Report on MLDC 1

"Global LISA Inference Group" (GLIG)*

Ed Bloomer(1), James Clark(1), Nelson Christensen(2), Martin Hendry(1), Chris Messenger(1), Renate Meyer(3) , Matt Pitkin(1), Christian Roever3) , Alexander Stroeer(4,5), Jennifer Toher(1), Richard Umstaetter(3) , John Veitch(1), Alberto Vecchio(4,5), Graham Woan(1)

(1)Department of Physics and Astronomy, University of Glasgow, Glasgow G12 8QQ, UK
(2) Physics and Astronomy, Carleton College, Northfield, MN, USA
(3) Department of Statistics, The University of Auckland, Auckland, New Zealand
(4) School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK
(5) Department of Physics and Astronomy, Northwestern University, Evanston, IL 60208, USA

December 3, 2006

Abstract

We have been developing an approach based on Bayesian inference to the analysis of LISA data sets. Here we report preliminary results obtained with the software that we have developed; in particular we present results for the training data sets 1.1.1a, 1.1.1c (single source galactic binary) and 1.2.1 (massive black holes) and the challenge data set 1.1.1c (single source galactic binary).

VERSION 2.0

*mldc-group@star.sr.bham.ac.uk

Contents

1 Introduction

Here we present our results of our analysis on some of the Mock LISA Data Challenge (MLDC) data sets. We approach this task by using Bayesian inference. Specifically, we apply Markov chain Monte Carlo (MCMC) methods as a means of estimating parameter values associated with the signals. While our results and analysis methods for LISA data analysis are still in their beginning stages, we believe that these MCMC methods offer potential. MCMC techniques are especially well suited for problems where the parameter number is large. Hence, as LISA captures complex signals simultaneously from numerous sources it will be imperative that these sophisticated parameter estimation techniques are applied.

Members of our research team have previously made a demonstration that MCMC methods could be utilized in characterizing a large, yet unknown know, number of signals [13]. Since then the popularity of MCMC methods applied to LISA data analysis has exploded. Members of our group have also demonstrated the use of MCMC methods for a wide range of signals (white dwarfs, massive black holes and extreme-mass ratio inspirals), though on somewhat simplified data sets [15, 9]. In this report we show our initial approach to the MLDC data; we present the preliminary results of our examination of data from galactic binary systems, and massive black hole binary systems. We also discuss the time delay interferometry (TDI) variables which we use in our analysis, and the comparison of our signal values to those of the LISA Simulator.

2 TDI variables

For ground based Michelson interferometers, equal arm lengths ensure the cancellation of laser noise when the light from each path is recombined. This, however, is not the case with LISA, where the compromise between a sustainable orbit and a rigid shape results in unequal arm lengths. Time Delay Interferometry (TDI) provides a solution to this problem by linearly combining and time-shifting the appropriate armlength data to cancel the laser noise component. The time delay corresponding to a particular light travel time from an arm (i) is denoted by $_{,i}$.

Following Armstrong *et al* [1], symmetrical antenna paths will acquire the same laser noise terms that can then be cancelled out during the laser interferometry. One combination of the data streams that satisfies the laser noise cancellation criteria is the *Sagnac* variable α, defined by equation 2.0.1, which is constructed from the difference of the Doppler variables y_{ij} for a clockwise and anti-clockwise path around the space-craft.

$$\alpha = y_{21} - y_{31} + y_{13,2} - y_{12,3} + y_{32,12} - y_{23,13} \tag{2.0.1}$$

Here the time delays ensure that the two light paths both arrive at spacecraft 1 at time t. Note that two other combinations (β, γ), corresponding to light paths that arrive at spacecraft 2 and 3 respectively at time t, can be derived by cyclic permutation of the indices through $1 \to 3$. A δ-gravitational wave signal will result in a six-pulse response; the signal appearing at six different times in the full data-stream, the exact time stamps dependent on the arrival direction of the wave and the configuration of the detector.

The above combinations will be free of laser noise; however, they require the full set of data-streams. Equation 2.0.2 describes a Michelson combination constructed from the data-streams from only two arms [12] and which is free from laser noise even if those arms are of unequal length. This restricts the available time delays but also the sources of laser noise, resulting in a eight-pulse response to gravitational waves.

$$X = y_{32,322} - y_{23,233} + y_{31,22} - y_{21,33} + y_{23,2} - y_{32,3} + y_{21} - y_{31} \tag{2.0.2}$$

Similar to equation 2.0.1, there are two corresponding combinations (Y, Z) found by permutation of the indices. For the first Mock LISA Data Challenge (MLDC), the raw data-sets consist of these three Unequal-Arm Michelson Combinations, X, Y and Z.

2.1 Uncorrelated TDI variables

The raw TDI variables X, Y and Z will be free from laser noise but in general will be strongly correlated with each other. This is a potentially serious issue for parameter estimation because it will result in inefficient MCMC sampling of the likelihood function. To optimise the performance of our MCMC codes, therefore, we define new TDI variables, A, E and T, which are mutually uncorrelated at any given time, t.

Our procedure for constructing A, E and T follows closely that of Prince *et al* [6], in which the authors present expressions for uncorrelated TDI variables, A_p, E_p and T_p, derived from the correlated Sagnac TDI variables α, β and γ[1]

Under the assumption of equal arm lengths, Prince *et al* establish that the noise covariance matrix for the Sagnac variables takes the simplified form of equation 2.1.1.

$$C = \begin{pmatrix} S_\alpha & S_{\alpha\beta} & S_{\alpha\beta} \\ S_{\alpha\beta} & S_\alpha & S_{\alpha\beta} \\ S_{\alpha\beta} & S_{\alpha\beta} & S_\alpha \end{pmatrix} \tag{2.1.1}$$

Uncorrelated TDI combinations are then constructed by diagonalising this matrix – i.e. by computing its eigenvalues and eigenvectors and thus identifying which linear combinations of the Sagnac variables will have a diagonal covariance matrix. These combinations are given by equation 2.1.2

$$A_p = \frac{1}{\sqrt{2}}(\tilde{\gamma} - \tilde{\alpha}) \tag{2.1.2}$$

$$E_p = \frac{1}{\sqrt{6}}(\tilde{\alpha} - 2\tilde{\beta} + \tilde{\gamma}) \tag{2.1.3}$$

$$T_p = \frac{1}{\sqrt{3}}(\tilde{\alpha} + \tilde{\beta} + \tilde{\gamma}) \tag{2.1.4}$$

Note that two of the eigenvalues of the matrix in equation 2.1.1 are equal, due to the fact that all off-diagonal elements are equal, and the expressions in equation 2.1.2 are orthonormalised.

In our case the structure of the noise covariance matrix for X, Y and Z is identical in form to that in equation 2.1.1. Hence, we can obtain uncorrelated TDI variables, A, E and T, via the same diagonalisation procedure. Note, however, that the expressions for A, E and T utilised by the LISA simulator do not correspond trivially to those in Prince *et al* – i.e. simply by replacing α by X, β by Y etc. – but can be easily obtained by interchanging the labels A and E and re-scaling. The expressions for A, E and T are therefore

$$A = \frac{1}{3}(2X - Y - Z) \tag{2.1.5}$$

$$E = -\frac{1}{\sqrt{3}}(Z - Y) \tag{2.1.6}$$

$$T = -\frac{\sqrt{2}}{3}(X + Y + Z) \tag{2.1.7}$$

In our software implementation, those variables are computed using functions derived from the LISA simulator (v2.1.1) [4, 3]; the LISA arm-length is always set to 5×10^9 cm (which introduces a negligible mis-match between signal and template). We always construct the TDI variables in the time domain and then FFT them to perform the computation of the likelihood in the Fourier domain.

[1] Note that, in fact, the derivation in Prince *et al* is for Sagnac variables computed in the frequency domain. However, if we work instead in the time domain then the algebraic relationship between the correlated and uncorrelated TDI variables is the same as in the frequency domain.

2.2 Verification of TDI generation

We carried out extensive tests to validate the TDI generation functions. In particular tests were performed on the noise-free Training sets 1.1.1a and 1.1.1c by computing the match between those data sets and the output of our TDI generation functions. The match is defined as the inner product between two signals maximized over the arbitrary reference phase, see *e. g.* [14] for a definition. Here's a brief summary of the results:

```
Trainingset 1.1.1a, 6 months data, full TDI simulation
match X 0.9999973
match Y 0.9999961
match Z 0.9999989

Trainingset 1.1.1c, 6 months data, full TDI simulation
match X 0.9999412
match Y 0.9999523
match Z 0.9999537

Trainingset 1.1.1a, 6 months data, fixing LISA's arm-length
match X 0.9999904
match Y 0.9999926
match Z 0.9999946

Trainingset 1.1.1a, 6 months data, fixing LISA's arm-length
match X 0.999902
match Y 0.999943
match Z 0.999967

Trainingset 1.1.1a, 12 months data, full TDI simulation
match X 0.9999834
match Y 0.9999901
match Z 0.9999829

Trainingset 1.1.1c, 12 months data, full TDI simulation
match X 0.999813
match Y 0.999856
match Z 0.999896

Trainingset 1.1.1a, 12 months data, fixing LISA's arm-length
match X 0.9999804
match Y 0.9999899
match Z 0.9999839

Trainingset 1.1.1c, 12 months data, shortcut
match X 0.999732
match Y 0.999798
match Z 0.999744
```

3 Galactic binary systems

3.1 Analysis method

We have developed an analysis pipeline to perform Bayesian inference based on Markov chain Monte Carlo (MCMC) methods to sample from the unknown posterior density function that we aim to compute. Our code is constructed to estimate the signal parameters *and* the unknown number of signals present in the data set through the technique known as Reversible Jump MCMC (RJMCMC, e.g. [5]) methods. However, in the present round of MLDCs we apply our approach only to data sets where there is a single source. This decision is based entirely on the computational burden required at present by our algorithm.

Our implementation is based on a hybrid sampler consisting of a random walk Metropolis (RWM) sampler, actively adapted using the adaptive acceptance probability techniques (AAP, [2]) during the burn-in of the sampler, and combined with genetic algorithms through differential evolution MCMC sampling (DEMCMC) with exponential crossover [11]. We implement this technique in a parallel fashion according to [7], in which uncorrelated MCMC chains running on different CPUs can be combined to increase the number of samples for each PDF. This parallelization is also useful for the DEMCMC method, in which parallel chains are needed to establish exponential crossover. In the frame of DEMCMC, the parallel chains become populations, which develop over generations, here expressed as sampling step of the underlying MCMC. We have communication between the populations using exponential crossover, a proposal which compares two populations in order to create a new generation for a third population by simple difference of the parameters within this chain. We setup this hybrid sampler in the following way: one RWM update with ongoing AAP adaption is followed by one DEMCMC update with underlying exponential crossover. This setup is designed to limit the CPU time, since proposals for the RWM stage get actively adapted towards optimum, and since DEMCMC ensures that initial bolder proposals quickly reduce towards smaller regions of interest around the mode of the posterior; equally ensuring, that the full posterior is evaluated at the beginning of the run, when bolder proposals are taken. This algorithm has been applied to a wide variety of signals (white dwarfs, massive black holes and EMRIs) in several test-cases for LISA data analysis [15, 9, 10].

3.2 Results for training data sets – 1.1.1a

Here we present the results for the analysis of the Training data set 1.1.1a. We show plots of the PDFs of each parameter. We performed the MCMC run on the TDI variable X over an observation time of 6 months. Each likelihood estimate for X takes 8 seconds to compute on a typical GHz CPU. The model waveform is generated in the time domain and then FFTW (as well as the data set), so that the computations are done in the Fourier domain. However, instabilities in the Beowulf cluster that we used for the analysis led to a much longer time, namely an average of ≈ 4 minutes. Our code in general considers the noise level as an unknown. Due to the hardware problems we have experience we decided to use a different strategy and estimate the noise level from the data set using a Welch method. Overlapping segments of the data are individually Fourier transformed, and the resulting spectra were averaged. The frequency region used for the noise estimate was XXXXXXX.

Priors were chosen as follows: for the amplitude, uniform prior within 1.76×10^{-22} and 1.82×10^{-22}, for the frequency uniform prior within 0.986 and 1.02 mHz, and for the angles uniform within periodic boundaries between 0 to π or 0 to $2\times\pi$ according to the nature and range of the variable.

Figure 1 shows the PDFs obtained in our analysis. The results are not as clean as one might have hoped, and the main reason (we believe) is to be attributed to the fact that we could not build sufficiently long chains for the PDFs. We were able to identify the signal in the data set and to derive posterior distributions sufficiently close to the true parameters, however several problems are apparent (which we are still investigating): the polar latitude shows a multimodal PDF far from the true value which indicates convergence problems. We further see a flat uniform posterior for the amplitude. Although the median of this distribution is close to the true value, we are either seeing (i) an MCMC output still in its burn-in phase, therefore following the uniform prior, or more likely, (ii) a very wide Gaussian-like distribution so that the mode of the Gaussian follows a plateau. The frequency also shows a peculiar shape. We also observe an offsets/bias of the mode; this can be explained by the fact that the MCMC sampling provides just an approximation to the true posterior, and, in our case, it was performed only for a short time. We close, that a) run-times are still too short in order to derive more conclusive answers about shape and behaviour of the MCMC sampler and sampling, b) that a run purely may not yield enough information to secure data analysis, although it is able to spot the main features of the signal quite accurately.

3.3 Results for training data sets – 1.1.1c

Here we present results for the Training data set 1.1.1c. The code and strategy is exactly the same as the one described for the training data set 1.1.1a, with the only (but important) exception that we worked on the combined TDI variables A and E. Each likelihood estimate for the combined A and E variables took 31 seconds, though problems with the Cluster yet again slowed down the computation to 10 min.

Priors were chosen uniform within the following values: between 6×10^{-23} and 1.00×10^{-22} (amplitude), between 10 and 11.4 mHz (frequency) and within the whole relevant range for the angles.

For this data set the results are much more consistent with what one would expect than for 1.1.1c: the use of both A and E outputs helps considerably in identifying the signal. The PDFs, see Fig 2, are centered around the value of the parameters that describe the signal present in the data, though we still observe a bias/offset in the mode. In some cases we also still see a multimodal distribution. Our results are once again degraded by the small length of the chain that we were able to build.

3.4 Results for blind data sets – 1.1.1c

Here we present results for our MCMC run on the Challenge data set 1.1.1c. Figure 3 and Table 1 contain the results of the analysis. Because of CPU limitations, we do know that the results that we are presenting here are only preliminary, and not accurate. The analysis strategy is the same as for the Training data set 1.1.1c. An important point to stress is that the chain started from the values of the true parameters for the Training data set 1.1.1c (these values are represented by the red vertical lines in the plot). The priors were chosen uniform within the following ranges: for the amplitude within 1×10^{-24} and 4×10^{-22}, for the frequency within 8 and 12 mHz, and for the angles uniform within periodic boundaries between 0 to π or 0 to $2\times\pi$ according to the nature and range of the variable.

The results show that the chain might still be in the burn-in phase; in fact several PDFs are still centered around the initial value of the chain. However, for the frequency, inclination angle, polar longitude (and to some extent) polar latitude we observe convergence to a PDF away from the initial values. Since these distributions have been generated independently in each of the separate parallel chains, we have some evidence that the chains have locked on the signal contained in the data set.

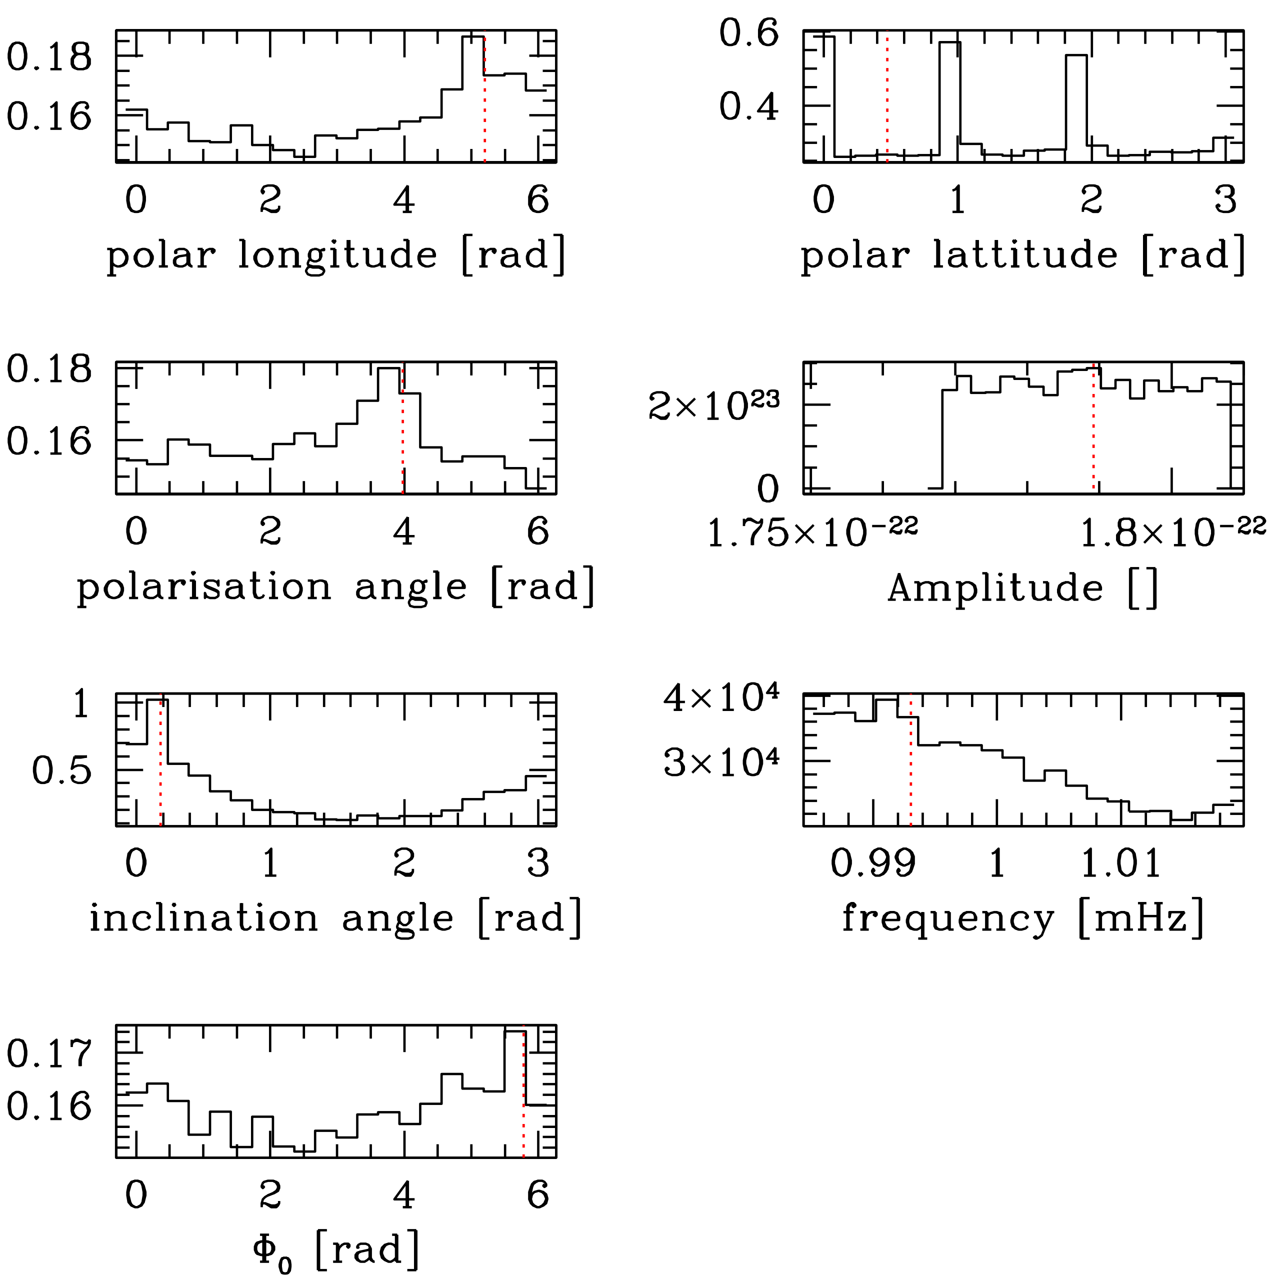

Figure 1: PDFs for the training data set 1.1.1a. The red dashed line represents the true value of the parameter of the signal present in the data set.

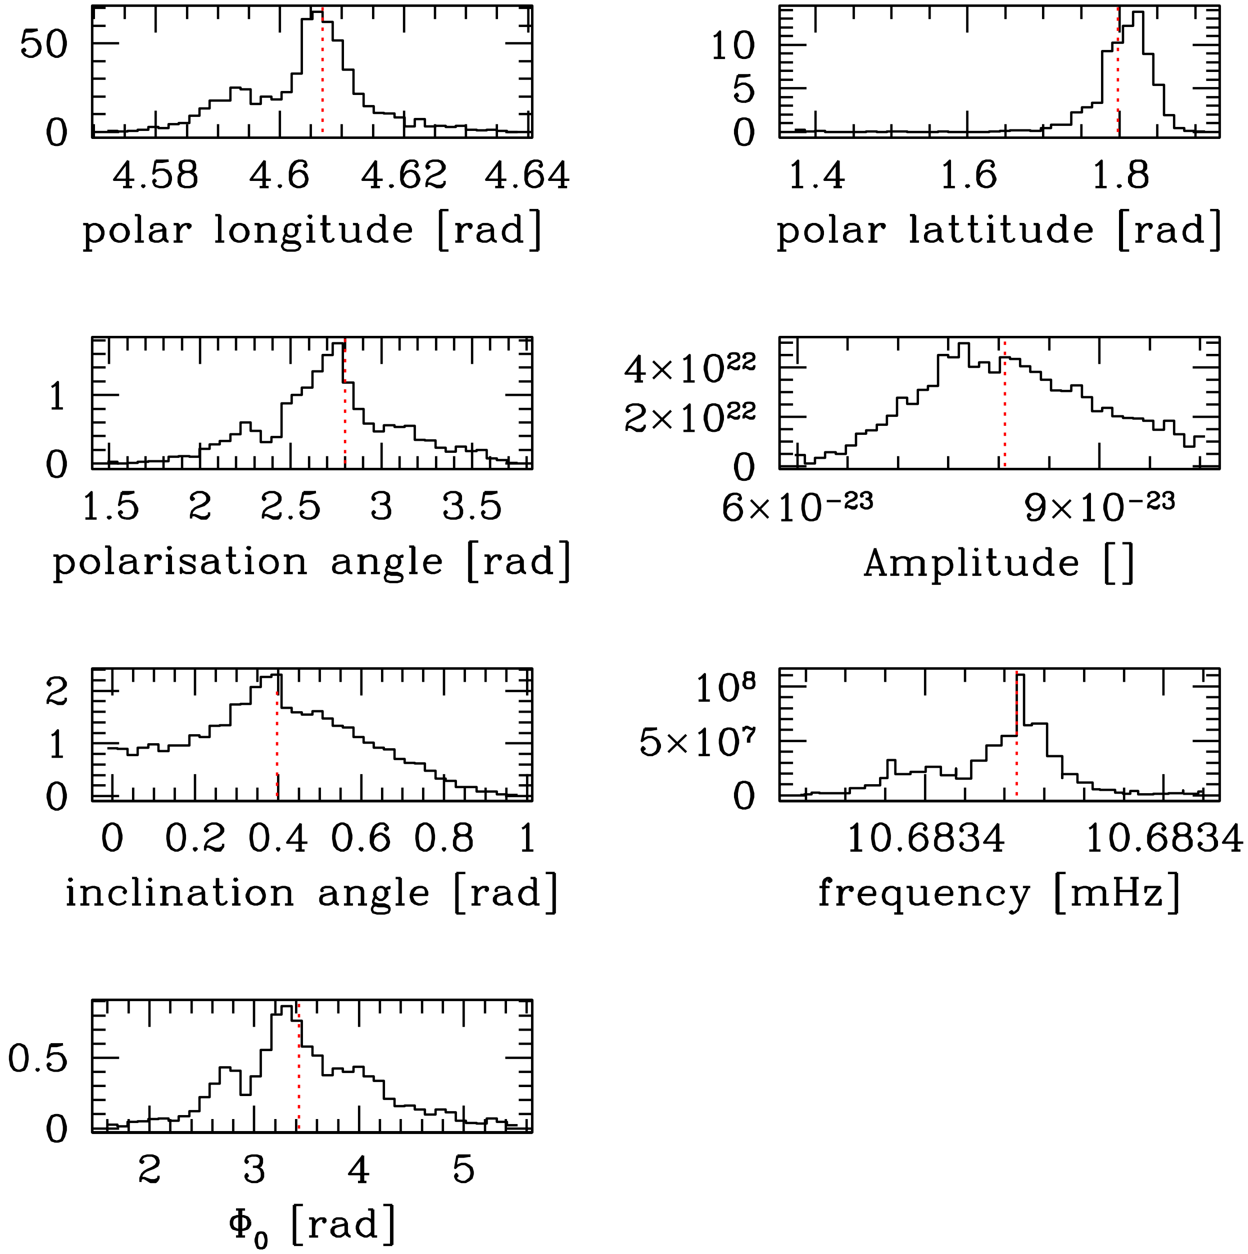

Figure 2: PDFs for the training data set 1.1.1c.

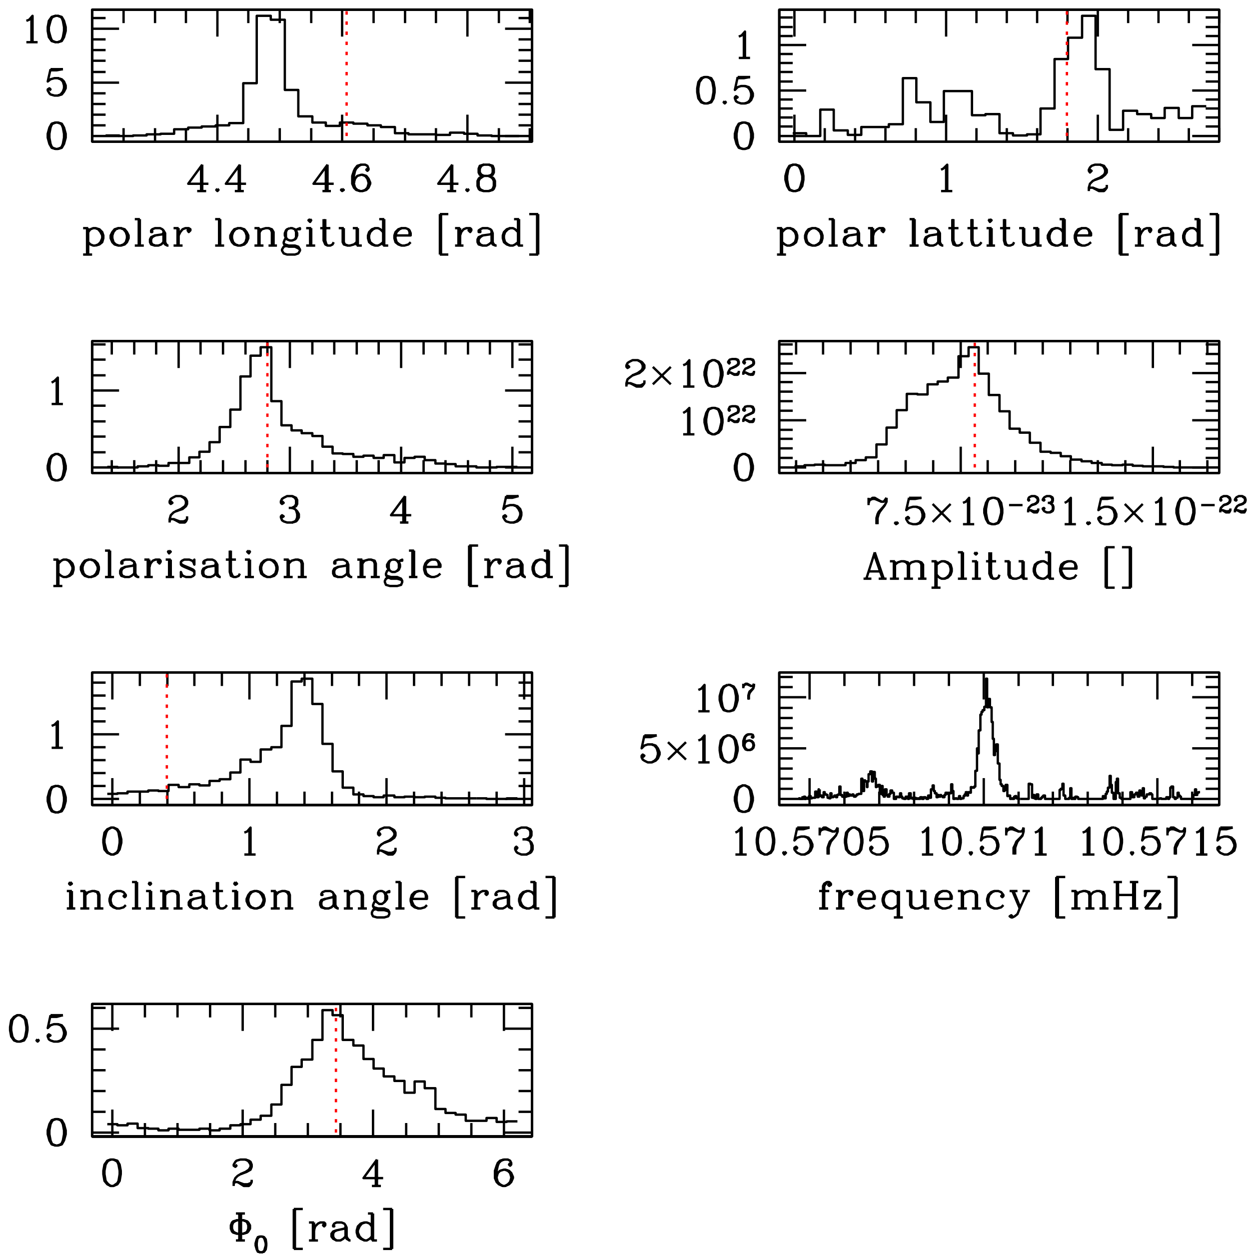

Figure 3: PDFs for the challenge data set 1.1.1c.

parameter	mode	median	mean	sampled standard deviation	unit
polar latitude (ϑ_N)	1.98740517	2.08807	1.91043	7.65262e-01	[rad]
polar longitude (φ_N)	4.48554686	4.51755	4.53682	7.24685e-02	[rad]
polarisation angle (Ψ)	2.83239415	2.83159	2.85234	1.85166e-01	[rad]
inclination (ι)	1.45692286	1.41366	1.41894	1.16190e-01	[rad]
amplitude (A)	8.21384359e-23	8.05668e-23	8.05418e-23	6.84563e-24	[]
frequency (f)	10.5710105	10.5710	10.5710	1.88855e-04	[mHz]
initial GW phase (Φ_0)	3.37761223	3.46168	3.48469	3.08036e-01	[rad]

Table 1: This table shows a summary of the results for Challenge set 1.1.1c.

4 Massive Black Hole Binary systems

4.1 Analysis method

We have developed a Markov chain Monte Carlo (MCMC) program for parameter estimation of massive black hole binary signals for the LISA mock data challenge. Specifically, our code is designed to analyze the signals from Challenge 1.2 - MBH Binaries. This code is a variation of one designed to analyze binary inspiral gravitational wave signals from multiple ground based interferometric detectors; this code is described in [8]. The main difference between our LISA code here and our code for ground based interferometers is that the LISA code is made to handle the form of the LISA data, and the incorporation of the *LISA Simulator* [4] to derive LISA's TDI response to given parameterised GW signals. In order to analyse data in terms of a Bayesian framework, we implemented an MCMC sampler in C that eventually is supposed to both *find* the global mode(s) of the posterior distribution and then '*explore*' the distribution, i.e. simulate random draws from the posterior.

Here we report on our progress in setting up an MCMC sampler to do inference on massive black hole binary inspiral signals in LISA data within a Bayesian framework. For the Challenge 1.2 - MBH Binaries we have adapted our code to take in the given data, and compare it with 2.0 post-Newtonian (PN) inspiral signal templates. In our program we read the data from the MLDC file, given in *lisaXML* format, compute A/E/T TDI variables, and then investigate the inspiral parameters' posterior distribution. The likelihood function is based on Fourier transforms of data and signals, which are both originally available in the time domain. We implemented a general scheme in which arbitrarily parameterised signal waveforms may be injected, and where the likelihood then is computed in the frequency domain. The results of the MCMC runs on MBH binary inspiral signals are based on 2.0 PN templates, characterised by 9 parameters.

The prior information for this Bayesian analysis is straightforward for the 'geometrical' parameters that determine the location and orientation of the inspiral event. A priori, the event is assumed to be equally likely across all possible directions; this leads to a uniform prior for the ecliptic longitude φ, and a prior density

$$f(\theta) = \tfrac{1}{2}\cos(\theta) \tag{4.1.1}$$

that is proportional to the circumference of the corresponding circle of latitude for the ecliptic latitude θ. Analogously, the prior density of the inclination ι is

$$f(\iota) = \tfrac{1}{2}\sin(\iota), \tag{4.1.2}$$

and the remaining parameters, polarisation ψ and coalescence phase ϕ_0, have uniform priors. The coalescence time t_c prior is uniform between 158 and 198 days. For the prior range for the masses we followed the limits set by this MLDC, and assumed the the m_1 prior to be uniform between $1 \times 10^6 \mathrm{M}_\odot$ and $5 \times 10^6 \mathrm{M}_\odot$, and that for m_2 is uniform between $0.25 \times 10^6 \mathrm{M}_\odot$ and $5 \times 10^6 \mathrm{M}_\odot$. Internally however we reparameterised the masses to chirp mass (m_c) and mass ratio (η). Our prior for the luminosity distance D for now is assumed flat as well, from zero to infinity.

4.2 Results for training data set 1.2.1

Presently our only results are with the Challenge 1.2.1 training data set. Table 4.2 shows a summary of the results. In Fig. 4 and 5 we see posterior probabilities generated by our MCMC code for this data. The marginal densities (1-dimensional and 2-dimensional) are shown, the dotted lines indicating the true parameter values. The 9 parameters are chirp mass (m_c), mass ratio (η), coalescence phase (ϕ_0), coalescence time (t_c), latitude (θ), longitude (φ), distance (D), polarisation (ψ) and inclination angle (ι). The estimates for most parameters are very good. But most notably the estimates for three of the parameters are slightly off: coalescence time, distance and phase. Our definition of the phase angle is probably different than that used to generate the data, and is not a major concern. The difference in coalescence time and distance is possibly due to numerical differences, but it could also be from another (yet undiscovered) discrepancy in our code or that used to generate the signals. The discrepancy in distance (which thereby effects the amplitude estimation) can be seen when we attempt to match our templates with the noiseless MLDC signals; see Fig. 6. The slight difference in the amplitude obviously affects our estimate of the signal amplitude. Note too that the SNR of this signal was ≈ 500, so it may not come as a surprise that the discrepancies in templates also turn out as ‘significant’ in the posterior, although numerically the differences may not be as relevant. The power spectral density estimate (necessary for the evaluation of the likelihood function) was computed from the second half of the same data set, after coalescence. For the initial test of our MCMC code we used the true parameter values as the starting point for our chain.

At this point in time our computation is limited by the speed of the likelihood calculation (the by far most time-consuming part being the derivation of LISA’s response to a given signal), and not the amount of data, which is actually comparable to what we were already working on in the application to ground-based measurements. Presently we analyze the final month of data before the inspiral coalescence. A likelihood evaluation takes about 8 seconds, which makes 450 likelihood values per hour. In order for this MCMC code to become practical we will need to increase this rate. The speed of the LISA Simulator is what is limiting the speed of the code, and what for now prevents us from conducting a truly blind search.

parameter	mean	median	95% central interval	true	unit
chirp mass (m_c)	1 024 127	1 024 123	(1 023 875, 1 024 419)	1 023 866	$\mathrm{M}_\odot$
mass ratio (η)	0.200 97	0.200 95	(0.199 82, 0.202 37)	0.199 50	
coalescence phase (ϕ_0)	0.161	0.160	(0.058, 0.290)	6.082*	rad
coalescence time (t_c)	16 545 532	16 545 531	(16 545 513, 16 545 557)	16 545 493	s
ecliptic latitude (θ)	−0.172	−0.176	(−0.278, −0.035)	−0.263	rad
ecliptic longitude (φ)	3.824	3.826	(3.784, 3.852)	3.795	rad
luminosity distance (D)	22 740	22 756	(21 187, 24 073)	27 000	Mpc
polarisation angle (ψ)	2.736	2.737	(2.718, 2.748)	2.730	rad
inclination angle (ι)	1.306	1.306	(1.297, 1.314)	1.317	rad

*orbital *initial* phase

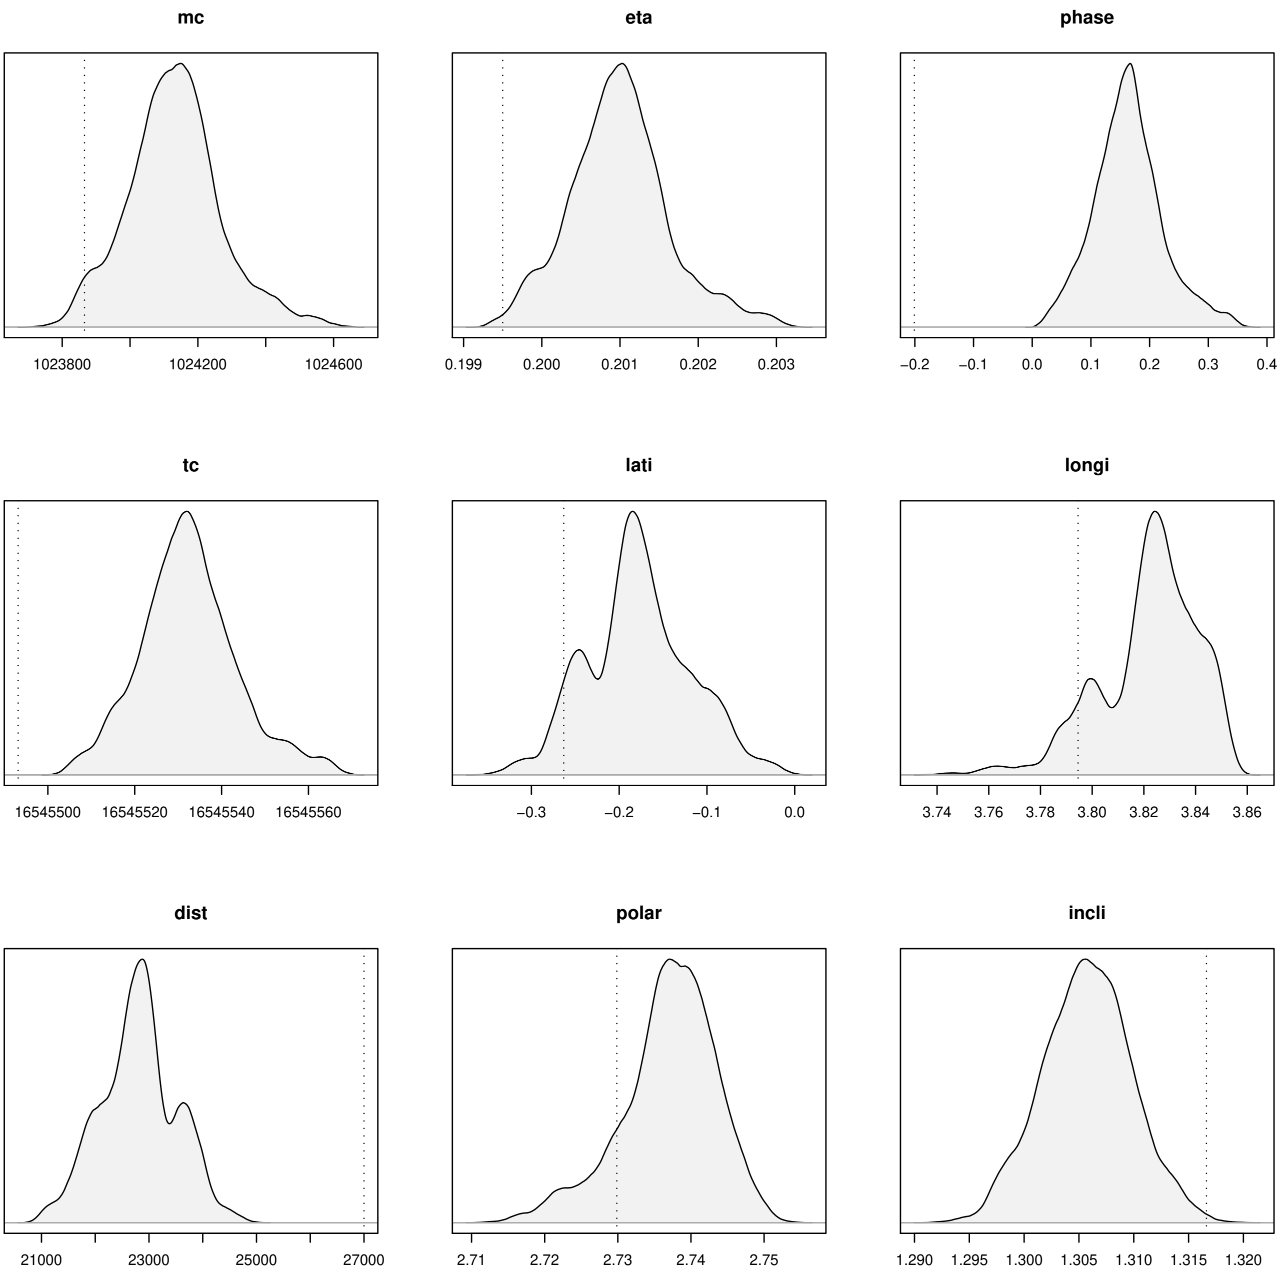

Figure 4: PDFs for the training data set of Challenge 1.2.1

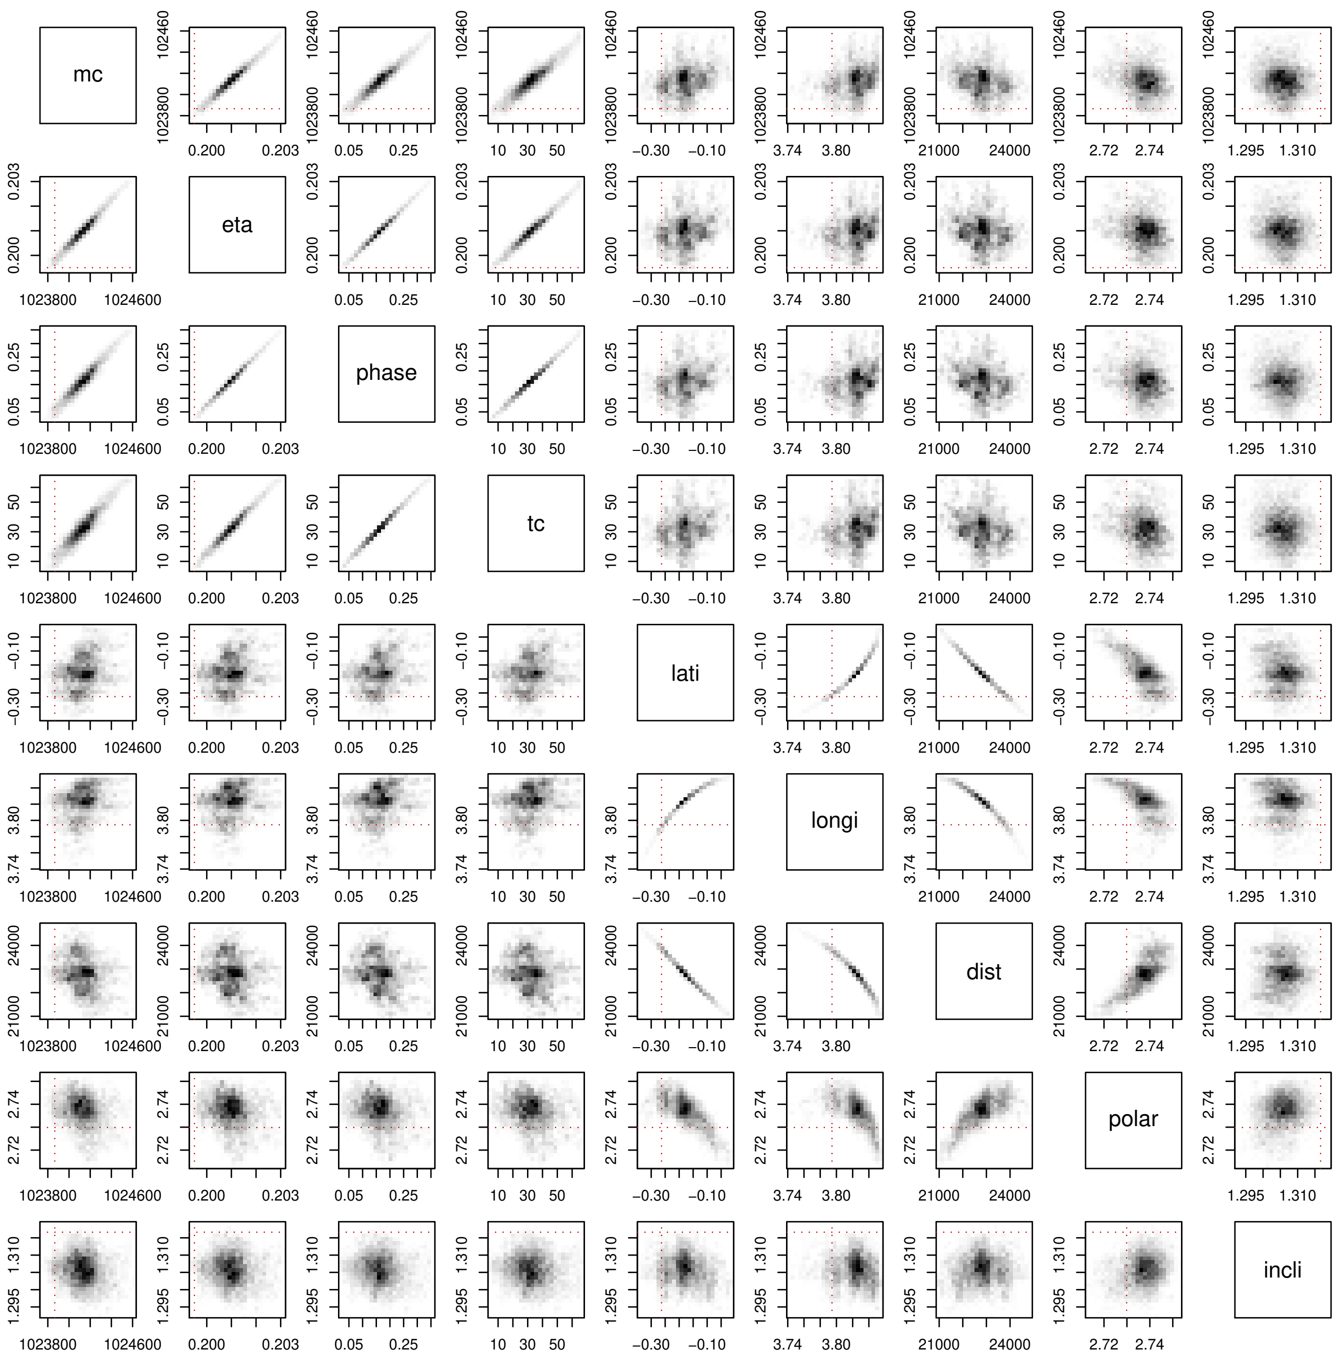

Figure 5: Two dimensional PDFs for the training data set of Challenge 1.2.1

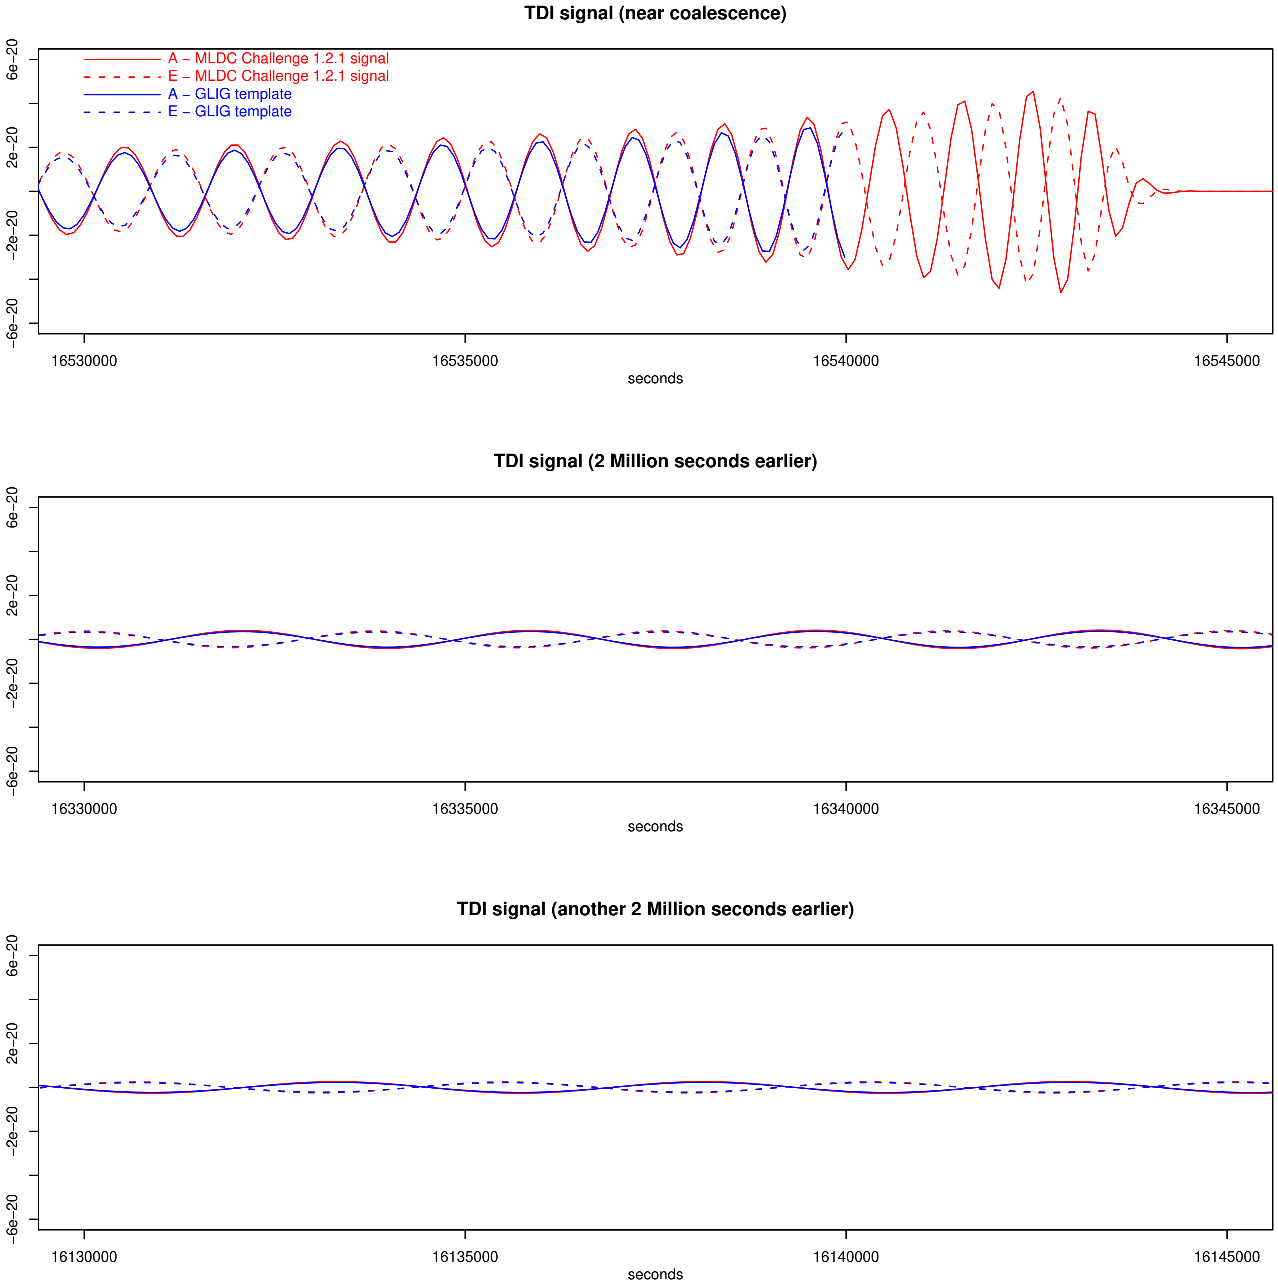

Figure 6: The match of one of our template massive black hole binary templates to that of an equivalent noiseless MLDC signal. This is done for the TDI A and E variables

5 REFERENCES

[1] J. W. Armstrong, F. B. Estabrook, and M. Tinto. Time-Delay Interferometry for Space-based Gravitational Wave Searches. *Astrophysical Journal*, 527:814–826, December 1999. 3

[2] Yves F. Atchade and Jeffrey S. Rosenthal. On Adaptive Markov Chain Monte Carlo Algorithms. Technical report, University of Montreal, July 2003. 6

[3] J. Cornish, N and L. J. Rubbo. LISA response function. *Physical Review D*, 67(2):022001, January 2003. 4

[4] N. J. Cornish, L.J. Rubbo, and O. Poujade. The LISA simulator. `www.physics.montana.edu/LISA/`. 4, 11

[5] P. Green. Trans-dimensional Markov chain Monte Carlo, 2003. 6

[6] T. A. Prince, M. Tinto, S. L. Larson, and J. W. Armstrong. LISA optimal sensitivity. *Physical Review D*, 66(12):122002, December 2002. 4

[7] Jeffrey S. Rosenthal. Parallel computing and monte carlo algorithms, 1999. 6

[8] C. Röver, R. Meyer, and N. Christensen. Coherent Bayesian inference on compact binary inspirals using a network of interferometric gravitational wave detectors. `http://arxiv.org/abs/gr-qc/0609131` (preprint), September 2006. 11

[9] A. Stroeer, J. Gair, and A. Vecchio. Automatic Bayesian inference for LISA data analysis strategies. *ArXiv General Relativity and Quantum Cosmology e-prints*, September 2006. 3, 6

[10] A. Stroeer and A. Vecchio. in preparation. 6

[11] C. J. F. ter Braak . Genetic algorithms and Markov Chain Monte Carlo: Differential Evolution Markov Chain makes Bayesian computing easy (revised). Technical report, August 2005. 6

[12] M. Tinto, F. B. Estabrook, and J. W. Armstrong. Time delay interferometry with moving spacecraft arrays. *Physical Review D*, 69(8):082001, April 2004. 3

[13] R. Umstätter, N. Christensen, M. Hendry, R. Meyer, V. Simha, J. Veitch, S. Vigeland, and G. Woan. Bayesian modeling of source confusion in LISA data. *Physical Review D*, 72:022001, 2005. 3

[14] A. Vecchio and E. D. Wickham. Effect of the LISA response function on observations of monochromatic sources. *Physical Review D*, 70(8):082002–+, October 2004. 5

[15] E. D. L. Wickham, A. Stroeer, and A. Vecchio. A Markov chain Monte Carlo approach to the study of massive black hole binary systems with LISA. *Classical and Quantum Gravity*, 23:819, October 2006. 3, 6