

Catching supermassive black hole binaries without a net

Neil J. Cornish and Edward K. Porter

Department of Physics, Montana State University, Bozeman, Montana 59717, USA

(Received 26 May 2006; published 3 January 2007)

The gravitational wave signals from coalescing Supermassive Black Hole Binaries are prime targets for the Laser Interferometer Space Antenna (LISA). With optimal data processing techniques, the LISA observatory should be able to detect black hole mergers anywhere in the Universe. The challenge is to find ways to dig the signals out of a combination of instrument noise and the large foreground from stellar mass binaries in our own galaxy. The standard procedure of matched filtering against a grid of templates can be computationally prohibitive, especially when the black holes are spinning or the mass ratio is large. Here we develop an alternative approach based on Metropolis-Hastings sampling and simulated annealing that is orders of magnitude cheaper than a grid search. For the first time, we show that it is possible to detect and characterize the signals from binary systems of Schwarzschild Black Holes that are embedded in instrument noise and a foreground containing millions of galactic binaries. Our technique is computationally efficient, robust, and applicable to both high and low signal-to-noise ratio systems.

DOI: [10.1103/PhysRevD.75.021301](https://doi.org/10.1103/PhysRevD.75.021301)

PACS numbers: 95.85.Sz, 04.30.-w

Supermassive Black Hole Binaries (SMBHBs) and Extreme Mass Ratio Inspirals (EMRIs) of a compact object into a supermassive black hole are two of the most exciting targets for the LISA observatory [1]. Studies of these objects will yield insights into the role played by black holes in structure formation and galactic dynamics. The signals will also encode information about strong field, dynamical gravity that can be used to perform precision tests of general relativity [2–4].

The SMBHB and EMRI signals contain a wealth of information that is encoded in a highly modulated time series composed of multiple harmonics of several distinct, evolving periods. The complexity of the signal is good news in terms of the science yield, but it poses a significant challenge to the data analyst. The signal from a binary system of structureless spinning objects, as described by general relativity and detected by LISA, is controlled by 17 parameters. In the case of SMBHBs the systems are expected to have circularized before entering the LISA band, thereby reducing the search to 15 parameters. In the case of EMRIs the orbits are expected to maintain significant eccentricity in the LISA band, but the spin of the smaller body can be neglected, thereby reducing the search to 14 parameters.

The large dimension of the search spaces and the high computational cost of generating the search templates make SMBHBs and EMRIs challenging targets for data analysis [5]. The problem only gets worse when one considers that we need to extract these signals from a timeseries that also contains the signals from millions of galactic binaries, and in the case of EMRIs, a possible self-confusion from hundreds of other EMRI systems [6].

It has been estimated that it would take 10^{40} templates to perform an optimal grid search for EMRI signals [5]. The numbers are less for SMBHBs, but still out of reach computationally. Several alternative approaches have been discussed, including nontemplate based strategies

that look for tracks in spectrograms [7], and hierarchical, semicoherent grid based searches [5]. Here we consider an alternative approach that uses Metropolis-Hastings sampling and simulated annealing to search through the space of templates. Our search method is closely related to the Markov Chain Monte Carlo (MCMC) [8] method that is used to explore the posterior distribution of the model parameters once the source has been located. In previous work the MCMC approach was used to test the Fisher matrix predictions for SMBHB parameter uncertainties by starting the chains off very close to true source parameters [9,10]. However, existing MCMC approaches, including sophisticated adaptive Reverse Jump implementations, perform poorly when searching large regions of parameter space. Here we demonstrate for the first time a robust and efficient *detection* algorithm that is many orders of magnitude faster than the MCMC search algorithms that have been investigated to-date. Advanced MCMC techniques that employ importance resampling and well designed priors have been used to study 5-parameter binary inspiral signals in the context of ground based gravitational wave detectors [11]. It would be very interesting to see how this algorithm performs in the LISA context. We apply our search algorithm to simulated LISA data streams that include the signals from a pair of nonrotating black holes and a foreground produced by galactic white dwarf binaries. While the SMBHB system we consider is simpler than the general case (the model is described by 9 parameters rather than 15), it serves to illustrate the relative economy of the gridless approach.

The gravitational waveform for a supermassive black hole system consisting of two Schwarzschild black holes is described by 9 parameters: the redshifted chirp mass, M_c ; the redshifted reduced-mass, μ ; the sky location, (θ, ϕ) ; the time-to-coalescence, t_c ; the inclination of the orbit of the binary, ι ; the phase of the wave at coalescence, φ_c ; the luminosity distance, D_L ; and the polarization angle, ψ . The

parameters M_c , μ are intrinsic to the system, while D_L , ι , φ_c , ψ are extrinsic as they depend on the perspective of the observer. The other three parameters, (θ, ϕ) and t_c , would be extrinsic if the LISA observatory were static, but the motion of the detector couples these parameters to the intrinsic evolution. The extrinsic parameters D_L , ι , φ_c , ψ can be analytically solved for using a generalized F -statistic [12], leaving a 5 dimensional search space.

We illustrate the performance of the search algorithm by considering two representative LISA sources—a $10^6 - 10^5 M_\odot$ binary system at $z = 1$ and a $10^5 - 5 \times 10^4 M_\odot$ binary system at $z = 5.5$. In each case the time of observation is 6 months, and the observations end ~ 1 week prior to merger. The early termination of the signal is designed to demonstrate LISA's ability to give early warning to other telescope facilities. The $z = 1$ example has parameters $(M_c/M_\odot, \mu/M_\odot, D_L/\text{Gpc}, t_c/\text{months}, \theta, \phi, \iota, \varphi_c, \psi) = (4.93 \times 10^5, 1.82 \times 10^5, 6.6, 6.23, 1.325, 2.04, 1.02, 0.95, 0.66)$ and the $z = 5.5$ example has parameters $(M_c/M_\odot, \mu/M_\odot, D_L/\text{Gpc}, t_c/\text{months}, \theta, \phi, \iota, \varphi_c, \psi) = (3.95 \times 10^5, 2.17 \times 10^5, 53, 6.25, 1.927, 0.351, 1.318, 2.0, 0.23)$. To make the searches more realistic, we add in a galactic foreground consisting of approximately 26×10^6 galactic sources. The galactic binary foreground is generated using a Nelemans, Yungelson and Zwart galaxy model [13,14]. The signal-to-noise ratio (SNR) for the sources is estimated using the combined instrument and galactic confusion noise quoted in Ref. [14]. The $z = 1$ example has SNR = 118.0 and the $z = 5.5$ example has SNR = 9.87. These SNR ratios are on the low side for typical LISA observations of SMBHBs as we terminate the observations a week before merger. The full inspiral signals would give SNRs of ~ 387 and ~ 182 for the two cases, and the merger and ringdown signals would further boost the SNRs by a factor of ~ 2 or more. In Fig. 1 we plot the detector response to the galactic foreground and instrumental noise, along with the noise-free response to the SMBHB signals. We use restricted post-Newtonian waveforms with 2-PN evolution of the phase and we employ the two independent Michelson-like interferometry channels that are available at low frequencies [15]. At low frequencies these Michelson-like channels are equivalent to the generation 1.5 A and E time delay interferometry channels [16]. As the equations that describe the phase evolution break down before we reach the last stable circular orbit at $R = 6M$, we terminate the search templates at a maximum value of $R = 7M$. For the sources in question, the observation period terminates a week from coalescence, so the maximum gravitational wave frequency reached is 0.28 mHz for the $z = 1$ example and 0.32 mHz for the $z = 5.5$ example. With this frequency range, the SMBHBs overlap with over 22.5×10^6 galactic binaries. To minimize the computational cost, the search templates were generated at a sample cadence of 4.2 mHz.

Our search algorithm uses Metropolis-Hastings rejection sampling, simulated annealing and algebraic extrem-




FIG. 1. The strain spectral density in a single LISA channel. The gray line is the LISA response to a galactic background of 26×10^6 sources plus simulated instrumental noise. The solid black lines show the LISA response to the SMBHB signals alone. The dashed black line indicates the RMS instrument plus galactic confusion noise level.

ization over extrinsic and quasiextrinsic parameters. The sampling proceeds as follows: Choose a random starting point \vec{x} in parameter space. Using a proposal distribution $q(\cdot|\vec{x})$, draw a new point \vec{y} . Evaluate the Hastings ratio

$$H = \frac{\pi(\vec{y})p(s|\vec{y})q(\vec{x}|\vec{y})}{\pi(\vec{x})p(s|\vec{x})q(\vec{y}|\vec{x})}. \quad (1)$$

Accept the candidate point \vec{y} with probability $\alpha = \min(1, H)$, otherwise remain at the current state \vec{x} . Here $\pi(\vec{x})$ are the priors on the parameters,

$$p(s|\vec{x}) = \text{const.} e^{-\langle s-h(\vec{x})|s-h(\vec{x}) \rangle/2}, \quad (2)$$

is the likelihood and $q(\vec{x}|\vec{y})$ is the proposal distribution. The angular brackets $\langle s-h(\vec{x})|s-h(\vec{x}) \rangle$ denote the standard noise weighted inner product of the signal s minus the template $h(\vec{x})$. We used the combined instrument and galactic confusion noise model of Timpano *et al.* [14]. We employ three different proposal distributions that are designed to give small, medium and large jumps. This mixture of jump sizes gives the search the flexibility to fully explore the parameter space and the ability to quickly hone in on promising regions. The small jumps are drawn from a multivariate Normal distribution, the medium sized jumps are given by a uniform draw of $\pm 10\sigma$ in each parameter and the large jumps come from a full range, uniform draw on all the parameters. We used a mixture of 20 small jump proposals for every medium or large jump proposal. Correlations between the parameters can seriously hurt the acceptance rate, so we use a multivariate Normal distribution that is the product of Normal distributions in each eigendirection of the Fisher information matrix,

$\Gamma_{ij}(\vec{x})$. The standard deviation in each eigendirection is set equal to $\sigma_i = 1/\sqrt{DE_i}$, where $D = 9$ and E_i is the corresponding eigenvalue of the Fisher matrix [17]. The Fisher matrix is also used to scale the medium size jumps.

The simulated annealing is done by multiplying the noise weighted inner product $\langle s|h \rangle$ by an inverse temperature β . We used a standard power-law cooling schedule:

$$\beta = \begin{cases} 10^{B(1-i/N_c)} & 0 \leq i \leq N_c, \\ 1 & i > N_c \end{cases}, \quad (3)$$

where i is the number of steps in the chain and N_c is the number of steps the chain takes to reach the normal temperature. We found that an initial heat factor of between 10 to 100 and a cooling schedule that lasted for $\sim 10^4$ steps worked well, but the performance was not particularly sensitive to these choices. For low SNR sources smaller initial heat factors and slightly longer cooling schedules yielded better results.

The F -statistic is used to automatically extremize over the four parameters (D_L , ι , ψ , φ_c), but the motion of the LISA detector sets a time reference, so the usual trick of using a fast Fourier transform to extremize over the time to coalescence, t_c , is not strictly permitted. However, the waveforms are much less sensitive to the sky location

than they are to t_c , so we employed t_c maximization during the annealing phase for the large and medium jump proposals. This procedure biases the solution, but the bias is erased by subsequent jumps.

In dozens of tests applied to many different examples, our search algorithm never failed to detect the SMBHB signals. On occasions the chain would lock onto a secondary maxima of the likelihood function, but this behavior can be heavily suppressed by using longer cooling schedules. Once the annealing phase is complete the t_c maximization is turned off and our search algorithm becomes a standard Markov Chain Monte Carlo (MCMC) algorithm for exploring the posterior distribution function. The MCMC method can be used to perform model comparisons, estimate instrument noise, and provide error estimates for the recovered parameters [8]. The method is now in widespread use in many fields, and is starting to be used by astronomers and cosmologists. MCMC techniques have been applied to ground based gravitational wave data analysis [18]; a toy LISA problem [19]; and the extraction of multiple overlapping galactic binaries from simulated LISA data [17].

For the example at $z = 1$ we use the following uniform priors in our search: we choose the mass ratio to lie between 5 and 15, the redshifted total mass between 5×10^5

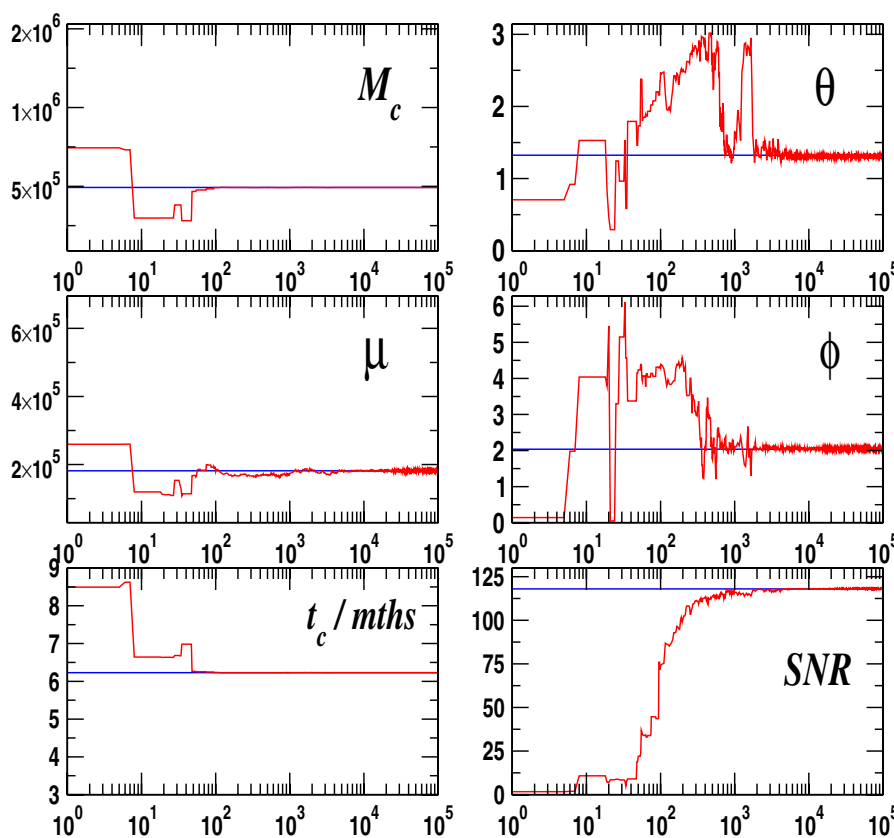


FIG. 2 (color online). A plot of the search chains for the five intrinsic parameters and the SNR. In all cases, the straight solid line represents the true values of the SMBHB parameters. After $N = N_c = 10000$ steps the search becomes a standard MCMC exploration of the posterior distribution function.

NEIL J. CORNISH AND EDWARD K. PORTER

and 5×10^6 solar masses, t_c is chosen to lie within 3 and 9 months, and θ and ϕ are drawn from a uniform sky distribution. The initial heat was set at 100 ($B = -2$) and the annealing lasted for $N_c = 10\,000$ steps. The search took three hours to run on a single 2 GHz processor. In Fig. 2 we plot a representative search chain. Because the search algorithm locks onto the source in $N \sim 1000$ steps, we use a logarithmic scale for the number of iterations, N . The t_c maximization allows the search to hone in on M_c and t_c very quickly. The reduced-mass μ is less well constrained and takes a little longer to lock in, and the sky location gets fixed last of all. The extrinsic parameters D_L , ι are recovered once the sky location is determined, while ψ and φ_c continue to explore their full range throughout the evolution. The failure to fix ψ and φ_c is consistent with the Fisher matrix predictions for the uncertainties in these parameters. The errors in the recovered search parameters were: $\Delta M_c = -42M_\odot = -0.154\sigma$; $\Delta \mu = 729M_\odot = 0.147\sigma$; $\Delta t_c = 76s = -0.171\sigma$; $\Delta \theta = 0.82^\circ = 1.06\sigma$; and $\Delta \phi = 0.65^\circ = -1.28\sigma$. The mean values for the recovered parameters and their standard deviations were determined from the post-annealing, MCMC portion of the chains.

We found that our gridless search algorithm is able to reliably identify the SMBHB signal within ~ 1000 steps. We have calculated that it would take 9.3×10^{12} templates to cover the same search range with an F -statistic based grid search at a minimal match level of 0.9 [20]. The comparison is not entirely fair since we also used an illegal maximization over t_c during the annealing phase, but we have verified that the annealed chains are able to find the SMBHB signal without this trick, it just takes 10 to 100 times longer. Either way, our search algorithm is significantly more economical than a naive grid based search.

For the example at $z = 5.5$ we use the following uniform priors in our search: we choose the mass ratio to lie between 1 and 5, the redshifted total mass between 2×10^5 and 2×10^6 solar masses, t_c is chosen to lie within 5 and 7 months and θ and ϕ are drawn from a uniform sky distribution. The initial heat was set at 10 ($B = -1$) and the annealing phase lasted for $N_c = 20\,000$ steps. In Fig. 3 we plot the SNR evolution for three runs. Two of these runs happened to lock onto an alternative solution for the sky location that exist because of the approximate symmetry $\phi \rightarrow \phi + \pi$ and $\theta \rightarrow \pi - \theta$ that holds for the low frequency LISA response function. Since the two solutions for the sky position have almost equal likelihood, the bimodality of the solution is a feature, rather than a flaw,




FIG. 3 (color online). Examples of the SNR evolution for three runs searching for the low SNR signal at redshift $z = 5.5$. The chains typically found, then lost the signal on several occasions early in the runs due to the high temperature and low SNR. The chains usually locked on for good at around iteration $N \sim 12\,000$.

of the search algorithm. As might be expected, the search algorithm takes longer to lock onto weak sources than strong sources, but the run times are still measured in hours, not days.

Here we have shown that it is possible to dig a SMBHB signal out from under instrument noise and the signals from foreground sources. The errors in the recovered parameters are consistent with a Fisher matrix prediction that treats the galactic foreground as an addition source of Gaussian noise. We will present a detailed study of detection threshold and the posterior distributions in the presence of galactic foregrounds in a future publication. The next step is to apply the same techniques to the more complicated signals from spinning SMBHB's and EMRIs. The larger parameter spaces are not expected to pose a problem as the search cost is expected to scale linearly with the search dimension. Indeed, it should be possible to simultaneously search for multiple, overlapping EMRI signals. We consider our current work as a proof-of-principle that the LISA data analysis challenge can be addressed with modest computational resources.

[1] P. Bender *et al.*, LISA prephase A report 1998 (unpublished).

[2] F. D. Ryan, Phys. Rev. D **56**, 1845 (1997).

[3] N. A. Collins and S. A. Hughes, Phys. Rev. D **69**, 124022

CATCHING SUPERMASSIVE BLACK HOLE BINARIES ...

PHYSICAL REVIEW D **75**, 021301(R) (2007)

- (2004).
- [4] E. Berti, V. Cardoso, and C. M. Will, *Phys. Rev. D* **73**, 064030 (2006).
- [5] J. R. Gair, L. Barack, T. Creighton, C. Cutler, S. L. Larson, E. S. Phinney, and M. Vallisneri, *Class. Quant. Grav.* **21**, S1595 (2004).
- [6] L. Barac and C. Cutler, *Phys. Rev. D* **70**, 122002 (2004).
- [7] L. Wen and J. R. Gair, *Class. Quant. Grav.* **22**, S445 (2005); **22**, S1359 (2005).
- [8] *Markov Chain Monte Carlo in Practice*, edited by W. R. Gilks, S. Richardson, and D. J. Spiegelhalter (Chapman & Hall, London, 1996).
- [9] E. D. L. Wickham, A. Stroeer, and A. Vecchio, *Class. Quant. Grav.* **23**, S819 (2006).
- [10] N. J. Cornish and E. K. Porter, *Class. Quant. Grav.* **23**, S761 (2006).
- [11] C. Röver, R. Meyer, and N. Christensen, *Class. Quant. Grav.* **23**, 4895 (2006).
- [12] P. Jaranowski, A. Królak, and B. F. Schutz, *Phys. Rev. D* **58**, 063001 (1998).
- [13] G. Nelemans, L. R. Yungelson, and S. F. Portegies Zwart, *Mon. Not. R. Astron. Soc.* **349**, 181 (2004).
- [14] S. Timpano, L. Rubbo, and N. J. Cornish, *Phys. Rev. D* **73**, 122001 (2006).
- [15] C. Cutler, *Phys. Rev. D* **57**, 7089 (1998).
- [16] T. A. Prince, M. Tinto, S. L. Larson, and J. W. Armstrong, *Phys. Rev. D* **66**, 122002 (2002).
- [17] N. J. Cornish and J. Crowder, *Phys. Rev. D* **72**, 043005 (2005).
- [18] N. Christensen and R. Meyer, *Phys. Rev. D* **64**, 022001 (2001); N. Christensen, A. Lisbon, and R. Meyer, *Class. Quant. Grav.* **21**, 317 (2004); N. Christensen, R. J. Dupuis, G. Woan, and R. Meyer, *Phys. Rev. D* **70**, 022001 (2004); R. Umstätter, R. Meyer, R. J. Dupuis, J. Veitch, G. Woan, and N. Christensen, *AIP Conf. Proc.* **735**, 336 (2004).
- [19] R. Umstätter, N. Christensen, M. Hendry, R. Meyer, V. Simha, J. Veitch, S. Vigeland, and G. Woan, *Phys. Rev. D* **72**, 022001 (2005).
- [20] B. J. Owen, *Phys. Rev. D* **53**, 6749 (1996); B. J. Owen and B. S. Sathyaprakash, *Phys. Rev. D* **60**, 022002 (1999).