The Astrophysical Journal, 743:136 (10pp), 2011 December 20 doi:10.1088/0004-637X/743/2/136


# BLINDLY DETECTING ORBITAL MODULATIONS OF JETS FROM MERGING SUPERMASSIVE BLACK HOLES

R. O'Shaughnessy[1], D. L. Kaplan[1], A. Sesana[2,3], and A. Kamble[1]
[1] Physics Department, University of Wisconsin–Milwaukee, Milwaukee, WI 53211, USA; oshaughn@gravity.phys.uwm.edu, kaplan@uwm.edu, kamble@uwm.edu
[2] Albert Einstein Institute, Am Muhlenberg 1, D-14476 Golm, Germany; alberto.sesana@aei.mpg.de
[3] Center for Gravitational Wave Physics, The Pennsylvania State University, University Park, PA 16802, USA
*Received 2011 July 29; accepted 2011 September 8; published 2011 November 30*

## ABSTRACT

In the last few years before merger, supermassive black hole (SMBH) binaries will rapidly inspiral and precess in a magnetic field imposed by a surrounding circumbinary disk. Multiple simulations suggest that this relative motion will convert some of the local energy to a Poynting-dominated outflow, with a luminosity $\sim 10^{43}\,\mathrm{erg\,s^{-1}}\,(B/10^4 G)^2 (M/10^8 M_\odot)^2 (v/0.4c)^2$, some of which may emerge as synchrotron emission at frequencies near 1 GHz where current and planned wide-field radio surveys will operate. On top of a secular increase in power (and $v$) on the gravitational wave inspiral timescale, orbital motion will produce significant, detectable modulations, both on orbital periods and (if black hole spins are not aligned with the binary's total angular momenta) spin–orbit precession timescales. Because the gravitational wave merger time increases rapidly with separation, we find that vast numbers of these transients are ubiquitously predicted, unless explicitly ruled out (by low efficiency $\epsilon$) or obscured (by accretion geometry $f_{\rm geo}$). If the fraction of Poynting flux converted to radio emission times the fraction of lines of sight accessible $f_{\rm geo}$ is sufficiently large ($f_{\rm geo}\epsilon > 2 \times 10^{-4}$ for a 1 year orbital period), at least one event is accessible to future blind surveys at a nominal $10^4$ deg$^2$ with 0.5 mJy sensitivity. Our procedure generalizes to other flux-limited surveys designed to investigate electromagnetic (EM) signatures associated with many modulations produced by merging SMBH binaries.

*Key words:* black hole physics – cosmology: observations – radio continuum: general – surveys

*Online-only material:* color figures

## 1. INTRODUCTION

Merging supermassive black hole (SMBH) binaries should naturally possess a circumbinary accretion disk whose magnetorotational (MRI)-driven turbulence generates and imposes a substantial external magnetic field. Recent simulations suggest that SMBHs moving through an imposed magnetic field should produce a Poynting-dominated outflow (Mösta et al. 2010), consisting of a jet (Palenzuela et al. 2010) and diffuse emission (Mösta et al. 2011). Thus, even in the absence of accretion, during their last years before merger, SMBH binaries should generally have faint emission, modulated by their orbital motion. Though the outflow (henceforth "jet") power increases during the inspiral, ending in a bright flare, as discussed in Kaplan et al. (2011), such flares are too faint or too rare to easily detect, unless almost all the Poynting flux is efficiently converted to low-frequency radiation. On the other hand, because SMBH binaries spiral in very slowly through gravitational wave emission, each merger flare is preceded by a long phase during which the jet is modulated and growing, at only slightly reduced efficiency. These modulations should be easily accessible to future radio surveys.

In this paper, we calculate how frequently modulations will occur and how often they can be detected. Our paper adopts and extends the assumptions used in Kaplan et al. (2011) for the jet power versus binary masses and for fiducial SMBH binary merger rates. In Section 2 we sum over all SMBH binaries to determine the number of binaries on our past light cone with detectable modulation. In Section 3 we describe how long, how frequent, and how bright modulations from each SMBH binary should be. Finally, in Section 4 we discuss that, in order to detect the signature presented here, targeted electromagnetic (EM) and gravitational wave (GW) surveys must overcome limitations of intrinsic active galactic nucleus (AGN) variability or small limiting distances, respectively. We also briefly explain how the many distinctive variations in the predicted light curve will distinguish this source from other candidate modulation.

To date, most 1 GHz radio transient surveys have surveyed at most a few thousand square degrees with rms sensitivities of $\sim$ mJy (Ofek et al. 2011), with many having smaller area, less sensitivity, and sporadic sampling. Upcoming surveys will be larger, more regular, and more sensitive. In particular, the Variables and Slow Transients (VAST; T. Murphy et al. 2011, in preparation) project using the Australian Square Kilometer Array Pathfinder (Johnston et al. 2007) will survey roughly $10^4$ deg$^2$ daily down to a nominal sensitivity of 0.5 mJy. We adopt these parameters to motivate our discussion, keeping in mind that survey plans often change and that unknown factors such as efficiencies or obscurations strongly impact our results. Rather than fix these choices, we leave in scalings, so that interested parties can make their own predictions.

### 1.1. Context

Several mechanisms for precursor, prompt, and delayed electromagnetic signatures of SMBH merger have been proposed, including Poynting-dominated jets (Kaplan et al. 2011; Palenzuela et al. 2010; Mösta et al. 2011), disk emission powered by post-merger perturbations and shocks (Schnittman & Krolik 2008; O'Neill et al. 2009; Megevand et al. 2009; Zanotti et al. 2010; Corrales et al. 2010; Rossi et al. 2010), the viscous-driven refilling of the post-merger disk (Milosavljević & Phinney 2005; Tanaka & Menou 2010), accretion of a fossil inner disk driven by the smaller body's inspiral (Chang et al. 2010), increased

tidal disruption rates (Chen et al. 2009; Wegg & Bode 2010; Stone & Loeb 2011) and perturbations of galactic-center stellar cores, and direct modulation of a close circumbinary disk (Haiman et al. 2009b; Bode et al. 2011; Roedig et al. 2011; see Schnittman 2010, and references therein). These merging binaries will also produce gravitational waves potentially accessible to pulsar timing arrays and *LISA* (Jenet et al. 2009; Demorest et al. 2009; Hughes 2002; Sesana et al. 2007, 2009). However measured, the SMBH merger rate would provide invaluable direct constraints on the merger rates of galaxies (versus direct methods, as discussed in Bertone & Conselice 2009, and references therein) and the formation and evolution of SMBHs (see, e.g., Volonteri 2010, and references therein). Depending on the identification epoch, mechanism, and follow-up observations, direct detection of merging SMBH binaries could provide information about accretion, stellar dynamics, and the evolution of galaxies (Volonteri 2010; Bloom et al 2009).

A number of binary SMBHs are known or suspected (e.g., Komossa 2006; Rodriguez et al. 2006; Smith et al. 2010; Burke-Spolaor 2010). While many dual AGN candidates (e.g., 3C75 and NGC 6240) have been discovered with kpc scale separation, very few candidate SMBH binaries are known with $<10$ pc separation. Objects with shifted broad-line regions like SDSS J092712.65+294344.0 (Komossa et al. 2003) or J153636.22+044127.0 (Boroson & Lauer 2009) may be SMBH binaries; see, e.g., Bogdanović et al. (2009), Dotti et al. (2009), and references therein. Rodriguez et al. (2006) demonstrate that O402+379 has two resolved cores with a projected 7 pc separation. More speculatively, OJ287 has precessing jets (Marscher & Jorstad 2011) and outbursts with an approximately 12 year period (Sillanpaa et al. 1988). These modulations have been fit both by binary motion (Valtonen et al. 2008) and by a warped accretion disk (see, e.g., Katz 1997). Since most SMBH binaries are too far away to be resolved by very long baseline interferometry, future candidates will likely be found by offset emission lines or modulated spectral features (e.g., Comerford et al. 2009; Shen & Loeb 2010; Tsalmantza et al. 2011). For example, recently, Eracleous et al. (2011) found 14 quasars with offset H$\beta$ emission lines, with offsets that shifted significantly between widely separated observations. These deviations could arise from line-of-sight orbital acceleration in an SMBH binary.

The search for binary SMBH is complicated by AGN variability. This variability is largely not periodic; detecting such periodic behavior in a radio light curve would be a strong indication of an inspiralling SMBH pair (e.g., Komossa 2006).

Our study complements the previous model-neutral investigation by Haiman et al. (2009a), who also calculate both the number of SMBH binaries per unit orbital period and the conditions under which (optical) surveys could detect their variability; see, e.g., their Figure 9. They examined both gravitational-wave-dominated and disk-driven inspiral. By contrast, in this paper we study the gravitational-wave-dominated phase only, adopt a concrete emission model, employ a state-of-the art population of merging SMBH binaries distributed over cosmic time, and address how AGN background variability and jet beaming would impact survey performance. We also discuss how GW searches with pulsar timing naturally complement EM surveys for massive, comparable-mass SMBH binaries.

## 2. POPULATION OF MODULATED OUTFLOWS FROM MERGING BINARIES

Following Palenzuela et al. (2010) and Kaplan et al. (2011), we assume that each merging binary black hole has a Poynting-dominated outflow, even in the absence of accretion.[4] For convenience, we scale the outflow's instantaneous luminosity to a fraction $\epsilon_{\rm Edd}$ of the Eddington luminosity at the merger event:

$$L_{\rm flare} = \epsilon_{\rm Edd} q^2 (v/v_{\rm max})^2 L_{\rm edd}, \quad (1)$$

where $q = m_2/m_1 < 1$ is the binary's mass ratio and $v = \beta c$ is the binary's coordinate relative circular velocity $(|\partial_t(\vec{r}_2 - \vec{r}_1)|$. Palenzuela et al. (2010) find that most of the emission is trapped into collimated jets, with $\epsilon_{\rm Edd} \approx 0.002$. Using a refined calculation of the outgoing flux, Mösta et al. (2011) calculate that the jets, while present, are associated with a much brighter diffuse quadrupolar emission; in total, Mösta et al. (2011) estimate that the same fields correspond to $\epsilon_{\rm Edd} \approx 0.02$. In either scenario, the Poynting-dominated outflow will naturally interact with the strong ambient magnetic field and coronal plasma to produce radiation. Lacking a detailed emission model, we assume that a fraction $\epsilon_{\rm radio}$ of the beam power is converted to radio as $F_\nu \propto \epsilon_{\rm radio} L/\nu$ and absorb all details of spectral model, $K$-correction, and most issues pertaining to anisotropic emission[5] into $\epsilon_{\rm radio}$. A survey with a flux threshold $F_{\nu,\rm min}$ will therefore be sensitive to all binaries inside a luminosity distance $d_{\rm Edd}(v, M, q)$:

$$d_{\rm Edd} \simeq \sqrt{L/4\pi F_{\rm min}} \simeq 14.2\,{\rm Gpc}(\beta/\beta_{\rm max})\sqrt{\frac{q^2(M/10^6\,M_\odot)(\epsilon/0.002)}{(F_{\nu,\rm min}/\,{\rm mJy})(\nu/{\rm GHz})}} \quad (2)$$

$$\epsilon \equiv \epsilon_{\rm radio}\epsilon_{\rm Edd}, \quad (3)$$

where $M$ is the total mass of the binary. All key results and figures are presented as functions of $\epsilon$ to allow efficient scaling to any physically motivated values of $\epsilon_{\rm Edd}$, $\epsilon_{\rm radio}$.

The jet power gradually increases over time as the binary spirals inward toward merger due to the influence of gravitational radiation, with circular velocity increasing as

$$\frac{d\beta}{dt} = \frac{c^3}{G}\frac{96}{15}\frac{\eta}{M}\beta^9, \quad (4)$$

where the symmetric mass ratio $\eta = m_1 m_2/M^2 = q/(1+q)^2$. As each outflow is tied to a single black hole, we anticipate modulation at the orbital period

$$P = \frac{G}{c^3} 2\pi M \beta^{-3}. \quad (5)$$

[4] On physical grounds we expect $L_{\rm jet} \propto B^2 M^2 q^2 v^2$, based on the black hole area, the energy density of the magnetic field, and the relative orbital velocity of the hole and ambient field. Self-consistent MRI amplification already produces a large enough $B$ field in the disk (see, e.g., Pessah et al. 2006) to power the instantaneous outflow assumed above. Following Beckwith et al. (2009), Rothstein & Lovelace (2008), and references therein, we anticipate that small loops will escape into the circumbinary region, expand rapidly in the absence of confining material, and form a fluctuating but partially coherent large-scale field in the interior that threads the circumbinary region and the hole. Ordered ambient flux can be advected inward, increasing the ordered field near the horizon. Neither process requires matter accretion. Still higher magnetic fluxes can be achieved in the presence of accreting matter, a small amount of which will flow inward from the nonrelaxed inner edge of the circumbinary disk. A detailed discussion of $B$ field evolution and flux advection onto the horizon is beyond the scope of this paper. Rather than adopt a model for $B$ fields near the hole, as described in Kaplan et al. (2011), we conservatively adopt a smaller magnetic field, such that the maximum jet luminosity just prior to merger is the Eddington luminosity if $\epsilon_{\rm Edd} = 1$.

[5] Later we will explicitly account for how orientation-dependent obscuration ($f_{\rm geo}$) or beaming ($f_{\rm beam}$) impacts the number versus flux distribution.

Finally, for black holes with spin $S = \chi m^2$ misaligned with the orbital angular momentum, the orbital plane should precess, leading to modulations on a timescale of order (Apostolatos et al. 1994)

$$\frac{2\pi}{T_{\rm prec}} \simeq \frac{G}{c^2}\frac{J}{r^3}(2+3q) \simeq \frac{G}{c^2} \times \begin{cases} \frac{S_1}{r^3}(2+3q) \simeq 2\frac{|\chi|}{Mc}\beta^6 & \text{(late)} \\ \frac{L}{r^3}(2+3q) \simeq 2\frac{q}{Mc}\beta^5 & \text{(early).} \end{cases} \quad (6)$$

In these expressions we differentiate between early in the inspiral, when the orbital angular momentum is much larger than the black hole spin, and late in the inspiral, when $L \lesssim S_1$. Though the latter case occurs only in the last few orbits before merger with significant spin–orbit, it ensures a significant change in the direction of $\hat{L}$ during each precession period and provides the best opportunity for jet modulation. On the contrary, early in the evolution, the orbital plane precesses through a small angle $\simeq \mathcal{O}(S/L)$; while many cycles occur, only for exceptionally bright jets with well-understood variability could we ascribe small modulations to precession. Spin–orbit misalignment may not occur in gas-rich mergers, as accretion may align spins to the orbital plane (Bogdanović et al. 2007).

Given these timescales and a birthrate $\partial_t n$ per unit comoving volume and time, we can determine the number of SMBH binaries on our past light cone whose modulation timescales lie in a desired range and whose jets are bright enough to detect:

$$N_I = \int dV_c \left[ \int_{\beta_-(z)}^{\beta_+(z)} d\beta \frac{\partial_t n}{d\beta/dt} \right] = \int dV_c \partial_t n \left[ T(\beta_+) - T(\beta_-) \right], \quad (7)$$

where $dV_c$ is the comoving volume element, where $T(v)$ is the time until merger for a binary of velocity $c\beta$

$$T(\beta) = -\frac{G}{c^3}\frac{5}{256}\frac{M}{\eta}\beta^{-8}, \quad (8)$$

and where the redshift-dependent expressions $v_\pm$ are defined using the smallest interval satisfying (1) the desired time–velocity relationship, (2) the constraint that $d < d_{\rm Edd}$, and (3) that $T(\beta_\pm)$ are less than the time since decoupling from the circumbinary disk $\simeq \mathcal{O}(1)\,{\rm Myr}\, q^{7/13}(M/10^8\,M_\odot)^{17/13}$ (see, e.g., Schnittman & Krolik 2008; O'Neill et al. 2009, and references therein).[6] For example, the total number of binaries whose observed periods lie in $[P_{\rm obs,-}, P_{\rm obs,+}]$ follows from

$$\beta_c \equiv \beta_{\rm max}\sqrt{4\pi d_L^2 F_{\rm min}\Delta\nu/L_{\rm edd}\epsilon q^2} \quad (9a)$$

$$\beta_* \equiv T^{-1}(0.86\,{\rm Myr} \times q^{7/13}(M/10^8 M_\odot)^{17/13}) \quad (9b)$$

$$\beta_\pm = \min[\beta_{\rm max}, \max[P^{-1}(P_{{\rm obs},\mp}/(1+z)), \beta_c, \beta_*]], \quad (9c)$$

where $T^{-1}$, $P^{-1}$ are the inverse functions of Equations (5) and (8); a similar expression with $P \to T_{\rm prec}$ holds for precession. This expression distinguishes between the source-frame period $P$ and the observed period $P_{\rm obs} = P(1+z)$. Figure 1 illustrates how these four critical surfaces define a range of orbital velocity $\beta$ consistent with maximum source distance, maximum lifetime, and allowed period range. For example, circumbinary disks limit the orbital period of *decoupled, GW-inspiral-dominated* binaries to $P \lesssim 3\,{\rm yr}\, q^{15/26}(M/10^8\,M_\odot)^{29/26}$. Similarly, circumbinary disks a priori limit the number of decoupled black hole binaries to less than $\simeq$1 Myr times the all-sky merger rate: less than $10^7$, or $<240\,{\rm deg}^{-2}$. This limit is reached only for very sensitive surveys that probe all black holes throughout the universe. As implied by Figure 1, for typical surveys and for the vast majority of decoupled binaries at moderate redshift, the flux limit or period limit bounds the space of detectable binaries; the circumbinary disk plays no role.

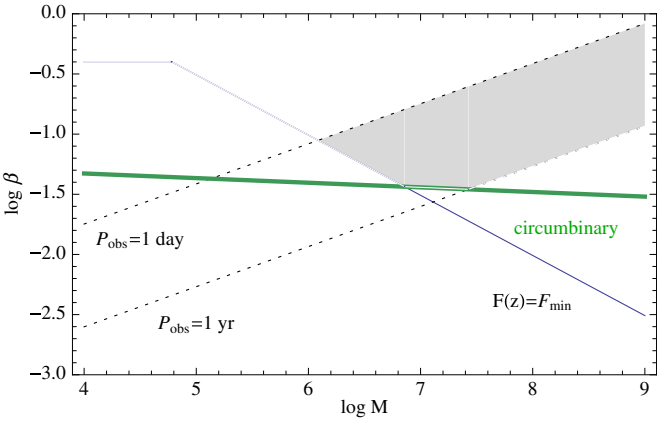


**Figure 1.** Critical orbital velocities vs. mass: for an equal-mass binary of mass $M$ at $z = 0.6$, a plot of orbital velocity vs. mass. The two dotted upward-sloping lines are the velocities associated with 1 year (bottom) and 1 day (top) orbital periods; a 1 year survey with 1 day cadence will be roughly sensitive to this range of time periods. The downward-sloping blue line shows the minimum velocity $\beta_c$, above which the jet is bright enough to be seen at $z = 0.6$, assuming $\epsilon = 0.002$ and $F_{\rm min} = 1$ mJy; this curve depends strongly on source distance, efficiency, and survey flux limit. Finally, the thick green line is $\beta_*$, the velocity at which the binary separates from the circumbinary disk (Equation (9)). The shaded region is bounded by $\beta_\pm$ and corresponds to the range of velocities where a binary is both isolated and detectable as a periodic source. The time $\Delta T$ a binary spends evolving through this region is $\Delta T = T(\beta_+) - T(\beta_-)$ for $T$ given by Equation (8).

(A color version of this figure is available in the online journal.)

Following Kaplan et al. (2011), we calculate the distribution of modulated flux at the earth using a distribution of merger rates as a function of black hole mass and cosmic time. The assembly of SMBHs is reconstructed through dedicated Monte Carlo merger simulations which are framed in the hierarchical structure formation paradigm. Briefly, these models evolve the black hole population starting from black hole "seeds," through accretion episodes triggered by galaxy mergers, and include the dynamical evolution of SMBH–SMBH binaries. The SMBH population is consistent with observational constraints, e.g., the luminosity function of quasars at $1 < z < 6$, the $M$–$\sigma$ relation and the black hole mass density at $z = 0$ (Volonteri et al. 2003, 2008; Volonteri & Begelman 2010). To illustrate the situation, we adopt two of the fiducial merger distributions as used in Arun et al. (2009). The two models used here are representative of two proposed seeding scenarios (see Sesana et al. 2011a for further discussion). In the notation of Arun et al. (2009), we compare models LE and SE, where S versus L refers to the seed size—large or small—and E refers to "efficient" accretion. The models described in that paper are representative of a range of plausible SMBH growth scenarios. As with uncertainties in

[6] For brevity, in our expressions we assume an accretion rate $\dot{m} = 0.3\dot{M}_{\rm edd}$ and disk viscosity parameter $\alpha = 0.3$, following King et al. (2007).

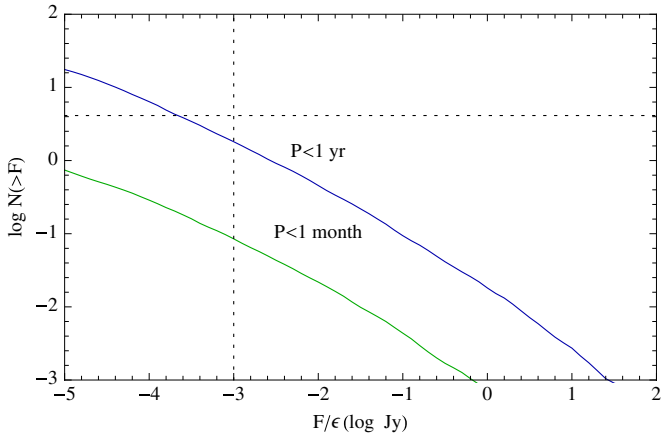


**Figure 2.** Sources with precession modulation: total number: assuming all binaries have significant spin–orbit misalignment, the number of sources with precession periods $P_{\rm min} = 1$ minute and $P_{\rm max} = 1$ year (blue) or 1 month (green) over the entire sky, with flux above $F/\epsilon$. To guide the eye to scales comparable to future radio surveys like VAST (T. Murphy et al. 2011, in preparation), a horizontal dashed line corresponds to $1/10^4$ deg$^2$ and a vertical dashed line corresponds to 1 mJy.

(A color version of this figure is available in the online journal.)

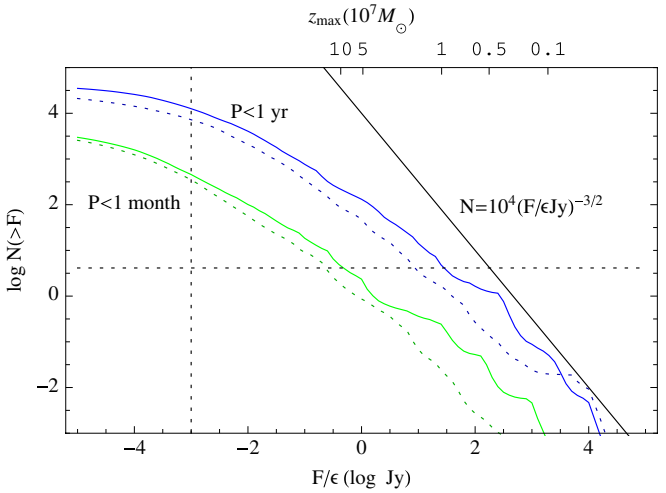


**Figure 3.** Sources with orbit period modulation I: total number: at any instant of time, the number of decoupled SMBH binaries vs. average jet flux, limiting to orbit periods within $P_{\rm min} = 1$ minute and $P_{\rm max} = 1$ year (blue and dotted blue) or 1 month (green and dotted green) over the entire sky, with flux above $F/\epsilon$ for $\epsilon = \epsilon_{\rm radio}\epsilon_{\rm Edd}$ the composite efficiency of the jet, using the SE (small seed, efficient accretion; dashed lines) and LE (large seed, efficient accretion; solid lines) model of Sesana et al. (2011a). For example, for an efficiency $\epsilon = 10^{-3}$ and target flux sensitivity 1 mJy, one to several hundred binaries are bright enough to see and have periods less than 1 year; of order one has a period less than 1 month. In the limit of infinite sensitivity, all rates must be less than the all-sky rate ($\simeq$30 yr$^{-1}$) times the longest time allowed (1 Myr, corresponding to the time since separation from a circumbinary disk). In the limit of poor sensitivity, the number vs. flux scales as $F^{-3/2}$, as usual for the nearby universe. For comparison, the solid black line is $10^4(F/\epsilon\,{\rm Jy})^{-3/2}$. The top axis shows ticks at $z_{\rm max}$, the redshift at which $L_{\rm edd}/4\pi d_L^2 = F/\epsilon$ for $M = 10^7 M_\odot$; as we assume $\epsilon < 1$, no source can be detected at higher redshift.

(A color version of this figure is available in the online journal.)

$\epsilon_{\rm radio}$, we attempt to make our conclusions robust to specific merger assumptions.

Any reasonably large sky patch should contain many bright SMBH binary jets whose plausible modulation timescales are within reach: minutes to months. Precessing sources, though the most unambiguous sources of modulation, are rare: only for the most optimistic efficiencies ($\epsilon \simeq 1$) will a typical-scale survey ($10^4$ deg$^2$) have sufficiently bright precessing jet in its sky patch (Figure 2). On the contrary, many binaries have orbital periods between minutes and months and are close enough to produce bright jets (Figure 3). The slow nature of gravitational wave losses ensures that most binaries are discovered near the widest orbits. By contrast, though the largest and closest binaries are most likely to be detected, any binary not inconsistent with the flux limit of the survey and that can plausibly be produced at that mass and redshift has a reasonable chance to be recovered (Figure 4). On the one hand, each value of $F/\epsilon$ defines a minimum mass versus redshift $M_c(z) \propto F/\epsilon$ below which no emission is visible even at merger;[7] these limiting curves are shown in Figure 4 for several choices of $F/\epsilon$. On the other hand, at each redshift the merger process does not efficiently produce SMBHs above a critical mass. In our two merger trees, this maximum mass $M_{\rm mgr}(z)$ is approximately

$$M_{\rm mgr}(z) \simeq 10^{10-z/2.5}\, M_\odot. \quad (10)$$

As a concrete example, in Figure 4 we show the mass–redshift distribution expected for a fiducial survey with $F = 0.5$ mJy sensitivity (e.g., T. Murphy et al. 2011, in preparation), adopting the most optimistic efficiency $\epsilon = 1$. Though concentrated at moderate mass and redshift, the distribution extends throughout the region bounded by the smallest flux (black) and largest mass (black dotted).[8] Results at other flux sensitivities are qualitatively similar, derived from this figure by suppressing binaries near and beyond the desired flux limit. Finally, the mass–redshift distribution implies a cumulative flux distribution $N(>F/\epsilon)$, with each value of $N(>F/\epsilon)$ corresponding to the integral in mass and redshift above a specific threshold curve for $F/\epsilon$ (blue). In particular, this relationship and these contours suggest that for every flux limit $F/\epsilon < 10^3$ Jy, a significant fraction of all detected binaries are always high mass and moderate- to high-redshift sources ($z > 1$).

Relatively few AGNs have unobscured lines of sight to their central engines, particularly at radio frequencies. Adopting a parameter $f_{\rm geo}$ to characterize the fraction of unobscured lines of sight,[9] the cumulative distribution of sources versus flux is $f_{\rm geo}N(>F/\epsilon)$. Fitting to the brightest sources in Figure 3, we estimate that the number of detectable binaries with periods less than 1 year is between[10]

$$N_{\rm obs} \simeq 10\text{–}30 f_{\rm geo}\left(\frac{\epsilon\,{\rm Jy}}{F}\right) \quad (11)$$

for a moderate range of flux ($F/\epsilon \in [10^{-2}, 10^3]$ Jy). Taking the lower limit, a fiducial survey with $10^4$ deg$^2$ coverage and 0.5 mJy sensitivity should have at least one binary with bright,

[7] This mass–redshift boundary is defined by $\epsilon L_{\rm edd}/4\pi d_L^2 = F_{\rm min}$ and would be attained only at merger of a comparable-mass binary.

[8] For numerical reasons, these boundaries also determine the range of fluxes we can reliably model. Our catalog consists of redshift bins $\Delta z \simeq 0.1$ out to $z \simeq 10$. For fluxes $F/\epsilon < 0.1$ mJy, binaries from $z > 10$ may contribute significantly. By contrast, for $F/\epsilon \gtrsim 10^4$ Jy, the limiting sensitivity passes inside the smallest redshift bins for a potentially significant proportion of mergers.

[9] This parameter absorbs all larger-scale obscuration effects, such as any obscuring torus or clouds that appear in AGN unification schemes; see Urry & Padovani (1995), Lawrence & Elvis (2010), and references therein. Standard AGN unification schemes require the torus subtend 65°, corresponding to a factor $f_{\rm geo}^{-1} \simeq 2$–4 in the radio (Lawrence & Elvis 2010).

[10] At the highest flux, the local universe dominates. Based on Equation (7) and a uniform local merger rate, we expect that $N(>F/\epsilon)$ at large flux must scale as $N \propto M^{5/6}P^{5/3}/(F/\epsilon)^{3/2}$.

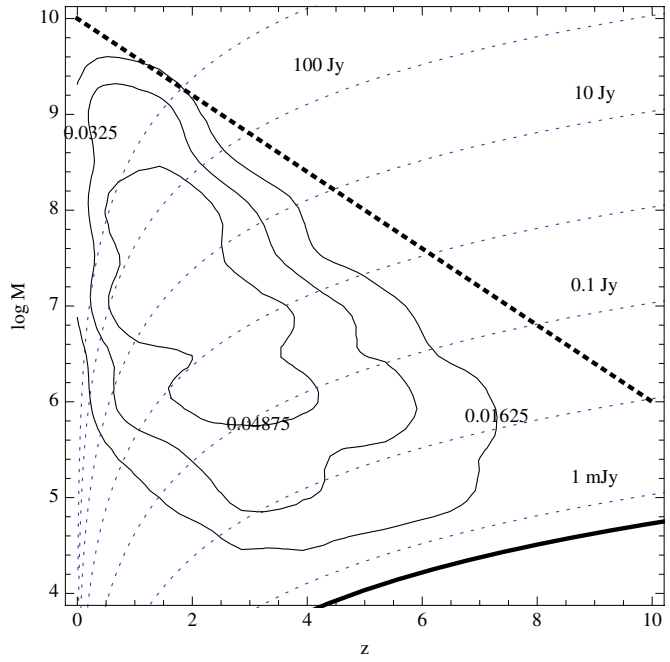


**Figure 4.** Sources with orbit period modulation Ib: mass–redshift distribution: contours of the mass–redshift distribution (solid black) of all binary black hole outflows with $F/\epsilon >= 0.5$ mJy (thick black curve) and observed periods $P_{\rm obs}$ between 1 year and 1 minute. This scaled flux limit corresponds to the fiducial VAST survey sensitivity ($F = 0.5$ mJy) for the most optimistic conversion of outflow to radio power ($\epsilon = 1$). Contours are shown at 1/4, 1/2, and 3/4 of the maximum value ($dN/d\log Mdz \simeq 0.06$). The distribution shown is smoothed, built by convolving a Gaussian kernel with the underlying merger tree, with smoothing lengths $\Delta \log M \simeq 0.3$ in mass and ($\Delta z = 0.3$) in redshift. Also shown are contours of the largest possible (Eddington-limited) flux at a given redshift: $F/\epsilon = L_{\rm edd}/4\pi d_L^2 \nu$ for $10^{-3}, 10^{-2}, \ldots, 10^3$ Jy (dotted blue, bottom to top). Finally, the thick dotted black line is an empirical relation for the maximum SMBH binary mass vs. redshift (Equation (10)).

(A color version of this figure is available in the online journal.)

detectable jets and observed period less than 1 year if the source obscuration and composite radiative efficiency satisfy

$$f_{\rm geo}\epsilon \gtrsim 2 \times 10^{-4}(F_{\rm min}/0.5\,{\rm mJy}). \quad (12)$$

In terms of Figure 3, this constraint corresponds to the number of orders of magnitude between (1) the flux at which the sky coverage limit (dotted) and the blue cumulative intersect and (2) the survey flux limit (just off scale to left). The number of sources in wide orbits is large; only extremely inefficient conversion of energy to radio can prevent them from being seen.

Conversely, Figure 2 shows that only a few binaries are both massive enough to produce a detectable radio jet by this mechanism and tight enough to undergo precession on an accessible timescale. Empirically, Figure 2 is approximated by $N(>F) \simeq 10^{-1.8}(F/\epsilon\ {\rm Jy})^{-2/3}$. Therefore, even an equally wide ($10^4$ deg$^2$) but much deeper (0.05 mJy) survey will have at least one binary with bright, detectable precessing jets only if jet power is efficiently converted to radio energy:

$$f_{\rm geo}^{3/2}\epsilon > 0.2(F_{\rm min}/0.05\,{\rm mJy}) \quad \text{(precession)}.$$

For brevity and clarity in this discussion we assume that the outflow emits roughly isotropically in all directions. If instead the same amount of power is beamed exclusively along a fraction $f_{\rm beam} < 1$ of all lines of sight, the number of detectable sources versus flux $N_{\rm beam}$ can be calculated from the previous expression via correcting for the fraction of the sources pointing toward us and the increased flux emitted along the remaining lines of sight:

$$N_{\rm beam}(>F/\epsilon) = f_{\rm beam} N\left(> \frac{F f_{\rm beam}}{\epsilon}\right). \quad (13)$$

Ignoring detectors' sensitivity, the fraction of jets pointing toward us and in the circumbinary phase is less than the total number of binaries that have decoupled from their circumbinary disk on our past light cone ($\lesssim 10^7$) times $f_{\rm beam} \equiv \Delta\Omega/4\pi$, the relative solid angle covered by a jet. By contrast, for experimentally accessible and observationally interesting flux ranges, the number versus flux trend $N(>F) \propto 1/F$ (Equation (11) and Figure 3), leading to $N_{\rm beam}$ independent of beaming.[11] While strictly true, beam modulation allows much longer orbits to be accessible and thereby can still significantly increase the number of accessible sources, as described below.

## 3. JET MODULATION

Throughout the inspiral, each black hole drives a Poynting-dominated outflow, powered by the strong ambient field provided by the circumbinary disk. As described in Kaplan et al. (2011; cf. their Footnote 1), our assumptions correspond to assuming that the disk provides a magnetic field such that the jet luminosity near merger $L_j(\beta = \beta_{\rm max}) \propto B^2M^2$ is the efficiency times the Eddington luminosity ($\epsilon_{\rm Edd}L_{\rm edd}$); the corresponding constant $B$ field is $B \simeq B_{\rm edd} \simeq 6 \times 10^4 \epsilon_{\rm Edd}^{1/2}(M/10^8\ M_\odot)^{-1/2}G$. A detailed treatment of the outflow kinematics and emission driven by this ideal-MHD jet is beyond the scope of this paper; see, e.g., Section 5 of Lyutikov (2011) for plausible emission microphysics and Palenzuela et al. (2010) for MHD. However, the large circumbinary field and coronal plasma provide a natural mechanism for dissipating the outflow, through strong synchrotron losses, inverse Compton scattering off of accelerated particles, an ambient medium to shock, and naturally relativistic particle velocity scales.[12] Following Kaplan et al. (2011), we assume that a fraction $\epsilon_{\rm radio}$ of the jet power is eventually emitted at radio frequencies, associated with a jet tied to the spin axes of each black hole.

As the black holes orbit and precess, the flux along our line of sight will be modulated by orbital effects (e.g., bolometric Doppler boosting, as in Loeb & Gaudi 2003), by precession of the spin axes changing the orientation of the jet, and by other relativistic effects associated with an ultracompact orbit. The observable modulation depends sensitively on jet dynamics and emission mechanisms, all beyond the scope of this paper. As an example with efficient conversion ($\epsilon \lesssim 1$) and simple modulation, in the rest of this paper we optimistically assume that Poynting flux is efficiently re-radiated by a sufficiently dense corona of weakly relativistic electrons as synchrotron radiation, in a short optically thin low Lorentz factor "jet." In this scenario, modulations at the orbit period arise from Doppler boosting during the orbit of each unresolved jet, with relative amplitude $\simeq \beta/\beta_{\rm max}$ times factors of order unity depending on the line of sight, spectrum, and mass ratio. Modulations at the precession period arises as the jet flow direction relative to the line of sight rotates, with line-of-sight flux fluctuations of order

[11] For surveys limited to very low redshift sources, $N(>F) \propto F^{-3/2}$. To the extent that all beamed sources are also nearby, beaming would indeed increase the detectable number: $N_{\rm beam} \propto f_{\rm beam}^{-1/2}$.

[12] For example, the Poynting flux in this ideal-MHD jet ($E \cdot B = 0$) has a strong electric field $|\vec{E}| \simeq \sqrt{L_j/r^2} \simeq B_{\rm edd}(M/r)$. Even in the absence of direct acceleration, this field produces strong transient drift currents $v \simeq c$.

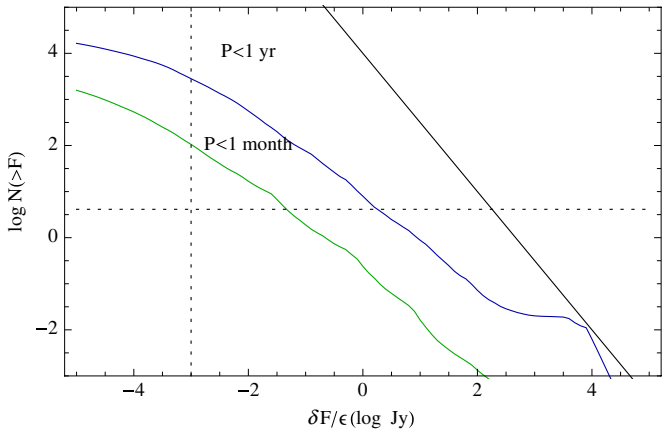


**Figure 5.** Sources with orbit period modulation II: significant modulation: the number of point sources with flux modulation greater than the limiting sensitivity vs. that sensitivity. Same as Figure 3, except the number is shown vs. $\delta F/\epsilon$, assuming a relative flux modulation $\delta$ per orbit period given by Equation (15). (A color version of this figure is available in the online journal.)

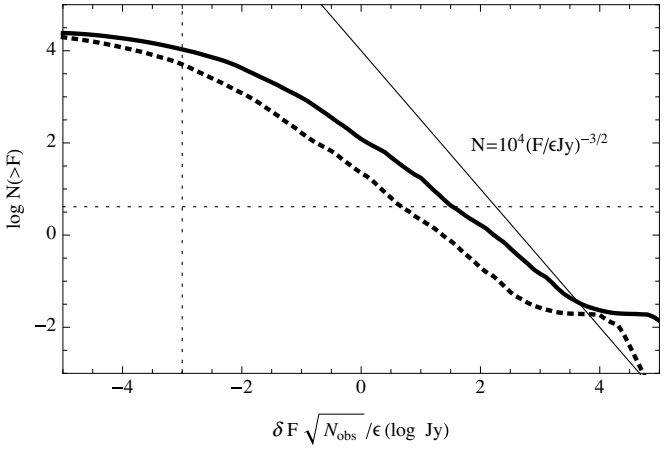


**Figure 6.** Sources with orbit period modulation III: stacking significant modulation: for the small seed (SE) model, the cumulative number of sources $N(> \delta F\sqrt{N_{\rm obs}}/\epsilon)$ at any instant which could be recovered by a survey with 1 day$^{-1}$ (thick) or 7 year$^{-1}$ (dotted), assuming 1 year observing time and stacking all $N_{\rm obs}$ flux measurements into a composite statistic. Each curve is directly comparable to the dark blue curve in Figure 5, after shifting the flux scale of the latter by $\sqrt{N_{\rm obs}}$.

unity for a low Lorentz factor jet. On the contrary, models with highly relativistic outflows have much stronger beaming, much less efficient conversion of jet to radio power, and a weak if any tie between the radio modulation timescale and any binary period: few sources will point toward us, while each will be radio-fainter and more randomly variable.

To identify modulation, the absolute change in flux from the modulation ($F \times \delta$ for $\delta$ the relative intensity change) should be larger than the detector noise. Equivalently, the per-measurement figure of merit derived from Fourier transforming the received flux

$$\mathrm{S/N} = \delta \frac{F}{F_{\rm min}} \tag{14}$$

should be greater than 1. Though realistic surveys will have a higher detection threshold, this factor can also be absorbed into $\epsilon$. For Doppler boosting at the orbit period, the relative change in power is small for bright and long-lived jets ($q \simeq 1$), owing to minimal contrast between the two holes' jets: to order of magnitude the relative power increase $\delta$ is (cf. Equation (1))

$$\delta \simeq \frac{\beta}{\beta_{\rm max}} \frac{1-q^2}{1+q^2}. \tag{15}$$

In this case, the analysis of Section 2 follows, replacing Equation (9) with

$$\beta_{c,\rm rel} = \beta_{\rm max} \left( \frac{4\pi d_L^2 F_{\rm min} \Delta\nu}{L_{\rm edd} \epsilon q^2} \times \frac{1+q^2}{1-q^2} \right)^{1/3}. \tag{16}$$

Out of the population of binaries whose jet power is bright enough to be seen and an orbit period in a testable range (Figure 3), Figure 5 shows a not insignificant fraction have jet power bright enough that Doppler boosting could be accessible. By contrast, for precession-induced changes in the jet direction, we expect a sinusoidal relative flux change of order unity: any precessing jet that is bright enough to be seen will have detectable modulation.

To this point we have considered the sensitivity limits for single measurements. By folding together multiple cycles we improve our sensitivity to faint sources. A priori, the flux sensitivity increases as the square root of the number of measurements $N_{\rm obs}$; the figure of merit becomes

$$\mathrm{S/N} = \delta \frac{F}{F_{\rm min}} \times \sqrt{N_{\rm obs}}. \tag{17}$$

Additionally, the survey duration $T_{\rm obs}$ and cadence $\Delta T = T_{\rm obs}/N_{\rm obs}$ naturally define a minimum and maximum range of periods to which the survey is sensitive: periods $P_{\rm obs}$ between approximately $2T_{\rm obs}/N_{\rm obs}$ and $T_{\rm obs}$. This subtlety has minimal impact, as binaries cluster near the longest orbital periods. As implied by the ratio of Equations (17) and (14), simply rescaling the distribution of all binaries whose peak-to-trough modulations are above the detection threshold (Figure 5) produces the distribution of sources with detectable periodic flux modulations (Figure 6).

In the above we have assumed that the flux is purely sinusoidally modulated, suitable to weakly Doppler-boosted synchrotron emission detected in the radio. By contrast, a highly relativistic outflow will generally be strongly Doppler boosted, tightly beamed, and emit most power preferentially at high energies. As noted earlier, beamed emission from the rare jets pointing toward us could be detected farther away (Equation (13)). Few binaries have bright jets and precession periods less than a year (Figure 2). However, any transverse structure $\propto \delta\theta$ in the narrow jet beam modulates the emitted flux on much shorter timescales $T_{\rm mod} \simeq \delta\theta/\Omega_p \simeq \sqrt{f_{\rm beam}} T_{\rm prec}$. The modulation in high-energy flux with time ($\delta(t)$) will generally contain a broad spectrum of frequencies, with frequency content depending on the details of the jet beam shape and its orientation relative to the line of sight. Because many orders of magnitude more binaries have much longer precession periods, these modulations in principle allow surveys of short cadence to identify sources with long natural periods. To order of magnitude and focusing only on bright jets in the local universe, the total number of sources increases as (solid angle of jet)$\times$ (inspiral time starting at a 1 yr/$\delta\theta$ precession period) $\times (d(\beta)/\delta\theta)^3$, the increased volume to which a beamed jet can be seen. Fixing $\beta$ in the center and final factors by requiring $T_{\rm prec} = 1\,{\rm yr}/\delta\theta \propto \beta^6$ (Equation (6)) and substituting for $d(\beta)$, *in principle* the number of detectable precessing sources increases relative to the small-

statistics tail (Figure 2) by a substantial factor:

$$\frac{N_{\rm prec,beam}}{N_{\rm prec}} \simeq \delta\theta^{-3/2} \simeq f_{\rm beam}^{-3/4}.$$

A much less favorable scaling applies at larger distances: following Equation (11), we expect $N \propto \delta\theta^{-1/2} \propto f_{\rm beam}^{-1/4}$. In practice, however, the long-timescale nonperiodic modulations that should arise from slowly precessing, beamed AGN jets should be difficult to distinguish from background AGN variability, as described in Section 4 below.

### 3.1. Ultrarelativistic Outflows

To this point we have implicitly assumed the radio flux to be weakly modulated on the orbit or precession period, as expected from prompt emission along a weakly relativistic orbit. By contrast, if these outflows are similar to conventional ultrarelativistic outflow models for AGN jets, then time of flight, generation, and reprocessing delays may modify any intrinsic structure, delaying and distorting the expected pattern from precession or the orbit.

Slowly precessing ultrarelativistic jets share many qualitative features with more familiar ultrarelativistic outflows like short gamma-ray bursts. In both cases, emission can be shock-driven, powered by interaction with the ambient medium. At best, emission from this shock will be delayed from prompt emission by time of flight and deceleration delays. Synchrotron radiation at low frequencies can be self-absorbed; this reprocessing delays and distorts any modulations imposed by the orbit. Finally, only observers over a range of angles $\theta_j + 1/\Gamma$ can see the jet (with $\theta_j$ being the jet opening angle and $\Gamma$ the bulk Lorentz factor); it may become visible only at late times. Just from outflow dynamics, the radio power expected from a relativistic outflow should plausibly peak on timescales of weeks to months after prompt emission.

Therefore, ultrarelativistic outflows if present can further bias us toward the longest timescales: sources that produce the shortest timescales will average away, leaving the multiple-year modulations expected among wide SMBH binaries.[13]

## 4. AGN BACKGROUNDS AND TARGETED SEARCHES

To this point we have assumed a blind survey, limited only by background and detector noise. In practice, additional sources of noise intrinsic and extrinsic to the binary system can severely complicate its detection or identification as a truly periodic source. On the contrary, highly relativistic SMBH binaries have unique kinematic and gravitational wave signatures that help distinguish true binary modulation from a background. Below we discuss two particular sources of noise (AGN variability and scintillation), and two methods to verify that a particular modulation has binary origin (multiple timescales and pulsar timing).

### 4.1. AGN Backgrounds and Targeted Searches

The host systems of binaries might be identified by residual accretion-powered emission. For example, accretion-fueled AGN jets persist $\mathcal{O}$(Myr) after the black holes stop accreting and separate from the disk. Magnetically powered jets may also already be identified. Conceivably, a targeted search could much more efficiently identify variability candidates.

On the other hand, AGN jets are well known to be strongly variable, with flux roughly performing a random walk on long timescales; see, e.g., Krolik (1998), and references therein. Likewise, these magnetically driven flows could be variable. A periodic modulation superimposed on a variable jet would be significantly more difficult to detect.

To illuminate the targeting problem, we parameterize the worst possible scenario: every merging binary has a bright residual radio jet, emitting at a significant fraction of the binary's Eddington luminosity ($\epsilon_{\rm old} L_{\rm edd}$, with $\epsilon_{\rm old} \gtrsim \epsilon$) and just as variable as active AGNs. For time changes less than a year, AGN fluxes are roughly random walks in amplitude (Hughes et al. 1992).[14] Specifically, if $w$ is a white random variable, a prototypical AGN flux and flux power spectrum is

$$F_{\rm AGN}(t) = \frac{\epsilon_{\rm old} L_{\rm edd}}{4\pi d_L^2} \int^t w dt \tag{18}$$

$$\begin{aligned} S_F(f) &\equiv \int dt e^{2\pi i f t} \langle F_{\rm AGN}(0) F_{\rm AGN}(t)\rangle \\ &= \delta_{\rm crit}^2 \left(\frac{\epsilon_{\rm old} L_{\rm edd}}{4\pi d_L^2}\right)^2 (f\ {\rm yr})^{-2}\ {\rm yr} \end{aligned} \tag{19}$$

normalizing to $\delta_{\rm crit}$ relative variations on 1 year timescales. In this worst case hypothesis, for any detectable jet, AGN variability is the dominant noise. As both the signal and dominant noise propagate from the source to our detector, the figure of merit for detectable sources depends only on ratios of properties at the source:[15]

$$\begin{aligned} {\rm S/N} &= \frac{\delta F}{\sqrt{S_F(1/P)/T}} \\ &= \frac{\epsilon q^2 (\beta/\beta_{\rm max})^3 (1-q^2)/(1+q^2)}{\epsilon_{\rm old}\delta_{\rm crit}} \sqrt{\frac{T \times 1\ {\rm yr}}{P^2}} \\ &= \frac{\epsilon}{\epsilon_{\rm old}\delta_{\rm crit}} \frac{\sqrt{T\,1\ {\rm yr}}}{P^2} (2\pi MG/(c\beta_{\rm max})^3) \frac{q^2(1-q^2)}{1+q^2}, \end{aligned} \tag{20, 21}$$

where for clarity we have replaced $\beta^3$ by the period (Equation (5)).

If indeed the residual jet radio power is brighter than modulation in the magnetically driven radio emission ($\epsilon < \epsilon_{\rm old}\delta_{\rm crit}$), the factors in this a priori expression are potentially unfavorable to targeted searches. As black hole light crossing times ($MG/c^3$) are less than a few minutes, targeted searches for periodic AGN variability from a jet superimposed on an existing source would only be sensitive to short timescales (i.e., $P < \sqrt{1\ {\rm yr}\ 500\ {\rm s}} \simeq 1$ day). Even then, detection could only occur if the binary jet and residual jet fluctuations are finely tuned to nearly the same flux (i.e., $\epsilon \simeq \epsilon_{\rm old}\delta$).

On the other hand, this Poynting-dominated outflow is not driven by accretion and is therefore likely far more stable than accretion-driven jets. Additionally, any residual jets should have dimmed and expanded in the thousands to million years since accretion terminated. Assuming that the transverse crossing time

[13] By contrast, an extended epoch of delayed emission will significantly *enhance* our prospects of detecting single brief events, such flares accompanying the merger of SMBH mergers; see Kamble et al (2011).

[14] We assume no correlation between AGN variability and the binary mass.

[15] Unlike the previous case, the AGN provides a unique noise realization for all detectors; repeated measurements at the same or multiple detectors does not improve our sensitivity.

sets the characteristic AGN jet variability timescale, a not-face-on residual jet of age $t_j$ should have variability timescales $\gtrsim \theta_j t_j$: as the jet ages, it loses the capacity to undergo rapid changes. As detailed treatment of this kind of residual AGN jet background is beyond the scope of this paper, we can neither endorse nor rule out the possibility that periodic variations might be accessible on top of likely jet variability backgrounds.

### 4.2. Scintillation

Emission from distant binaries propagates through plasma in their host, our galaxy, and the intergalactic medium. The motion of this intervening medium also introduces significant relative amplitude fluctuations ("scintillation"). Scintillation reprocesses both our target modulated signal and any colocated background. The amount of scintillation depends strongly on the line of sight to the source, the plasma content and motion near the merging binary, the observing frequency, the angular size of the source, and the modulation timescale of interest; see Walker (1998), Rickett (1990), and references therein. Because of the wide range in possible parameters that are unrelated to the intrinsic nature of the source, we do not model scintillation-induced backgrounds in detail here.

To some degree, modeling is unnecessary. To the extent that background AGN activity contaminates a signal, a phenomenological approach to AGN activity automatically includes all noise sources, scintillation included. Likewise, scintillation has characteristic power-law scalings of amplitude and timescale with frequency (Walker 1998), allowing identification and partial subtraction of these modulations.

That said, scintillation naturally produces $O(1)$ flux variations on an unknown but long timescale comparable to any periodicity of interest. For extended sources experiencing strong refractive scattering, scintillation naturally occurs on long timescales, roughly the time an intervening plasma element $v$ would need to cross the source size $R$. While the size of the emission region is poorly known, the lower bound of the horizon size ensures that this crossing time is large:

$$t_{\rm scin} \propto R/v \simeq 20\,{\rm day}(M/10^8\,M_\odot)(R/M)(50\,{\rm km\,s^{-1}}/v),$$

where we scale $v$ to typical galactic plasma velocities. In this case, typical scintillation flux variations should be large, with $\delta F/F \propto (\nu/\nu_o)^{17/30} \simeq O(1)$; based on Figure 1 from Walker (1998) we expect a transition frequency $\nu_o$ between 1 and 10 times a fiducial 1 GHz observing frequency, depending on the line of sight. For binary candidates found by radio surveys, multiband follow-up observations will be required to disentangle the impact of scintillation from any intrinsic modulation.

### 4.3. Unique Time Signatures Versus Confusing Variability Backgrounds

As exemplified by OJ287 (Katz 1997), an observed periodicity could be fit by other processes produced by single black holes, such as a warped disk causing the black hole to precess (Katz 1997) or periodic accretion from or excited by a companion (see, e.g., Section 1.1; Artymowicz & Lubow 1996; Hayasaki et al. 2007; Farris et al. 2011, and references therein). Binary black holes can likewise produce periodic emission through other mechanisms, including orbiting accretion flows and modes excited in the circumbinary disk. A detailed treatment of all possible backgrounds is beyond the scope of this paper. However, the mere presence of an unmodeled periodicity alone does not uniquely determine this mechanism.

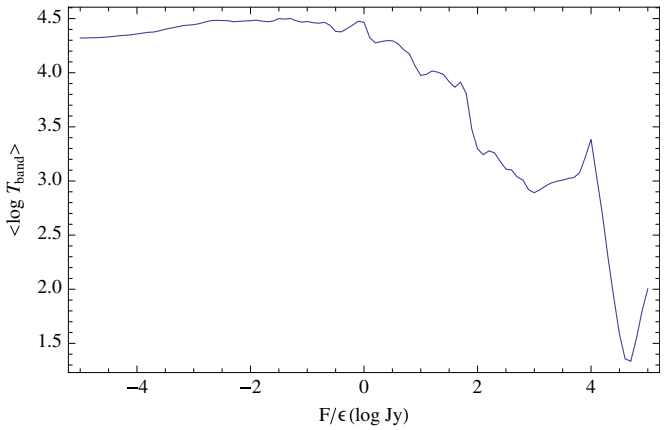


**Figure 7.** Typical lifetime: the average log lifetime $\langle \log T \rangle$ vs. scaled flux $F/\epsilon$, including sources with observed orbital period $<1$ yr. The brightest sources have short typical lifetimes, as they consist of the most massive nearby SMBH binaries in tight orbits.

(A color version of this figure is available in the online journal.)

On the other hand, extremely sensitive surveys that recover many merging events may find a few with distinctive signatures: either (1) a distinctive ultrarelativistic (i.e., Doppler-boosted and shapiro-delay-distorted) light curve, (2) multiple periodic timescales, or (3) a noticeable chirp. The most distinctive signature, a noticeable chirp, occurs surprisingly often among detectable binaries. For example, an all-sky survey with $F/\epsilon = 0.1$ Jy can potentially observe between 300 and 1000 binaries with less than 1 year observed periods (Figure 3). That population has a $\simeq$10% probability of one member undergoing a chirp from 1 year to 1 minute within 10 years. Moreover, the most massive and brightest members generally have shorter chirp times; see, e.g., the average log lifetime in Figure 7.

### 4.4. Pulsar Timing GW Signature

The massive, wide binaries which produce the most numerous and brightest jets are precisely the same sources whose gravitational wave emission is targeted by pulsar timing arrays (Sesana et al. 2009; Hobbs et al. 2010). Pulsar timing arrays are very likely to have a significantly shorter range than direct electromagnetic searches. To order of magnitude, the ratio of radio to gravitational wave power is $\simeq M^2 r^4 \omega^6 / L_{\rm edd}$, or

$$\frac{L_{\rm GW}}{L_j} = \frac{M\sigma_T}{m_p c^2}(2\pi/P)^2\beta^2 \simeq 3\times 10^6 \epsilon^{-1}(P/\,{\rm yr})^2(M/10^8\,M_\odot)\beta^2. \quad (22)$$

On the other hand, the relative energy flux sensitivity of a gravitational wave survey with characteristic strain sensitivity $h_s$

$$\frac{F_{\rm gw,min}}{F_{\rm min}\nu} = \frac{h_s^2(2\pi/P)^2}{F_{\rm em}\nu}. \quad (23)$$

As a result, the relative luminosity distance sensitivity of pulsar timing arrays and EM surveys goes as

$$\frac{d_{L,\rm gw}}{d_{L,\rm em}} = \sqrt{\frac{F_{\rm radio}\nu}{F_{\rm GW}}\frac{L_{\rm GW}}{L_{\rm EM}}} \quad (24)$$

$$= h_s^{-1}\beta\sqrt{F_{\rm radio}\nu M\sigma_T/m_p} \quad (25)$$

$$\simeq \frac{3\times10^{-4}}{(h_s/10^{-13})}\beta\sqrt{\frac{(F_{\rm radio}/\,{\rm mJy})(M/10^8\,M_\odot)}{(\epsilon/0.002)}}, \quad (26)$$

where we have scaled $h_s$ to existing pulsar timing array sensitivities (Yardley et al. 2010). Future pulsar timing arrays will have enough reach to identify sources up to their confusion limit. Conservatively assuming only one source per frequency can be distinguished (Sesana et al. 2009), at least a few sources over the entire sky could be individually resolved. Depending on how well the network localizes sources on the sky, the confusion limit could be significantly lower (cf. Boyle & Pen 2010). Nonetheless, barring extremely inefficient conversion of jet power to radio, direct electromagnetic surveys will have substantially greater sensitivity than pulsar timing. A pulsar timing survey by itself will only be more effective if the efficiency $\epsilon$ and flux $F$ imply that $N(>F/\epsilon)$ is small—in other words, only if radio surveys can find at best a handful of sources over the entire sky.

Nonetheless, electromagnetic and gravitational wave surveys naturally complement each other. Owing to considerable uncertainties in the emission geometry and model, only a gravitational wave signature can confirm that a jet candidate corresponds to an SMBH binary. Conversely, starting with a well-localized binary candidate from pulsar timing (Sesana & Vecchio 2010), electromagnetic observations can tightly localize a suitably modulated companion (Sesana et al. 2011b) with suitable overall circumbinary spectra (Tanaka et al. 2011; Milosavljević & Phinney 2005). Subsequently, as observations provide a candidate light curve corresponding to a known binary, archival and directed EM surveys can find more, at higher confidence. Eventually, in the limit of perfect coupling, a joint EM and gravitational wave survey could construct joint confidence in a proposed reconstruction of its data as a superposition of nearby and distant binaries plus background variability.

## 5. DISCUSSION

Each SMBH binary should possess a pair of magnetically powered bright jets attached to each hole during the last stages of gravitational-wave-driven inspiral. SMBHs therefore ubiquitously produce long-lived, modulated, and moderately bright emission even in (local) vacuum, without direct accretion. Though each jet is faint and though mergers are rare, gravitational-wave-driven inspiral is so slow that many modulation cycles exist per binary. Unlike rare merger events (Kaplan et al. 2011), these modulations should easily be seen with future blind radio surveys like VAST (T. Murphy et al. 2011, in preparation), only excepting pessimistic choices for the efficiencies $\epsilon_{\rm radio}$, $\epsilon_{\rm Edd}$, and obscuration $f_{\rm geo}$. For example, following Palenzuela et al. (2010) in adopting the conservative value $\epsilon_{\rm Edd}=0.002$, we predict that 30–60 observable binaries with observable orbital periods <1 yr could be seen by the VAST survey ($10^4$ deg sky coverage, and either 1 day$^{-1}$ or 7 year$^{-1}$ cadence; see Figure 6). Alternatively, for the same conditions, Mösta et al. (2011) predict significantly more diffuse emission; adopting their result $\epsilon_{\rm Edd}=0.04$, roughly 10 times more systems could be detected. Our results are insensitive to the precise merger rate we assume.

At any time, the most accessible sources are the largest black hole binaries in the widest (≃year) orbits. Nonetheless, surveys with shorter cadence can achieve comparable sensitivity by stacking cycles. The gravitational wave emission from these systems is also the target of pulsar timing arrays (Sesana et al. 2009). Unless very little jet power escapes in modulated radio jets, radio surveys will generally be more efficient than pulsar timing surveys for the same binaries.

The region surrounding an SMBH offers many opportunities to produce variability on comparable timescales and comparable bands. In special but not exceptionally rare circumstances, this mechanism can be identified by multiple timescales—depending on the scale, some combination of effects from its ultrarelativistic orbit (Doppler and Shapiro), spin–orbit precession, and in exceptional circumstances the gravitational wave chirp.

If most merging binaries accreted and produced long-lived AGN jets before decoupling from the circumbinary disk, then this long-lived remnant likely has significant variability which could severely limit the efficacy of a radio survey. Targeted searches toward known jets must overcome similar variability challenges.

While we emphasize and scale fiducial results to radio-band measurements, our expressions involve only flux sensitivity limits and efficiencies and scale to any band.[16] On the other hand, our preferred model of sinusoidal modulations visible equally well in all directions likely scales to low-energy emission. High-energy emission naturally arises from an ultrarelativistic and tightly beamed outflow, with dominant modulations from precession and subdominant modulations from orbital motion (e.g., from curvature near the source, for marginally detected lines of sight into tight orbits). In that case, the probability of detecting prompt emission from any one jet will be reduced by $\delta\theta^2/4\pi$. On the other hand, the internal structure of ultrarelativistic jets versus angle could magnify the effect of small changes in angle. A handful of tightly beamed jets could exhibit a distinctive rise and fall on the longest observing timescales.

In this paper we only address the prospects for detecting a periodically modulated outflow in the radio, after the binary decouples completely from its circumbinary disk. We assume no accretion. However, so long as gravitational radiation dominates the inspiral, our calculations can be easily generalized to any partially or totally accretion-powered EM signal, just by replacing Equation (1) for $L_j(\beta)$ and by modeling the modulation $\delta(P)$ in Equation (14). For example, Chang et al. (2010) propose that an inner circumbinary disk will be accreted onto a more massive companion, with disk luminosity $L\propto T^{-5/4}\propto\beta^{10}$ peaking at $L\simeq 0.1L_{\rm Edd}$. Observations will easily distinguish between these generalizations: each generically predicts different scalings of number versus flux, period, and modulation shape. For example, Haiman et al. (2009a) discuss how the flux distribution $N(>F)$ could be used to extract information about the relative scaling of inspiral and emission with orbit period.

In our paper, we assume that emission tracks the underlying SMBH dynamics, the outflow provides enough power for potentially significant reach in a quiet band (e.g., $\epsilon>10^{-4}$), and reprocessing or relativistic kinematics at most weakly distorts

[16] At sufficiently high energies, less than one photon arrives per orbit period in a reasonable-scale detector area. This effect is significant only for extremely low mass binaries in ultrarelativistic orbits, exceptionally high photon energies, or exceptionally low efficiencies: $M<0.5\times10^3\,M_\odot(d/\,{\rm Gpc})(E_\gamma/\,{\rm keV})^{1/2}v^{3/2}/\epsilon^{1/2}$. Even for $\epsilon\simeq10^{-6}$, the most interesting sources would still produce photon modulations in the X-ray.

the signal. The details of the emission mechanism remain unknown and merit further detailed investigations.

D.L.K. and A.K. are partially supported by NSF award AST-1008353. R.O.S. is supported by NSF award PHY-0970074, the Bradley Program Fellowship (BCS), and the UWM Research Growth Initiative. A.S. is supported by the Max Planck Society. The authors thank M. Eracleous, J. Krolik, H. Bignall, C. Palenzuela, L. Rezzolla, P. Mösta, D. Alic, O. Zanotti, and J. Lazio for helpful feedback, as well as the anonymous referee for thoughtful suggestions. R.O.S. and A.S. also thank Aspen Center for Physics, where this work was completed, and all participants of its 2011 program on Galaxy and Supermassive Black Hole Coevolution.

## REFERENCES

Apostolatos, T. A., Cutler, C., Sussman, G. J., & Thorne, K. S. 1994, Phys. Rev. D, 49, 6274
Artymowicz, P., & Lubow, S. H. 1996, ApJ, 467, L77
Arun, K. G., Babak, S., Berti, E., et al. 2009, Class. Quantum Grav., 26, 094027
Beckwith, K., Hawley, J. F., & Krolik, J. H. 2009, ApJ, 707, 428
Bertone, S., & Conselice, C. J. 2009, MNRAS, 396, 2345
Bloom, J. S., Holz, D. E., Hughes, S. A., et al. 2009, arXiv:0902.1527
Bode, T., Bogdanovic, T., Haas, R., et al. 2011, arXiv:1101.4684
Bogdanović, T., Eracleous, M., & Sigurdsson, S. 2009, ApJ, 697, 288
Bogdanović, T., Reynolds, C. S., & Miller, M. C. 2007, ApJ, 661, L147
Boroson, T. A., & Lauer, T. R. 2009, Nature, 458, 53
Boyle, L., & Pen, U.-L. 2010, arXiv:1010.4337
Burke-Spolaor, S. 2011, MNRAS, 410, 2113B
Chang, P., Strubbe, L. E., Menou, K., & Quataert, E. 2010, MNRAS, 407, 2007
Chen, X., Madau, P., Sesana, A., & Liu, F. K. 2009, ApJ, 697, L149
Comerford, J. M., Gerke, B. F., Newman, J. A., et al. 2009, ApJ, 698, 956
Corrales, L. R., Haiman, Z., & MacFayden, A. 2010, MNRAS, 404, 947
Demorest, P., Lazio, J., Lommen, A., et al. 2009, in Astronomy, Vol. 2010, AGB Stars and Related Phenomena, Astro2010: The Astronomy and Astrophysics Decadal Survey, 2010, 64
Dotti, M., Montuori, C., Decarli, R., et al. 2009, MNRAS, 398, L73
Eracleous, M., Boroson, T. A., Halpern, J. P., & Liu, J. 2011, arXiv:1106.2952
Farris, B. D., Liu, Y. T., & Shapiro, S. L. 2011, Phys. Rev. D, 84, 024024
Haiman, Z., Kocsis, B., & Menou, K. 2009a, ApJ, 700, 1952
Haiman, Z., Kocsis, B., Menou, K., Lippai, Z., & Frei, Z. 2009b, Class. Quantum Grav., 26, 094032
Hayasaki, K., Mineshige, S., & Sudou, H. 2007, PASJ, 59, 427
Hobbs, G., Archibald, A., Arzoumanian, Z., et al. 2010, Class. Quantum Grav., 27, 084013
Hughes, P. A., Aller, H. D., & Aller, M. F. 1992, ApJ, 396, 469
Hughes, S. A. 2002, MNRAS, 331, 805
Jenet, F., et al. 2009, arXiv:0909.1058
Johnston, S., Bailes, M., Bartel, N., et al. 2007, PASA, 24, 174
Kaplan, D. L., O'Shaughnessy, R., Sesana, A., & Volonteri, M. 2011, ApJ, 734, L37
Katz, J. I. 1997, ApJ, 478, 527
King, A. R., Pringle, J. E., & Livio, M. 2007, MNRAS, 376, 1740
Komossa, S. 2006, Mem. Soc. Astron. Ital., 77, 733
Komossa, S., Burwitz, V., Hasinger, G., et al. 2003, ApJ, 582, L15
Krolik, J. 1998, in Active Galactic Nuclei, ed. J. H. Krolick (Princeton Series in Astrophysics; Princeton, NJ: Princeton Univ. Press)
Lawrence, A., & Elvis, M. 2010, ApJ, 714, 561
Loeb, A., & Gaudi, B. S. 2003, ApJ, 588, L117
Lyutikov, M. 2011, Phys. Rev. D, 83, 064001
Marscher, A. P., & Jorstad, S. G. 2011, ApJ, 729, 26
Megevand, M., Anderson, M., Frank, J., et al. 2009, Phys. Rev. D, 80, 024012
Milosavljević, M., & Phinney, E. S. 2005, ApJ, 622, L93
Mösta, P., Palenzuelá, C., Rezzalla, L., et al. 2010, Phys. Rev. D, 81, 064017
Mösta, P., Alic, D., Rezzolla, L., & Zanotti, O. 2011, arXiv:1109.1177
Ofek, E. O., Frail, D. A., Breslauer, B., et al. 2011, ApJ (arXiv:1103.3010)
O'Neill, S. M., Miller, M. C., Bogdanović, T., Reynolds, C. S., & Schnittman, J. D. 2009, ApJ, 700, 859
Palenzuela, C., Lehner, L., & Liebling, S. L. 2010, Science, 329, 927
Pessah, M. E., Chan, C., & Psaltis, D. 2006, Phys. Rev. Lett., 97, 221103
Rickett, B. J. 1990, ARA&A, 28, 561
Rodriguez, C., Taylor, G. B., Zavala, R. T., et al. 2006, ApJ, 646, 49
Roedig, C., Dotti, M., Sesana, A., Cuadra, J., & Colpi, M. 2011, MNRAS, 415, 3033R
Rossi, E. M., Lodato, G., Armitage, P. J., Pringle, J. E., & King, A. R. 2010, MNRAS, 401, 2021
Rothstein, D. M., & Lovelace, R. V. E. 2008, ApJ, 677, 1221
Schnittman, J. D. 2010, arXiv:1010.3250
Schnittman, J. D., & Krolik, J. H. 2008, ApJ, 684, 835
Sesana, A., Gair, J., Berti, E., & Volonteri, M. 2011a, Phys. Rev. D, 83, 044036
Sesana, A., Roedig, C., Reynolds, M. T., & Dotti, M. 2011b, arXiv:1107.2927
Sesana, A., & Vecchio, A. 2010, Phys. Rev. D, 81, 104008
Sesana, A., Vecchio, A., & Volonteri, M. 2009, MNRAS, 394, 2255
Sesana, A., Volonteri, M., & Haardt, F. 2007, MNRAS, 377, 1711
Shen, Y., & Loeb, A. 2010, ApJ, 725, 249
Sillanpaa, A., Haarala, S., Valtonen, M. J., Sundelius, B., & Byrd, G. G. 1988, ApJ, 325, 628
Smith, K. L., Shields, G. A., Bonning, E. W., et al. 2010, ApJ, 716, 866
Stone, N., & Loeb, A. 2011, MNRAS, 412, 75S
Tanaka, T., & Menou, K. 2010, ApJ, 714, 404
Tanaka, T., Menou, K., & Haiman, Z. 2011, arXiv:1107.2937
Tsalmantza, P., Decarli, R., Dotti, M., & Hogg, D. W. 2011, arXiv:1106.1180
Urry, C. M., & Padovani, P. 1995, PASP, 107, 803
Valtonen, M. J., et al. 2008, Nature, 452, 851
Volonteri, M. 2010, A&AR, 18, 279
Volonteri, M., & Begelman, M. C. 2010, MNRAS, 409, 1022
Volonteri, M., Haardt, F., & Madau, P. 2003, ApJ, 582, 559
Volonteri, M., Lodato, G., & Natarajan, P. 2008, MNRAS, 383, 1079
Walker, M. A. 1998, MNRAS, 294, 307
Wegg, C., & Bode, J. N. 2010, arXiv:1011.5874
Yardley, D. R. B., Hobbs, G. B., Jenet, F. A., et al. 2010, MNRAS, 407, 669
Zanotti, O., Rezzolla, L., Del Zanna, L., et al. 2010, A&A, 523, A8