

Effect of magnetic fields on the dynamics and gravitational wave emission of Papaloizou-Pringle instability-saturated self-gravitating accretion disks: Simulations in full GR

Erik Wessel¹, Vasileios Paschalidis², Antonios Tsokaros^{3,4}, Milton Ruiz⁵, and Stuart L. Shapiro⁶

¹*Department of Physics, University of Arizona, Tucson, 85721 Arizona, USA*

²*Departments of Astronomy and Physics, University of Arizona, Tucson, 85721 Arizona, USA*

³*Department of Physics, University of Illinois, Urbana-Champaign, 61801 Illinois, USA*

⁴*National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign, Urbana, 61801 Illinois, USA*

⁵*Departamento de Astronomía y Astrofísica, Universitat de València, Dr. Moliner 50, 46100 Burjassot (València), Spain*

⁶*Departments of Physics and Astronomy, University of Illinois, Urbana-Champaign, 61801 Illinois, USA*



(Received 13 March 2023; accepted 30 May 2023; published 28 June 2023)

We explore the effect magnetic fields have on self-gravitating accretion disks around spinning black holes via numerical evolutions in full dynamical magnetohydrodynamic spacetimes. The configurations we study are unstable to the Papaloizou-Pringle instability (PPI). PPI-saturated accretion tori have been shown to produce gravitational waves, detectable to cosmological distances by third-generation gravitational wave (GW) observatories. While the PPI operates strongly for purely hydrodynamic disks, the situation can be different for disks hosting initially small magnetic fields. Evolutions of disks without self-gravity in fixed black hole (BH) spacetimes have shown that small seed fields can initiate the rapid growth of the magneto-rotational instability (MRI), which then strongly suppresses the PPI. Since realistic astrophysical disks are expected to be magnetized, PPI-generated GW signals may be suppressed as well. However, it is unclear what happens when the disk self-gravity is restored. Here, we study the impact of magnetic fields on the PPI-saturated state of a self-gravitating accretion disk around a spinning BH ($\chi = 0.7$) aligned with the disk angular momentum, as well as one around a nonspinning BH. We find the MRI is effective at reducing the amplitude of PPI modes and their associated GWs, but the systems still generate GWs. Estimating the detectability of these systems across a wide range of masses, we show that magnetic fields reduce the maximum detection distance by Cosmic Explorer from 300 Mpc (in the pure hydrodynamic case) to 45 Mpc for a $10M_{\odot}$ system, by LISA from 11500 to 2700 Mpc for a $2 \times 10^5 M_{\odot}$ system, and by DECIGO from $z \approx 5$ down to $z \approx 2$ for a $1000M_{\odot}$ system.

DOI: [10.1103/PhysRevD.107.123031](https://doi.org/10.1103/PhysRevD.107.123031)

I. INTRODUCTION

Thick gaseous tori that orbit and accrete onto black holes have been a major topic of study in astrophysics for decades. Early theoretical work on these systems was motivated by the need to model x-ray binaries and active galactic nuclei [1]. Supermassive star collapse [2], tidal disruption events by black holes, mergers of binary neutron stars, and of binary black hole-neutron stars [3,4] can all form short-lived (massive) accretion tori. Thus, black hole (BH) disk systems are known to be ubiquitous in the Universe, and much work has been done to model the electromagnetic signatures of these objects. However, with the rise of gravitational wave (GW) astronomy [5] and the promise of third generation GW observatories to open up new frequency ranges and achieve greater sensitivity [6–8], interest has arisen in BH-disk systems as potential GW sources.

In order to produce GW radiation, an astrophysical object must have significant time variation of its quadrupole moment.¹ For BH-disk systems, the simplest way for this to occur is for the disk to be nonaxisymmetric, with a significant concentration of mass in one or more orbiting lumps that persist for many orbits. The well-known hydrodynamic Papaloizou-Pringle instability (PPI) [9] saturates to such a quasi-steady-state configuration [10]. The PPI is a global instability with modes that grow exponentially in amplitude from initially small perturbations in axisymmetric disks. The growth of PPI modes can be understood intuitively as a runaway feedback loop resulting from the exchange of a conserved—but *not* lower-bounded—quantity between

¹GWs may also be produced by changes in higher-order mass and angular momentum moments, but in practice these are typically subdominant contributions.

wavelike disturbances on the inner edge of the disk which propagate opposite the flow, and wavelike disturbances on the outer edge of the disk which propagate with the flow [11–16]. The fastest-growing modes of the PPI are simple nonaxisymmetric perturbations of the rest-mass density ρ_0 of the form

$$\rho_0 \propto e^{i(m\phi - \sigma t)}, \quad (1)$$

where ϕ is the azimuthal angular coordinate and m is the integer mode number, which determines the number of overdense lumps produced in the disk, σ is a different complex number for each m , with a real component that controls the pattern's orbital frequency, and an imaginary component that controls its exponential growth rate. The low- m modes are fastest growing, with patterns that rotate at roughly the orbital frequency of the maximum-density region of the disk.

Susceptibility to the PPI is determined by the specific angular momentum profile of the initially axisymmetric disk. In a Newtonian context, the specific angular momentum is defined as

$$j \equiv r v_\phi, \quad (2)$$

where r is the cylindrical radius and v_ϕ the azimuthal component of the velocity. For axisymmetric disks j typically is nearly a power law in terms of r , and the criteria for PPI instability is [17]

$$j \propto r^{2-q}, \quad \sqrt{3} < q \leq 2. \quad (3)$$

Disks with ($q > 2$) have radially decreasing specific angular momentum, and are quickly accreted due to a runaway axisymmetric instability first described by Rayleigh [18]: we will not concern ourselves with such disks. While the criteria above were originally derived in the Newtonian regime, it turns out that these criteria can still fruitfully be applied to the PPI in relativistic disks [19], provided we use a relativistic definition for the specific angular momentum,

$$j \equiv u^t u_\phi, \quad (4)$$

where u^μ is the fluid four-velocity in a polar coordinate system.

The PPI was first studied for a self-gravitating disk orbiting a nonspinning BH in full, dynamical spacetime numerical relativity in [20]. Shortly thereafter longer evolutions of a similar BH disk system were published by [21], who were able to extract long GW waveforms produced by the PPI, and assessed the potential detectability of such signals. Smoothed-particle hydrodynamics simulations of tori produced by tidal disruption events demonstrated one astrophysically plausible pathway for producing PPI-unstable disks [22], although the estimates

of GWs produced by those disks [23] show that the amplitude would be much lower than the configuration studied in [21]. In [24] PPI-unstable self-gravitating disks were evolved around tilted, spinning black holes, but starting from constraint-violating initial data.

In [25] we studied the PPI in self-gravitating disks around spinning BHs with constraint-satisfying initial data for the first time. Three nearly identical BH disk systems were evolved, differing primarily in the spin of the central BH, which in one case had no spin, and in the other two had a dimensionless spin magnitude of $\chi = 0.7$ either aligned or antialigned with the disk orbital angular momentum. We found that BH spin only altered the PPI dynamics through the changes induced in disk orbital angular frequency, but that the behavior of the instability was generally unchanged. However, in the aligned-spin case, the reduction in the innermost stable circular orbit (ISCO) radius slowed accretion and improved signal duration. We also carried out a thorough assessment of detectability, showing that such a system could potentially be detected out to cosmological distances by next-generation space-based GW observatories. In [26] we followed up with hydrodynamic simulations of self-gravitating tori tilted with respect to the BH spin, where we found that PPI saturates earlier and these systems generate GWs beyond the dominant (2,2) mode. In [27] GWs produced by the growth of low- m nonaxisymmetries in thick tori with shallow j profiles around spinning BHs were investigated as a plausible alternative explanation for GW190521.

While these results are exciting from a GW astronomy point of view, these studies did not account for magnetic fields. Magnetic fields are important because the magneto-rotational instability (MRI) grows rapidly in weakly magnetized accretion disks [28]. MRI quickly amplifies small initial fields and drives turbulence that transports angular momentum which allows accretion to proceed [29].

Since the quadrupole moment of the BH-disk system is primarily dictated by the large-scale bulk motion of the disk, an end-state configuration produced by the MRI could be unfavorable to detecting GWs. Both the PPI and the MRI grow at rates near the orbital frequency of the tori, so it is an open question which instability would prevail in PPI-unstable disks hosting small seed magnetic fields. In the context of neglecting the disk self-gravity, this question was answered in [30], where simulations of accretion tori with weak toroidal seed fields (the configuration with the slowest MRI growth rate) revealed that the MRI overpowers the PPI, leading to a final state with no hint of the large nonaxisymmetries favored by the PPI. For disks that start out with nearly axisymmetric configurations, this has effectively ruled out the potential for the PPI to grow and produce detectable gravitational waves, unless the initial field is very weak [22,30].

However, when considering astrophysical processes that can produce massive accretion disks around black holes it is plausible that disk self-gravity could make a difference.

Furthermore, the dynamical events that produce massive tori generically lead to configurations that can be significantly nonaxisymmetric from the outset, e.g., tidal disruption of a neutron star by a black hole. The disk self-gravity was not treated in [30], and no prior work has determined what final configuration such disks will evolve to, or whether such disks can retain significant nonaxisymmetry long enough to produce a detectable gravitational wave signal.

In this paper, we take the first step in exploring such configurations. We take advantage of the evolutions performed in our previous study of the PPI [25], in which self-gravitating disks evolved into highly nonaxisymmetric stable configurations. In this study, we resume two of these prior simulations from near the time of maximum $m = 1$ mode amplitude, but with the introduction of small seed poloidal magnetic fields. In this way, these configurations can stand for a generic case of an initially nonaxisymmetric massive torus with weak initial magnetic fields. This way we also give the PPI the best chance at surviving the magnetic fields. In other words, if the PPI is suppressed by MRI when then PPI is already saturated, then it is unlikely that PPI will be able to dominate over the MRI in any other scenario.

In general we find that, for the configurations we simulated, the MRI reduces the amplitude of the PPI-favored $m = 1$ modes by a factor of 5 within a few orbits. At late times, our simulations still show an elevated $m = 1$ amplitude in their settled states, which results in weak but discernable GW signals. Given that the resolution of our simulations is limited we take our evolutions as characteristic of the impact of magnetic fields on persistent nonaxisymmetric modes in accretion disks, and calculate the effect on GW detectability. Since our simulations end before the GW signals stop, we fit our extracted GW waveforms to a simple model of the full duration signal, as we did in [25], and then exploit the scale freedom in our evolution equations to scale the signals from our simulations to a range of masses. These are chosen to span a range including BH disks left over from BH-NS mergers ($\sim 10M_\odot$), and supermassive star (SMS) collapse of ($\sim 1000M_\odot$) and ($\sim 2 \times 10^5 M_\odot$). We find that the maximum detection luminosity distances are reduced by factors of 3.16–6.67 for the magnetized configurations. This limits Cosmic Explorer (CE) to detecting $10M_\odot$ systems within 45 Mpc, LISA to detecting $2 \times 10^5 M_\odot$ systems within 2700 Mpc, and DECIGO, which for nonmagnetized BH disks could detect GWs from a $1000M_\odot$ system out to $z \approx 5$, to within just $z \approx 2$. Overall, we find the optimistic model still sharply reduces the expectation of detection for the magnetized disk configurations studied here, despite the large headstart given to the PPI and the presence of strong self-gravitation.

This paper is organized as follows. In Sec. II we describe our methods. This covers the generation of initial data, the evolution of hydrodynamic and magnetohydrodynamic disks, and the definitions of all diagnostics we compute to track the behavior of our BH disk systems. In Sec. III we

use these diagnostics to quantify the dynamics of our disks and compare the new magnetized evolutions to our prior unmagnetized ones. We also extract GW signals and repeat the analysis of [25] to assess the effect magnetization has on the maximum detection distance of such sources by LIGO, Cosmic Explorer (CE), DECIGO, and LISA. We conclude in Sec. IV with a discussion of our results and their broader implications for multimessenger astronomy. In this work we use geometrized units where ($G = c = 1$), except where stated otherwise. M designates the Christodoulou mass [31] of the central BH.

II. METHODS

Here we summarize the methods used to produce initial data and evolve it in this study. These techniques have all been discussed in detail elsewhere, and we provide appropriate references accordingly. We also describe in detail the diagnostics used in our analysis. Analysis and visualization of data produced by Cactus-based codes was carried out with the help of the KUIBIT Python module [32].

A. Initial data

The initial configurations evolved in this study were previously generated and evolved in [25]. The only significant difference between the two configurations is the spin angular momentum of the central BH, which is zero in the case we refer to as S_0 , but is parallel to the disk orbital angular momentum and possesses a dimensionless magnitude of $\chi = 0.7$ in the case we refer to as $S_\uparrow$. Using the techniques in [33], the COCAL code solves the initial data problem in general relativity (GR) for each configuration. Starting from an initially Kerr configuration, an elliptic form of the Einstein equations is solved via the Komatsu-Eriguchi-Hachisu scheme [34] as applied to BHs [35], producing spacetime and fluid initial data for equilibrium nonmagnetized BH-disk configurations with the disk self-gravity included. Both BH-disk configurations have equal BH Christodoulou masses [31], M , and the inner edges of each disk are placed at the same radial coordinate, which is well outside the respective innermost stable circular orbit (ISCO). Disk material is modeled as a perfect fluid with a polytropic equation of state,

$$P = k\rho_0^\Gamma, \quad (5)$$

where $\Gamma = 4/3$, ρ_0 is the rest-mass density, and k is the polytropic constant, which scales out of the problem.

To be PPI unstable, these disk configurations are designed to possess shallow specific angular momentum profiles of the form,

$$j(\Omega) = A^2(B_0 - \Omega), \quad (6)$$

as in [33], with $A = 0.1$. B_0 and the other disk parameters are determined by an iteration procedure that repeatedly




FIG. 1. The equatorial specific orbital angular momentum (j) as a function of radial coordinate (r). The black dotted line shows the steepest specific angular momentum profile that is PPI unstable, configurations with steeper profiles will be stable to the PPI according to the criterion in Eq. (3).

generates configurations until finding ones that satisfy the requirements that the disk rest masses be approximately $\sim 0.1M$.

We also make sure that the disk rest masses differ from one another by less than 10%. The result is self-gravitating disks on a nearly-Kerr spacetime, where $j \sim r^{0.01}$, as shown in Fig. 1, which corresponds to $q \approx 2$ thereby satisfying the criteria of Eq. (3) for PPI susceptibility.² The parameters for both configurations are presented in Table I.

The COCAL code solves for initial data configurations on a polar coordinate system in which the interior of the BH is excised. In order to evolve this data with a moving-puncture scheme as described in Sec. II C, it is necessary to provide initial data for the excised region as well, since it includes parts of the BH interior that are still covered by the puncture coordinates. Fortunately, the additional initial data can be constraint violating without influencing the BH exterior as long as it is sufficiently smooth, and so we use the so-called “smooth junk-filling” technique to fill the BH interior [37].

B. Evolution grid structure

For time evolution the initial data is interpolated from the spherical coordinates used by COCAL onto Cartesian grids. Our grids are a nested hierarchy of concentric cubes of half-side lengths which are powers of 2 starting at the finest level $2.19M$ and extending to $2250M$ for a total of 11 levels of refinement, with each successive level having half the resolution of the coarsest one that it contains. This allows the stringent requirements for resolving the event horizon near the BH to be met, as well as the more relaxed requirements for resolving disk dynamics, while avoiding excessive resource use in the outlying regions of the domain where there is little spatial or temporal variation.

²The disks are not gravitationally unstable: the specific angular momentum profiles are steep enough that the Toomre stability criterion [36] is robustly satisfied, even ignoring the stabilizing effect of finite thickness.

TABLE I. Parameters of initial data. All masses, distances, and timescales expressed in terms of BH Christodoulou mass [31] M . From left to right we have: the dimensionless spin parameter of central BH (χ); the ADM mass of the spacetime (M_{ADM}); coordinate radii of the innermost stable circular orbit (r_{ISCO}) inner edge of the disk (r_{inner}), maximum density point of the disk (r_c), outer edge of the disk (r_{outer}); the disk orbital period at r_c (t_{orb}); and the rest mass of the disk (M_{disk}).

Model	χ	$\frac{M_{\text{ADM}}}{M}$	$\frac{r_{\text{ISCO}}}{M}$	$\frac{r_{\text{inner}}}{M}$	$\frac{r_c}{M}$	$\frac{r_{\text{outer}}}{M}$	$\frac{t_{\text{orb}}}{M}$	$\frac{M_{\text{disk}}}{M}$
$S_{\uparrow}$	0.7	1.13	3.39	9.00	15.6	31.7	390	0.12
S_0	0.0	1.14	6.00	9.00	16.9	35.0	427	0.135

In our experience, it is not uncommon for evolutions of the Baumgarte-Shapiro-Shibata-Nakamura (BSSN) formulation of general relativity in our code to slowly accumulate secular drifts of the coordinates, which, over long evolutions, result in the center of mass of the system gradually meandering. Indeed, we encounter just such a drift in these simulations as well as those of [25]. The exact cause of this phenomena in our code has not been closely studied. In our evolution of the shift we use up-winding stencils which could potentially be sensitive to a minor numerical instability leading to an error buildup which could cause this coordinate motion. While the drift starts too early for any physical causal influence from the boundary to have propagated into center of the domain, BSSN exhibits some gauge modes which are superluminal, so this could also be a boundary effect. There are other conceivable explanations, however, regardless of the cause, the secular coordinate drift only impacts the coordinate gauge and does not directly influence the physical degrees of freedom of the system.

On the other hand, the coordinate drift is a hazard to the mesh-refinement scheme, since it might cause the BH to stray across refinement boundaries and become under-resolved. To prevent this, the simulation is periodically regridded every $2.5M$, with all data interpolated onto a new set of grids centered on the current BH apparent horizon centroid.³ Table II lists the grid parameters. The jump in size between the $n = 3$ and $n = 4$ grids is to ensure that the disk is well resolved and to avoid a mesh-refinement boundary intersecting with the bulk of the disk, which could introduce artifacts. This compares favorably to the grids used to study the PPI in [20,21], and in our prior study [25] this exact structure was used and shown to successfully capture the dynamics of the PPI.

C. Nonmagnetized evolution

Initially, we evolve our disks purely hydrodynamically. These are the same evolutions reported previously in [25]

³Although the grid structure is moved, we do not alter the coordinates themselves, so this procedure does not repeatedly move the BH to the center of the coordinate system.

for the $S_\uparrow$ and S_0 cases. We use the Illinois GRMHD adaptive-mesh-refinement (AMR) code [38–40], which was built within the open-source Cactus-Carpet infrastructure [41,42] and is itself the basis of the publicly available counterpart in the Einstein Toolkit [43]. This code evolves the metric via the BSSN formulation of the Einstein equations [44,45]. Moving puncture gauge conditions are used [46,47] with $1 + \log$ time slicing [48] and the “gamma driver” shift condition [49] cast into a first-order form as in Eq. (18) of [50].

The disk material is treated as a perfect fluid with a Γ -law equation of state (EOS),

$$P = (\Gamma - 1)\rho_0\epsilon, \quad (7)$$

where ϵ is the internal specific energy, and $\Gamma = 4/3$ as before. This fluid model is consistent with the polytropic fluid model used in the initial data [Eq. (5)] when the disk material is isentropic, as it is assumed to be in the initial data. The choice of $\Gamma = 4/3$ is realistic for an optically thick radiation-pressure dominated gas, and allows also close comparison to the work of [30].

D. Magnetic field insertion

Using checkpoints of our original hydrodynamic evolutions we resume our simulations from shortly after the PPI growth saturates—when the $m = 1$ PPI nonaxisymmetric mode is near its maximum, and after a short period of hydrodynamic evolution when the time of all refinement levels is the same, we seed an initial magnetic field to render the disk MRI unstable.

The vector potential and magnetic field, here defined with respect to coordinate-stationary observers, are A_i and B_i , respectively, and they are related by

$$B_i = \epsilon^{ijk}\partial_j A_k, \quad (8)$$

where ϵ^{ijk} is the Levi-Civita tensor associated with spatial slice metric γ_{ij} . In terms of the vector potential, our initial field is

$$A_\phi = A_b \max(P - P_{\text{cut}}, 0). \quad (9)$$

A_b is a strength parameter of the field, which indirectly determines the ratio of magnetic to gas pressure at the disk center. The cutoff term P_{cut} is chosen to be 10% of the maximum pressure to ensure that the field is supported only within the high-density region of the disk, and the purely ϕ -directed vector potential results in a B field that is initially poloidal. Critically, the configurations we study here have pressure below P_{cut} all along the azimuthal axis, and in our hydrodynamics scheme all shocks are smoothed over a few grid cells, so this choice of vector potential has no discontinuities. Additionally, although the max function is not formally differentiable at one point, this is not a

problem for derivatives calculated by finite difference methods.

To choose the field strength A_b , we consider the fact that in order for the disk to be MRI unstable, λ_{MRI} must fit within the vertical height of the high-density region of the disk, otherwise the growth of the MRI modes would be incompatible with the effective boundary conditions at the disk edge. The MRI wavelength is

$$\lambda_{\text{MRI}} \equiv 2\pi \frac{|v_A|}{|\Omega|}, \quad (10)$$

where Ω is the orbital frequency for local material. The Alfvén velocity is

$$v_A = \sqrt{b^2/(b^2 + \rho_0(1 + \epsilon) + P)}, \quad (11)$$

where $b^i = B_u^i/\sqrt{4\pi}$ and $B_{(u)}^i$ is the magnetic field observed in the fluid rest frame [39]. In the limit of a weak field $b^2 \ll \rho_0(1 + \epsilon) + P$, so $v_A \propto |b|$ in general, which implies $\lambda_{\text{MRI}} \propto |b|$. Therefore, the constraint that the MRI mode fit within the disk provides an upper bound on the value of b^2 , or equivalently B^2 , and thus on the field strength parameter A_b .

For the lower limit on A_b , we follow the empirical guidance of [51], who found that an average of $\lambda_{\text{MRI}}/\Delta(n) \approx 10$ in the vertical direction and ≈ 20 in the toroidal direction was sufficient to resolve the MRI turbulent state. In our case, the dense region of the disks sits between about $10M$ and $35M$, and so $\Delta(4)$ from Table II is the grid spacing that determines the minimum resolvable λ_{MRI} .

For both of the disks we evolved, we seeded fields with parameters that met all these criteria. Figure 2 compares the disk rest-mass configuration to the vertical MRI wavelength at the time of field insertion.

We can also compare the magnetic energy of the seed field to the internal energy of the gas at the moment of field insertion, which is shown in Fig. 3. As can be seen from the figure, the inserted field is weak in the sense that the energy is still dominated by the gas. This means that the gas will

TABLE II. Parameters of the nested Cartesian grids. From left to right: the refinement level (n), the box half-side length (or “minimum radius”) for the grid of level n [$r_{\text{box}}(n)$], the grid spacing for the grid of level n [$\Delta(n)$]. The grids are all perfect cubes initially centered at the origin but periodically updated to follow the BH centroid. Grid spacings are equal to $\Delta(n)$ in the x , y , and z directions.

n	$r_{\text{box}}(n)$	$\Delta(n)$
1	$2.19M$	$(M/25.6)$
2–3	$2^{(n-1)}r_{\text{box}}(1)$	$2^{(n-1)}\Delta(1)$
4–11	$2^n r_{\text{box}}(1)$	$2^{(n-1)}\Delta(1)$




FIG. 2. Cross sections of the baryon rest-mass density in both accretion disks at the time of field insertion, superimposed with a plot of half the vertical MRI wavelength ($\lambda_{\text{MRI}}/2$) at the midplane. In both cases, the MRI modes fit within the disks on the overdense side, while the MRI wavelength on the rarified side is too large for modes to grow there. For $S_{\uparrow}^{\text{Mag}}$ the yz cross section is used, while for S_0^{Mag} the xz cross section is used, with the x axis flipped to align the high-density sides of the disks.

not be significantly disrupted by the field, and that the resulting magnetized configuration is still constraint satisfying to within numerical fidelity. Thus, we can think of the moment of field insertion as specifying a new initial configuration, which can be evolved and meaningfully compared to our previous nonmagnetized evolutions.

E. Magnetized evolution

With magnetic fields now inserted, we continue the evolution using the Illinois GRMHD AMR code (as described in Sec. II C) adopting ideal magnetohydrodynamics. The magnetic induction equation is evolved using a vector potential method to ensure perfect satisfaction of the $\nabla \cdot B = 0$ condition [39]. This code has previously been used for numerous studies of magnetohydrodynamics in extreme dynamical spacetimes, such as magnetized BH-NS mergers [4,40], MRI-driven jet launching from binary neutron star merger remnant disks [52], MRI-driven accretion disks around binary black holes [53,54], etc. For our electromagnetic gauge choice, we adopt the generalized Lorenz gauge condition developed in [55], which avoids the development of spurious magnetic fields across the AMR levels, setting the generalized Lorenz gauge damping parameter to $\xi = 0.5/M$ for the nonspinning case and $\xi = 0.65/M$ for the spinning case. A slightly higher value was used in the spinning case because an early test run suffered an instability triggered by magnetic flux buildup on the refinement boundary nearest the BH. Since the dynamics of magnetic fields near the spinning BH are different than for




FIG. 3. Ratio of the magnetic energy density ($b^2/2$) to the gas internal energy (ϵ) in the midplane of the disks at the moment of field insertion, on a log scale. Nowhere is the ratio greater than 10^{-4} , so the fields are extremely weak and do not disrupt the gas or introduce non-negligible constraint violations.

the nonspinning case, and the curvature at the horizon is slightly greater, the need for a slightly different level of damping was not totally unexpected. As this parameter only impacts the gauge evolution, it should have no effect on the physical dynamics of the simulations besides keeping them numerically stable. Separately from the second-order dissipation in the generalized Lorenz gauge, we apply higher-order Kreiss-Oliger dissipation to the vector potential and fluid quantities inside the apparent horizon. This suppresses any poorly resolved short-wavelength modes that might appear there, which could otherwise potentially leak out of the horizon due to finite-difference truncation error.

F. Diagnostics

To ensure the quality of the spacetime evolution, we monitor the normalized Hamiltonian and momentum constraints as defined in [56] [Eqs. (40)–(43)]. The output of these diagnostics is discussed in the Appendix.

The evolution of nonaxisymmetric density modes is tracked in two ways. Using 3D grids of the rest-mass density ρ_0 output every $80M$ during the simulation, we

define a simple nonaxisymmetric *mode amplitude* diagnostic,

$$C_m = \int d^3x \rho_0 e^{im\phi}, \quad (12)$$

where ϕ is the azimuthal coordinate of a polar coordinate system centered on the BH centroid. Because the BH dominates the mass of the system, the center of mass is always near the origin of this coordinate system, so this diagnostic compensates for the slow secular drifting of our coordinates described in Sec. II B, which would otherwise induce mode mixing between different m and undermine the interpretation of our results.

The definition of mode amplitudes used here is slightly different from that used in prior studies [25,57–59],

$$C_m^{\text{old}} = \int \sqrt{-g} d^3x u^0 \rho_0 e^{im\phi}, \quad (13)$$

where g is the determinant of the full spacetime metric and u^0 is the time component of the fluid four-velocity. Unfortunately, while we did extract Eq. (13) every $10M$, the routine that computed it was not set up to evaluate the integral in a BH-tracking coordinate system, so at late times when the center of mass is far from the coordinate origin the quality of this diagnostic suffers noticeably. In practice, Eq. (12) agrees closely with Eq. (13) during the times when the system center of mass has not yet drifted too far from the origin (despite the difference in integration measure), and although the frequency of sampling is lower, Eq. (12) is still computed multiple times per orbit, capturing dynamics down to the orbital frequency.

While our mode amplitudes can be informative, they do not exactly match the diagnostics used by [30], who instead extract the *mode power*, normalized with respect to the $m = 0$ mode,

$$\frac{P_m}{P_0} = \frac{\int_{r_{\text{in}}}^{r_{\text{out}}} \int_0^\pi |\int_0^{2\pi} \rho_0 e^{im\phi} d\phi|^2 \sqrt{\gamma} r^2 \cos \theta d\theta dr}{\int_{r_{\text{in}}}^{r_{\text{out}}} \int_0^\pi |\int_0^{2\pi} \rho_0 d\phi|^2 \sqrt{\gamma} r^2 \cos \theta d\theta dr}, \quad (14)$$

where r, θ, ϕ are polar coordinates with their origin at the BH centroid (in [30] this coincides with the COM), and γ is the determinant of the spatial metric in Cartesian coordinates. To compute this measure of the m -mode strengths for our own evolutions we do a linear interpolation of our Cartesian 3D grid data onto a set of grid points in a polar coordinate system centered on the BH. The polar grid has points at 150 evenly spaced radial coordinate values spanning $M \leq r \leq 100M$, 100 evenly spaced θ values between $-\pi/2$ and $\pi/2$, and 200 evenly spaced ϕ values between 0 and 2π . Once the quantities have been interpolated onto this grid we evaluate the integrals via the trapezoidal method.




FIG. 4. The difference between the square root of the spatial metric determinant ($\sqrt{\gamma}$) and the Cartesian flat-space value (1), on a log scale, for both configurations at the time of field insertion. The dashed black lines show the $r = 9M$ coordinate radius. Because most disk material lies outside of $9M$ in our evolutions, approximating $\sqrt{\gamma} \approx 1$ is appropriate when evaluating diagnostic integrals for these simulations.

The spatial metric determinant γ was not regularly saved in our evolutions; however, it was recorded at a few key times, allowing us to determine that by the time of field insertion γ has settled to a static state. Figure 4 shows the difference between $\sqrt{\gamma}$ and 1 along a line through the BH at the time of field insertion, for both configurations. The initial inner edge of the disks is at $r = 9M$ in both cases, and we found that the integrals are dominated by material further than $9M$ from the BHs throughout the evolutions as well. Outside this distance, it can be seen that $\sqrt{\gamma}$ differs from 1 (the flat-space Cartesian value) at the $\sim 10\%$ level.

Therefore, even though we lack the true values of $\sqrt{\gamma}$ at all times, we are able to approximate $\sqrt{\gamma} \approx 1$ when computing Eq. (14) and still make rough comparisons with the results of [30], since our values are not likely to be more than a few percent different on account of the approximation.

To keep track of the magnetic state of the disk, we output 3D grids of B_i every $80M$ as well. From these data, we can calculate many useful diagnostics. As in [51,60], we also compute the ρ_0 -weighted average B -field square magnitude in the $i = \hat{\phi}, R, z$ directions as

$$\langle B_i^2 \rangle \equiv \frac{\int B_i^2 \rho_0 d^3x}{\int \rho_0 d^3x}, \quad (15)$$

where R is the radial component of the cylindrical coordinate basis (rather than the r component, which points in the spherical radial direction). The index $\hat{\phi}$ is meant to indicate the orthonormal ϕ basis component. Therefore $B_{\hat{\phi}}$

will have units of $1/M$ just like B_z and B_R . This convention is used by [51,60].

Gravitational waveforms were extracted using the Ψ_4 Weyl scalar of the Newman-Penrose formalism on the surfaces of concentric coordinate spherical shells centered on the origin. The extraction radii were chosen to be distant enough from the source to be safely in the far field, while being near enough that the grids still resolved the GW wavelengths well. This was achieved by choosing extraction radii at which the signals fall off as $1/r$, a behavior from which the signals deviate if r is too close to or too far from the BH (and hence in an under-resolved region). The complex value of Ψ_4 over each extraction sphere was decomposed into $s = -2$ spin-weighted spherical harmonics for all modes with $\ell \leq 3$. The strain h is the double time integral of Ψ_4 , and was computed via the fixed frequency integration technique of [61], with the real and complex parts providing the $+$ and $\times$ GW polarizations in the natural basis: $h = h_+ + ih_\times$.

G. Resolution dependence

We lacked the computational resources to conduct a formal resolution study as a part of this work. In addition to the $O(\Delta^{-4})$ scaling of resource requirements typical of 3D

time evolutions, the resources needed to perform a resolution study in our case were further inflated by the requirement that we re-evolve the hydrodynamic PPI to the point of saturation before inserting magnetic fields. Therefore, assessing the impact of resolution on the turbulent magnetic field state would have been especially expensive.

Nonetheless, we can indirectly assess the impact resolution is likely to have given what is known of our code's properties as well as the results of prior studies evolving similar systems.

The Illinois GRMHD AMR code is designed to be second-order convergent for hydrodynamic quantities, and tests have shown it to be successful [38–40,43]. In prior work it has performed as expected in similar situations, both for pure relativistic hydrodynamics and magnetohydrodynamics [53,54,62].

The resolution of our grids, as detailed in Sec. II B, easily matches or exceeds that of prior numerical evolutions of the hydrodynamic PPI in general relativity [20,21]. In addition, as reported in [25], we performed a set of incomplete runs with a different grid structure, where a refinement boundary crossed into the dense regions of the disk, and most of the disk evolved with only half the resolution. In these runs the




FIG. 5. Perspective renders of both disks at different times after magnetic field insertion. The left column shows $S_{\uparrow}^{\text{Mag}}$, the right shows S_0^{Mag} ; the top panels show the disks shortly after magnetic field insertion, the bottom panels show the disks later on. Volumetric shading indicates baryon number density, and the surfaces within the disks are isodensity contours to further emphasize the disk structure. Magnetic field lines are traced by the shaded white filaments, the black hole apparent horizon is black, and in all panels the disks are viewed from a coordinate distance of $70M$ from the BH centroid at an angle of 35° above the xy plane. The difference in size of the apparent horizons is a known effect of the chosen coordinates around spinning black holes. While the $m = 1$ nonaxisymmetries are present with a large amplitude shortly after magnetic field insertion, at later times the action of the MRI has significantly reduced their amplitudes, while the field configuration has become much more turbulent and toroidal, as expected from the MRI, and even features some magnetic field lines radiating from the BH poles in a manner reminiscent of jets (although we do not study jets here, and these features are seen in both the spinning and nonspinning cases).

early hydrodynamic evolution of the PPI was nearly identical to our production runs with the higher resolution grid, so we are confident that the spacetime and hydrodynamic evolutions are adequately resolved.

For magnetized accretion tori, capturing the turbulent state induced by the MRI is crucial for sustaining the disk dynamo [29,63]. This turbulence is challenging to resolve, requiring many decades of length scales to be captured in order to reach full convergence. What level of resolution is adequate to obtain realistic bulk behavior for MRI-driven accretion disks in GRMHD has been the subject of numerous studies [51,60,64]. While there is reason to question whether any global accretion disk simulations are yet fully convergent [64], simulations where $\lambda_{\text{MRI}}/\Delta \approx 10$ in the vertical direction and ≈ 20 in the azimuthal direction in the high-density regions of the disk are thought to capture MRI-driven turbulence decently [60]. As mentioned previously in Sec. II D, our disks meet this criteria upon field insertion, and for a period of time after. Therefore, while the qualitative behavior of our results is on solid footing, making quantitative error estimates requires a resolution study which goes beyond the scope of the current work.

III. RESULTS

The two BH-disk systems we studied, the spinning BH case $S_{\uparrow}$ ($\chi = 0.7$) and the nonspinning case S_0 , were previously simulated in full GR without magnetic fields [25]. Of the three configurations in our previous study, these were the two that produced the most powerful and sustained GW signals. In this study we do not evolve the configurations from their initial axisymmetric states, but instead we resume the evolutions near the PPI saturation time, when the $m = 1$ mode is at its maximum, and insert a seed poloidal magnetic field to study its effect on the subsequent dynamics. As can be seen from the 3D renders in Fig. 5, the $m = 1$ nonaxisymmetries are reduced significantly by the introduction of magnetic fields and the growth of the MRI. In the following sections, we make this assessment more quantitative, and determine what it implies for potential detectability.

A. MRI growth and saturation

We can track the MRI growth and saturation through averages of the magnetic field components. Figure 6 shows time series of the average squared values of B -field components, weighted by rest-mass density ρ_0 [Eq. (15)]. The dynamics in both the $S_{\uparrow}$ and S_0 are similar. Initially, the R and z components of the field have a finite value, while the toroidal component is near zero. Subsequently differential rotation rotates nontoroidal components into the toroidal direction, reducing the former and increasing the latter. Quickly the decline of the nontoroidal components halts and reverses, consistent with the expected behavior of




FIG. 6. The density-weighted square averages of magnetic field components [as defined by Eq. (15)]. Top: sum of the non-azimuthal field components R and z . Middle: azimuthal field component $\hat{\phi}$. Bottom: ratio of the square averages of the R and $\hat{\phi}$ components, on a log scale.

the MRI where backreaction of the increasing field line tension on the orbiting fluid leads to runaway growth of the field strength (see Sec. IV B of [29] for an intuitive treatment). Within two orbits (as measured by the orbital period of the maximum density material in the initial data) the MRI appears to have saturated, with both the $S_{\uparrow}$ and S_0 simulations reaching their maximum magnetic strengths.

After this initial growth plateaus, the magnitudes of the magnetic field averages begin to decay. This is commonly seen in accretion disk simulations, see for example Fig. 1(b) of [60] (although note our initial field configuration is different), and is associated with a general decline in magnetic energy [51,60]. This decline could be due to the depletion of the disk. We can look at diagnostics insensitive to the overall flux: the bottom panel of Fig. 6 shows the ratio of the R and $\hat{\phi}$ B -field components. The value this ratio relaxes to, about 0.05, is typical of slightly under-resolved MRI turbulence seen in previous studies (Fig. 4 of [51] and Fig. 5 of [60]).




FIG. 7. Comparison of accretion rates between the nonmagnetized BH-disk systems (dashed lines) and same simulations rerun with magnetic fields seeded at the indicated times (solid lines). The dotted vertical lines indicate the time of magnetic field insertion.

B. Accretion rate

The ultimate source of power for BH-disk systems is the release of gravitational binding energy as material is accreted onto the central black hole. The accretion rate therefore determines the maximum power that such systems can emit in any form, including EM and GW emission. In combination with the total mass of the tori, the rate of accretion also determines the lifetime of these objects, and therefore bounds the lifetimes of any potential observable signatures. The conservation of net angular momentum means that accretion can only be enabled by transport of angular momentum from inner to outer disk regions, and the PPI and MRI instabilities have long been known to produce conditions that can make this transport efficient in astrophysical disks [29]. In Fig. 7, the accretion rate is tracked as a function of time for the original hydrodynamic evolutions—subject only to the PPI—and the new MHD evolutions, where both the PPI and MRI are present. We see that the magnetized disks accrete faster than their nonmagnetized counterparts. The main impact of the increased accretion rate for this study will be the resulting shortening of the disk lifetime, which we will account for when placing an upper bound on the GW detectability.

C. Decay of low- m density modes

We can directly track the mode amplitudes via the diagnostic described in II F [Eq. (12)]. As shown in Figs. 8 and 9, the introduction of the seed magnetic field




FIG. 8. Amplitude of the normalized $m = 1$ density mode after magnetic field insertion for the spin-up case. Note that the components of the $m = 1$ nonaxisymmetric mode take on a slightly nonzero average value at late times in this case. This seems to be due to slight eccentricity that develops due to the nonaxisymmetric variations in accretion that occur intermittently as a result of both the disk nonaxisymmetry and the non-axisymmetry of the seed magnetic field. Our diagnostic measures just the angular $m = 1$ mode, it does not account for radial distribution of matter and bears no relation to the disk center of mass moment. There is therefore no reason to expect it to always oscillate around zero, and this detail can be ignored as irrelevant to the dynamics of interest.

significantly reduces the amplitude of the dominant $m = 1$ mode within the first four orbits. After this decline, the $m = 1$ mode amplitude of the magnetized cases has an average amplitude of roughly $\sim 1/5$ of their original value. Our prior results for nonmagnetized disks show that the saturation amplitudes of $m = 1$ modes remain constant after saturation, and therefore at late times (specifically, at times more than four orbits after field insertion) the magnetized disks have lost $\sim 1/5$ of their amplitudes compared to the nonmagnetized cases. Table III lists the exact average amplitudes compared to their initial values for both cases.

D. Comparison to prior fixed-spacetime study

To make closer contact with the fixed-spacetime (Cowling approximation) results of [30], we use the normalized *mode power* diagnostic [Eq. (14)].




FIG. 9. Amplitude of the normalized $m = 1$ density mode after magnetic field insertion for the nonspinning case.

TABLE III. Average amplitudes of the $m = 1$ mode ($\langle C_1/C_0 \rangle_{\text{late}}$) compared to their values shortly after PPI saturation and MRI seed-field insertion (initial C_1/C_0). The last column shows the ratio between the late time average for the magnetized disks and the initial values. Late times are those four or more orbits after magnetic field insertion in both cases.

Model	Initial C_1/C_0	$\langle C_1/C_0 \rangle_{\text{late}}$	Ratio
$S_{\uparrow}^{\text{Mag}}$	0.29	5.9×10^{-2}	0.21
S_0^{Mag}	0.40	9.3×10^{-2}	0.23

A key point of comparison is Fig. 11 of [30], which shows the $m = 1$ mode power over time for a purely hydrodynamic disk and a magnetized disk with the $m = 1$ mode perturbed (jump starting its early growth), as well as initially axisymmetric magnetized disks. Both the hydrodynamic disks and the perturbed magnetized disks reach $P_1/P_0 \approx 4 - 7 \times 10^{-2}$, but for the latter this is only a transient feature, and the $m = 1$ power quickly settles to match that of the initially axisymmetric magnetized disks. This is one of the most consequential findings of [30]: the presence of weak magnetic fields, through amplification by the MRI and subsequent effect on disk dynamics, does not merely *suppress growth* of the PPI-unstable $m = 1$ modes; they also quickly *erase all traces* of any powerful $m = 1$ modes that once existed. By comparing these results to our own, we can begin to address one of the outstanding questions raised in the conclusion of [30]: whether this is also true for *self-gravitating* disks.

Figures 10 and 11 show the mode powers for the first four modes for each of our disks. Our self-gravitating disks share the preference for lower- m modes seen in, e.g., Fig. 6 of [30]. We start out with slightly stronger $P_1/P_0 \approx 2 \times 10^{-1}$, and yet within a few orbits P_1/P_0 has been reduced by nearly an order of magnitude in both our evolutions. Self-gravity and BH-disk interactions in our dynamic spacetime simulations do not seem to have protected the $m = 1$ modes from being weakened, nor have they resulted in any significant increases to the $m = 2, 3, 4$ modes that might preserve GW detectability.




FIG. 10. The normalized density mode powers for $m = 1, 2, 3, 4$ after magnetic field insertion for the spin-up case.




FIG. 11. The normalized density mode powers for $m = 1, 2, 3, 4$ after magnetic field insertion for the nonspinning case.

Are all traces of the initially powerful $m = 1$ modes *completely* erased? This question is more challenging to answer completely in the absence of a resolution study. Our simulations show a persistent $m = 1$ mode in the MRI-saturated steady state of our disks, and this mode has about 10 times the power of those seen at late times in the magnetized disks of [30], and is much more coherent as well. However, while our azimuthal resolution is roughly comparable to that of [30] in the region of the disk, our vertical and radial resolution is just $\sim 1/3$ of theirs, notwithstanding the differences in convergence between the different numerical methods. These differences could be important in light of the well-known difficulties of adequately resolving steady-state MRI turbulence [51,60,63–65]. To understand what might be supporting the $m = 1$ mode at late times, we can look at the specific angular momentum profile. Figure 12 shows that the profile has steepened substantially toward the PPI-stable limit over the evolution, however in the region within about $20M$ where most of the disk material is the disks remain marginally PPI unstable, suggesting that the PPI might be responsible for the persistence of these modes.




FIG. 12. Comparison of the equatorial radial profile of the specific angular momentum (j) between the initial configurations (solid lines) and magnetized configurations at late times (dashed lines). The black dotted line shows the steepest specific angular momentum profile that is PPI unstable, configurations with steeper profiles will be stable to the PPI.

Therefore, it is not straightforward for us to conclude with certainty whether the relatively strong $m = 1$ mode we observe even with magnetic fields is attributable to self-gravity, and/or other differences with the disk models used in [30]. Yet another important difference between our evolutions and those of [30] is that our disks are seeded with a *poloidal* magnetic field configuration while in [30] disks with initially *toroidal* fields are evolved.

What is clear is that self-gravity is not a panacea for the ills the MRI causes to PPI-supported $m = 1$ modes: even disks with significant self-gravity, starting from a highly asymmetric, $m = 1$ dominated state can be quickly smoothed out by the presence of a weak poloidal magnetic field.

E. Gravitational wave signal

Figures 13 and 14 show the strain of the (2, 2) mode extracted for each configuration. The solid lines indicate the case where magnetic fields are inserted, the dashed lines show the strain from the original, nonmagnetized evolutions. In both cases it can be seen that the reduced $m = 1$ amplitude caused by the magnetization of the disks translates into a reduced GW amplitude.

To gain better insight to the effect this will have on detection, we look instead at the frequency-domain quantity of characteristic strain. This quantity is useful for assessing detectability, as the area under a characteristic strain curve and above a given noise floor roughly corresponds to the signal-to-noise ratio (SNR) for the given optimal matched filtering (see [66] for details). Characteristic strain is defined for positive frequencies only and can be computed as

$$h_c = 2f|\tilde{h}_{\text{res}}|, \quad (16)$$




FIG. 13. Strain for the spinning BH case. The solid line is the magnetized case; the dashed line is the nonmagnetized case. The top time axis shows orbits since magnetization; the bottom uses geometrized units and measures time from peak signal amplitude.




FIG. 14. Strain for the nonspinning BH case. The solid line is the magnetized case; the dashed line is the nonmagnetized case. The top time axis shows orbits since magnetization; the bottom uses geometrized units and measures time from peak signal amplitude.

where $\tilde{h}_{\text{res}}$ is the Fourier transform of the interferometer response. Here we will assume the response to be

$$\tilde{h}_{\text{res}} = \sqrt{\frac{|\tilde{h}_+|^2 + |\tilde{h}_\times|^2}{2}}. \quad (17)$$

This averages the detector response over polarizations, which is appropriate for assessing average SNR of sources whose arrival time and sky location are unknown. We will therefore take h_c computed from (17) and (16) to be the *polarization-averaged characteristic strain*. Figures 15 and 16 show this quantity for the portion of the signals after the




FIG. 15. Characteristic strain for the nonspinning BH case, for the GW signal emitted after the magnetic field insertion. The solid line is the magnetized case; the dashed shows the signal for the same times in the nonmagnetized case.




FIG. 16. Characteristic strain for the spinning BH case, for the GW signal emitted after the magnetic field insertion. The solid line is the magnetized case; the dashed shows the signal for the same times in the nonmagnetized case.

magnetic field insertion. In both cases, the peak amplitude of the signal is reduced by a factor of about 2 from the nonmagnetized case (which will generally imply a similar reduction to expected SNR). While the amplitude is reduced, the frequency of the $m = 1$ modes remains roughly double the orbital frequency, strongly suggesting that the PPI is still responsible for the maintenance of these modes, rather than them being the result of unrelated random fluctuations and turbulence in the disk.

The time-series plots show that the loudest portions of the gravitational wave signals for the magnetized disks occur early on, as the amplitude of the $m = 1$ mode is still in the process of being reduced. This early time is a period of transition, and transient dynamics tend to depend sensitively on the initial conditions. Since we mean for our disks to act as stand-ins for the potential end states of a wide range astrophysical events, it is wise to avoid paying too much attention to the transient dynamics, and instead focus on the late time steady state, which we can hope to be more generic. We define the settled, late-time part of the signal to start $2500M$ (roughly six orbits) after the maximum GW amplitude for the magnetized disks. The magnetized disks have their seed fields inserted shortly after PPI saturation, so this cutoff time also ensures that the signal starts nearly two orbits after MRI saturation and does not capture any transients of MRI growth in addition to avoiding transients of PPI saturation. These signals are obviously shorter than the full duration GW signals, with durations of $2000M$ and $2500M$ for S_0^{Mag} and $S_{\uparrow}^{\text{Mag}}$, respectively, so they are compared with the last $2000M$ of S_0 and the last $2500M$ of $S_{\uparrow}$, to ensure a meaningful comparison of the polarization-averaged characteristic strains in Figs. 17 and 18. As is apparent from these figures, the effect of magnetization is even more extreme




FIG. 17. Characteristic strain for the nonspinning BH case, for the GW signal emitted during the period when the disk has entered a settled steady state. The solid line is the magnetized case; the dashed shows the signal for an equivalent period in the nonmagnetized case.

for this part of the GW signal, with the amplitudes reduced by factors closer to ($\sim 4-5$).

F. Detectability estimates

Finally, we assess what the results of our numerical evolutions mean for the prospects of detecting gravitational waves from real astrophysical disks with similar configurations. To do this we take the GW signals from the configurations we simulated, scale them to a range of different masses, and then compute their amplitude and redshift for a range of source distances in order to estimate the SNR that can be expected from next-generation GW




FIG. 18. Characteristic strain for the spinning BH case, for the GW signal emitted during the period when the disk has entered a settled steady state. The solid line is the magnetized case; the dashed shows the signal for an equivalent period in the nonmagnetized case.

observatories. Because our simulations do not extend for the full lifetime of the disks, we also construct a model that extrapolates the signal to its estimated full duration, and perform the same analysis with the extended waveforms. We did this previously in [25] for the nonmagnetized disks, here we follow the exact same procedure for both magnetized and nonmagnetized disks in order to determine the impact magnetization might have on detectability.

Given that we do not perform a resolution study, we regard the estimates provided here as a *rough upper bound* on detectability for real BH-disk systems of comparable configurations, in accordance with the discussion in Sec. II G.

1. Computing rescaled signals at detector

Because we modeled disk material as a perfect fluid with a Γ -law EOS (Sec. II C), the magnetohydrodynamic and GR equations are invariant under uniform rescaling of mass, which serves as the unit of length and time in geometrized units. Thus, by simply adjusting the mass scales of the relevant observables, we are able to use each of our simulations to represent an entire family of BH-disk systems that share the same dimensionless properties of BH spin magnitude, orientation, and BH-to-disk mass ratio. For GWs, the strain amplitude is proportional to system mass, while the frequency is inversely proportional, so applying these rescalings to our GW model for $rh(t/M, \theta, \phi)$ is straightforward. To go from $rh(t, \theta, \phi)$ in the source frame to the observed strain at a detector, we first choose a value r that will be the luminosity distance to the source in the detector frame. Then the ratio of M/r controls the overall amplitude of the strain.

For example, if we chose $M = 10M_\odot$ and $r = 300$ Mpc, then we would plug $M[s] = 10M_\odot [\text{kg}] \cdot \frac{G}{c^3}$, and $M[\text{Mpc}] = 10M_\odot [\text{kg}] \cdot \frac{G}{c^2} \cdot \frac{\text{Mpc}}{3 \times 10^{22} \text{ m}}$, into the below formula to compute h ,

$$h(t[s], \theta, \phi) = rh(t[s]/M[s], \theta, \phi) \frac{M[\text{Mpc}]}{300 \text{ Mpc}}. \quad (18)$$

The above procedure accounts for the mass scale and luminosity distance only. To accommodate large cosmological distances, we also shift the observed frequencies by an amount determined by the redshift z , which we calculate from the luminosity distance in a standard Λ CDM cosmology with $H_0 = 67.7$ km/s/Mpc and $\Omega_m = 0.3089$, $\Omega_\Lambda = 0.6911$. The dominant mode of the GW signal is the $(2, \pm 2)$ mode, so we will only consider it in this analysis. To get a sense of average detectability, we compute the strain for observers viewing the source from an angle of $\theta = \pi/2.34$ relative to the disk orbital plane, for which the amplitude of the $(2, 2)$ mode equals its θ -averaged value. From this we compute the polarization-averaged response at the detector $\tilde{h}_{\text{res}}$ via (17), and then the

characteristic strain h_c via (16), for which we may then compute the sky-averaged SNR (assuming an optimal matched filter) given detector parameters for Advanced LIGO [5], Cosmic Explorer [6], DECIGO [8], and LISA [7]. The sky-averaged sensitivity factor for a 90° interferometer [Eq. (51) of [66]] along with the sensitivity curves provided by [67] were used for LIGO and CE, while the approximate analytic sky-averaged sensitivities⁴ in [68,69] were used for DECIGO and LISA, respectively.

2. Masses and distances of potential sources

We motivate the choice of mass scale in our detectability analysis by considering three potential formation channels for BH-disk systems similar to those we simulated. At the low end, BH-disk systems of $3 - 20M_\odot$ could result following NSNS and BH-NS mergers [70], and are expected to be reasonably common in the nearby universe. Systems massing $25 - 14M_\odot$ and $250M_\odot$ or greater could be the end results of population III stars ending their lives as collapsars [71–74]. These stars are expected to have peak formation rates in the $z \sim 5-8$ range [75,76], and there is observational evidence for them at $z \approx 6$ [77]. The highest-mass systems of $10^3M_\odot - 10^6M_\odot$ could conceivably result from the collapse of supermassive stars (SMSs), simulations of which have been shown to produce BHs surrounded by disks of about 10% the BH mass [2,78–81]. These events have been proposed to occur in the early Universe and provide the seeds for the formation of supermassive BHs that emerge at $z \approx 7$ [82–88] (see reviews [89–91]).

In addition to considering a range of masses that encompasses all the scales listed above, we also choose a set of three distinct masses to act as representatives for each of the three channels above. These scales, as labeled in Figs. 19 and 20 in the subsequent subsections, are

- (a) $M = 10M_\odot$, representing the possible aftermath of a BH-NS merger.
- (b) $M = 1000M_\odot$, representing the possible aftermath of a low-mass SMS collapse.
- (c) $M = 2 \times 10^5M_\odot$, representing the possible aftermath of an high-mass SMS collapse.

3. Detectability horizons of settled GW signals

Since we aim to gain insight into the GW detectability of a wide class of phenomena that may lead to nonaxisymmetric disks, including the early time transient GW signals

⁴When comparing ground and space-based detectors, a clash of conventions is encountered. For ground-based detectors, averaging over polarization is assumed to be part of the definition of characteristic strain h_c ; for space-based detectors, it is included as part of the sky-averaged response. Here we compromise by multiplying the space-based sensitivities by a factor of $\sqrt{2}$ to account for the $1/\sqrt{2}$ in h_c , effectively removing the contribution of polarization averaging. This rescues the relationship between SNR and the area between the h_c and sensitivity curves.




FIG. 19. Comparison of the detectability by GW observatories LIGO, CE, DECIGO, and LISA, for optimally oriented BH disks that are purely hydrodynamic ($S_{\uparrow}$, S_0) versus those that were seeded with magnetic fields ($S_{\uparrow}^{\text{Mag}}$, S_0^{Mag}), for spinning and nonspinning central BHs, respectively. Labels indicate features corresponding to three mass scales of astrophysical interest: (a) $10M_{\odot}$, (b) $1000M_{\odot}$, (c) $2 \times 10^5 M_{\odot}$. Top: horizons of maximum detectable distance for each system mass for each observatory. The threshold for a system at a given mass and distance on the chart being considered detectable is that its GW signal achieves a sky and polarization averaged $\text{SNR} \geq 8$ (given an optimal matched filter). Line styles denote observatories, while colors indicate which of the four BH-disk configurations studied the GW signal originates from. The + (×) labels show the furthest detectable distance of nonmagnetized (magnetized) configurations for each of the labeled mass scales: (a) 50 Mpc (10 Mpc) by CE, (b) 19000 Mpc (7000 Mpc) by DECIGO, (c) 2600 Mpc (1000 Mpc) by LISA. Bottom: characteristic strain curves for signals at the labeled mass scales and distances marked by the + symbols on the top plot. These strain curves are overlaid on top of noise curves for each of the detectors, and the area between curves is a proxy for the SNR.

is inappropriate. Therefore, we only assess the detectability of the *settled* late-time dynamics, defined as the portion of the signal 2500*M* after the maximum GW amplitude for the magnetized cases, as in Sec. III E. To enable a meaningful comparison between all signals, here they are all truncated to the same end time, which is set to be 4500*M* by our shortest simulation, and therefore all the settled late-time signals in this section have a 2000*M* duration. This is notably shorter than the duration of the signals analyzed in [25], and results in lower maximum detection distances. However, since our simulations do not capture the full signal duration anyway, the absolute detectability distances are less important than relative differences between the configurations.

In Fig. 19 we show the maximum detectability distances for these signals across a range of masses and propagated from a range of distances. That is, for each simulated BH-disk system we choose a set of mass scales, and for each of those masses we compute the GW strain in the detector

frame for a range of luminosity distances (following the methods of Sec. III F 1). We then use the detector-frame strain to compute the SNR for each observatory, assuming an optimal matched filter. The lines in the top panel of the figure show the maximum distances systems at each mass were found to be detectable, where we have chosen an SNR of 8 as the threshold of detectability for this analysis. We see that the spinning cases $S_{\uparrow}$ and $S_{\uparrow}^{\text{Mag}}$ are more detectable than their nonspinning counterparts S_0 and S_0^{Mag} . In [25] we showed this was accounted for by the increased orbital frequency of the disk in the $S_{\uparrow}$ case, due to the smaller ISCO radius of the cospinning BH, so it is not surprising that this carries over to $S_{\uparrow}^{\text{Mag}}$. The difference between the nonmagnetized and magnetized configurations is stark: across all masses the detectability horizon has been reduced by factors of ~ 3 – 5 for the magnetized disks. While the accretion rate is somewhat larger in the magnetized cases, for the times represented here little accretion has had time




FIG. 20. The same quantities as in Fig. 19, but this time for the model waveforms produced by extrapolating our simulated GWs to their expected full duration. Three mass scales of astrophysical interest are again labeled: (a) $10M_\odot$ (b) $1000M_\odot$ (c) $2 \times 10^5 M_\odot$. In the top panel, sources with these masses are marked by + (×) to denote the maximum distances nonmagnetized (*magnetized*) BH-disk configurations of each mass could be detected: (a) 300 Mpc (45 Mpc) by CE, (b) 60,000 Mpc (19,000 Mpc) by DECIGO, (c) 11,500 Mpc (2700 Mpc) by LISA.

to occur and the disks are nearly the same masses. The difference in detectability is therefore predominantly due to the suppression of the $m = 1$ modes that occurs during the growth of the MRI in the magnetized disks.

4. Detectability horizons of extrapolated GW models

The actual detectability of GW signals depends on both their amplitude and duration. Therefore, we need to extrapolate beyond the short duration we simulated to get a model of the full signal for each case. To do this, we use the same simple signal model as [25]: motivated by the quadrupole formula, the signal at late times is assumed to be of the form

$$rh = B \exp((i\omega_0 - \gamma)t + i\phi_0). \quad (19)$$

The amplitude, B , of the waveform is set by the amplitude of the extracted GWs, but the exponent of amplitude falloff, γ , is extracted by fitting the disk mass at late times to a decaying exponential. The frequency ω_0 and initial phase ϕ_0 are set by matching to the unrolled phase profile of the GW signal, and extrapolating it past the final time with a linear best fit. The model fitting is shown in Fig. 21.

We cannot know for certain how long the GW-producing $m = 1$ modes will persist after the end of our evolutions, so we conservatively estimate that the GW signal terminates when only 10% of the original disk mass remains, once again using a smooth falloff to avoid artifacts from the imposed cutoff.

We now take the settled late-time GW signals (after the cutoff of $2500M$), and instead of truncating them at $4500M$ as we did in the previous section, we use their full duration. In the last $500M$, we smoothly match onto the model waveform fitted for that case, which then extends the signal past the end of our simulations.

We perform the same analysis as in the previous section to these extended model waveforms, and the results are shown in Fig. 20. Here, the differences in detectability are even more noticeable: the maximum detection distances are reduced by factors of ~ 3.1 – 6.6 for the magnetized disks. This can at least partially be attributed to the increased accretion rate caused by the magnetic shear, which dramatically shortens the lifetime of the disk extrapolated by the model. However, although the change in signal lifetime is significant, its impact on detectability is not that large because most of the GW signal power comes from the early-time emission, so the reduction of the $m = 1$ amplitude is the dominant effect.




FIG. 21. Comparison of analytic model to quantities extracted from simulations for both magnetized and nonmagnetized accretion disks (indicated by line colors). The dashed curves show analytic model fit to the magnetized disks, the dotted curves show model fit to the nonmagnetized disks, and the solid lines show the quantities being fit: strain amplitude in the top row, disk rest mass in the bottom row. Analytic models are fit using only data after the time marked by the vertical dash-dotted lines, to avoid influence of early-time transients on the extrapolated late-time signal model.

Now that the GW signals are not cut off prematurely by the end of the simulation, the estimates of maximum detectability distances are more realistic. In all cases the spinning BH configurations are most detectable, and in Table IV we list the estimated maximum detection distances for each of the mass scales of interest introduced in Sec. III F 2.

TABLE IV. Maximum detection luminosity distances for the $S_{\uparrow}$ and $S_{\uparrow}^{\text{Mag}}$ BH-disk configurations for a set of masses motivated by plausible formation channels (see Sec. III F 2). In each case, we consider only the detector that achieves its greatest detection distance for the given source. In all cases the signal is not powerful enough for detection by Advanced LIGO, so it is not included here.

Label	BH mass ($M_{\odot}$)	Maximum distance (Mpc)			Observatory
		$S_{\uparrow}$	$S_{\uparrow}^{\text{Mag}}$	$S_{\uparrow}/S_{\uparrow}^{\text{Mag}}$	
(a)	10	300	45	6.67	CE
(b)	1000	60000	19000	3.16	DECIGO
(c)	2×10^5	11500	2700	4.26	LISA

IV. DISCUSSION

In this study, we conducted magnetohydrodynamic simulations in dynamical spacetime comparing the dynamics of self-gravitating accretion tori around spinning and nonspinning BHs for both magnetized and unmagnetized configurations. Our goal was to estimate a rough upper bound on the resilience of the nonlinear saturated state of the PPI in self-gravitating disks to the effects of the MRI on the disk dynamics, so MRI-susceptible field configurations were only seeded once PPI modes reached their maximum amplitude, giving the PPI its best chance of survival. Despite this, these disks, which were previously shown to be powerful sources of gravitational radiation [21,25] due to their persistent $m = 1$ nonaxisymmetries, reacted to the introduction of magnetic fields by exhibiting quick MRI growth and strong suppression of the $m = 1$ mode and the associated gravitational radiation. Witnessing such strong and sudden suppression of these modes, even in the fully nonlinear self-gravitating saturation state, suggests that suppression of the $m = 1$ mode will generally occur for a wide range of conceivable configurations with less developed PPI modes, whenever weak magnetic fields of similar magnitude are present.

These results are similar to the findings of [30] for disks without self-gravity around nonspinning black holes. By constructing a mode-power diagnostic, we are able to approximately compare the power of the low- m nonaxisymmetric modes in our disks to those reported by [30]. We find that the PPI supported $m = 1$ mode is quickly reduced in amplitude, dropping nearly an order of magnitude in four orbits in both $S_{\uparrow}^{\text{Mag}}$ and S_0^{Mag} configurations. Tracking the components of magnetic field flux, we see fast increases in the azimuthal and radial components for the first two orbits, characteristic of MRI growth [29]. After this, the fields saturate, then decay. It is notable (but not surprising) that the majority of the $m = 1$ amplitude reduction happens during the first few orbits, when the magnetic fields are strongest. At late times, the strength of the fields is determined by the equilibrium between dissipation (purely numerical in our simulations) and MRI-driven turbulence [29,63]. In our simulations, the field strength peaks before it relaxes to weaker magnitudes not high above those at initialization, and this seems to allow the PPI to support the $m = 1$ modes against further disruption, resulting in $m = 1$ modes of about 10 times the power of those seen at late times by [30].

When interpreting this result, it is important to consider the known difficulties in adequately resolving steady-state MRI turbulence [51,60,64]. As detailed in Sec. III D, our simulations are slightly lower in resolution in some directions than those of [30], so a careful resolution study would need to be carried out for error estimates in our results. However, the lack of a resolution study does not prevent the results reported here from providing a correct qualitative description of the ability of saturated $m = 1$ PPI modes in massive, self-gravitating disks to survive the

transient growth of the MRI. This is because greater resolution is expected to make the average magnetic field stronger, and the dynamical influence of the MRI and associated turbulence should be even more pronounced [60,64], while the PPI is already well resolved.

Our key finding is therefore that the MRI's ability to reduce the amplitude of large nonaxisymmetric features is robust even in the nonlinear self-gravitating case, and even when starting from extremely nonaxisymmetric disks. Because high-amplitude orbiting nonaxisymmetric patterns are necessary to produce powerful GWs, this result reduces our expectation that such systems will act as detectable GW sources at the large distances determined by unmagnetized models.

To assess detectability, we fit a simple GW model to the extracted gravitational waveforms, and used our study of the observed disk dynamics to extrapolate an optimistic case where the MRI reduces but does not eliminate the disk's $m = 1$ mode, allowing GW production at a reduced amplitude, and for a duration shortened by the more rapid accretion of the magnetized configurations. This optimistic case predicts reductions of the maximum luminosity distances of detection by factors ranging from 3.16 to 6.67. For nonmagnetized disks, the spinning BH configuration $S_{\uparrow}$ was detectable by DECIGO out to redshifts of $z \approx 5$ for systems of mass $1000M_{\odot}$, however the most detectable magnetized configuration $S_{\uparrow}^{\text{Mag}}$ produces a signal too weak to be detected by DECIGO beyond $z \approx 2$, a 76% reduction in luminosity distance.

Finally, we caution that our study does not rule out the potential of accretion disks as engines of detectable gravitational waves. We have shown that a particular nonaxisymmetric disk configuration, which acts as a powerful GW source when not magnetized, quickly becomes significantly more axisymmetric when the MRI is allowed to develop, greatly reducing GW emission. This does not prove that there are no other magnetized disk configurations that could be sustained sources of GW production. The task of finding such configurations, or ruling out their existence, remains to be done.

ACKNOWLEDGMENTS

The following grants supported this work: National Science Foundation (NSF) Grants No. PHY-1912619 and No. PHY-2145421, and National Aeronautics and Space Administration (NASA) Grant No. 80NSSC22K1605 to the University of Arizona; NSF Grant No. PHY-2006066 and NASA Grant No. 80NSSC17K0070 to the University of Illinois at Urbana-Champaign. M.R. acknowledges also support by the Generalitat Valenciana Grant No. CIDEGENT/2021/046 and by the Spanish Agencia Estatal de Investigación (Grant No. PID2021-125485NB-C21). Extreme Science and Engineering Discovery Environment (XSEDE) Grant No. TG-PHY190020

provided the high performance computing resources. XSEDE itself is supported by NSF Grant No. ACI-1548562. A.T. acknowledges support from the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign through the NCSA Fellows program. Simulations were run on the Stampede2 cluster, maintained and operated by the Texas Advanced Computing Center (TACC) at the University of Texas at Austin, supported by NSF Grant No. ACI-1540931. Additional simulations and data analyses were run on the University of Arizona's (UArizona) Ocelote cluster, which is maintained by the UArizona Research Technologies department and supported by UArizona Technology and Research Initiative Fund, UArizona University Information Technology Services, and the UArizona Office of Research, Innovation, and Impact.

APPENDIX: EVOLUTION OF CONSTRAINT VIOLATIONS

Throughout the evolutions we track the constraint violations to ensure the physical validity. In terms of the BSSN variables, the Hamiltonian constraint is

$$\mathcal{H} = \tilde{\gamma}^{ij} \tilde{D}_i \tilde{D}_j e^{\phi} - \frac{e^{\phi}}{8} \tilde{R} + \frac{e^{5\phi}}{8} \tilde{A}_{ij} \tilde{A}^{ij} - \frac{e^{5\phi}}{12} K^2 + 2\pi e^{5\phi} \rho. \quad (\text{A1})$$

We monitor the L_2 norm of $\mathcal{H}$, defined as the square root of the integral of $|\mathcal{H}|^2$ over the volume, which we here




FIG. 22. Evolution of the L_2 norm of the Hamiltonian constraint violations for both configurations, with the magnetized disk evolutions overlaid on the nonmagnetized evolutions.




FIG. 23. Evolution of the L_2 norm of the momentum constraint violations for both configurations, with the magnetized disk evolutions overlayed on the nonmagnetized evolutions.




FIG. 25. Evolution of the differences between the L_2 norms of the momentum constraint violations between the magnetized and nonmagnetized evolutions, for both configurations.




FIG. 24. Evolution of the differences between the L_2 norms of the Hamiltonian constraint violations between the magnetized and nonmagnetized evolutions, for both configurations.

denote $||\mathcal{H}||$. Figure 22 shows the values of this constraint for the purely hydrodynamic and magnetized runs.

The momentum constraint is

$$\mathcal{M}^i = \tilde{D}_j (e^{6\phi} \tilde{A}^{ji}) - \frac{2}{3} e^{6\phi} \tilde{D}^i K - 8\pi e^{6\phi} S^i. \quad (\text{A2})$$

In this case, we monitor the same L_2 norm as above for each of the three components of $\mathcal{M}^i$. We then take the norm of this vector: $||\mathcal{M}|| \equiv \sqrt{||\mathcal{M}^x||^2 + ||\mathcal{M}^y||^2 + ||\mathcal{M}^z||^2}$. This is plotted in Fig. 23.

As can be seen, the constraint values are nearly identical after field insertion, indicating that the initial field has a small enough energy density to avoid introducing any noticeable constraint violations. It is also clear that the differing dynamics between the two simulations have little impact on the constraint evolution. To verify that the near-indistinguishability of the two constraint evolutions between the magnetized and unmagnetized cases is not a mistake in our plot, Figs. 24 and 25 show the differences between the two sets of constraint violations for each configuration.

- [1] M. J. Rees, *Annu. Rev. Astron. Astrophys.* **22**, 471 (1984).
- [2] M. Shibata and S. L. Shapiro, *Astrophys. J.* **572**, L39 (2002).
- [3] G. Lovelace, M. D. Duez, F. Foucart, L. E. Kidder, H. P. Pfeiffer, M. A. Scheel, and B. Szilágyi, *Classical Quantum Gravity* **30**, 135004 (2013).
- [4] V. Paschalidis, M. Ruiz, and S. L. Shapiro, *Astrophys. J. Lett.* **806**, L14 (2015).
- [5] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams *et al.* (The LIGO Scientific Collaboration, The Virgo Collaboration, and The KAGRA Collaboration), *Living Rev. Relativity* **23**, 3 (2020).
- [6] D. Reitze, R. X. Adhikari, S. Ballmer, B. Barish, L. Barsotti, G. Billingsley, D. A. Brown, Y. Chen, D. Coyne, R. Eisenstein *et al.*, *Bull. Am. Astron. Soc.* **51**, 035 (2019).
- [7] P. Amaro-Seoane, H. Audley, S. Babak, J. Baker, E. Barausse, P. Bender, E. Berti, P. Binetruy, M. Born, D. Bortoluzzi *et al.*, [arXiv:1702.00786](https://arxiv.org/abs/1702.00786).
- [8] S. Sato, S. Kawamura, M. Ando, T. Nakamura, K. Tsubono, A. Araya, I. Funaki, K. Ioka, N. Kanda, S. Moriwaki *et al.*, *J. Phys. Conf. Ser.* **840**, 012010 (2017).
- [9] J. C. B. Papaloizou and J. E. Pringle, *Mon. Not. R. Astron. Soc.* **208**, 721 (1984).
- [10] J. F. Hawley, *Mon. Not. R. Astron. Soc.* **225**, 677 (1987).
- [11] O. M. Blaes, *Mon. Not. R. Astron. Soc.* **216**, 553 (1985).
- [12] O. M. Blaes and W. Glatzel, *Mon. Not. R. Astron. Soc.* **220**, 253 (1986).
- [13] P. Goldreich, J. Goodman, and R. Narayan, *Mon. Not. R. Astron. Soc.* **221**, 339 (1986).
- [14] R. Narayan, P. Goldreich, and J. Goodman, *Mon. Not. R. Astron. Soc.* **228**, 1 (1987).
- [15] J. Goodman and R. Narayan, *Mon. Not. R. Astron. Soc.* **231**, 97 (1988).
- [16] D. M. Christodoulou and R. Narayan, *Astrophys. J.* **388**, 451 (1992).
- [17] J. C. B. Papaloizou and J. E. Pringle, *Mon. Not. R. Astron. Soc.* **213**, 799 (1985).
- [18] J. W. Strutt, *Proc. R. Soc. A* **93**, 148 (1917).
- [19] Y. Kojima, *Prog. Theor. Phys.* **75**, 1464 (1986).
- [20] O. Korobkin, E. B. Abdikamalov, E. Schnetter, N. Stergioulas, and B. Zink, *Phys. Rev. D* **83**, 043007 (2011).
- [21] K. Kiuchi, M. Shibata, P. J. Montero, and J. A. Font, *Phys. Rev. Lett.* **106**, 251102 (2011).
- [22] R. Nealon, D. J. Price, C. Bonnerot, and G. Lodato, *Mon. Not. R. Astron. Soc.* **474**, 1737 (2018).
- [23] M. Toscani, G. Lodato, and R. Nealon, *Mon. Not. R. Astron. Soc.* **489**, 699 (2019).
- [24] V. Mewes, J. A. Font, F. Galeazzi, P. J. Montero, and N. Stergioulas, *Phys. Rev. D* **93**, 064055 (2016).
- [25] E. Wessel, V. Paschalidis, A. Tsokaros, M. Ruiz, and S. L. Shapiro, *Phys. Rev. D* **103**, 043013 (2021).
- [26] A. Tsokaros, M. Ruiz, S. L. Shapiro, and V. Paschalidis, *Phys. Rev. D* **106**, 104010 (2022).
- [27] M. Shibata, K. Kiuchi, S. Fujibayashi, and Y. Sekiguchi, *Phys. Rev. D* **103**, 063037 (2021).
- [28] S. A. Balbus and J. F. Hawley, *Astrophys. J.* **400**, 610 (1992).
- [29] S. A. Balbus and J. F. Hawley, *Rev. Mod. Phys.* **70**, 1 (1998).
- [30] M. Bugli, J. Guilet, E. Müller, L. Del Zanna, N. Bucciantini, and P. J. Montero, *Mon. Not. R. Astron. Soc.* **475**, 108 (2018).
- [31] D. Christodoulou, *Phys. Rev. Lett.* **25**, 1596 (1970).
- [32] G. Bozzola, *J. Open Source Software* **6**, 3099 (2021).
- [33] A. Tsokaros, K. Uryū, and S. L. Shapiro, *Phys. Rev. D* **99**, 041501 (2019).
- [34] H. Komatsu, Y. Eriguchi, and I. Hachisu, *Mon. Not. R. Astron. Soc.* **237**, 355 (1989).
- [35] A. A. Tsokaros and K. Uryū, *Phys. Rev. D* **75**, 044026 (2007).
- [36] A. Toomre, *Astrophys. J.* **139**, 1217 (1964).
- [37] Z. B. Etienne, J. A. Faber, Y. T. Liu, S. L. Shapiro, and T. W. Baumgarte, *Phys. Rev. D* **76**, 101503 (2007).
- [38] M. D. Duez, Y. T. Liu, S. L. Shapiro, and B. C. Stephens, *Phys. Rev. D* **72**, 024028 (2005).
- [39] Z. B. Etienne, Y. T. Liu, and S. L. Shapiro, *Phys. Rev. D* **82**, 084031 (2010).
- [40] Z. B. Etienne, Y. T. Liu, V. Paschalidis, and S. L. Shapiro, *Phys. Rev. D* **85**, 064029 (2012).
- [41] Cactus, <https://www.cactuscode.org/>.
- [42] E. Schnetter, S. H. Hawley, and I. Hawke, *Classical Quantum Gravity* **21**, 1465 (2004).
- [43] Z. B. Etienne, V. Paschalidis, R. Haas, P. Mösta, and S. L. Shapiro, *Classical Quantum Gravity* **32**, 175009 (2015).
- [44] M. Shibata and T. Nakamura, *Phys. Rev. D* **52**, 5428 (1995).
- [45] T. W. Baumgarte and S. L. Shapiro, *Phys. Rev. D* **59**, 024007 (1998).
- [46] M. Campanelli, C. O. Lousto, P. Marronetti, and Y. Zlochower, *Phys. Rev. Lett.* **96**, 111101 (2006).
- [47] J. G. Baker, J. Centrella, D.-I. Choi, M. Koppitz, and J. van Meter, *Phys. Rev. Lett.* **96**, 111102 (2006).
- [48] C. Bona, J. Massó, E. Seidel, and J. Stela, *Phys. Rev. Lett.* **75**, 600 (1995).
- [49] M. Alcubierre, B. Brügmann, P. Diener, M. Koppitz, D. Pollney, E. Seidel, and R. Takahashi, *Phys. Rev. D* **67**, 084023 (2003).
- [50] I. Hinder, A. Buonanno, M. Boyle, Z. B. Etienne, J. Healy, N. K. Johnson-McDaniel, A. Nagar, H. Nakano, Y. Pan, H. P. Pfeiffer *et al.*, *Classical Quantum Gravity* **31**, 025012 (2014).
- [51] J. F. Hawley, X. Guan, and J. H. Krolik, *Astrophys. J.* **738**, 84 (2011).
- [52] V. Paschalidis, *Classical Quantum Gravity* **34**, 084002 (2017).
- [53] R. Gold, V. Paschalidis, Z. B. Etienne, S. L. Shapiro, and H. P. Pfeiffer, *Phys. Rev. D* **89**, 064060 (2014).
- [54] V. Paschalidis, J. Bright, M. Ruiz, and R. Gold, *Astrophys. J. Lett.* **910**, L26 (2021).
- [55] B. D. Farris, R. Gold, V. Paschalidis, Z. B. Etienne, and S. L. Shapiro, *Phys. Rev. Lett.* **109**, 221102 (2012).
- [56] Z. B. Etienne, J. A. Faber, Y. T. Liu, S. L. Shapiro, K. Taniguchi, and T. W. Baumgarte, *Phys. Rev. D* **77**, 084002 (2008).
- [57] V. Paschalidis, W. E. East, F. Pretorius, and S. L. Shapiro, *Phys. Rev. D* **92**, 121502 (2015).
- [58] W. E. East, V. Paschalidis, and F. Pretorius, *Classical Quantum Gravity* **33**, 244004 (2016).
- [59] W. E. East, V. Paschalidis, F. Pretorius, and S. L. Shapiro, *Phys. Rev. D* **93**, 024011 (2016).

- [60] J. F. Hawley, S. A. Richers, X. Guan, and J. H. Krolik, *Astrophys. J.* **772**, 102 (2013).
- [61] C. Reisswig and D. Pollney, *Classical Quantum Gravity* **28**, 195015 (2011).
- [62] P. L. Espino, V. Paschalidis, T. W. Baumgarte, and S. L. Shapiro, *Phys. Rev. D* **100**, 043014 (2019).
- [63] O. Blaes, *Space Sci. Rev.* **183**, 21 (2014).
- [64] B. R. Ryan, C. F. Gammie, S. Fromang, and P. Kestener, *Astrophys. J.* **840**, 6 (2017).
- [65] H. Shiokawa, J. C. Dolence, C. F. Gammie, and S. C. Noble, *Astrophys. J.* **744**, 187 (2012).
- [66] C. J. Moore, R. H. Cole, and C. P. L. Berry, *Classical Quantum Gravity* **32**, 015014 (2015).
- [67] M. Evans, J. Harms, and S. Vitale (2016), <https://dcc.ligo.org/LIGO-P1600143/public>.
- [68] K. Yagi and N. Seto, *Phys. Rev. D* **83**, 044011 (2011).
- [69] T. Robson, N. J. Cornish, and C. Liu, *Classical Quantum Gravity* **36**, 105011 (2019).
- [70] R. Voss and T. M. Tauris, *Mon. Not. R. Astron. Soc.* **342**, 1169 (2003).
- [71] A. I. MacFadyen and S. E. Woosley, *Astrophys. J.* **524**, 262 (1999).
- [72] A. I. MacFadyen, S. E. Woosley, and A. Heger, *Astrophys. J.* **550**, 410 (2001).
- [73] A. Heger and S. E. Woosley, *Astrophys. J.* **567**, 532 (2002).
- [74] A. Heger, C. L. Fryer, S. E. Woosley, N. Langer, and D. H. Hartmann, *Astrophys. J.* **591**, 288 (2003).
- [75] L. Tornatore, A. Ferrara, and R. Schneider, *Mon. Not. R. Astron. Soc.* **382**, 945 (2007).
- [76] J. L. Johnson, V. C. Dalla, and S. Khochfar, *Mon. Not. R. Astron. Soc.* **428**, 1857 (2013).
- [77] D. Sobral, J. Matthee, B. Darvish, D. Schaerer, B. Mobasher, H. J. A. Röttgering, S. Santos, and S. Hemmati, *Astrophys. J.* **808**, 139 (2015).
- [78] S. L. Shapiro and M. Shibata, *Astrophys. J.* **577**, 904 (2002).
- [79] S. L. Shapiro, *Astrophys. J.* **610**, 913 (2004).
- [80] L. Sun, V. Paschalidis, M. Ruiz, and S. L. Shapiro, *Phys. Rev. D* **96**, 043006 (2017).
- [81] H. Uchida, M. Shibata, T. Yoshida, Y. Sekiguchi, and H. Umeda, *Phys. Rev. D* **96**, 083016 (2017).
- [82] A. Loeb and F. A. Rasio, *Astrophys. J.* **432**, 52 (1994).
- [83] S. L. Shapiro, *AIP Conf. Proc.* **686**, 50 (2003).
- [84] S. M. Koushiappas, J. S. Bullock, and A. Dekel, *Mon. Not. R. Astron. Soc.* **354**, 292 (2004).
- [85] S. L. Shapiro, *Astrophys. J.* **620**, 59 (2005).
- [86] M. C. Begelman, M. Volonteri, and M. J. Rees, *Mon. Not. R. Astron. Soc.* **370**, 289 (2006).
- [87] G. Lodato and P. Natarajan, *Mon. Not. R. Astron. Soc.* **371**, 1813 (2006).
- [88] M. C. Begelman, *Mon. Not. R. Astron. Soc.* **402**, 673 (2010).
- [89] Z. Haiman, *Astrophys. Space Sci. Libr.* **396**, 293 (2013).
- [90] M. A. Latif and A. Ferrara, *Publ. Astron. Soc. Aust.* **33**, E051 (2016).
- [91] A. Smith, V. Bromm, and A. Loeb, *Astron. Geophys.* **58**, 3.22 (2017).