

Gas-driven massive black hole binaries: signatures in the nHz gravitational wave background

B. Kocsis^{1,2,3★†} and A. Sesana^{4,5★}

¹Harvard–Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA

²Institute for Advanced Study, Einstein Dr, Princeton, NJ 08540, USA

³Eötvös University, Pázmány P. st. 1/A, Budapest H-1117, Hungary

⁴Albert Einstein Institute, Am Mühlenberg 1 D-14476 Golm, Germany

⁵Center for Gravitational Wave Physics, The Pennsylvania State University, University Park, PA 16802, USA

Accepted 2010 September 28. Received 2010 September 2; in original form 2010 January 27

ABSTRACT

Pulsar timing arrays (PTAs) measure nHz frequency gravitational waves (GWs) generated by orbiting massive black hole binaries (MBHBs) with periods between 0.1 and 10 yr. Previous studies on the nHz GW background assumed that the inspiral is purely driven by GWs. However, torques generated by a gaseous disc can shrink the binary much more efficiently than GW emission, reducing the number of binaries at these separations. We use simple disc models for the circumbinary gas and for the binary–disc interaction to follow the orbital decay of MBHBs through physically distinct regions of the disc, until GWs take over their evolution. We extract MBHB cosmological merger rates from the Millennium simulation, generate Monte Carlo realizations of a population of gas-driven binaries and calculate the corresponding GW amplitudes of the most luminous individual binaries and the stochastic GW background. For steady state α -discs with $\alpha > 0.1$, we find that the nHz GW background can be significantly modified. The number of resolvable binaries is however not changed by the presence of gas; we predict 1–10 individually resolvable sources to stand above the noise for a 1–50 ns timing precision. Gas-driven migration reduces predominantly the number of small total mass or unequal mass ratio binaries, which leads to the attenuation of the mean stochastic GW background, but increases the detection significance of individually resolvable binaries. The results are sensitive to the model of binary–disc interaction. The GW background is not attenuated significantly for time-dependent models of Ivanov, Papaloizou & Polnarev.

Key words: black hole physics – gravitational waves – planet–disc interactions – pulsars: general – cosmology: theory.

1 INTRODUCTION

Inspiring massive black hole binaries (MBHBs) with masses in the range $\sim 10^4$ – $10^{10} M_{\odot}$ are expected to be the dominant source of gravitational waves (GWs) at \sim nHz–mHz frequencies (Haehnelt 1994; Jaffe & Backer 2003; Wyithe & Loeb 2003; Sesana et al. 2004, 2005). The frequency band $\sim 10^{-5}$ –1 Hz will be probed by the *Laser Interferometer Space Antenna* (LISA; Bender et al. 1998), a space-borne GW laser interferometer developed by European Space Agency and National Aeronautics and Space Administration. The observational window 10^{-9} – 10^{-6} Hz, corresponding roughly to orbital periods 0.03–30 yr, is already accessible using pulsar timing arrays (PTAs, e.g. the Parkes radio-telescope; Manchester

2008). The complete Parkes PTA (PPTA; Manchester 2008), the European PTA (EPTA; Janssen et al. 2008) and NanoGrav (Jenet et al. 2009) are expected to improve considerably on the capabilities of these surveys, eventually joining their efforts in the international PTA project (IPTA; Hobbs et al. 2010); and the planned Square Kilometer Array (SKA; Lazio 2009) will produce a major leap in sensitivity.

Radio pulses generated by rotating neutron stars travel through the Galactic interstellar medium and are detected by radio telescopes on the Earth. The arrival times of pulses are fitted for a model including all the known and measured systematic effects affecting the signal generation, propagation and detection (Edwards, Hobbs & Manchester 2006). Timing residuals between the observed pulses and the best-fitting model carry information on additional unmodelled effects, including the presence of GWs. Indeed, GWs modify the propagation of radio signals from the pulsar to the Earth (Sazhin 1978; Detweiler 1979; Bertotti, Carr & Rees 1983; Hellings

★E-mail: bkocsis@cfa.harvard.edu (BK); alberto.sesana@aei.mpg.de (AS)

†Einstein Fellow.

& Downs 1983; Jenet et al. 2005) and PTAs measure the direction-dependent systematic variations in the arrival times of signals from a sample of nearly stationary pulsars in the Galaxy distributed over the sky.

PTAs provide a direct observational window on to the MBH binary population, and can contribute to address a number of open astrophysical questions, such as the shape of the bright end of the MBH mass function, the nature of the MBH–bulge relation at high masses and the dynamical evolution at subpc scales of the most massive binaries in the Universe (particularly relevant to the so-called ‘final pc problem’, Milosavljevic & Merritt 2003) PTAs can detect gravitational radiation of two forms: (i) the stochastic GW background produced by the incoherent superposition of radiation from the whole cosmic population of MBHBs and (ii) individual sources that are sufficiently bright in GWs to outshine the background (typically massive, $M \gtrsim 10^9 M_\odot$, and ‘cosmologically nearby’, $d_L \lesssim 3$ Gpc). Both classes of signals are of great interest, and PTAs could lead to the discovery of systems difficult to detect with other techniques [for alternatives using active galactic nuclei (AGNs), see Haiman, Kocsis & Menou 2009 and references therein].

Popular scenarios of MBH formation and evolution (e.g. Volonteri, Haardt & Madau 2003; Koushiappas & Zentner 2006; Malbon et al. 2007; Yoo et al. 2007) predict frequent MBH mergers (up to several hundreds per year), implying the existence of a large number of subpc MBHBs. The prospect for detecting GW signals using PTAs depends on the number and cosmological distribution of MBHBs with orbital periods of 0.03–30, or separations typically in the range 0.001–0.1 pc. The three main ingredients for calculating the GW background are

- (i) the merger rate of MBHBs as a function of mass and redshift,
- (ii) the relative time each binary spends at these separations during a merger episode and
- (iii) the amplitude of the GW signal produced by each individual stationary system.

Recently, Sesana, Vecchio & Colacino (2008, hereafter SVC08) and Sesana, Vecchio & Volonteri (2009, hereafter SVV09) carried out a detailed study of the expected signals (stochastic and individual), focusing on the uncertainties related to (i). They found that the background is affected by the galaxy merger rate evolution along the cosmic history, the massive black hole mass function and the accretion history of the MBHB during a galaxy merger, and they predict a factor of ~ 10 uncertainty for the characteristic strain amplitude in the range (2×10^{-15}) – (2×10^{-16}) , at $f = 1 \text{ yr}^{-1}$, within the expected detection capabilities of the complete PPTA and of the SKA. They pointed out that the GW signal can be separated into individually resolvable sources and a stochastic background, and found the number of individually resolvable sources for a 1-ns timing precision level to be between 5 and 15, depending on the considered model.

In this paper, we examine for the first time how predictions relevant for PTA observations are modified by the presence of ambient gas, affecting the inspiral rate of binaries during a merger episode, ingredient (ii) above. A gaseous envelope is expected to surround the binary because MBHBs are produced in galaxy mergers, which are known to trigger inflows of large quantities of gas into the central region, as shown by hydrodynamic simulations (Springel et al. 2005). This gas, accreted on to the MBHs, is responsible for luminous AGN activity, and is also expected to catalyze the coalescence of the new-formed MBH pair (e.g. Escala et al. 2004; Dotti et al. 2007), as described below. The forming MBHB spirals inward initially as a result of dynamical friction on dark matter, ambient stars

and gas (Begelman, Blandford & Rees 1980). As the binary separation shrinks to subpc scales, the supply rate of stars crossing the orbit decreases, and the interaction with stars becomes less and less efficient to shrink the binary. In gas rich mergers, the dense nuclear gas is expected to cool rapidly and settle into a geometrically thin circumbinary accretion disc (e.g. Barnes 2002; Escala et al. 2004). Torques from the tidal field of the binary clear a gap in the gas with a radius less than twice the separation of the binary, and generates a spiral density wave in the gaseous disc, which in turn drains angular momentum away from the binary on a relatively short time-scale within the last pc, $\lesssim 10^7$ yr (Armitage & Natarajan 2002, 2005; Escala et al. 2005; Dotti et al. 2007; Hayasaki 2009, see however Lodato et al. 2009). Ultimately, at even smaller separations, corresponding to an orbital time-scale of \sim yr, the emission of GWs becomes the dominant mechanism driving the binary to the final coalescence. The main point of this paper is to notice that the most sensitive PTA frequency band corresponds to orbital separations near the transition between gas and GW-dominated evolution. *As binaries shrink more quickly inwards in the gas-driven phase, the number of binaries emitting at each given separation is decreased compared to the purely GW-driven case.* The subject of this work is to explore how various gas-driven models modify the expectations on the GW signal potentially observable by PTAs.

Recently, Haiman et al. (2009, hereafter HKM09) examined the evolution of MBHBs in the gas-driven regime for simple models of geometrically thin circumbinary discs (see also Syer & Clarke 1995; Armitage & Natarajan 2005). The interaction between the binary and the gaseous disc is analogous to type-II planetary migration, and evolves through two main phases. First, the inspiral is analogous to the disc-dominated type-II migration of planetary dynamics, where the binary migrates inwards with a radial velocity equal to that of the gas accreting towards the centre. Later, as the mass of the gas within a few binary separations becomes less than the reduced mass of the binary, the evolution slows down, and it is analogous to the planet-dominated (or secondary-dominated) type-II migration. In both cases, the radial inspiral rate is still much faster than in the purely GW driven case at orbital separations beyond a few 100 Schwarzschild radii. For standard Shakura–Sunyaev α -disc models, the viscosity is assumed to be proportional to the total pressure, and is consequently very large in the radiation pressure dominated phase at small radii, increasing the migration rate in the radiation pressure dominated phase. On the other hand, for β -disc models, where the viscosity is proportional to the gas pressure only, the increase of radiation pressure does not impact the viscous time-scale, and the migration rate is relatively slower in this regime. Finally, we note that the binary–gas interaction is also significantly different for non-steady models of accretion (Ivanov et al. 1999; HKM09). In the typical secondary-dominated phase, gas flows in more quickly than how the binary separation shrinks, and is repelled close to the outer edge of the gap by the torques of the binary. This causes gas to accumulate near the gap, and delays the merger of the binary relative to the steady-state models.

In this paper, we couple the HKM09 models for the migration of MBHBs in the presence of a steady gaseous disc, to the population models derived in SVV09, and we compute the effects on GWs at nHz frequencies. Here, we restrict to nearly circular inspirals for simplicity, using the corresponding GW spectrum [ingredient (iii) above]. This assumption might be violated in gas-driven inspirals (Armitage & Natarajan 2005; Cuadra et al. 2009), and is the subject of a future paper (Sesana & Kocsis, in preparation).

The paper is organized as follows. In Section 2 we introduce the theory of the GW signal from an MBHB population, describing its

characterization in terms of its *stochastic level* and of the statistics of *individually resolvable sources*. In Section 3, we describe our MBHB population model, coupling models of coalescing binaries derived by cosmological N -body simulations to a scheme for the dynamical evolution of the binaries in massive circumbinary discs. We present in detail our results in Section 4, and in Section 5 we briefly summarize our main findings. Throughout the paper we use geometric units with $G = c = 1$.

2 DESCRIPTION OF THE GRAVITATIONAL WAVE SIGNAL

The theory of the GW signal produced by the superposition of radiation from a large number of individual sources, was extensively presented in SVC08 and SVV09, here we review the basic concepts, deriving the GW signal for gas-driven mergers.

The dimensionless characteristic amplitude of the GW background produced by a population of binaries with a range of masses m_1 and m_2 and redshifts z is given by (Phinney 2001)

$$h_c^2(f) = \frac{4}{\pi f^2} \int dz dm_1 dm_2 \frac{\partial^3 n}{\partial z \partial m_1 \partial m_2} \frac{1}{1+z} \frac{dE_{\text{gw}}}{d \ln f_r}, \quad (1)$$

where $dE_{\text{gw}}/d \ln f_r$ is the total emitted GW energy per logarithmic frequency interval in the comoving binary rest frame, $f_r = (1+z)f$ is the rest-frame frequency, f is the observed frequency, n is the comoving number density of sources and the $1/(1+z)$ factor accounts for the redshift of the observed GW energy. The characteristic GW amplitude h_c is related to the present-day total energy in GWs as $\rho_{\text{gw}} = (\pi/4) \int f h_c^2(f) df$.

2.1 GW-driven inspirals

To the leading quadrupole order, for circular purely GW-driven binaries orbiting far outside the innermost stable circular orbit (ISCO), equation (1) can be evaluated assuming that $E_{\text{gw}} = E_{\text{pot}} = -m_1 m_2 / a$ is equal to the Newtonian potential energy of the binary, and that f_r is equal to twice the Keplerian orbital frequency (Phinney 2001). In this case,

$$\frac{dE_{\text{gw}}}{d \ln f_r} = \frac{1}{3} \mu (\pi M f_r)^{2/3} = \frac{1}{3} (\pi f_r)^{2/3} \mathcal{M}^{5/3} \quad (2)$$

for $f_r < f_{\text{ISCO}} \equiv 1/(6^{3/2} \pi M)$. Here $\mu = m_1 m_2 / M$, $M = m_1 + m_2$ and $\mathcal{M}^{5/3} = \mu M^{2/3}$ are the reduced, total and chirp masses for a binary, and f_{ISCO} is the GW frequency at ISCO. Substituting into equation (1),

$$h_c^2(f) = \frac{4 f^{-4/3}}{3 \pi^{1/3}} \int dz dm_1 dm_2 \frac{\partial^3 n}{\partial z \partial m_1 \partial m_2} \frac{\mu M^{2/3}}{(1+z)^{1/3}}. \quad (3)$$

It is also useful to examine the number of MBHBs and their respective contributions to the total signal (Phinney 2001; SVC08). Equation (3) can be rewritten as

$$h_c^2(f) = \int dz dm_1 dm_2 \frac{\partial^4 N}{\partial m_1 \partial m_2 \partial z \partial \ln f_r} h^2(\mathcal{M}, z, f_r), \quad (4)$$

where N is the number of sources, which we can calculate from a comoving merger rate density as explained below, and

$$h(\mathcal{M}, z, f_r) = \frac{8}{10^{1/2}} \frac{\mathcal{M}}{d_L(z)} (\pi \mathcal{M} f_r)^{2/3}, \quad (5)$$

is the sky and polarization-averaged characteristic GW strain amplitude of a single binary with chirp mass \mathcal{M} , at the particular orbital radius corresponding to f_r .

We generate the distribution $\partial^4 N / (\partial m_1 \partial m_2 \partial z \partial \ln f_r)$ corresponding to a comoving merger rate density¹ $\partial^4 N / (\partial m_1 \partial m_2 \partial t_r \partial V_c)$ (see Section 3.1), assuming that the number of binaries emitting in the interval $\ln f_r$ is proportional to the time the binary spends at that frequency:

$$\frac{\partial^4 N}{\partial m_1 \partial m_2 \partial z \partial \ln f_r} = \frac{\partial^4 N}{\partial m_1 \partial m_2 \partial t_r \partial V_c} \frac{dV_c}{dz} \frac{dz}{dt_r} \frac{dt_r}{d \ln f_r}, \quad (6)$$

where dV_c/dz and dz/dt_r are given by the standard cosmological relations between comoving volume, redshift and time, given in e.g. Phinney (2001). The last factor can be expressed using the residence time $t_{\text{res}} = a(da/dt_r)^{-1}$ the binary spends at a particular semimajor axis as

$$\left| \frac{dt_r}{d \ln f_r} \right| = \left| \frac{dt_r}{d \ln a} \frac{d \ln a}{d \ln f_r} \right| = \frac{2}{3} t_{\text{res}}, \quad (7)$$

where Kepler's law, $a = M(\pi M f_r)^{-2/3}$, was used to obtain $d \ln a / d \ln f_r = 2/3$, and the residence time for a purely GW-driven evolution is

$$t_{\text{res}} = t_{\text{res}}^{\text{gw}} \equiv \frac{dt_r}{d \ln a} = \frac{5}{64} \mathcal{M} (\pi \mathcal{M} f_r)^{-8/3}. \quad (8)$$

In summary, the distribution of sources in equation (4) becomes

$$\frac{\partial^4 N}{\partial m_1 \partial m_2 \partial z \partial \ln f_r} = \frac{2}{3} \frac{\partial^4 N}{\partial m_1 \partial m_2 \partial t_r \partial V_c} \frac{dV_c}{dz} \frac{dz}{dt_r} t_{\text{res}}. \quad (9)$$

Equations (3) and (4–9) are equivalent, but (4–9) are practical to generate discrete Monte Carlo realizations of a given source population. Moreover, equations (4–9) provide a transparent interpretation of (3). The total rms background, $h_c \propto \sqrt{N} h \propto f^{-2/3}$, comes about because the mean number of binaries per frequency bin is $N \propto t_{\text{res}}^{\text{gw}} \propto f_r^{-8/3}$ and each binary generates an rms strain $h \propto f_r^{2/3}$. The scaling $h_c \propto f^{-2/3}$ is a consequence of averaging over the local inspiral episodes of merging binaries in the GW driven regime, but is completely independent on the overall cosmological merger history or on the involved MBH masses. The latter affects only the overall constant of proportionality. Equation (3) also shows that this scaling constant is insensitive to the cosmological redshift distribution of mergers, $h_c \propto (1+z)^{1/6}$, as well as the number of minor mergers with $\mu \ll M$, once the total mass in satellites merging with BHs of mass M is fixed (Phinney 2001). However, the background is sensitive to the assembly scenarios of major mergers (SVV09). More importantly, as shown in SVC08, the actual GW signal in any single realization of inspiralling binaries is qualitatively different from $h_c \propto f^{-2/3}$, as a small discrete number of individual massive binaries dominate the nHz GW background, creating a very spiky GW spectrum. We further discuss the discrete nature of the signal in Section 2.3 below.

2.2 Gas-driven inspirals

Let us now derive the GW background for an arbitrary model of binary evolution. We can derive the background using the residence time t_{res} the binary spends at each semimajor axis a , where $t_{\text{res}} < t_{\text{res}}^{\text{gw}}$ in the gas-driven phase (cf. equation 8 for the purely GW-driven case). We define t_{res} for various accretion-disc models in Section 3.2. Generally, the emitted GW spectrum is

$$\frac{dE_{\text{gw}}}{d \ln f_r} = \frac{dE_{\text{gw}}}{dt_r} \frac{dt_r}{d \ln f_r} = \frac{2}{3} \frac{dE_{\text{gw}}}{dt_r} t_{\text{res}}. \quad (10)$$

¹ In practice, $\partial^4 N / (\partial m_1 \partial m_2 \partial t_r \partial V_c)$ is a function of m_1 , m_2 and z .

The second equality follows the definition of t_{res} and Kepler's law (see Equation 7). The emitted power dE_{gw}/dt_r depends only on the masses and the geometry of the orbit, but is independent of the global migration rate of the binary, and therefore it is the same as in the pure GW-driven case. The effects of migration is fully encoded in t_{res} . Plugging back in equation (1), the mean square signal in the gas-driven phase is

$$h_c^2(f) = \frac{4f^{-4/3}}{3\pi^{1/3}} \int dz dm_1 dm_2 \frac{\partial^3 n}{\partial z \partial m_1 \partial m_2} \frac{\mu M^{2/3}}{(1+z)^{1/3}} \times \frac{t_{\text{res}}(M, \mu, f_r)}{t_{\text{res}}^{\text{gw}}(M, f_r)}. \quad (11)$$

A comparison with (3) shows that the rms signal is attenuated in the gas-driven phase by $\sqrt{t_{\text{res}}(M, \mu, f_r)/t_{\text{res}}^{\text{gw}}(M, f_r)}$. This factor is a complicated function of f_r , that behaves differently for different masses and mass ratios of the binaries; the overall spectrum is no longer a power law.

We generate Monte Carlo realizations of the GW signal by sampling the population of the inspiralling systems. To do this, it is sufficient to recognize that the derivation given by equations (4–9) remains valid in the gas-driven phase if using the appropriate $t_{\text{res}}(M, \mu, f_r)$, since the contribution given by each individual source to the signal is the same. *The net GW spectrum changes because of a reduction in the number of sources in the gas-driven case.*

Note that the individual GW signal given by equation (5) depends on three parameters only: \mathcal{M} , z and f_r . This implies that signals from sources with the same observed frequency, redshift and chirp mass, (f_r, z, \mathcal{M}) , but different mass ratios, q , are totally indistinguishable.² We can make use of this property and reduce the number of independent parameters in the distribution by integrating over the mass ratio in equation (4):

$$\frac{\partial^3 N}{\partial \mathcal{M} \partial z \partial \ln f_r} = \int_0^1 dq \frac{\partial^4 N}{\partial m_1 \partial m_2 \partial z \partial \ln f_r} \left| \frac{\partial(m_1, m_2)}{\partial(\mathcal{M}, q)} \right|, \quad (12)$$

Note, that this step is different for the GW- and gas-driven cases, because $\partial^4 N/(\partial m_1 \partial m_2 \partial z \partial \ln f_r)$ is proportional to t_{res} in equation (9). Here $|\partial(m_1, m_2)/\partial(\mathcal{M}, q)|$ is the determinant of the Jacobian matrix corresponding to the variable change from (m_1, m_2) to (\mathcal{M}, q) . With equations (9) and (12), we derive $\partial^3 N/(\partial \mathcal{M} \partial z \partial \ln f_r)$ for any gas-driven model given by $t_{\text{res}}(M, \mu, f_r)$, and draw Monte Carlo samples of inspiralling binaries from this distribution when generating the GW signal.

2.3 Statistical characterization of the signal

In observations with PTAs, radio pulsars are monitored weekly for total periods of several years. Assuming a repeated observation in uniform Δ time intervals for a total time T , the maximum and minimum resolvable frequencies are $f_{\text{max}} = 1/(2\Delta)$, corresponding to the Nyquist frequency, and $f_{\text{min}} = 1/T$. The observed GW spectrum is therefore discretely sampled in bins of $\Delta = f_{\text{min}}$. For circular orbits, the frequency of the GWs is twice the orbital frequency.

Let us examine whether the sources' GW frequency evolves during the observation relative to the size of the frequency bins. Writing the frequency shift during an observation time T as $\Delta f_{\text{evol}} \approx \dot{f}T = (d \ln f / d \ln a)(d \ln a / dt) f T = (3/2) f T / [(1+z)t_{\text{res}}]$ and considering the frequency resolution bin to be $\Delta f_{\text{bin}} = 1/T$,

² This is true only in the angular-averaged approximation. We neglect the directional sensitivity of PTAs.




Figure 1. Residence time, t_{res} , as a function of the binary period. Different panels correspond to selected model parameters as labelled. In each panel, the thick curves represent the residence time for the gas-driven dynamics according to HKM09. Solid, long-dashed and short-dashed curves are for mass ratios $q = 1, 0.1, 0.01$, respectively. The two thin dotted vertical lines approximately bracket the PTA observable window. The extrapolated pure GW-driven evolution ($t_{\text{res}}^{\text{gw}} \propto t_{\text{orb}}^{8/3}$) is shown as thin lines, for comparison. The ratio $t_{\text{res}}/t_{\text{res}}^{\text{gw}}$ gives the relative decrease in the number of binaries in the gas-dominated case compared to the GW driven case.

then the frequency evolution relative to the frequency resolution bin is

$$\frac{\Delta f_{\text{evol}}}{\Delta f_{\text{bin}}} \approx \frac{3}{2} \frac{f T^2}{(1+z)t_{\text{res}}} = \frac{0.015}{1+z} \mathcal{M}_{8.5}^{5/3} f_{50}^{11/3} T_{10}^2 \frac{t_{\text{res}}^{\text{gw}}}{t_{\text{res}}}, \quad (13)$$

where $\mathcal{M}_{8.5}$ is the chirp mass in units of $10^{8.5} M_{\odot}$, f_{50} is the frequency in units of 50 nHz, T_{10} is the observation time in unit of 10 yr, and for the second equality we have used equation (8). Equation (13) shows that typical binaries contribute to a single frequency bin as stationary sources in the GW-driven regime.³ We shall demonstrate that this is also true in the gas-driven case (see Fig. 1 below).

We generate N_k different realizations of the signal (usually $N_k = 1000$), i.e. N_k realizations of the MBHB population, consistent with equation (9) (see Section 3.1). Each of those consists of $N_b \sim 10^3 - 10^4$ binaries producing a relevant contribution to the signal, which we label by $(\mathcal{M}_i, z_i, f_{ri})$, $i = 1, 2, \dots, N_b$. The total signal (equation 4) in each frequency resolution bin Δf is evaluated as the sum of the contributions of each individual source [see Amaro-Seoane et al. (2010) for the detailed numerical procedure]:

$$h_c^2(f) = \sum_i h_{c,i}^2(\mathcal{M}_i, z_i, f_{ri}), \quad (14)$$

where for each f the sum is over the inspiralling sources emitting in the corresponding blueshifted (i.e. rest frame) Δf_r frequency resolution bin and $h_{c,i} = h_i \sqrt{f_i T}$ is the angle and polarization-averaged GW strain given by equation (5), multiplied by the square root of the number of cycles completed in the observation time. For comparison, we also evaluate the continuous integral equation (4),

³ The detection of the frequency shift of an $\mathcal{M}_{8.5} = f_{50} = z = 1$ source would require an extended observation with $T \gtrsim 35$ yr.

which represents the rms average of equation (14) over different realizations of MBHB populations in the limit $N_k \rightarrow \infty$.

Since the mass function of merging binaries is in general quite steep, the relative contributions of the few most massive binaries turn out to dominate the background in each frequency bin. The total GW signal depends very sensitively on these rare binaries, and the inferred spectrum is very spiky. It is useful to separate the total signal h_c into a part generated by a population of GW-bright *individually resolvable sources*, and a *stochastic level* h_s , which includes the contribution of all the unresolvable, dimmer sources. More precisely, in each frequency resolution bin, we find the MBHB with the largest $h_{c,i}^2(\mathcal{M}_i, z_i, f_{r_i})$, and we define it *individually resolvable* if its signal is stronger than the total contribution of all the other sources in that particular frequency bin. The *stochastic level* is consequently defined by adding up only the unresolvable sources in equation (14). Since the signal is dominated by few individual sources in each frequency bin, the $h_s(f)$ distribution obtained over the N_k realizations is far from being Gaussian or just even symmetric. To give an idea of the uncertainty range of $h_s(f)$, we calculate the 10, 50 and 90 per cent percentile levels of the $h_s(f)$ distribution of the N_k different realizations.

The separation of individually resolvable sources is useful for several reasons (see SVC; SVV09). First, it is useful from a statistical point of view for understanding the variance of the expected GW spectrum among various realizations of the inspiralling MBHBs. The discrete nature of the resolvable sources allows a different statistical analysis than for the smooth background level corresponding to the stochastic level, $h_s(f)$. The individually resolvable signals could also be important observationally. A sufficiently GW-bright resolvable binary allows to measure the GW polarization using PTAs, and give information on the sky position of the binary (Sesana & Vecchio 2010), which might be used to search for direct electromagnetic signatures like periodically variable AGN activity (HKM09). A coincident detection of GWs and electromagnetic emission of the same binary would have far reaching consequences in fundamental physics, cosmology and black hole physics (Kocsis, Haiman & Menou 2008).

2.4 Timing residuals

In general, the characteristic GW amplitudes (either of a stochastic background or of a resolvable source) can be translated into the pulsar timing language by converting $h_c(f)$ into a ‘characteristic timing residual’ $\delta t_c(f)$ corresponding to the sky position and polarization-averaged delay in the time of arrivals of consequent pulses due to GWs,

$$\delta t_c(f) = \frac{h_c(f)}{2\pi f}. \quad (15)$$

The pulsar timing residuals expected from an individual stationary GW source are derived in section 3 of SVV09 in detail. The corresponding measurement can be represented, in the time domain, with a residual:

$$\delta t(t) = r(t) + \delta t_N(t), \quad (16)$$

where $r(t)$ is the contribution due to the GW source (which accumulates continuously with observing time t , see below), and $\delta t_N(t)$ represents random fluctuations due to noise. The latter is the superposition of the intrinsic noise in the measurements and the GW stochastic level from the whole population of MBHBs, with an rms value

$$\delta t_{N,\text{rms}}^2(f) = \langle \delta t_N^2(f) \rangle = \delta t_p^2(f) + \delta t_s^2(f), \quad (17)$$

where $t_p(f)$ is the rms instrumental and astrophysical noise corresponding to the given pulsar, and $\delta t_s(f) = h_s(f)/(2\pi f)$ is due to the rms stochastic GW background of unresolved MBHBs, as defined in the previous section.

The sky angle and orbital orientation-averaged signal-to-noise ratio (SNR) at which one MBHB, radiating at (GW) frequency f , can be detected using a *single* pulsar with matched filtering is

$$\text{SNR}^2 = \frac{\delta t_{\text{gw}}^2(f)}{\delta t_{N,\text{rms}}^2(f)}. \quad (18)$$

Here $\delta t_{\text{gw}}(f)$ is the rms timing residual signal resulting from GWs emitted by the individual stationary source over the observation time T defined as

$$\delta t_{\text{gw}}(f) = \sqrt{\frac{8}{15}} \frac{h(\mathcal{M}, z, f)}{2\pi f} \sqrt{fT}, \quad (19)$$

where $h(\mathcal{M}, z, f)$ is the angle and polarization-averaged GW strain amplitude given by equation (5), the pre-factor $\sqrt{8/15}$ averages the observed signal over the ‘antenna beam pattern’ of the array (equation 21 in SVV09)⁴, and the \sqrt{fT} term accounts for the residual build-up with the number of cycles. For N_p number of pulsars, the total detection SNR of an individually resolvable MBHB is the rms of the contributions of individual pulsars given by (18). For N_p identical pulsars, the effective noise level is therefore attenuated by $N_p^{-1/2}$.

In the following, we will represent the overall GW signal and the stochastic background by using either their characteristic amplitudes, $h_c(f)$ and $h_s(f)$, or the corresponding characteristic timing residuals, $\delta t_c(f)$ and $\delta t_s(f)$, according to equation (15). We study the detection significance of individually resolvable sources and the distribution of their numbers as a function of the induced δt_{gw} . For each Monte Carlo realization of the emitting MBHB population, we count the cumulative number of all (N_t) and resolvable (N_r) sources above δt_{gw} as a function of δt_{gw} :

$$N_{t/r}(\delta t_{\text{gw}}) = \int_{\delta t_{\text{gw}}}^{\infty} \frac{\partial N_{t/r}}{\partial \delta t'_{\text{gw}}} \delta t'_{\text{gw}}, \quad (20)$$

where the integral is either over all sources or only the individually resolvable sources (i.e. restricted to those that produce residuals above the rms stochastic level, see Section 2.3).

3 THE EMITTING BINARY POPULATION

We calculate the GW signal by generating a catalogue of binaries consistent with equation (9). This requires (i) a model for the comoving merger rate density of coalescing MBHBs, $\partial^4 N / (\partial m_1 \partial m_2 \partial t_r \partial V_c)$, and (ii) a model for the evolution of individual inspiralling binaries $t_{\text{res}}(\mathcal{M}, \mu, f_r)$. These two items are the subjects of the next two subsections below.

3.1 Population of coalescing massive black hole binaries

We use the population models described in section 2 of SVV09, the reader is referred to that paper for full details. We extract catalogues of merging binaries from the semi-analytical model of Bertone, De Lucia & Thomas (2007) applied to the Millennium run (Springel et al. 2005). We then associate a central MBH to each merging

⁴ Note that that the square root in the pre-factor $\sqrt{8/15}$ is missing in equation (20) of SVV09, because of a typo there.

galaxy in our catalogue. We explored a total of 12 models, combining four M_{BH} -bulge prescriptions found in the literature with three different accretion scenarios during mergers. The 12 models are listed in table 1 of SVV09. In the present study, we shall use the Tu-SA population model as our default case. In this model, the MBH masses in the merging galaxies correlate with the masses of the bulges following the relation reported in Tundo et al. (2007), and accretion is efficient on to the more massive black hole, *before* the final coalescence of the binary. Since the comparison among different MBHB population models is not the main purpose of this study, we will present results only for this model. However, we tested our dynamic scenario on other population models presented in SVV09, finding no major differences for alternative models.

Assigning an MBH to each galaxy, we obtain a catalogue of mergers labelled by MBH masses and redshift. From this, we generate the merger rate per comoving volume, $\partial^4 N / (\partial m_1 \partial m_2 \partial t_r \partial V_c)$. In practice, due to the large number of mergers in the simulation, this is a finely resolved continuous function, describing the merger rate density as a function of m_1 , m_2 and z . After plugging into equations (9) and (12), we can obtain the continuous distribution $\partial^3 N / (\partial \mathcal{M} \partial z \partial \ln f_r)$ one would observe in an ‘ideal snapshot’ of the whole sky. We then sample this distribution to generate random Monte Carlo realizations of the GW signal. In summary, the chosen MBHB population model fixes the cosmological merger history, i.e. the function $\partial^4 N / (\partial m_1 \partial m_2 \partial t_r \partial V_c)$, and ‘Monte Carlo sampling’ refers to first choosing a realization of the cosmological merger history (i.e. generate N_b number of binary masses and redshifts)⁵ and then assigning an orbital frequency (or time to merger) to each binary. In addition to our fiducial merger rate, we also examine for the first time the situation where minor mergers do not contribute to the coalescence rate of the central MBHs. This is motivated by recent numerical simulations indicating that minor mergers with mass ratios $q < 0.1$ lead to tidal stripping of the merging satellite, and the resulting core does not sink efficiently to the centre of the host galaxy (Callegari et al. 2009). We identify the mass ratio of merging galaxies using the mass of the stellar components, to avoid complications due to the tidal stripping of merging dark matter haloes. In practice, we find that suppressing all the minor mergers does not affect the resulting GW signal, implying that the contribution of mergers involving dwarf galaxies is negligible.

3.2 Binary evolution in massive circumbinary discs

We adopt the simple analytical models of HKM09 to describe the dynamical evolution of MBHBs in a geometrically thin circumbinary accretion disc. Here we provide only a short summary of the gas-driven models, highlight their main assumptions and refer the reader to section 2 of HKM09 and references therein for more details.

For the typical MBH masses and separations we are considering, the tidal torque from the binary dominates over the viscous torque in the disc, opening a gap in the gas distribution. A spiral density wave is excited in the disc which torques the binary and pushes it inward. First, when the binary separation is relatively large, and the local disc mass is larger than the mass of the secondary, then the secondary migrates inward with the radial velocity of the accreting

gas,

$$t_{\text{res}}^{\nu} = -\frac{a}{\dot{a}_{\nu}} = \frac{2\pi r_0^2 \Sigma_0}{\dot{M}}, \quad (21)$$

where r_0 is the radius of the gap (typically $r_0 \sim 1.5a$, where a is the semimajor axis of the binary, Artymowicz & Lubow 1994), Σ_0 is the local surface density of the disc with no secondary and \dot{M} is the accretion rate far from the binary (the ν index denotes viscous evolution). This is analogous to disc-dominated type-II planetary migration. More typically, however, the binary is more massive than the local disc mass. In this case, analogous to secondary-dominated type-II migration, the angular momentum of the binary is absorbed less efficiently by the gas outside the gap, and the evolution slows down according to (Syer & Clarke 1995)

$$t_{\text{res}}^{\text{SC}} = -\frac{a}{\dot{a}_{\text{SC}}} = q_B^{-k_1} t_{\nu}, \quad (22)$$

where q_B is a measure of the lack of local disc mass dominance

$$q_B = \frac{4\pi r_0^2 \Sigma_0}{\mu} = \frac{2\dot{M}}{\mu} t_{\nu}, \quad (23)$$

which is less than unity in this case; μ is the reduced mass of the binary and k_1 is a constant that depends on the surface density–accretion rate power-law index. Here $k_1 = 0$ if $q_B > 1$; otherwise, $k_1 = 7/17$ at large separations, if the disc opacity at the gap boundary is dominated by free–free absorption, and $k_1 = 3/8$ closer in, if the opacity is dominated by electron scattering.

In the secondary-dominated regime, equation (22) assumes that the migration is very slow, so that the binary can influence the surface density very far upstream, and an approximate steady state is reached where the gas density is enhanced outside the gap by a factor q_B^{-k} relative to the gas density with no secondary (Syer & Clarke 1995). To examine the sensitivity of our conclusions to these assumptions, we consider an additional non-steady model of Ivanov et al. (1999). In their model, which is also axisymmetric by construction, they assumed that the binary torques are concentrated in a narrow ring near the edge of the cavity, which results in a time-dependent pile up of material near the cavity. The migration slows down further relative to $t_{\text{res}}^{\text{SC}}$ as

$$t_{\text{res}}^{\text{IPP}} = \frac{\mu}{2\dot{M}} \left(\frac{a}{a_0} \right)^{1/2} \left\{ \frac{1}{1 + \delta (1 - \sqrt{a/a_0})} \right\}^{k_2}, \quad (24)$$

where \dot{M} is the constant accretion rate far outside the binary, a_0 is the initial semimajor axis where the local disc mass is just equal to the secondary (i.e. $q_B = 1$), δ and k_2 are constants which depend on the density–accretion rate profile of the disc ($\delta = 4.1$ and 6.1 , $k_2 = 0.29$ and 0.26 for free–free and electron scattering opacity, respectively). Note that as the separation decreases well below its initial value, $a \ll a_0$, the curly bracket is a -independent, implying that $t_{\text{res}}^{\text{IPP}} \propto a^{1/2} \propto t_{\text{orb}}^{1/3}$.

Since the migration rate is proportional to the gas surface density and accretion rate, it is sensitive to the structure of the accretion disc. Following Syer & Clarke (1995), we estimate Σ_0 with that of a steady accretion disc of a single accreting BH, but consider two different models for that, α - and β -discs (see HKM09 for explicit formulae). For the classic Shakura & Sunyaev (1973) α -disc, the viscosity is proportional to the total (gas+radiation) pressure of the disc. Until very recently, this model, if radiation pressure dominated, has been thought to be thermally and viscously unstable (Lightman & Eardley 1974; Piran 1978). In the alternative β -model,

⁵ Here N_b is chosen randomly for each distribution, it can vary for different realizations of the population according to a Gaussian distribution with $\sigma = 1/\sqrt{\langle N_b \rangle}$ around the mean $\langle N_b \rangle$.

the viscosity is proportional to the gas pressure only⁶, and it is stable in both sense. The nature of viscosity is not well understood to predict which of these prescriptions lies closer to reality. Recent numerical magnetohydrodynamic simulations (Hirose, Krolik & Blaes 2009) suggest that the thermal instability is avoided in radiation pressure dominated situations because stress fluctuations lead the associated pressure fluctuations, and seem to favour the α prescription over the β -model. We carry out all calculations for both models, but consider the α prescription as our fiducial disc model. In both cases, the model is uniquely determined by three parameters: the central BH mass, the accretion rate far from the binary \dot{M} and the α viscosity parameter. The exact value of these parameters is not well known. Observations of luminous AGN imply an accretion rate around $\dot{m} = \dot{M}/\dot{M}_{\text{Edd}} = (0.1-1)$ with a statistical increase towards higher quasar luminosities (Kollmeier et al. 2006; Trump et al. 2009). Here $\dot{M}_{\text{Edd}} = L_{\text{Edd}}/(\eta c^2)$ is the Eddington accretion rate for $\eta = 10$ per cent radiative efficiency, where L_{Edd} is the Eddington luminosity. Observations of outbursts in binaries with an accreting white dwarf, neutron star or stellar black hole imply $\alpha = 0.2-0.4$ (Dubus, Hameury & Lasota 2001; King, Pringle & Livio 2007, and references therein). Theoretical limits based on simulations of magnetohydrodynamic turbulence around black holes are inconclusive, but are consistent with α in the range 0.01–1 (Pessah, Chan & Psaltis 2007). It is however unclear whether these numbers are directly applicable to circumbinary MBH systems. The binary exerts a torque that pushes the gas away on average, and consequently might reduce the accretion rate (see Lubow, Seibert & Artymowicz 1999; Lubow & D’Angelo 2006, in the planetary context, and MacFadyen & Milosavljević 2008; Cuadra et al. 2009, for MBH binaries). We explore several choices covering all 6 combinations with $\dot{m} = \{0.1, 0.3\}$ and $\alpha = \{0.01, 0.1, 0.3\}$ for both α - and β -discs, respectively. Motivated by the considerations outlined above, we highlight the α -disc with $\dot{m} = 0.3$ and $\alpha = 0.3$ as our default model.

We regard $t_{\text{res}}^{\text{SC}}$ and $t_{\text{res}}^{\text{IPP}}$ as lower and upper limits of the true residence time during secondary-dominated gas-driven migration. However, we caution that these estimates are subject to many uncertainties related to the complexity of binary accretion disc and migration physics. Other models for the binary–disc interaction (i.e. Hayasaki 2009; Hayasaki, Ueda & Isobe 2010) lead to a slower migration rate. Two-dimensional hydrodynamic simulations show that the flow into the gap is generally non-axisymmetric, and the gap and the binary orbit become eccentric (Lubow & D’Angelo 2006; MacFadyen & Milosavljević 2008; Cuadra et al. 2009). The accretion rate and inspiral velocity in reality may be higher than in the axisymmetric approximation. Resonant interactions lead to an enhanced angular momentum transport (Goldreich & Tremaine 1980). For unequal masses or thick discs, there may be significant inflow into the gap (Artymowicz & Lubow 1996). In this case, corotation torques need to be considered, and the migration is expected to be much faster inward or outward, analogous to type-I planetary migration (Tanaka, Takeuchi & Ward 2002). Disc thermodynamics and turbulence may also significantly influence the migration rate (Nelson & Papaloizou 2004; Paardekooper & Papaloizou 2009; Paardekooper et al. 2010).

These steady-state accretion disc models also assume that the self-gravity of the disc is negligible, the disc is optically thick, and the temperature of the gas is large enough ($T \gtrsim 10^4$ K) that the

opacity corresponds to that of ionized gas. These assumptions break down at radial distances corresponding to orbital periods of around $t_{\text{orb}} \gtrsim (5-10)$ yr for BHs with masses $M \gtrsim (10^8-10^9) M_{\odot}$ (see fig. 1 of HKM09). Thus, these disc models are self-consistent for the relatively short-period binaries, but become suspect for binaries emitting at the low-frequency edge of the PTA window.

Another shortcoming of these gas-driven migration models is that they were derived assuming the disc surface density is a decreasing function of radius. This affects the specific values of k_1 and k_2 in equations (22) and (24). This assumption is safe for β -discs in general, as well as for α -discs if the gas pressure dominates over radiation pressure, but it is violated for radiation pressure dominated α -discs. However, gas-driven models with binary separation in the PTA frequency band and large masses $M > 10^8 M_{\odot}$ are in fact radiation pressure dominated near the gap.⁷ It is unclear how k_1 and k_2 change for radiation pressure dominated α -discs. As an approximation, following Armitage & Natarajan (2005) and HKM09, we extrapolate the values of the gas pressure dominated outer regions.

Substituting the surface density profiles $\Sigma_0(r)$ of α - and β -discs in equations (21–22), we obtain the inspiral rate of binaries as they evolve from large to small radii in a gaseous medium. Fig. 1 shows the corresponding residence times for selected MBH masses and disc models. The binary migration history can be divided in three phases.

- (i) Initially, migration occurs on the viscous time-scale $t_{\text{res}}^{\nu} \propto t_{\text{orb}}^{5/6}$.
- (ii) Then as the disc mass decreases below the secondary mass, the gas-driven migration slows down according to $t_{\text{res}}^{\text{SC}}$ or $t_{\text{res}}^{\text{IPP}}$. The migration rate in this case is non-uniform, it changes as the local accretion disc structure varies due to changes in the dominant source of opacity (free–free for large separation to electron scattering for smaller separations) and the source of pressure (thermal gas pressure at large a to radiation pressure at small a). In particular, for the steady Syer & Clarke (1995) models, $t_{\text{res}}^{\text{SC}} \propto t_{\text{orb}}^{25/51} - t_{\text{orb}}^{7/12}$ with free–free to electron scattering opacity for thermal pressure support, and $t_{\text{res}}^{\text{SC}} \propto t_{\text{orb}}^{35/24} - t_{\text{orb}}^{7/12}$, for radiation pressure support for α - to β -discs, respectively; while for the time-dependent models of Ivanov et al. (1999), the accretion rate approaches $t_{\text{res}}^{\text{IPP}} \propto t_{\text{orb}}^{1/3}$. Note that the residence time in this regime is generally smaller for α -discs than for β -discs. This is explained by the fact that the overall surface density is smaller for radiation pressure dominated α -discs, and so to achieve the same net accretion rate, the radial gas inflow velocity is larger for α -discs. The migration rate is a monotonic function of the radial inflow velocity (see equations 21 and 22), so that the binary is pushed in more quickly for α -discs.
- (iii) Finally, when the separation is sufficiently reduced, the emission of GWs becomes very efficient and determines the inspiral rate according to $t_{\text{res}}^{\text{GW}} = a/\dot{a}_{\text{gw}} \propto a^4 \propto t_{\text{orb}}^{8/3}$ (see equation 8).

Note that for circular orbits, the GW frequency (in the binary rest frame) is simply $f_r = 2/t_{\text{orb}}$. The number of binaries at any given t_{orb} is proportional to the residence time t_{res} . Therefore, the decrease in t_{res} in the gas-driven regime compared to the GW-driven case implies a decrease in the population of MBHs, which ultimately leads to the attenuation of the low-frequency end of the observable total GW spectrum. The rms GW spectrum averaged over the whole population of binary inspiral episodes is no longer a power law. It is interesting to examine the contributions of various evolutionary phases according to equation (11), $h_c \propto f^{-2/3} \sqrt{t_{\text{res}}^{\text{GW}}/t_{\text{res}}}$. If all

⁶ The name comes from the definition of viscosity $\nu \propto \alpha p_{\text{gas}}^{\beta} p_{\text{tot}}^{1-\beta}$, where $\beta = 1$ for the β -model, while $\beta = 0$ for the α -model. In both cases, α is a free model parameter.

⁷ Note that the most massive binaries, however, are GW driven, and are not sensitive to the accretion disc.

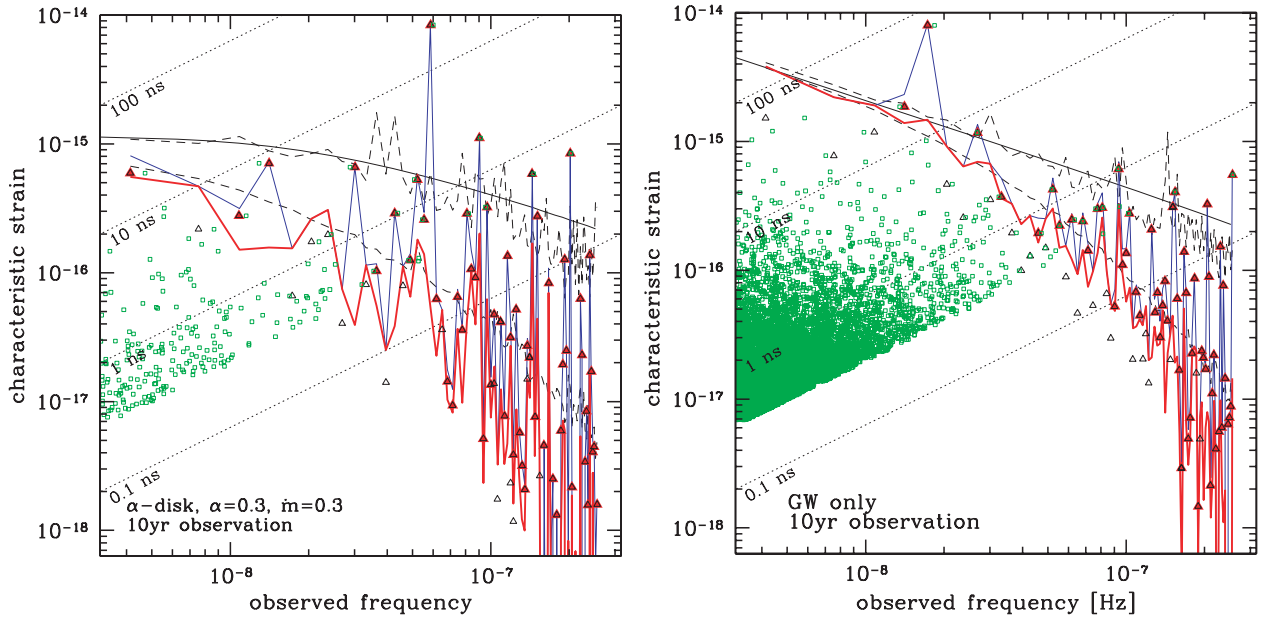


Figure 2. Components of the GW signal from a population of inspiralling MBHBs. In the left-hand panel, we consider all binaries embedded in a gaseous α -disc for our default model (all mergers wet), while in the right-hand panel all binaries are purely GW driven (all mergers dry). In each panel, the smooth solid line is the rms total characteristic GW strain h_c using the integral expressions (4–9) (which correspond to an average over $N_k \rightarrow \infty$ Monte Carlo realizations), the two dashed black lines represent the rms total signal (upper) and the rms stochastic background level (lower) averaged over $N_k = 1000$ Monte Carlo realizations, respectively. The jagged blue line displays a random Monte Carlo realization of the GW signal. The small black and red triangles show the contributions of the brightest and resolvable sources in each frequency bin, respectively. The jagged red line is the stochastic GW background for this realization, i.e. once the resolvable sources in each frequency bin are subtracted. The green dots label all the systems producing an rms residual $t_{\text{gw}} > 0.3$ ns over $T = 10$ yr. The dotted diagonal lines show constant t_{gw} levels as a function of frequency. An observation time of 10 yr is assumed.

binaries were in the GW-driven phase then $h_c^{\text{GW}} \propto f^{-2/3}$. If all were in the secondary-dominated type-II migration regime with a steady radiation pressure dominated disc⁸ $h_c^{\text{SC}} \propto f^{-1/16} f^{3/8}$ for α - and β -discs, respectively, further out $h_c^{\text{SC}} \propto f^{3/8} f^{43/102}$ in the gas pressure dominated regime (with electron scattering versus free-free opacity), while $h_c^{\text{PP}} \propto f^{1/2}$ asymptotically for the non-steady models, and finally $h_c^{\text{v}} \propto f^{1/4}$ in the disc-dominated type-II migration regime. In any case, the GW spectrum $h_c(f)$ is much shallower in the gas-driven phase. Note the gas-driven phase contributes a nearly flat or an *increasing* spectrum $h_c(f)$, very different from the nominal $f^{-2/3}$ GW-driven case. In general, the orbital separation for more massive objects at a given f_r is smaller in terms of Schwarzschild radii. Since the transition to gas-driven migration, when expressed in Schwarzschild radii, is roughly independent of mass, it follows then at any given frequency bin the more massive objects are typically GW driven and lighter objects are gas driven. The total average spectrum⁹ assuming only wet mergers¹⁰, is between $f^{-2/3}$ and $f^{1/2}$ depending on the ratio of GW-driven to gas-driven binaries. Note, however, that the individually resolvable sources are typically very massive and are in the GW-driven phase for the relevant range

of binary separations, and therefore their properties should not be modified by gas effects.

4 RESULTS

4.1 Description of the signal

As stated in Section 2.3, the relevant frequency band for pulsar timing observations, assuming a temporal observation baseline T and a time interval between subsequent observations Δ , is between $f_{\text{min}} \approx 1/T$ and $f_{\text{max}} \approx 1/(2\Delta)$ with a resolution $\Delta f \approx 1/T$. In our calculations, we assume a default duration of $T = 10$ yr for the PTA campaign with $\Delta \approx 1$ week. This gives $f_{\text{min}} \approx 3 \times 10^{-9}$ Hz, $f_{\text{max}} \approx 10^{-6}$ Hz and $\Delta f \approx 1/T \approx 3 \times 10^{-9}$ Hz. The simulated signal is computed by doing a Monte Carlo sampling of the distribution $\partial^3 N / (\partial z \partial \mathcal{M} \partial \ln f_r)$, and adding the GW contribution of each individual source. In each frequency bin, we identify the individually resolvable sources and the stochastic background. We repeat this exercise $N_k = 1000$ times for the 12 steady disc models defined in Section 3.2 and for the purely GW-driven case. In addition, we also calculate the GW signal using the integral expressions (4–9) using the continuous distribution function, which corresponds to the rms average of the GW signal in the $N_k \rightarrow \infty$ limit.

The GW signal and its most important ingredients are plotted in Fig. 2. The left-hand panel shows results for gas-driven inspirals using our default α -disc (i.e. ‘gas on’ – all mergers wet); the right-hand panel shows results for purely GW-driven inspirals for comparison (i.e. ‘gas off’ – all mergers dry) for the same underlying cosmological MBHB coalescence rate. A randomly selected Monte Carlo realization of the signal is depicted as a dotted blue jagged line. The

⁸ Most of the gas-driven binaries contributing to the background at large frequencies are in this regime.

⁹ If averaging over each binary episode but not over the cosmological merger tree.

¹⁰ Here ‘wet’ refers to gas-rich mergers where the dynamics of the binary is driven by both the circumbinary disc and GW emission depending on the binary separation, while ‘dry’ refers to the case where the binary dynamics is driven by GW emission only at all radii. Note that in both cases we neglect interactions with stars.

green dots represent the contributions of individual binaries; systems producing $\delta t_{\text{gw}}(f) > 0.3$ ns timing residual are shown. In each frequency bin, the brightest source is marked by a black triangle. The individually resolvable sources are marked by superposed red triangles, and the stochastic level from all unresolvable sources is shown with a solid red jagged line. Clearly, the GW signal for any single realization is far from being smooth. The noisy nature of the signal is due to rare massive binaries rising well above the stochastic level. Solid black curves in Fig. 2 show $h_c(f)$, the rms value of the GW signal averaged over the merger episodes, using the integral expressions (4–9). The upper black dashed curve is the GW signal, averaged over $N_k = 1000$ Monte Carlo realizations of the merging systems. This is still noisy, due to the finite number of realizations used, but is consistent with the integrated average shown by the solid black curve. The lower black dashed curve marks the rms stochastic component, $h_{s,\text{rms}}(f)$, averaged over the same $N_k = 1000$ realizations. Fig. 2 shows that the gaseous disc greatly reduces the number of binaries at small frequencies compared to the GW-only case, as gas drives the binaries in quickly towards the final coalescence. Consequently, the signal is more spiky in the presence of gas. There is a clear flattening of the rms spectrum at low frequencies ($f < 1/\text{yr}$) in the gas-driven case compared to the purely GW-driven case. The stochastic level is more suppressed in the wet-only case and is much steeper than $f^{-2/3}$ at high frequencies. Despite the spiky signal in a given Monte Carlo realization, the overall shape of the background is well recognizable in the GW driven model, while the characterization of the global shape of the signal in the gas-driven case appears to be less viable.

Gas-driven migration becomes more and more prominent at large binary separations, corresponding to large orbital times or small GW frequencies. Therefore, to check the ultimate maximum impact of gas effects on future PTA detections, we simulated the spectrum for a hypothetical $T = 100$ yr observation baseline. Fig. 3 shows a realization of the spectrum for such an extended observation, assuming our default gas model (i.e. all mergers wet). Given the extended temporal baseline, the minimum observable frequency is




Figure 3. Same as left-hand panel of Fig. 2, but for an observation lasting 100 yr (averages over the 1000 realizations are not shown in this case for clarity).

pushed down to $\sim 3 \times 10^{-10}$ Hz, where the spectrum of gas-driven mergers is considerably flatter. Moreover, the frequency resolution bin is then narrower, the number of sources per bin is much smaller, and more sources become resolvable. At the smallest frequencies, the induced timing residual can be higher than $1 \mu\text{s}$ and more than 50 sources will be individually resolvable at 1 ns precision level, some of them with an SNR as high as 100. These numbers are not severely modified if considering purely GW-driven dry mergers only (see Fig. 5 below). Interestingly, due to the high-frequency resolution of such an extended observation, the GW frequency of some of the resolvable binaries may evolve significantly during the observation (e.g. a binary with masses $m_1 = m_2 = 5 \times 10^8 M_\odot$ or higher are non-stationary relative to the bin size at frequencies above $f \gtrsim 40$ nHz, see equation 13). Therefore, it may be possible to detect the frequency evolution of individually resolvable binaries during such an extended monitoring campaign. However, this computation is idealized: it is questionable if millisecond pulsars can maintain an ns timing stability over such a long time-scale; none the less, it points out the enormous capabilities of long-term PTA campaigns.

4.2 Individually resolvable sources

Let us now examine the prospects for detecting individual sources using PTAs. How many sources are expected to be individually resolvable? How significant is their detection?

As explained in Section 2.4, the signal from a specific binary can be detected using a PTA if the corresponding timing residual δt_{gw}^2 is above the rms noise level $\delta t_{N,\text{rms}}^2$ characterizing the PTA given by equation (18). We should notice, however, that this estimate of the number of individually resolvable sources is conservative and is likely to provide only a *lower limit* for the following reasons. First, we average over the sky position of the binaries and pulsars, while in reality, it may be possible to take advantage of the different GW polarization and amplitude generated by sources in different sky positions to deconvolve their signal (even if they have similar strength and frequency). Secondly, the brightest source identification algorithm can be implemented recursively, after an accurate subtraction of the identified sources. However, we find that, especially at low frequencies, the distribution of GW source amplitudes for various binaries in a single frequency bin is not strongly hierarchical, so that a recursive brightest source finding algorithm should not increase significantly the number of resolvable systems. We present here results both in terms of the total number (N_t) and resolvable systems (N_r , see equation 20).

In Fig. 4, we plot the distribution of the number of sources (total and resolvable) as a function of timing residual, detection frequency, redshift and chirp mass, found in two particular realizations of our default α -disc (all mergers wet) and in the purely GW-driven models (all mergers dry). The figure shows that even though there are much more sources in the purely GW-driven case, the number of sources rising above the stochastic level is almost the same in the purely dry and wet cases. Fig. 4 also shows the chirp mass, redshift and frequency distribution for sources above 1-ns timing level. As was previously shown in SVV09, the bulk of the sources are cosmologically nearby ($z \lesssim 1$) with masses peaking around $\mathcal{M} \sim 2 \times 10^8 M_\odot$. Fig. 4 shows that gas dynamics does not introduce a major systematic change in the shape of the redshift and the chirp mass distributions. The lower-left panel shows that gas removes a systematically larger fraction of sources at small frequencies.

Fig. 5 shows the cumulative number of binaries (total and resolvable) as a function of timing residual. The upper panel shows results




Figure 4. Top-left panel: characteristic amplitude of the timing residuals δt_{gw} (equation 19) as a function of frequency; the dots are the residuals generated by individual sources and the solid line is the estimated *stochastic level* of the GW signal. Top-right panel: distribution of the number of total (dotted lines) and resolvable (solid lines) sources per logarithmic redshift interval as a function of redshift, generating a $\delta t_{\text{gw}} > 1$ ns. Bottom-left panel: distribution of the number of total (dotted lines) and resolvable (solid lines) sources per logarithmic frequency interval as a function of the GW frequency, generating a $\delta t_{\text{gw}} > 1$ ns. Bottom-right panel: distribution of the total (dotted lines) and individually resolvable (solid lines) number of sources per logarithmic chirp mass interval as a function of chirp mass, generating a $\delta t_{\text{gw}} > 1$ ns. All the black elements refer to our default disc model, the red elements are for a GW-driven MBHB population. Distributions are averaged over $N_k = 1000$ realizations of the MBHB population.

for the α -disc models with $\dot{m} = 0.3$ and different α . The statistics of resolvable sources is almost unaffected by the large suppression of the total number of sources at a fixed timing residual. For example, in all of our models, we expect ~ 2 resolvable sources at a timing level of $\delta t_{\text{gw}} = 10$ ns, even though the total number of sources contributing to the signal at that level spans about an order of magnitude among the different models (~ 5 for $\alpha = 0.3$ to ~ 50 for binaries driven by GW only). The same is true for β -disc models (central panel) and for non-steady models (lower panel), even though in these cases the total number of sources at a particular δt_{gw} is not reduced dramatically by gas effects. This result can be understood with a closer inspection of Fig. 1. Let us focus on the α -disc model. As explained in Section 3.2, the impact of gas-driven dynamics in the PTA window is more significant for lower binary masses and unequal mass ratios. These are the binaries that build-up the bulk of the signal, and its stochastic level is consequently greatly reduced in the gas-driven case. On the other hand, the population of high equal mass binaries, which constitute most of the individually resolvable sources, is almost unaffected by the presence of the circumbinary disc, as they are already in the GW-driven regime in the relevant range of binary periods.

Fig. 6 shows expectations on the detection significance of resolvable sources. The SNR distribution of resolvable sources (equation 18) are shown as thin lines, accounting for the astrophysical GW noise from the unresolved binaries, but neglecting the intrinsic noise of the array (i.e. assuming an ideal detector with infinite sensitivity, $\delta t_p^2(f) = 0$, in equation 17). This calculation represents an *upper limit* of the SNR. Thick lines, in Fig. 6 plot the SNR considering a total detector noise of 1 ns (appropriate for SKA). The figure shows that the expected detection significance of resolvable sources is systematically higher for gas-driven models with larger α . In general, in an array with 1 ns sensitivity, we may expect a couple of individually resolvable sources with $\text{SNR} > 5$. For near future instruments with much worse sensitivity, the identification of resolvable sources might be more challenging.

4.3 Stochastic background

Fig. 7 shows the rms stochastic level (i.e. after subtracting off the individually resolvable sources, see Section 2.3) for selected steady state gas disc models, averaged over $N_k = 1000$ realizations. The top-left panel highlights the difference between wet and dry models,




Figure 5. Cumulative number of total $[N_t(\delta t_{\text{gw}})$, thin lines] and individually resolvable $[N_r(\delta t_{\text{gw}})$, thick lines] sources emitting above a given δt_{gw} threshold as a function of δt_{gw} . Upper panel: α -disc model with $\dot{m} = 0.3$. Lines are for $\alpha = 0.3$ (solid), 0.1 (long-dashed) and 0.01 (short-dashed). Central panel: same as upper panel but for a β -disc model. Lower panel: non-steady IPP model, assuming $\alpha = 0.3$ and $\dot{m} = 0.3$ (solid) and 0.1 (long-dashed). In all the panels, the dotted lines refer, for comparison, to the GW-driven model. All the distributions refer to the ensemble mean computed over all $N_k = 1000$ realizations of the MBHB population.




Figure 6. SNR distribution of individually resolvable sources. Thin and thick curves refer to neglecting the instrumental noise or using a $\delta t_{\text{gw}} = 1$ ns timing precision, respectively. Line styles as in the upper panel of Fig. 5.

the top-right and bottom-left panels collect different α -disc models and the lower right panel is for selected β -disc models (non-steady models, not shown here, give results basically identical to β -disc models). The different line styles show the effect of changing the




Figure 7. Influence of the gas-driven dynamics on h_c (thin lines) and on h_s (thick lines). Upper left panel: GW-driven dynamics versus gas-driven dynamics for our default disc model. Upper right panel: $\alpha = 0.3$ (solid), 0.1 (long-dashed), 0.01 (short-dashed), for an α -disc with $\dot{m} = 0.3$. Lower left panel: same as the upper right panel, considering an α -disc with $\dot{m} = 0.1$. Lower right panel: same as the upper right panel, considering a β -disc with $\dot{m} = 0.3$. The two dotted lines in each panel represent the sensitivity of the complete PPTA survey and an indicative sensitivity of 1 ns for the SKA.

α parameter of the disc. We also show the rms total signal level, which is exactly proportional to $f^{-2/3}$ in the dry case. In general, the stochastic level matches the total signal level at low frequencies, but is increasingly suppressed for frequencies above $f \gtrsim 10^{-8}$ Hz for all of our models. At sufficiently large frequencies, GW emission dominates even for wet mergers, and both the rms total signal and the stochastic level approaches the purely GW-driven case. However at small frequencies, a significant fraction of binaries is driven by gas, and the signal is attenuated and the spectrum is less steep compared to the dry case. Fig. 7 shows that gas-driven migration suppresses the stochastic background significantly, by a factor of 5 for our standard disc model below 10^{-8} Hz. The suppression of the total and stochastic levels is a strong function of the model parameters. Interestingly, there is almost no suppression for β -discs or for α -discs with a small accretion rate and/or a small α value. In these cases, the local disc mass is smaller, resulting in longer viscous time-scales, and hence the population of widely separated binaries is not reduced significantly.

Fig. 8 quantifies how the stochastic background changes among different Monte Carlo realizations, showing the range attained by 10–90 per cent of all $N_k = 1000$ realizations. The variance is not the product of uncertainties related to the parameters of the adopted cosmological and dynamical model, but is purely determined by the small number statistics of sources per frequency bin, intrinsic to the source distribution. At $f = 10^{-8}$ Hz, the variance of the signal produced by our default disc model is ~ 0.5 dex. Decreasing α to 0.1 and 0.01, the variance drops to ~ 0.35 and ~ 0.25 dex, respectively. In the case of all mergers driven by GWs, the variance at the same frequency is only ~ 0.15 dex. This means that, contrary to the GW-driven model (SVC08), it is impossible to predict the stochastic level of the signal accurately for our default gas-driven model. For a




Figure 8. Variance of the expected stochastic level of the signal as a function of α . In all the panels, we assume the α -disc model with $\dot{m} = 0.3$ (except for the lower right panel, referring to the GW driven model). Solid lines represent the median h_s over 1000 Monte Carlo realizations, while the shaded area enclosed within the two dashed lines is the 10–90 per cent confidence region. The two dotted lines in each panel represent the sensitivity of the complete PPTA survey and an indicative sensitivity of 1 ns for the SKA.

linear fit, any power law in the range $f^{-0.3}$ – $f^{-1.1}$ is acceptable within the level of variance of the signal in the frequency range 3–30 nHz.

5 DISCUSSION AND CONCLUSIONS

The presence of a strong nHz GW signal from a cosmological population of MBHBs is a clear prediction of hierarchical models of structure formation, where galaxy evolution proceeds through a sequence of merger events. The detailed nature of the signal depends however on a number of uncertain factors: the MBHB mass function, the cosmological merger rate, the detailed evolution of binaries and so on. In particular, MBHB dynamics determines the number of sources emitting at any given frequency, and thus the overall shape and strength of the signal. Previous works on the subject (e.g. SVC08 and SVV09) considered the case of GW-driven binaries only. In this paper, we studied the impact of gas-driven MBHB dynamics on the nHz GW signal detectable with PTAs. This is relevant because in any merger event, cold gas is efficiently funnelled towards the centre of the merger remnant, providing a large supply of gas to the MBHB formed following the galaxy interaction. Even a per cent of the galaxy mass in cold gas funnelled towards the centre is much larger than the masses of the putative MBHBs involved in the merger so that the newly formed binary evolution may be driven by gas until a few thousand years before final coalescence.

To conduct our study, we coupled models for gas-driven inspirals of HKM09 to the MBHB population models presented in SVV09. Our simulations cover a large variety of steady state and quasi-stationary one-zone disc models for the MBHB–disc dynamical interactions (with an extensive exploration of the α viscosity parameter and \dot{m} for α - and β -discs), along with several different prescriptions for the merging MBHB population (four different black hole mass–galaxy bulge relations coupled with three different accre-

tion recipes). The differences with respect to the purely GW-driven models (presented in SVV09) are qualitatively similar for all the considered MBHB populations, we thus presented the results for the Tu-SA population model only, focusing on the impact of the different disc models. Our main findings can be summarized as follows.

(i) The effect of gas-driven dynamics may or may not be important depending on the properties of the circumbinary disc. A robust result is that if the viscosity is proportional to the gas pressure only (β -disc models), there is basically no effect on the GW signal, independently of the other disc parameters. Similarly, there is a tiny effect for time-dependent migration models (i.e. Ivanov et al. 1999), where gas piles up and the accretion rate decreases as the binary hardens. However, if the viscosity is proportional to the total (gas+radiation) pressure (α -disc; Shakura & Sunyaev 1973), then the GW signal can be significantly affected for certain steady state circumbinary disc models (Syer & Clarke 1995) with $\alpha \gtrsim 0.1$ and $\dot{m} \gtrsim 0.3$. This difference is explained by the fact that the gas-driven inspiral rate is determined by the radial gas inflow velocity, which is significantly faster for radiation pressure dominated α -discs for a fixed accretion rate, as they are more dilute. Therefore gas effects can dominate over the GW-driven inspiral rate for α -discs at relatively small binary separations corresponding to the PTA frequency band.

(ii) With respect to the GW-driven case, the presence of massive circumbinary discs affects the population of low-unequal mass binaries predominantly ($M < 10^8 M_\odot$; $q < 0.1$), causing a significant suppression of the stochastic level of the signal, but leaving the number and strength of massive *individually resolvable sources* basically unaffected. In our default model ($\alpha = 0.3$, $\dot{m} = 0.3$), the stochastic background is suppressed by a factor of ~ 5 at $f < 10^{-8}$ Hz. This suppression factor decreases for smaller \dot{m} and α . The stochastic level is *not* suppressed significantly for non-steady discs, for β -discs (arbitrary \dot{m} and α), and for α -discs with either $\dot{m} \lesssim 0.1$ and arbitrary α or $\alpha \lesssim 0.01$ with arbitrary \dot{m} . *About 10 individual sources are resolvable at 1-ns timing level, independently of the adopted disc model.*

(iii) All the results shown here for the Tu-SA model, hold for every other MBHB population model we tested. There is a certain level of degeneracy between disc dynamics and MBHB mass function: the stochastic level given by a population of heavy binaries evolving by gas dynamics can mimic that of a population of lighter binaries that are driven by GWs only. However, the variance of the signal would be much bigger in the former case, because the signal is produced by fewer massive sources.

(iv) The detection of GWs emitted by MBHBs embedded in gaseous discs with high viscosity and accretion rate may be very challenging for relatively short-term PTA campaigns like the PPTA. In fact, we find that most of the 12 MBHB population models tested in SVV09 would *not* produce a stochastic signal detectable by the PPTA (i.e. the signal is a factor of 3 below the PPTA capabilities for the Tu-SA model). However, long-term projects like the SKA, which aim to nanosecond sensitivities, are expected to be able to detect the GW signal, resolving a handful of individual sources with high significance.

A word of caution should be spent to stress the fact that our models are idealized in many ways. We considered circular binaries only. If most systems were significantly eccentric, then the overall signal would be modified by the multiharmonic emission of each individual source. Moreover, we only considered radiatively efficient, geometrically thin, one-zone, steady-state and quasi-stationary

accretion discs. The most massive but gas-driven binaries, emitting at low-GW frequencies $f_{\text{gw}} \lesssim 5$ nHz, are radiation pressure dominated, but the migration rate estimates had to be extrapolated using the scaling exponents for gas pressure dominated discs. For even wider massive binaries (f_{gw} near the low-frequency observation limit) the discs are marginally Toomre unstable. We used simple models for the binary–disc interaction, scaling the type-II planetary migration formulae to MBHBs. The results are sensitive to the models: the steady-state models of Syer & Clarke (1995) lead to a decrease in the stochastic nHz background, while the more sophisticated time-dependent models of Ivanov et al. (1999) as well as the (Hayasaki 2009; Hayasaki et al. 2010) models lead to almost no effect (see also Section 3.2, for a list of caveats). Further studies should examine the accuracy of these approximations for gas-driven migration in circumbinary discs around MBHBs. Finally, it is likely that not all the binaries are gas driven on their way to the coalescence, and the efficiency of the disc–binary coupling may vary from merger to merger depending on the environmental conditions, which may modify the properties of the expected signal as well. None the less, our calculations provide clear predictions for the possible attenuation of the stochastic GW background, which may be confirmed or discarded by ongoing and forthcoming PTAs.

ACKNOWLEDGMENTS

AS is grateful to M. ‘penguin’ Giustini and V. Cappa for enlightening discussions. BK acknowledges support by NASA through Einstein Postdoctoral Fellowship grant no. PF9-00063 awarded by the Chandra X-ray Center, which is operated by the Smithsonian Astrophysical Observatory for NASA under contract NAS8-03060, and partial support by OTKA grant 68228.

REFERENCES

- Amaro-Seoane P., Sesana A., Hoffman L., Benacquista M., Eichhorn C., Makino J., Spurzem R., 2010, *MNRAS*, 402, 2308
- Armitage P. J., Natarajan P., 2002, *ApJ*, 567, L9
- Armitage P. J., Natarajan P., 2005, *ApJ*, 634, 921
- Artymowicz P., Lubow S. H., 1994, *ApJ*, 421, 651
- Artymowicz P., Lubow S. H., 1996, *ApJ*, 467, L77
- Barnes J. E., 2002, *MNRAS*, 333, 481
- Begelman M. C., Blandford R. D., Rees M. J., 1980, *Nat*, 287, 307
- Bertone S., De Lucia G., Thomas P. A., 2007, *MNRAS*, 379, 1143
- Bertotti B., Carr B. J., Rees M. J., 1983, *MNRAS*, 203, 945
- Callegari S., Mayer L., Kazantzidis S., Colpi M., Governato F., Quinn T., Wadsley J., 2009, *ApJ*, 696, 89
- Cuadra J., Armitage P. J., Alexander R. D., Begelman M. C., 2009, *MNRAS*, 393, 1423
- Detweiler S., 1979, *ApJ*, 234, 1100
- Dotti M., Colpi M., Haardt F., Mayer L., 2007, *MNRAS*, 379, 956
- Dubus G., Hameury J.-M., Lasota J.-P., 2001, *A&A*, 373, 251
- Edwards R. T., Hobbs G. B., Manchester R. N., 2006, *MNRAS*, 372, 1549
- Escala A., Larson R. B., Coppi P. S., Mardones D., 2004, *ApJ*, 607, 765
- Escala A., Larson R. B., Coppi P. S., Mardones D., 2005, *ApJ*, 630, 152
- Goldreich P., Tremaine S., 1979, *ApJ*, 233, 857
- Haehnelt M. G., 1994, *MNRAS*, 269, 199
- Haiman Z., Kocsis B., Menou K., 2009, *ApJ*, 700, 1952 (HKM09)
- Hayasaki K., 2009, *PASJ*, 61, 65
- Hayasaki K., Ueda Y., Isobe N., 2010, *PASJ*, 62, 1351
- Hellings R. W., Downs G. S., 1983, *ApJ*, 265, 39
- Hirose S., Krolik J. H., Blaes O., 2009, *ApJ*, 691, 16
- Hobbs G. et al., 2010, *Classical Quantum Gravity*, 27, 4013
- Ivanov P. B., Papaloizou J. C. B., Polnarev A. G., 1999, *MNRAS*, 307, 79
- Jaffe A. H., Backer D. C., 2003, *ApJ*, 583, 616
- Janssen G. H., Stappers B. W., Kramer M., Purver M., Jessner A., Cognard I., 2008, in Bassa C., Wang Z., Cumming A., Kaspi V. M., eds, *AIP Conf. Proc. Vol. 983, 40 YEARS OF PULSARS: Millisecond Pulsars, Magnetars and More*. Am. Inst. Phys., New York, p. 633
- Jetten F. A., Hobbs G. B., Lee K. J., Manchester R. N., 2005, *ApJ*, 625, 123
- Jetten F. et al., 2009, preprint (arXiv:0909.1058)
- King A. R., Pringle J. E., Livio M., 2007, *MNRAS*, 376, 1740
- Kocsis B., Haiman Z., Menou K., 2008, *ApJ*, 684, 870
- Kollmeier J. A. et al., 2006, *ApJ*, 648, 128
- Koushiappas S. M., Zentner A. R., 2006, *ApJ*, 639, 7
- Lazio J., 2009, preprint (arXiv:0910.0632)
- Lightman A. P., Eardley D. M., 1974, *ApJ*, 187, L1
- Lodato G., Nayakshin S., King A. R., Pringle J. E., 2009, *MNRAS*, 398, 1392
- Lubow S. H., D’Angelo G., 2006, *ApJ*, 641, 526
- Lubow S. H., Seibert M., Artymowicz P., 1999, *ApJ*, 526, 1001
- MacFadyen A., Milosavljević M., 2008, *ApJ*, 672, 83
- Malbon R. K., Baugh C. M., Frenk C. S., Lacey C. G., 2007, *MNRAS*, 382, 1394
- Manchester R. N., 2008, in Bassa C., Wang Z., Cumming A., Kaspi V. M., eds, *AIP Conf. Proc. Vol. 983, 40 YEARS OF PULSARS: Millisecond Pulsars, Magnetars and More*. Am. Inst. Phys., New York, p. 584
- Milosavljević M., Merritt D., 2003, *ApJ*, 596, 860
- Paardekooper S.-J., Papaloizou J. C. B., 2009, *MNRAS*, 394, 2283
- Paardekooper S.-J., Baruteau C., Crida A., Kley W., 2010, *MNRAS*, 401, 1950
- Pessah M. E., Chan C.-K., Psaltis D., 2007, *ApJ*, 668, L51
- Phinney E. S., 2001, preprint (astro-ph/0108028)
- Piran T., 1978, *ApJ*, 221, 652
- Sazhin M. V., 1978, *SvA*, 22, 36
- Sesana A., Vecchio A., 2010, *Phys. Rev. D*, 81, 4008
- Sesana A., Haardt F., Madau P., Volonteri M., 2004, *ApJ*, 611, 623
- Sesana A., Haardt F., Madau P., Volonteri M., 2005, *ApJ*, 623, 23
- Sesana A., Vecchio A., Colacino C. N., 2008, *MNRAS*, 390, 192 (SVC08)
- Sesana A., Vecchio A., Volonteri M., 2009, *MNRAS*, 384, 2255 (SVV09)
- Shakura N. I., Syunyaev R. A., 1973, *A&A*, 24, 337
- Springel V. et al., 2005, *Nat*, 435, 629
- Syer D., Clarke C. J., 1995, *MNRAS*, 277, 758
- Tanaka H., Takeuchi T., Ward W. R., 2002, *ApJ*, 565, 1257
- Trump J. R. et al., 2009, *ApJ*, 700, 49
- Tundo E., Bernardi M., Hyde J. B., Sheth R. K., Pizzella A., 2007, *ApJ*, 663, 57
- Volonteri M., Haardt F., Madau P., 2003, *ApJ*, 582, 599
- Wyithe J. S. B., Loeb A., 2003, *ApJ*, 590, 691
- Yoo J., Miralda Escudé J., Weinberg D. H., Zheng Z., Morgan C. W., 2007, *ApJ*, 667, 813

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