

# A three-stage search for supermassive black-hole binaries in LISA data

Duncan A Brown<sup>1,2</sup>, Jeff Crowder<sup>3</sup>, Curt Cutler<sup>2,3</sup>, Ilya Mandel<sup>2</sup> and Michele Vallisneri<sup>2,3</sup>

<sup>1</sup> LIGO Laboratory, California Institute of Technology, Pasadena, CA 91125, USA

<sup>2</sup> Theoretical Astrophysics, California Institute of Technology, Pasadena, CA 91125, USA

<sup>3</sup> Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA

Received 19 April 2007, in final form 22 April 2007

Published 19 September 2007

Online at [stacks.iop.org/CQG/24/S595](http://stacks.iop.org/CQG/24/S595)

## Abstract

Gravitational waves from the inspiral and coalescence of supermassive black-hole (SMBH) binaries with masses  $m_1 \sim m_2 \sim 10^6 M_\odot$  are likely to be among the strongest sources for the Laser Interferometer Space Antenna (LISA). We describe a three-stage data-analysis pipeline designed to search for and measure the parameters of SMBH binaries in LISA data. The first stage uses a time–frequency track-search method to search for inspiral signals and provide a coarse estimate of the black-hole masses  $m_1, m_2$  and the coalescence time of the binary  $t_c$ . The second stage uses a sequence of matched-filter template banks, seeded by the first stage, to improve the measurement accuracy of the masses and coalescence time. Finally, a Markov chain Monte Carlo search is used to estimate all nine physical parameters of the binary (masses, coalescence time, distance, initial phase, sky position and orientation). Using results from the second stage substantially shortens the Markov chain burn-in time and allows us to determine the number of SMBH-binary signals in the data before starting parameter estimation. We demonstrate our analysis pipeline using simulated data from the first Mock LISA Data Challenge. We discuss our plan for improving this pipeline and the challenges that will be faced in real LISA data analysis.

PACS numbers: 04.25.Nx, 04.30.Nb, 04.80.Nn, 95.75.Wx, 95.85.Sz

(Some figures in this article are in colour only in the electronic version)

## 1. Introduction

There is compelling evidence from electromagnetic observations that the cores of galaxies contain supermassive black holes (SMBHs) [1]. SMBH binaries can form after galactic mergers as the black holes from the individual galaxies fall to the center of the merged system

and form a bound pair. Hierarchical-merger models of galaxy formation predict that SMBH binaries will be common in galaxies [2, 3] and the presence of one such binary has been inferred from x-ray measurements of the core of the galaxy NGC 6240 [4]. The evolution of an SMBH binary will eventually be driven by radiation reaction from the emission of gravitational waves (GWs) and the binary will inspiral and merge to form a single SMBH. The GWs from inspirals of SMBH binaries with component masses in the range  $m \sim 10^4 - 10^7 M_\odot$  will be among the strongest sources for LISA, the planned space-based GW detector [5, 6]. The direct detection of SMBH binaries will be of wide astrophysical relevance, for example by probing the merger rates and histories of galaxies [7] or by providing cosmological standard candles [8].

Searching for SMBH binary inspiral signals is expected to be one of the more straightforward tasks in LISA data analysis. The velocities of the black holes during the inspiral are  $v/c \ll 1$  until the final  $\sim 10^1 - 10^2$  cycles, and so existing post-Newtonian waveforms [9, 10] will describe the inspiral gravitational waveforms with sufficient accuracy for use as templates in a matched-filter search [11]. As such, searches for SMBH binaries in LISA data will be similar in nature to existing searches for binary-neutron-star (BNS) inspirals in ground-based GW detectors, such as the Laser Interferometer Gravitational-Wave Observatory (LIGO) [12]. However, there are several key differences between LIGO and LISA binary inspiral searches. First, the LIGO pipelines are designed to search for signals with expected signal-to-noise ratios (SNRs)  $\lesssim 10$ , whereas the SNR of LISA SMBH binaries at distances  $z \lesssim 2$  is expected to be several hundred or more. Second, the BNS signals sweep through the sensitive frequency band of ground-based detectors on timescales of order a minute, during which detector velocities and orientations can be considered as fixed to high accuracy. By contrast, LISA will be able to observe a single SMBH inspiral for weeks to months. During that time, the LISA velocity and orientation change appreciably, inducing modulations in the recorded signal. Indeed, almost all the information about an SMBH binary's sky location and orientation is encoded in these modulations. (In the ground-based case, a network of three or more widely separated detectors is required to determine a binary's sky location by triangulation between the times of arrival of the GW signals at the different detector locations.) Finally, whereas the rate of BNS inspirals in ground-based detectors makes it unlikely that multiple signals will be observed concurrently, LISA data may contain simultaneous signals from a few different SMBH binaries.

Existing search pipelines developed for ground-based observations of stellar-mass binary inspirals can achieve high detection efficiency already at SNRs  $\sim 10$  [13–16], so the task of detecting SMBH inspirals with LISA seems easy in comparison. Furthermore, since SMBH binaries at  $z \sim 1$  have such high SNR, and because of LISA's relatively wider frequency band (roughly three orders of magnitude for LISA, compared to two for LIGO), it should also be possible to determine the masses and spins of the binaries with significantly higher accuracy in the LISA case than for ground-based detections. Fisher-matrix calculations suggest that, for SMBHs detected at  $z \sim 1$ , LISA should be able to determine the chirp mass to relative accuracy  $\sim 10^{-5}$ , both individual masses to  $\sim 10^{-3}$  and the SMBH spins to  $\sim 10^{-3} - 10^{-2}$  [17]. Indeed, the goal of our data-analysis pipeline is not only to detect the SMBH signals, but also to provide accurate measurements of the binary parameters.

Based partly on the considerations discussed above, our group has adopted the following three-stage search method. Low- $z$  SMBH binary inspirals are so bright that they are easily visible as tracks in time–frequency (TF) spectrograms. Therefore our first stage consists of a search for such TF tracks; the shape and location of the track yields a first estimate of the two masses,  $m_1$  and  $m_2$ , and the coalescence time,  $t_c$ . The second stage is a set of more refined grid-based matched-filter searches that start in a neighborhood of the best-fit parameters found in the first stage; these searches home in on more accurate values for the three parameters

$m_1$ ,  $m_2$  and  $t_c$ . The final stage is currently a straightforward implementation of a Markov chain Monte Carlo (MCMC) simulated-annealing search for the best-fit parameters in the full nine-dimensional parameter space (including also the binary's luminosity distance, initial phase, inclination, polarization, ecliptic latitude and longitude).

There are a few reasons for adopting such a complicated algorithm. First, we believe that the capability of looking for TF tracks is a very useful one to develop in the LISA context: it could possibly reveal tracks that do not correspond closely to signals from merging binaries, and so might not be found by more sophisticated (grid-based or MCMC) methods, even though they are visible to the eye in a TF plot. The track-search method also allows us to count the number of SMBH binary signals present in the data before attempting parameter estimation. Second, the grid search is useful to make sure that we do not miss any binary sources, by examining the entire parameter space. In the pipeline described here, however, we did not cover the entire parameter space in our grid search; rather, we seeded the second-stage search using the parameters obtained from the first stage. In future implementations, we intend to compare the full grid search to this method. Finally, the MCMC approach is clearly very adept at obtaining the final parameter estimates.

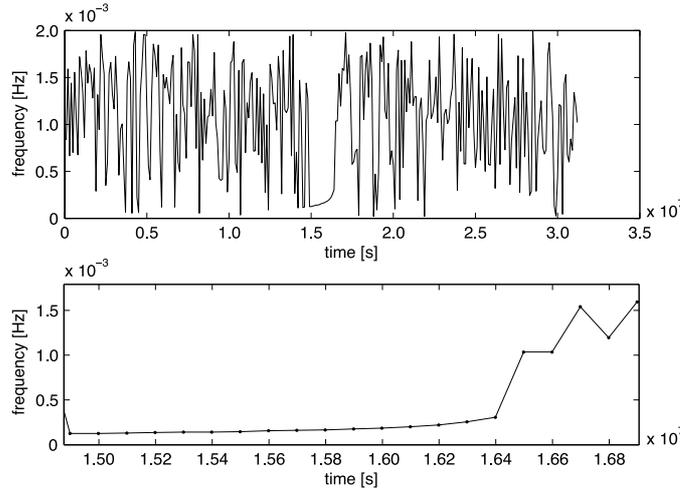
We have tested the performance of our SMBH binary search pipeline using data from the Mock LISA Data Challenges (MLDCs) [18, 19]. The MLDCs are a program sponsored by the LISA International Science Team to foster the development of LISA data-analysis methods and tools, and to demonstrate already acquired milestones in the extraction of science information from the LISA data output. In the MLDCs, GW signals whose parameter values are unknown to the challenge participants are embedded in synthetic LISA noise; participants are challenged to identify the signals and extract their parameters. Challenges of increasing difficulty are being issued roughly every 6 months. The results from the first challenge are summarized by Arnaud and colleagues in this volume [20]. Challenge 1 included two datasets with signals from isolated SMBH systems with differing SNRs (SNR < 500 for challenge 1.2.1 and SNR < 100 for challenge 1.2.2); we analyzed challenge 1.2.1. One of the goals of the MLDCs is to demonstrate that data-analysis pipelines can actually achieve the remarkable parameter measurement accuracy predicted by the Fisher-matrix analysis.

For simplicity, the challenge 1 problems omitted two important complications of space-based binary-inspiral searches that are not present in ground-based searches. First, SMBH binaries may enter the LISA band with considerable eccentricity, but they were modeled as circular in challenge 1 (by contrast, the BNSs observed by ground-based detectors will have become essentially circular by the time they enter the observation band). Second, through much of LISA's sensitivity band the dominant noise comes from unresolved Galactic white-dwarf binaries, which were not included in challenge 1 (by contrast, in the ground-based case the relevant noise originates almost entirely from the instrument). Our initial search pipeline took advantage of these simplifications.

The rest of this paper is organized as follows. In sections 2–4, we describe the three stages in our SMBH binary data-analysis pipeline: a track search in the time–frequency plane, a grid-based matched filtering search and a Markov chain Monte Carlo search; in section 5, we present the results of analyzing the MLDC dataset 1.2.1; and in section 6, we discuss our plans for improving the pipeline to cope with issues such as binary eccentricity and the noise sources likely to be observed in real LISA data.

## 2. Stage 1: search for tracks in the time–frequency plane

The TF spectrogram contains enough information to identify an SMBH binary inspiral at a high SNR. The techniques described below make it possible to quickly search for the presence



**Figure 1.** Time–frequency plot of the *brightest pixel* in each time bin, as computed for the *X* channel of challenge 1 training set 1.2.1. In the top plot, one can see the chirp beginning at roughly  $1.5 \times 10^7$  s. The bottom plot is a blown-up version of the top plot, for the  $\sim 2 \times 10^6$  s preceding the coalescence time, showing the presumed track found on the first pass through the data.

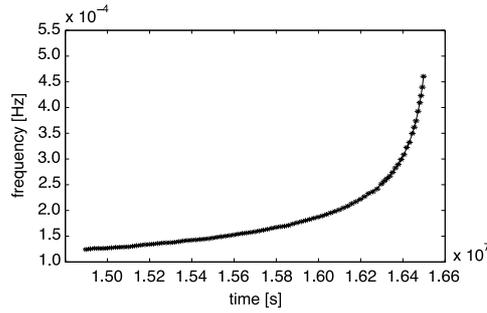
of an SMBH binary inspiral in the signal and to get rough estimates for the coalescence time and the two masses.

Challenge 1 includes signals from the adiabatic inspiral of a circular binary system of nonspinning SMBHs. The frequency evolution of these inspirals is given by (7.11a) of [21] in terms of the time of coalescence  $t_c$  and the two SMBH masses  $m_1$  and  $m_2$ . We write it here as a function of the symmetric mass ratio  $\eta = m_1 m_2 / (m_1 + m_2)^2$  and the chirp mass  $M_c = (m_1 + m_2) \eta^{3/5}$ , using the second-order post-Newtonian (2PN) approximation,

$$f_{\text{GW}}(t) = \frac{\eta^{3/5}}{8\pi M_c} (T_c - T)^{-3/8} \left\{ 1 + \left[ \frac{743}{2688} + \frac{11}{32} \eta \right] (T_c - T)^{-1/4} - \frac{3\pi}{10} (T_c - T)^{-3/8} \right. \\ \left. + \left[ \frac{1855099}{14450688} + \frac{56975}{258048} \eta + \frac{371}{2048} \eta^2 \right] (T_c - T)^{-1/2} + \mathcal{O}[(T_c - T)^{-5/8}] \right\}. \quad (1)$$

Here  $f_{\text{GW}}$  is the GW frequency in Hertz,  $M_c$  is expressed in seconds and  $T$  is the dimensionless time variable related to coordinate time  $t$  by  $T = t \eta^{8/5} / (5M_c)$ .

We create a TF map of the noisy data stream  $s(t) = h(t) + n(t)$  (in fact, one of the Time-delay interferometry (TDI) channels  $X(t)$ ,  $Y(t)$  and  $Z(t)$  provided in the MLDC datasets), sampled with timestep  $\delta t$ , in two passes. On the first pass, we split up the data stream into time bins of equal duration  $\Delta t$ . The TF spectrogram will then consist of pixels of size  $\Delta t \times \Delta f$ , where  $\Delta f = 1/(\Delta t)$ . We determine the normalized power contained in each pixel with a fast Fourier transform (FFT), normalizing by the power spectral density of the noise, and then find the peak frequency in each bin by searching for the loudest pixel (see below for details). The resulting set of {time, frequency} pairs allows us to search for an inspiral track on the TF map (see figure 1). Once such a track is identified, we make a second pass through the data, iterating through the track region with time bins of varying duration to create an improved TF map. Earlier in the track, a larger  $\Delta t$  helps to detect a weak signal and achieve greater



**Figure 2.** The stars represent individual points on the TF map obtained during the second pass through the data in the  $X(t)$  channel of challenge 1 training set 1.2.1. The curve is the result of fitting these points to the model (1).

frequency resolution; closer to coalescence, a smaller  $\Delta t$  reduces the error in estimating the rapidly chirping GW frequency.

In fact, we have made several improvements to the general approach outlined in the previous paragraph. The first set of improvements concerns the determination of the peak frequency in a given time bin. Simply searching for the loudest pixel would give frequency-determination errors of order  $1/(\Delta f)$ , even for a noiseless signal. Instead, we achieve higher accuracy by modeling the bleeding of frequency into neighboring pixels: specifically, we determine the peak frequency by fitting the logarithm of power in the pixels nearest to the brightest pixel to a parabola, using zero padding in the time domain to achieve better frequency resolution when necessary. We also apply a Hanning window to the signal prior to taking the FFT, and we overlap time bins to avoid information loss from windowing.

Another improvement concerns the variable timestep and the identification of outliers on the second pass through the data. If the peak frequencies of neighboring time bins differ by more than  $2\Delta f$ , we decrease  $\Delta t$  by a pre-set factor (say 1.5) to reduce the sweep of frequency in each bin. If this operation fails to bring the peak frequencies closer together, we declare the data point an outlier, and skip to the next bin.

The {time, frequency} data points obtained on the second pass serve as inputs to a MATLAB least-squares fitting algorithm that extracts the inspiral parameters  $t_c$ ,  $M_c$  and  $\eta$  by fitting the data points to the model of (1) (see figure 2). Specifically, we find the values  $t_c$ ,  $M_c$  and  $\eta$  that minimize the sum

$$\Sigma = \sum_{i=1}^N [f(t_i) - f_{\text{GW}}(t_i; t_c, M_c, \eta)]^2, \quad (2)$$

where  $t_i$  are the centers of the output time bins,  $f(t_i)$  are the associated frequencies and  $f_{\text{GW}}(t_i; t_c, M_c, \eta)$  is the model from (1).

Although one could weight the data points on the basis of the signal amplitude, such a weighting seems to carry little benefit: late in the inspiral, the increased amplitude offers greater SNR, which is however substantially offset by poorer frequency determination (due either to frequency drift within each time bin if  $\Delta t$  is not properly adjusted, or to low-frequency resolution if it is).

Table 1 shows the results of the TF search on the blind challenge dataset 1.2.1. After averaging results from the three TDI streams, we found  $M_c = 1.208 \times 10^6 M_\odot$ ,  $\eta = 0.17$  and  $t_c = 1.3372 \times 10^7$  s. The accuracy of these estimates is discussed in section 5; suffice it to

**Table 1.** Parameters extracted via TF searches from the  $X$ ,  $Y$  and  $Z$  channels of blind challenge dataset 1.2.1.  $N$  is the number of data points obtained during the second pass through the data and  $\Sigma$  is the sum of the squares of the residuals, as defined in (2).

	$N$	$\Sigma/10^{-11}$	$M_c/(10^6 M_\odot)$	$\eta$	$t_c/(10^7 \text{ s})$
$X$	156	9.2	1.2096	0.182	1.3373
$Y$	190	9.21	1.2033	0.139	1.3370
$Z$	192	11.5	1.2099	0.183	1.3373

say that these first-stage results were certainly accurate enough for our purpose. Finally, while challenge 1 tested our method for the case of SMBH signals in synthetic instrumental noise only, we have verified that our TF search also works well in the presence of confusion noise from Galactic binaries.

### 3. Stage 2: grid-based search

The grid-based part of the search relies on the template placement algorithm of Babak *et al* [13] and the `FINDCHIRP` matched filtering algorithm of Allen *et al* [14], both of which were developed for the LIGO binary neutron star searches. The basic algorithm is as follows: a grid of templates is constructed in the  $(m_1, m_2)$  plane using the metric-based square-grid placement algorithm [13, 22] implemented in the LIGO Algorithm Library (LAL) [23]<sup>4</sup>. The fineness of the grid is specified by its minimum-match parameter  $MM$ , which is the minimum overlap between any point in the parameter space and its nearest grid point. We have written C code that implements the matched filtering and template generation algorithms described in [14]. These C functions are then ‘wrapped’ by the Simplified Wrapper Interface Generator (SWIG), which allows them to be called from the Python high-level programming language. This approach allowed us to rapidly prototype and develop the procedure described below.

For each mass pair in the grid, we compute a (Fourier-transformed) waveform  $\tilde{h}(f)$  (corresponding to coalescence at  $t = 0$ ), using 2PN waveforms and the stationary phase approximation (SPA) [24]. Note that this  $\tilde{h}(f)$  effectively assumes that LISA’s arms have fixed orientation; i.e., it neglects the modulation due to LISA’s rotational and translational motion. (This is reasonable as a first pass, since most of the SNR accumulates in only a few days.) Using the long-wavelength approximation, we transform from  $\tilde{h}(f)$  to the LISA TDI variable  $\tilde{X}_h(f)$  using

$$\tilde{X}_h(f) = \sin^2(2\pi f L) \tilde{h}(f), \quad (3)$$

where  $L$  is the LISA arm length. Let the (Fourier-transformed) data be  $\tilde{X}_s(f)$ : then for each template waveform  $\tilde{X}_h(f)$  in our grid we use the FFT to compute the inverse Fourier transform

$$z(t) = \int \frac{\tilde{X}_s(f) \tilde{X}_h^*(f)}{S_X(f)} e^{2\pi i f t} df, \quad (4)$$

and we maximize  $|z(t)|$  over  $t$  to estimate the time of coalescence. We identify the best-fit point in the  $(m_1, m_2)$  plane, and then repeat the search in a neighborhood of that point with a finer grid. (More specifically, we initially generate a grid on a rectangular section of  $(m_1, m_2)$  space, but then ‘throw away’ points that are outside a narrow range in  $M_c$ .) We do this four times with a final minimum-match parameter  $MM = 0.995$ . For challenge 1.2.1, on the basis of results from the TF stage ( $m_1 \approx 2.9 \times 10^6 M_\odot$  and  $m_2 \approx 7.3 \times 10^5 M_\odot$ ), we chose our initial

<sup>4</sup> Babak *et al* also describe a more efficient hexagonal placement algorithm, but we were unable to place templates for LISA SMBH binaries using its LAL implementation. We intend to work with the authors of the LAL code to resolve this.

**Table 2.** True values and estimates from three steps for the challenge parameters. In stages 1 and 2 estimates were made only for parameters  $M_c$  and  $\eta$  (and therefore  $m_1$  and  $m_2$ ) and  $t_c$ .

Parameter	Unit	True value	Stage 1	Stage 2	Stage 3
$M_c$	$10^6 M_\odot$	1.2086	1.208	1.2108	1.2077
$\eta$		0.160	0.17	0.163	0.156
$m_1$	$10^6 M_\odot$	2.8972	2.74	2.8536	2.9652
$m_2$	$10^6 M_\odot$	0.7270	0.76	0.7381	0.7130
$t_c$	$10^7$ s	1.337 4027	1.3372	1.337 4149	1.337 4072
Ecl. lat. $\theta$	rad	-0.492	-	-	0.536
Ecl. long. $\phi$	rad	0.866	-	-	4.039
Pol. angle $\psi$	rad	3.234	-	-	5.886
Init. phase $\varphi_0$	rad	3.527	-	-	0.233
Distance $D$	$10^9$ pc	8.000	-	-	16.811
Incl. angle $\iota$	rad	1.944	-	-	0.617

grid to cover the portion of the  $(m_1, m_2)$  plane satisfying  $6 \times 10^5 < m_2 < m_1 < 3.2 \times 10^6 M_\odot$ , with initial MM = 0.30.

Now, our parameter-estimation errors are dominated not by the coarseness of the grid, but by the fact that our 2PN SPA waveforms are not identical to BBH waveforms injected into the Mock LISA data, even for the same parameter values. Besides neglecting the modulations due to LISA's motion, our 2PN SPA waveforms differ from the MLDC versions by higher-order PN terms. They are also simply cut off at the frequency of the innermost stable circular orbit (ISCO) of a test mass in the Schwarzschild spacetime, whereas the MLDC waveforms end with a very particular choice of taper. Therefore we do one final grid search using MLDC waveforms (again with MM = 0.995), for a sky position  $(\theta, \phi)$  at the North pole in ecliptic coordinates, and for some arbitrary choice of the three angles  $(\iota, \psi, \varphi_0)$ . Although all five of these angles are wrong, in this step the *other* features of the templates (e.g., the 2PN frequency evolution and the amplitude taper) do match those of the injected MLDC binary waveforms, and so presumably yield improved parameter estimates. Our best parameters at the end of this stage are shown in table 2.

#### 4. Stage 3: Markov chain Monte Carlo

So far, the first two stages have given estimates only of the two masses and coalescence time; in addition, the stage-2 analysis was based only on the  $X$  channel. Thus, we rely on the MCMC stage to find the distance, sky location, and the polarization and inclination angles of our source. A more efficient way to do this would be to use the  $\mathcal{F}$ -statistic [25, 26] to automatically optimize over four amplitude parameters that are functions of distance, polarization, inclination and initial phase; however, we did not have time to implement this procedure for challenge 1. Therefore our MCMC code does a brute force search over all parameters—but with the advantage that it starts in the right vicinity for the masses and coalescence time, as estimated in the first two stages.

MCMC approaches have shown promise in the extraction of GW-source parameters with LISA [27–32]. Nevertheless, it has been suggested that, for SMBH binaries, MCMC searches over a full parameter set need to be started in a neighborhood of the correct source parameters to efficiently characterize the posterior probability density functions [28]. Since the initial search grid provided a good estimate of three parameters (the constituent masses and coalescence time  $t_c$ ), and since it is trivial to extremize analytically over the luminosity distance, we were hopeful that we could determine the values of the sky location and binary orientation with

a straightforward implementation of the Metropolis–Hastings algorithm (MHA). Since time was limited and posterior distributions were not required for challenge 1, we chose not to estimate these, but rather to use the MHA to locate the best-fit parameters.

In the MHA, a Markov chain is built by accepting a new proposed point with the probability  $\alpha = \min(1, H)$ ;  $H$  is the *Hastings ratio* for a jump from position  $\vec{x}$  to  $\vec{y}$  in parameter space, and is given by

$$H = \frac{p(\vec{y})p(s|\vec{y})q(\vec{x}|\vec{y})}{p(\vec{x})p(s|\vec{x})q(\vec{y}|\vec{x})}, \quad (5)$$

where  $p(\vec{x})$  is the prior distribution,  $p(s|\vec{y})$  is the likelihood of the parameter set  $\vec{y}$  producing the signal  $s$  and  $q(\vec{x}|\vec{y})$  is the proposal distribution used to generate the move from  $\vec{x}$  to  $\vec{y}$ . If the noise is a normal process with zero mean, the likelihood is given by

$$p(s|\vec{\lambda}) \propto \exp\{-(s - h(\vec{\lambda})|s - h(\vec{\lambda}))^2/2\}, \quad (6)$$

with ‘ $(\cdot|\cdot)$ ’ the standard inner product computed with respect to the LISA instrument noise. In this work the priors for the masses and time of coalescence for the binary system were given by the MLDC Taskforce. Uniform priors were chosen for the angular parameters.

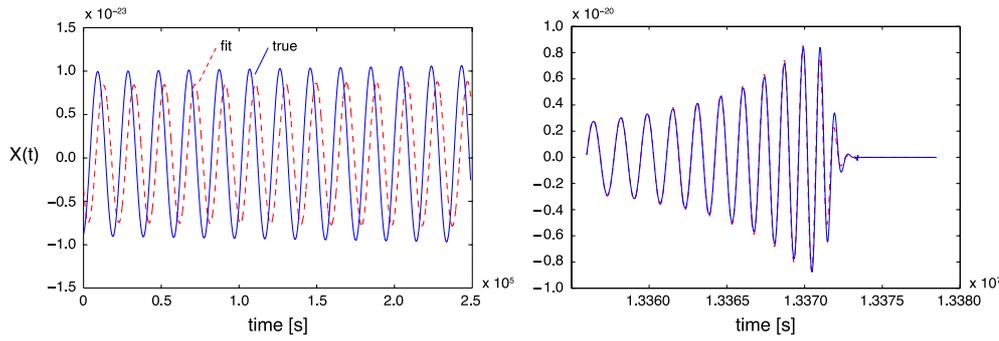
The Markov chain process is guaranteed to converge to the posterior probability distribution if the proposal distribution is nontrivial; however, the speed of convergence does depend on its choice. In this search we adopted two types of proposals: the first consisted of a multivariate normal distribution with jumps directed along the eigendirections of the Fisher information matrix, computed locally; the second amounted to drawing parameters from uniform distributions. For the angular parameters, both timid and bold draws (from small or large ranges) were made to ensure we were fully exploring parameter space; for the component masses, only timid draws (<1%) were used. The initial heating factor [27] for the runs was chosen empirically (through pilot runs) so that there was an initial acceptance rate of ~50% in the chains, to enhance the exploration of the parameter space. The cooling schedule was given by

$$\beta = \begin{cases} \beta_0 \left(\frac{1}{\beta_0}\right)^{i/N_a} & 0 \leq i \leq N_a \\ 1 & i > N_a \end{cases}, \quad (7)$$

where  $\beta_0$  is the initial heating factor,  $N_a$  is the length of the annealing phase and  $i$  is the step counter.

Multiple concurrent chains were started using the parameter estimates obtained in stage 2. These were run on a supercomputing cluster with 3.2 GHz Intel Pentium 4 processors, using Synthetic LISA [33] to reproduce the LISA response to the SMBH binary waveforms. Each run was limited to 12 h, providing ~3500 steps in each of the chains. The most promising candidates at the end of the first run were used as the starting locations of a second run. At the end of the first run the best candidates had reached log likelihood values in the neighborhood of 200 000; the second run saw them increase to ~205 000. The chains converged around two points in parameter space, differing by their locations on opposite sides of the sky. This was not unexpected: dual maxima at antipodal sky positions are a well-known near degeneracy for LISA sources (arising from the near degeneracy of the beam pattern functions). Our choice between the two final parameter sets was based on a visual comparison of the putative signals with the challenge dataset.

In future implementations of the pipeline, we plan to incorporate the  $\mathcal{F}$ -statistic in the MCMC stage to reduce the size of parameter space. This will increase search efficiency and relax the need to begin the search in a neighborhood of the best-fit parameters (something that will be necessary when searching for the dimmer SMBH binaries of challenge 2). Another



**Figure 3.** Comparison of our best-fit  $X(t)$  to the true  $X(t)$  for (a) a short stretch of time near  $t_c$  and (b) a short stretch near the beginning of the dataset. Clearly, our fit is excellent near  $t_c$ , where most of the SNR accumulates, but much poorer at early times.

time-saving measure will be to start the search on a limited portion of the data stream, and then steadily increase its size. This process, called frequency annealing [34], allows a quick initial exploration of parameter space, and a careful later investigation of the exquisitely sharp likelihood peaks close to bright SMBH binaries.

## 5. Results for MLDC challenge

As was the case for many challenge-1 participants, the 3 December 2006 submission deadline arrived before our pipeline was fully ready; nevertheless we decided to submit our best estimates for the parameters of the blind dataset 1.2.1. This dataset consisted of the three TDI unequal-Michelson channels  $X(t)$ ,  $Y(t)$  and  $Z(t)$ . In stage 1 of our search, we analyzed each of these channels separately, and simply averaged the three results to arrive at the stage-one parameter estimates shown in the fourth column of table 2. In stage 2, only the  $X(t)$  data were analyzed (partly because of time pressure). In stage 3, we analyzed two TDI channels given by  $X$  and  $(X + 2Y)/\sqrt{3}$ , which are noise orthogonal at moderately low frequencies.

The true signal parameters were made publicly available on 4 December, and here we briefly describe how our search fared in their recovery. The injected signal had a combined<sup>5</sup> ( $A + E$ ) SNR of 667.734; its true physical parameters are listed in the third column of table 2. Our best-fit waveform matched the true waveform rather well: it had an SNR of 664.47 and its cross-correlation with the true waveform was 0.994 for the  $A$  channel and 0.996 for the  $E$  channel [35]. The quality of the fit is illustrated in figure 3, which compares the true  $X(t)$  (produced by us from the key file) with our best-fit  $X(t)$ , for short time stretches near the coalescence time  $t_c$  and near the beginning of the dataset. Clearly our fit is excellent near  $t_c$ , where most of the SNR accumulates, but is much poorer at early times, when the contribution to the SNR is much lower. The lesson from the other two Michelson variables is qualitatively the same.

Our best-fit parameters are listed in the last column of table 2: our inferred chirp mass  $M_c$  was correct to within  $\Delta M_c/M_c < 10^{-3}$ , our inferred symmetric mass ratio  $\eta$  to within  $\Delta \eta \approx 4 \times 10^{-3}$ , and the error in our coalescence time was  $\Delta t_c \approx 45$  s, corresponding to approximately 0.05 GW periods just before the plunge. Nevertheless, it is clear from our estimates for the *other* parameters that, instead of converging on a neighborhood of the true

<sup>5</sup> In this context,  $A$  and  $E$  are the orthogonal, optimal TDI observables given by  $(2X - Y - Z)/3$  and  $(Z - Y)/\sqrt{3}$ , as used in [20]. The third orthogonal, optimal TDI observable,  $T$ , contributes only a tiny fraction of the total SNR for these sources.

maximum, our MCMC code locked into a high but secondary maximum of the posterior probability distribution. Our inferred sky position is almost at the antipodes of the actual location (i.e., our ecliptic latitude is approximately the negative of the true value, and our ecliptic longitude is off by nearly  $\pi$ ). This was not due to a mismatch of conventions or a bug in our code; rather, it reflects the above-mentioned degeneracy between antipodal sky locations (the degeneracy becomes perfect in the low-frequency limit). The four parameters ( $D, t, \psi, \varphi_0$ ) that determine the overall complex amplitudes of the GW polarizations  $h_+$  and  $h_\times$  were also off by factors of order one, except for our overall phase  $\varphi_0$ , which was correct to within 0.004 radians (modulo  $\pi$ ).

It is also instructive (and reassuring) to contemplate the performance of the first two stages of our search. Stage 1 returned  $M_c$  with a fractional error  $\Delta M_c/M_c < 10^{-3}$ ,  $\eta$  to within  $\sim 6\%$ , and  $t_c$  to within  $\sim 2 \times 10^3$  s. After stage 2, the estimated  $M_c$  was in fact slightly worse, but the errors in  $\eta$  and  $t_c$  were significantly reduced, to  $\Delta\eta \approx 0.003$  and  $\Delta t_c \approx 120$  s. This gratifying level of accuracy indicates that the coarser stages 1 and 2 were indeed accomplishing the job required of them.

## 6. Future directions

As explained above, the most obvious improvement to our pipeline will be to recast the MCMC stage so that it maximizes the  $\mathcal{F}$ -statistic on the five-dimensional space  $(M_c, \eta, t_c, \theta, \phi)$ , reducing the search-space dimensionality by 3. In addition, we will extend our grid search to handle the case where the merger occurs *after* the end of the dataset (we did not compete on dataset 1.2.2 because our current grid search could not handle such mergers). This generalization should be fairly straightforward.

In the second round of challenges (see the proceeding by Arnaud and colleagues in this volume [35]), dataset 2.2 contains signals from an entire Galaxy's worth of white-dwarf binaries, four to six SMBH binary inspirals (the exact number is not specified) with SNRs ranging from  $\sim 10$  to  $\sim 2000$ , and five EMRIs. Our plan is to first run our pipeline as a standalone search for the SMBH binaries, and then to join forces with Crowder and Cornish's WD binary search [32] to iteratively improve the fits provided by the two searches. Beyond that, we plan to extend the SMBH binaries search to include: (1) merger and ringdown waveforms; (2) spin-precession effects; and (3) the effects of nonzero eccentricity. For the first two items, we intend to make use of the technology already developed by the ground-based GW community. For instance, Buonanno, Chen, and Vallisneri [36] have shown how searches for binaries of spinning BHs can be made considerably more efficient by dividing the parameters into intrinsic (such as the masses) and extrinsic (such as the orientation of the orbital plane at a fiducial time), and optimizing over the extrinsic parameters semi-analytically. (This can be viewed as a generalization to spinning binaries of the  $\mathcal{F}$ -statistic analysis mentioned above.) We shall endeavor to generalize this strategy to LISA searches for SMBH binaries.

## Acknowledgments

JC's, CC's and MV's work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract to the National Aeronautics and Space Administration. JC would like to thank the staff of the supercomputing facilities of the Jet Propulsion Laboratory, California Institute of Technology, where the MCMC stage of the searches described herein was carried out. IM would like to acknowledge the Brinson Foundation, NASA grant NNG04GK98G and NSF grant PHY-0601459 for financial support. MV is grateful for support from the Human Resources Development Fund program at JPL. DB acknowledges

funding from NSF grant PHY-0601459 and the LIGO Laboratory. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY-0107417. This paper carries LIGO document number LIGO-P070018-00-Z.

© US Govt

## References

- [1] Krolik J H 1999 *Active Galactic Nuclei: From the Central Black Hole to the Galactic Environment* (Princeton, NJ: Princeton University Press)
- [2] Haehnelt M G and Kauffmann G 2002 *Mon. Not. R. Astron. Soc.* **336** L61–4 (Preprint [astro-ph/0208215](#))
- [3] Volonteri M, Haardt F and Madau P 2003 *Astrophys. J.* **582** 559–73 (Preprint [astro-ph/0207276](#))
- [4] Komossa S *et al* 2003 *Astrophys. J.* **582** L15–20 (Preprint [astro-ph/0212099](#))
- [5] Bender P L *et al* 1998 *LISA Pre-Phase A Report MPQ 233* (Garching, Germany: Max-Planck-Institut für Quantenoptik)
- [6] Haehnelt M G 1998 *Laser Interferometer Space Antenna, 2nd Int. LISA Symp. on the Detection and Observation of Gravitational Waves in Space (AIP Conf. Proc. vol 456)* ed W M Folkner (Preprint [astro-ph/9809328](#))
- [7] Hughes S A 2002 *Mon. Not. R. Astron. Soc.* **331** 805–16 (Preprint [astro-ph/0108483](#))
- [8] Holz D E and Hughes S A 2003 *Class. Quantum Grav.* **20** S65–72 (Preprint [astro-ph/0212218](#))
- [9] Arun K G, Blanchet L, Iyer B R and Qusailah M S S 2004 *Class. Quantum Grav.* **21** 3771–802 (Preprint [gr-qc/0404085](#))
- [10] Blanchet L, Faye G, Iyer B R and Joguet B 2002 *Phys. Rev. D* **65** 061501 (Preprint [gr-qc/0105099](#))
- [11] Wainstein L A and Zubakov V D 1962 *Extraction of Signals from Noise* (Englewood Cliffs, NJ: Prentice-Hall)
- [12] Barish B C and Weiss R 1999 *Phys. Today* **52** 44–50
- [13] Babak S, Balasubramanian R, Churches D, Cokelaer T and Sathyaprakash B S 2006 *Class. Quantum Grav.* **23** 5477–504 (Preprint [gr-qc/0604037](#))
- [14] Allen B, Anderson W G, Brady P R, Brown D A and Creighton J D E 2005 Preprint [gr-qc/0509116](#)
- [15] Brown D A *et al* 2004 *Class. Quantum Grav.* **21** S1625–33
- [16] Abbott B *et al* (LIGO Scientific Collaboration) 2007 Preprint [arXiv:0704.3368](#)
- [17] Lang R N and Hughes S A 2006 *Phys. Rev. D* **74** 122001 (Preprint [gr-qc/0608062](#))
- [18] Arnaud K A *et al* 2006 *Laser Interferometer Space Antenna, 6th Int. LISA Symp. (AIP Conf. Proc. vol 873)* ed S M Merkowitz and J C Livas pp 619–24 (Preprint [gr-qc/0609105](#))
- [19] Arnaud K A *et al* (Mock LISA Data Challenge Task Force) 2006 *Laser Interferometer Space Antenna, 6th Int. LISA Symp. (AIP Conf. Proc. vol 873)* ed S M Merkowitz and J C Livas pp 625–32 (Preprint [gr-qc/0609106](#))
- [20] Arnaud K A *et al* (Mock LISA Data Challenge Task Force) 2007 *Class. Quantum Grav.* **24** S529 (Preprint [gr-qc/0701139](#))
- [21] Will C M and Wiseman A G 1996 *Phys. Rev. D* **54** 4813–4848 (Preprint [gr-qc/9608012](#))
- [22] Owen B J and Sathyaprakash B S 1999 *Phys. Rev. D* **60** 022002
- [23] LSC Algorithm Library software packages LAL, LALWRAPPER, and LALAPPS; the CVS tag version iulgroup\_s1\_20030420 of LAL, version iulgroup\_s1\_20030416 of LALWRAPPER, and iulgroup\_030526 of LALAPPS were used in this analysis. <http://www.lsc-group.phys.uwm.edu/lal>
- [24] Droz S, Knapp D J, Poisson E and Owen B J 1999 *Phys. Rev. D* **59** 124016
- [25] Jaranowski P, Krolak A and Schutz B F 1998 *Phys. Rev. D* **58** 063001 (Preprint [gr-qc/9804014](#))
- [26] Królak A, Tinto M and Vallisneri M 2004 *Phys. Rev. D* **70** 022003 (Preprint [gr-qc/0401108](#))
- [27] Cornish N J and Crowder J 2005 *Phys. Rev. D* **72** 043005 (Preprint [gr-qc/0506059](#))
- [28] Cornish N J and Porter E K 2006 *Class. Quantum Grav.* **23** S761–8 (Preprint [gr-qc/0605085](#))
- [29] Wickham E D L, Stroeer A and Vecchio A 2006 *Class. Quantum Grav.* **23** S819–28 (Preprint [gr-qc/0605071](#))
- [30] Cornish N J and Porter E K 2007 *Phys. Rev. D* **75** 021301 (Preprint [gr-qc/0605135](#))
- [31] Stroeer A, Gair J and Vecchio A 2006 *Laser Interferometer Space Antenna, 6th Int. LISA Symp. (AIP Conf. Proc. vol 873)* ed S M Merkowitz and J C Livas pp 444–51 (Preprint [gr-qc/0609010](#))
- [32] Crowder J and Cornish N 2007 *Phys. Rev. D* **75** 043008 (Preprint [astro-ph/0611546](#))
- [33] Vallisneri M 2005 *Phys. Rev. D* **71** 022001 (Preprint [gr-qc/0407102](#))
- [34] Cornish N J and Porter E K 2007 *Class. Quantum Grav.* **24** S501 (Preprint [gr-qc/0701167](#))
- [35] Arnaud K A *et al* (Mock LISA Data Challenge Task Force) 2007 *Class. Quantum Grav.* **24** S551 (Preprint [gr-qc/0701170](#))
- [36] Buonanno A, Chen Y and Vallisneri M 2003 *Phys. Rev. D* **67** 104025 (Preprint [gr-qc/0211087](#))