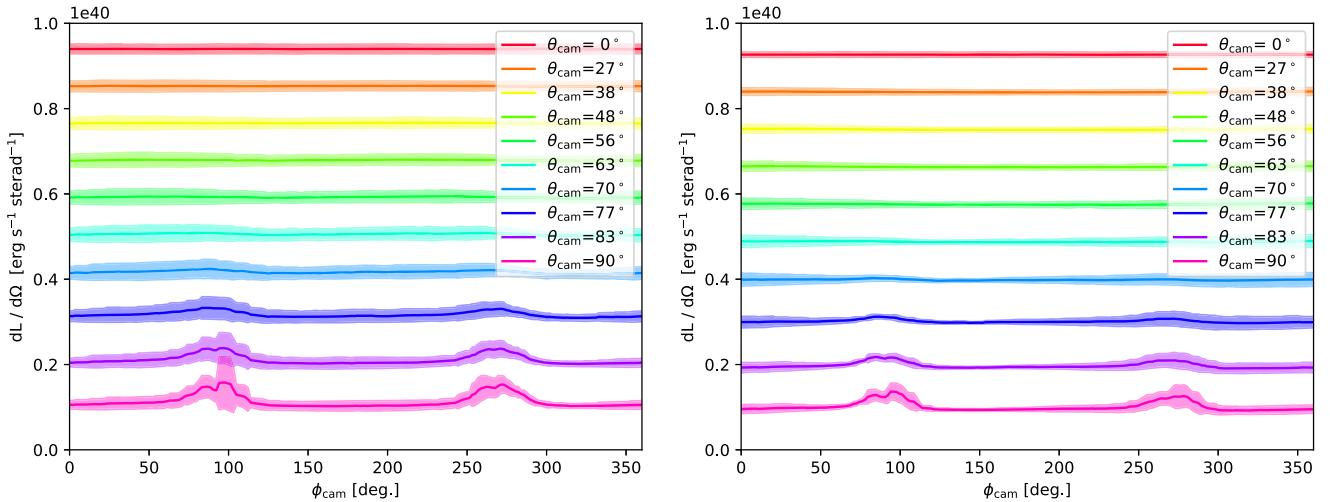


Figure 11. Dependence of the bolometric luminosity on the observer's azimuthal position (ϕ_{cam}) at a variety of inclination angles (θ_{cam}) for $\dot{m} = 8 \times 10^{-4}$. Solid curves represent time-averaged quantities over the second orbit (left) and the third orbit (right), while shaded regions show the standard deviation of the set of 12 time levels of data used to make the averages. The i th curve from the bottom is vertically offset by $9(i-1) \times 10^{38} \text{ erg s}^{-1} \text{ sr}^{-1}$, meaning the $\theta_{\text{cam}} = 0^\circ$ curve should really lie at $\sim 1.1 \times 10^{39} \text{ erg s}^{-1} \text{ sr}^{-1}$.

specific angular momentum only slightly greater than that at the ISCO, gas can accrete with comparatively little in the way of dissipation. When this occurs, matter is accreted with greater orbital energy, and therefore radiates almost a factor of two less efficiently (Beloborodov & Illarionov 2001).

Yet another relevant consideration is an artifact of the brief duration of our simulation. To travel from the circumbinary disk to the binary, matter must go inward once, be torqued to greater angular momentum, shock against the circumbinary disk, lose angular momentum, and then fall toward the binary. This process takes roughly as long as an orbital period at the circumbinary disk edge, approximately three binary orbital periods. Because our simulation ran only three binary orbital periods, it did not run long enough for the mass-supply rate to the binary to equilibrate with the mass accretion rate in the inner portion of the circumbinary disk.

4.3. Angular Dependence and Time Dependence

The accretion rate from a circumbinary disk to mini-disks surrounding the individual members of a binary system is in general modulated on frequencies comparable to the binary orbital frequency (MacFadyen & Milosavljević 2008; Roedig et al. 2011; Noble et al. 2012; Shi et al. 2012; Farris et al. 2014). Whether this translates into time dependence of radiation depends on how swiftly the cross-gap accretion rate translates into photon emission. When the inflow through the mini-disks is governed by conventional disk mechanics (i.e., turbulent MHD stresses), the inflow time is nearly always considerably longer than the binary orbital period. In the situation treated here, however, the mini-disk inflow time is comparable to or shorter than the modulation period. As mentioned in the previous subsection, this may in part be due to the fact that the mini-disks' outer edges are not far from the ISCO, permitting accretion to be driven by pressure gradients at least as much as by internal stresses (Beloborodov & Illarionov 2001). Additionally, because the mini-disks are in a binary, tidal forces induce spiral shocks that can also transport angular momentum, particularly if the disk is comparatively hot (Lynden-Bell & Pringle 1974; Ju et al. 2016), and these are indeed present in our simulation data (Bowen et al. 2018). As a

result of the accelerated inflow rate, the switching of the accretion from one mini-disk to the other that accompanies the overall modulation in accretion rate is reflected in a strong modulation of the accretion rate in each mini-disk; if the inflow rate were slower, the accretion rate in the individual mini-disks would track the longer-term accretion rate, not the modulated version.

In this context, relativistic effects can create time variability in a number of ways. As demonstrated in Section 3.3 with regard to the low accretion rate case, Doppler beaming and gravitational lensing can work together to induce periodic variation at twice the orbital frequency when the system is viewed within $\approx 20^\circ$ of the orbital plane. Although Figure 11 is posed in terms of azimuthal angle dependence in the corotating frame, this translates to time dependence in an inertial frame.

Similar effects can also modulate the light observed from a more rapidly accreting system viewed not too far from the plane, but they work somewhat differently. The key change is that these disks are much more optically thick. This means, for example, that they are more effective at blocking nearly edge-on views. In addition, and of greater interest, the time for photons to escape from within a mini-disk can be comparable to or longer than the inflow time within the mini-disk. When the photon diffusion time is longer than the inflow time, absorption can both smooth out the accretion rate modulation and also suppress the emitted luminosity. Quantitative evaluation of these effects demands more powerful simulation tools.

4.4. Caveats

Our methods are based upon some approximations that are often, but not always, valid. One is the assumption of “fast light.” This assumption essentially translates to ray-tracing within single snapshots corresponding to a single value of coordinate time, even though light requires a finite time to travel across the source region. This approximation would not support predictions of time variability on scales comparable to or shorter than the light crossing time. However, because the photon diffusion time through optically thick disks is generically an order of magnitude or more longer than the local dynamical time, and the dynamical time is $\sim (r/M)^{3/2} M$,

such rapid variability is unlikely to be important for the thermal portion of the spectrum. For the optically thin case, the relevant considerations are different. Our model for the Comptonized spectrum short-circuits any account of the time required for the spectrum to reach a steady state. Because we set the temperature in the optically thin corona component to a fixed value, emulating the corona temperature seen in observed AGNs, we neglect the time variability associated with temperature fluctuations in the corona that may be present in natural systems. The variability in the corona from our calculations is therefore due solely to the cooling timescale, set to be the local orbital period—a timescale much larger than the local light crossing time. Thus, mainly because of the assumptions built into our coronal emission model, our fast-light approximation is expected to accurately capture the time variability of our system.

Since the fast-light approximation required us to use the same metric time slice used by the simulation data snapshot, the geodesic calculation also assumed the fast-light approximation. The spacetime dynamics are tied to the orbit of BHs, so if the ratio of their orbital velocity to the speed of light is small then we would expect that the fast-light approximation is justified. At the $a \simeq 20M$ separation considered here, this ratio is $\simeq 0.1$, which is small enough for the approximation to be reasonable, but only marginally so. We will measure the error of this assumption in future work in which we propagate the photons in time with the spacetime and simulation data.

Another concern arises from the short duration of the underlying simulation (three binary orbits). Over such a short time, the system has not reached inflow equilibrium in any sense; put another way, the accretion rate varies substantially as a function of radius. This is likely true even if the accretion rates onto the mini-disks are summed. By contrast, because astrophysical BBHs have evolved for a number of orbits many orders of magnitude larger than the three orbits of our GRMHD simulation, in most instances they can be expected to have reached a very close approximation to inflow equilibrium (in terms of the total accretion rate onto the mini-disks). Departures from a radially uniform accretion rate can lead to distortions in the predicted spectrum by overemphasizing some temperature regions and underemphasizing others. In addition, as has been emphasized in both this paper and Bowen et al. (2018), it is unclear to what degree the mini-disk regions ever settle into a steady state given the possibility that they may follow a sequence of depletion/refilling episodes if their inflow rates are comparable to or faster than the binary orbital frequency. If so, a time-dependent analysis of their emission properties (as illustrated here) will be necessary.

A third concern is that our simulation determines the local scale height of the disk by a rather ad hoc mechanism. In real disks, it is the result of pressure support whose blend of radiation and gas pressure, as well as its temperature, are determined by a balance between turbulent dissipation and heat transport. By contrast, in our simulation pressure is assumed to be entirely gas pressure, and the local temperature is held close to an imposed target temperature by forcing the gas to lose heat at a rate comparable to the dynamical frequency. Although these gross approximations may be very significant for determination of the disk shape, they are likely less so for time-averaged predictions of photon output. One reason is that

the local rate of heat dissipation is dependent upon the mass accretion rate and the gravitational potential, rather than the shape of the disk. Another is that thermal equilibrium fixes the local spectrum once the surface brightness (the dissipation rate per unit area) is known. Our very approximate description of the cooling rate has the virtue of tying the local bolometric luminosity very closely to the dissipation required by the local accretion rate. It may, however, be unreliable in regions (e.g., the streams or very optically thick regions of the mini-disks) where the physical cooling time may be long compared to the time required to change the radius.

A further concern, closely related to the ad hoc cooling rate, is that although we describe one case in which the accretion rate is almost Eddington, we ignore both radiation forces and photon trapping. It is probably best to think of our $\dot{m} = 0.5$ example as an initial exploration of the properties of such disks.

It is also possible that by setting the criterion for emitting a locally thermal spectrum to be the existence of a Thomson scattering photosphere, we overestimated the thermal luminosity and underestimated the unthermalized radiation. Regions where the Thomson depth is only $< O(10)$ may, in fact, produce spectra rather harder than the Planckians we assigned them.

Our method (shooting rays from a distant camera to the source) describes photon propagation well only when the entire path is transparent. It is essentially for this reason that we restricted our high accretion rate predictions to face-on views. For predictions taking into account the shape of the photosphere, as well as Compton scattering in coronal regions, a method that follows photons from source to observer (e.g., *Pandurata*: Schnittman & Krolik 2013) is preferable.

5. Conclusions

In this paper we presented a first step toward estimating the radiative properties of SMBBHs in the stage immediately before merger.

When the accretion rate is great enough to make most of the accretion flow optically thick, our model produces thermal radiation with a spectrum that differs only modestly from ordinary single BH systems. The contrasts may be greater, however, for binaries with greater separations or containing more rapidly spinning BHs or if regions of modest optical depth achieve only partial thermodynamic equilibrium between gas and photons.

Outside thermalized regions, inverse Compton scattering between photons and high-energy electrons produces hard X-ray emission. The hard X-ray flux may also be subject to modulation on frequencies comparable to the binary orbital frequency, particularly when the system is viewed from a position near the orbital plane and the accretion rate is comparatively low. Both Doppler beaming and gravitational lensing can modulate the observed light flux seen by near-plane observers. Additional X-ray variability may arise from refilling/depletion episodes caused by periodic passage of the BHs near the overdensity feature at the edge of the circumbinary disk.

Some of our predictions are robust; others are subject to cautions we have enumerated in Section 4.4. However, the post-processing tool we have created has considerable flexibility and potential

power, one that can be re-used—employing more realistic assumptions—on data from future simulations.

S. D. was supported through the *Center for Computational Relativity and Gravitation* and *Frontier in Gravitational Astrophysics* program through RIT's office of research. D. B., M. C., and V. M. received support from NSF grants AST-1028087, AST-1516150, PHY-1305730, PHY-1707946, OAC-1550436 and OAC-1516125. S. C. N. was supported by AST-1028087, AST-1515982 and OAC-1515969, and by an appointment to the NASA Postdoctoral Program at the Goddard Space Flight Center administered by USRA through a contract with NASA. J. H. K. was partially supported by NSF grants AST-1516299, PHYS-1707826 and OAC-1516247 and the Simons Foundation (grant 559794, JHU). V. M. also acknowledges partial support from AYA2015-66899-C2-1-P.

Analysis and ray-tracing were performed on the Blue Waters system at the University of Illinois at Urbana-Champaign and its National Center for Supercomputing Applications, and the NewHorizons and BlueSky Clusters at Rochester Institute of Technology. The Blue Waters sustained-petascale computing project is supported by the National Science Foundation (awards OAC-0725070 and OAC-1238993) and the state of Illinois. This work is also part of the “Predicting the Transient Signals from Galactic Centers: Circumbinary Disks and Tidal Disruptions around Black Holes” PRAC allocation support by the National Science Foundation (award number OAC-151596). The NewHorizons and BlueSky Clusters were supported by NSF grant No. PHY-0722703, DMS-0820923, AST-1028087, and PHY-1229173. This work was performed in part at Aspen Center for Physics, which is supported by National Science Foundation grant PHY-1607611.

Appendix A The Lorentz-invariant Radiative Transfer Equation

Let I , α , and j be the Lorentz-invariant intensity, absorption coefficient, and emissivity, respectively, and let I_ν , α_ν , and j_ν be the values measured by a local observer that measures the photon to have frequency ν . The relationship between these two sets of quantities is

$$I = \frac{I_\nu}{\nu^3}, \quad j = \frac{j_\nu}{\nu^2}, \quad \alpha = \alpha_\nu \nu. \quad (15)$$

The radiative transfer equation in this observer's frame is

$$\frac{dI_\nu}{ds} = j_\nu - \alpha_\nu I_\nu, \quad (16)$$

where ds is the incremental distance along the geodesic as seen by the local observer. Inserting Equations (15) into (16) yields

$$\frac{dI}{ds} = \frac{1}{\nu}(j - \alpha I). \quad (17)$$

We wish to use the affine geodesic parameter to evolve this differential equation. If we denote the local observer's 4-velocity by v^μ and the photon's wavevector by k^μ , we have

the following relations:

$$\nu = -\frac{v^\mu k_\mu}{2\pi}, \quad (18)$$

$$ds = c dt_{\text{obs}} = -\frac{v^\mu dx_\mu}{c}. \quad (19)$$

Using the freedom we have to choose λ up to an affine transformation, we define the normalized wavevector $N_\mu = \frac{\partial x^\mu}{\partial \lambda} \equiv \frac{c}{2\pi} k^\mu$. Equations (18) and (19) yield $ds = \nu d\lambda$, and we may rewrite (17) in the simple form:

$$\frac{dI}{d\lambda} = j - \alpha I. \quad (20)$$

Appendix B Conversion Factors

In order to convert variables used in HARM3D and BOTHROS to cgs units, we need to match numerical scales (i.e., those used in the code) to physical scales in cgs units. In the following, we subscript with an “ n ” the variables in numerical units, with a “ c ” those in cgs units, and with a “ c/n ” the conversion factors between the two system of units, $X_{c/n} \equiv X_c/X_n$. Let us now discuss how ν_c is calculated from ν_n . We have:

$$\nu_{c/n} = -\frac{(u_\mu N^\mu)_{c/n}}{c}. \quad (21)$$

Since the conversion factors should all be constants along the geodesics, we can evaluate them at the camera. We denote by C^μ the 4-velocity of the camera, which we set to $(1, 0, 0, 0)$, and the subscript ∞ denotes a quantity evaluated at the camera. If the camera observes the photon at a given frequency $\nu_c \infty$, we then get:

$$(u_\mu N^\mu)_{c/n} = \frac{(C_\mu N^\mu)_{c \infty}}{(C_\mu N^\mu)_{n \infty}} = -\frac{c \nu_c \infty}{(C_\mu N^\mu)_{n \infty}}. \quad (22)$$

As we mentioned in Section 2.3.1, N_∞^μ is set such that the geodesics are null ($N^\mu N_\mu = 0$), point to the camera from the field, and move forward in time. Putting this into (21), we obtain:

$$\nu_{c/n} = -\frac{(u_\mu N^\mu)_{c/n}}{c} = -\frac{\nu_c \infty}{(C_\mu N^\mu)_{n \infty}}. \quad (23)$$

The density scale of the gas is set by choosing an average accretion rate into the system normalized by the numerical one measured from the simulation data. Because our circumbinary disk has reached inflow equilibrium only out to $r \sim 4a$ we choose $2a < r < 4a$ as the region within which to measure the simulation's accretion rate in code units, i.e., \dot{M}_n . Measuring this from the Noble et al. (2012) simulation, we find $\dot{M}_n \approx 0.03$ at $t = 50,000M$, the time from which we start the Bowen et al. (2018) simulation.

The other quantities are more straightforward to transform, and the scales used are shown in Table 1.

Table 1
Scales Used in BOTHROS to Determine the Conversion Factors in Translating Numerical Units to cgs Units

| Quantity (Q) | Numerical (Q_n) | CGS (Q_c) | Conversion ($Q_{c/n}$) |
|--------------------------------|--------------------------------|---|--|
| Length (L) | M_n [1] | GM_c/c^2 | L_c/L_n |
| Time (T) | M_n [1] | cL_c | cL_c/n |
| Mass (M) | M_n [1] | M_c [$10^6 M_\odot$] | M_c/M_n |
| Accretion rate (\dot{M}) | \dot{M}_n [0.03] | \dot{M}_c [$8 \times 10^{-4} - 0.5 \dot{M}_{\text{Edd}}$] | \dot{M}_c/\dot{M}_n |
| Mass density (ρ) | $\dot{M}_n/(4\pi L_n^2)$ | $\dot{M}_c/(4\pi c L_c^2)$ | $\dot{M}_c/n/(c L_c^2)$ |
| Frequency (ν) | $-(C_\mu N^\mu)_n \propto [1]$ | $\nu_c \propto [10^{14} - 10^{21} \text{ Hz}]$ | $-\nu_c \propto (C_\mu N^\mu)_n \propto$ |
| Affine parameter (λ) | L_n/ν_n | L_c/ν_c | $L_c/n/\nu_c$ |

Note. Refer to Equation (23) regarding ν_n . \dot{M}_n is measured at the inner edge of the circumbinary. We see that all conversion factors can be derived from the free parameters M_c , \dot{M}_c and ν_c . Actual scaling parameters used are specified in square brackets “[].”

ORCID iDs

Scott C. Noble <https://orcid.org/0000-0003-3547-8306>
 Dennis B. Bowen <https://orcid.org/0000-0002-7447-1142>
 Julian H. Krolik <https://orcid.org/0000-0002-2995-7717>
 Vassilios Mewes <https://orcid.org/0000-0001-5869-8542>

References

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016a, *PhRvX*, **6**, 041015
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016b, *PhRvD*, **93**, 122003
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016c, *PhRvL*, **116**, 241103
 Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, *PhRvL*, **118**, 221101
 Amaro-Seoane, P., Audley, H., Babak, S., et al. 2017, arXiv:1702.00786
 Balbus, S. A., & Hawley, J. F. 1998, *RvMP*, **70**, 1
 Bansal, K., Taylor, G. B., Peck, A. B., Zavala, R. T., & Romani, R. W. 2017, *ApJ*, **843**, 14
 Beloborodov, A. M., & Illarionov, A. F. 2001, *MNRAS*, **323**, 167
 Blanchet, L. 2014, *LRR*, **17**, 2
 Bode, T., Bogdanović, T., Haas, R., et al. 2012, *ApJ*, **744**, 45
 Bode, T., Haas, R., Bogdanović, T., Laguna, P., & Shoemaker, D. 2010, *ApJ*, **715**, 1117
 Bowen, D. B., Campanelli, M., Krolik, J. H., Mewes, V., & Noble, S. C. 2017, *ApJ*, **838**, 42
 Bowen, D. B., Mewes, V., Campanelli, M., et al. 2018, *ApJL*, **853**, L17
 Detweiler, S. L. 2005, *CQGra*, **22**, S681
 Dolence, J. C., Gammie, C. F., Mościbrodzka, M., & Leung, P. K. 2009, *ApJS*, **184**, 387
 D'Orazio, D. J., Haiman, Z., & MacFadyen, A. 2013, *MNRAS*, **436**, 2997
 Dotti, M., Ruszkowski, M., Paredi, L., et al. 2009, *MNRAS*, **396**, 1640
 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*, **447**, L80
 Farris, B. D., Duffell, P., MacFadyen, A. I., & Haiman, Z. 2015b, *MNRAS*, **446**, L36
 Farris, B. D., Gold, R., Paschalidis, V., Etienne, Z. B., & Shapiro, S. L. 2012, *PhRvL*, **109**, 221102
 Farris, B. D., Liu, Y. T., & Shapiro, S. L. 2011, *PhRvD*, **84**, 024024
 Giacomazzo, B., Baker, J. G., Miller, M. C., Reynolds, C. S., & van Meter, J. R. 2012, *ApJL*, **752**, L15
 Gold, R., Paschalidis, V., Etienne, Z. B., Shapiro, S. L., & Pfeiffer, H. P. 2014a, *PhRvD*, **89**, 064060
 Gold, R., Paschalidis, V., Ruiz, M., et al. 2014b, *PhRvD*, **90**, 104030
 Goulding, A. D., Greene, J. E., Bezanson, R., et al. 2018, *PASJ*, **70**, S37
 Graham, M. J., Djorgovski, S. G., Stern, D., et al. 2015, *MNRAS*, **453**, 1562
 Gültekin, K., & Miller, J. M. 2012, *ApJ*, **761**, 90
 Heyrovský, D. 2005, *ApJ*, **624**, 28
 Ireland, B., Mundim, B. C., Nakano, H., & Campanelli, M. 2016, *PhRvD*, **93**, 104057
 Johnson-McDaniel, N. K., Yunes, N., Tichy, W., & Owen, B. J. 2009, *PhRvD*, **80**, 124039
 Ju, W., Stone, J. M., & Zhu, Z. 2016, *ApJ*, **823**, 81
 Kelly, B. J., Baker, J. G., Etienne, Z. B., Giacomazzo, B., & Schnittman, J. 2017, *PhRvD*, **96**, 123003
 Kocsis, B., Haiman, Z., & Loeb, A. 2012, *MNRAS*, **427**, 2680
 Krolik, J. H. 1999, *Active Galactic Nuclei: From the Central Black Hole to the Galactic Environment* (Princeton, NJ: Princeton Univ. Press)
 Krolik, J. H., Hawley, J. F., & Hirose, S. 2005, *ApJ*, **622**, 1008
 Liu, T., Gezari, S., Burgett, W., et al. 2016, *ApJ*, **833**, 6
 Liu, T., Gezari, S., Heinis, S., et al. 2015, *ApJL*, **803**, L16
 Lynden-Bell, D., & Pringle, J. E. 1974, *MNRAS*, **168**, 603
 MacFadyen, A. I., & Milosavljević, M. 2008, *ApJ*, **672**, 83
 Milosavljević, M., & Phinney, E. S. 2005, *ApJL*, **622**, L93
 Mundim, B. C., Nakano, H., Yunes, N., et al. 2014, *PhRvD*, **89**, 084008
 Noble, S. C., & Krolik, J. H. 2009, *ApJ*, **703**, 964
 Noble, S. C., Krolik, J. H., & Hawley, J. F. 2009, *ApJ*, **692**, 411
 Noble, S. C., Krolik, J. H., Schnittman, J. D., & Hawley, J. F. 2011, *ApJ*, **743**, 115
 Noble, S. C., Leung, P. K., Gammie, C. F., & Book, L. G. 2007, *CQGra*, **24**, 259
 Noble, S. C., Mundim, B. C., Nakano, H., et al. 2012, *ApJ*, **755**, 51
 Palenzuela, C., Lehner, L., & Yoshida, S. 2010, *PhRvD*, **81**, 084007
 Pfister, H., Lupi, A., Capelo, P. R., et al. 2017, *MNRAS*, **471**, 3646
 Poisson, E. 2005, *PhRvL*, **94**, 161103
 Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, *Numerical Recipes in C. The Art of Scientific Computing* (2nd ed.; Cambridge: Cambridge Univ. Press)
 Pringle, J. E. 1991, *MNRAS*, **248**, 754
 Roedig, C., Dotti, M., Sesana, A., Cuadra, J., & Colpi, M. 2011, *MNRAS*, **415**, 3033
 Roedig, C., Krolik, J. H., & Miller, M. C. 2014, *ApJ*, **785**, 115
 Roedig, C., Sesana, A., Dotti, M., et al. 2012, *A&A*, **545**, A127
 Ryan, G., & MacFadyen, A. 2017, *ApJ*, **835**, 199
 Schnittman, J. D., & Krolik, J. H. 2013, *ApJ*, **777**, 11
 Schnittman, J. D., Krolik, J. H., & Noble, S. C. 2013, *ApJ*, **769**, 156
 Schnittman, J. D., Krolik, J. H., & Noble, S. C. 2016, *ApJ*, **819**, 48
 Shannon, R. M., Ravi, V., Lentati, L. T., et al. 2015, *Sci*, **349**, 1522
 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
 Tanaka, T., Menou, K., & Haiman, Z. 2012, *MNRAS*, **420**, 705
 Tanaka, T. L., & Haiman, Z. 2013, *CQGra*, **30**, 224012
 Tang, Y., Haiman, Z., & MacFadyen, A. 2018, *MNRAS*, **476**, 2249
 Tang, Y., MacFadyen, A., & Haiman, Z. 2017, *MNRAS*, **469**, 4258
 Trakhtenbrot, B., Ricci, C., Koss, M. J., et al. 2017, *MNRAS*, **470**, 800
 Tremblay, S. E., Taylor, G. B., Ortiz, A. A., et al. 2016, *MNRAS*, **459**, 820
 Zhu, Y., Davis, S. W., Narayan, R., et al. 2012, *MNRAS*, **424**, 2504
 Zilhão, M., & Noble, S. C. 2014, *CQGra*, **31**, 065013