

Detecting extreme mass ratio inspirals with LISA using time–frequency methods: II. Search characterization

Jonathan Gair¹ and Linqing Wen²

¹ Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge, CB3 0HA, UK

² Max Planck Institut fuer Gravitationsphysik, Albert-Einstein-Institut Am Muehlenberg 1, D-14476 Golm, Germany

E-mail: jgair@ast.cam.ac.uk

Received 8 April 2005, in final form 23 June 2005

Published 6 September 2005

Online at stacks.iop.org/CQG/22/S1359

Abstract

The inspirals of stellar-mass compact objects into supermassive black holes constitute some of the most important sources for LISA. Detection of these sources using fully coherent matched filtering is computationally intractable, so alternative approaches are required. In a previous paper (Wen L and Gair J R 2005 *Class. Quantum Grav.* **22** S445), we outlined a detection method based on looking for excess power in a time–frequency spectrogram of the LISA data. The performance of the algorithm was assessed using a single ‘typical’ trial waveform and approximations to the noise statistics. In this paper we present results of Monte Carlo simulations of the search noise statistics and examine its performance in detecting a wider range of trial waveforms. We show that typical extreme mass ratio inspirals can be detected at distances of up to 1–3 Gpc, depending on the source parameters. We also discuss some remaining issues with the technique and possible ways in which the algorithm can be improved.

PACS numbers: 04.25.Nx, 04.30.Db, 04.80.Cc, 04.80.Nn, 95.55.Ym, 95.85.Sz

1. Introduction

Astronomical observations indicate that many galaxies host a supermassive black hole (SMBH) in their centre. The inspirals of stellar-mass compact objects into such SMBHs with mass $M \sim \text{few} \times 10^5 M_\odot - 10^7 M_\odot$ constitute one of the most important gravitational wave (GW) sources for the planned space-based GW observatory LISA. The large parameter space of possible signals and their long duration makes detection using fully coherent matched filtering computationally intractable, so alternative search techniques are required. One option is to use a semi-coherent matched filtering algorithm and a preliminary analysis of that approach [2] suggests that the LISA extreme mass ratio inspiral (EMRI) detection rate will be dominated

by inspirals of $\sim 10M_{\odot}$ BHs onto $\sim 10^6M_{\odot}$ SMBHs, and could be as high as ~ 1000 in 3–5 years within ~ 3.5 Gpc.

Matched filtering algorithms are computationally very expensive and it is therefore valuable to scope out alternatives. In a previous paper [1], we described a strategy to detect GWs from EMRIs by accumulating the signal power in the time–frequency (t–f) domain. The t–f power spectrum is produced by dividing the data into segments of equal duration and carrying out a fast Fourier transform (FFT) on each. The length of the data segments affects the resolution of the spectrogram and must be chosen to allow the evolution of the source to be followed in the t–f plane. This is a search parameter that should eventually be optimized, probably in a source specific way. For now we choose ~ 2 week segments, as these provide appropriate resolution for typical EMRIs. The choice of segment duration does not significantly affect the computational requirements of the search, in contrast to the semi-coherent approach [2] which is limited to 2–3 week segments by computational cost.

The power spectrum is defined for each segment i and frequency bin k as

$$P(i, k) = \sum_{\alpha=I,II} \left[\frac{2|(h_k^{\alpha i} + n_k^{\alpha i})|^2}{\sigma_{n_k^{\alpha}}^2} \right] = \sum_{\alpha=I,II} \left[\frac{2(h_k^{\alpha i})^2}{\sigma_{n_k^{\alpha}}^2} + 4 \frac{\text{Re}[h_k^{\alpha i} (n_k^{\alpha i})^*]}{\sigma_{n_k^{\alpha}}^2} + \frac{2(n_k^{\alpha i})^2}{\sigma_{n_k^{\alpha}}^2} \right], \quad (1)$$

where h_k denotes the Fourier amplitude of the signal, $\sigma_{n_k}^2 = 0.5S_h(f)/(dt^2 df)$ is the expected variance of the noise component n_k at frequency bin k , characterized by $S_h(f)$, the strain spectral density of the noise, dt is the sampling cadence and df is the Fourier bin width. We assume that the LISA noise components, n_k , can be approximated by a stationary and Gaussian distribution, with spectral density $S_h(f)$. In addition to the instrumental noise, the LISA data stream will contain an unresolvable astrophysical foreground from galactic white dwarf (WD) binaries. We include the white dwarfs as a noise contribution in S_h in the usual way ([3], equations (52) and (54)). The instrumental noise spectral density is also taken from [3], equation (48). Including the WD foreground in this way assumes it is isotropic on the sky, which is not a good assumption in this case. For the same reason, the WD foreground will not be a stationary noise source on long timescales, but this is probably a reasonable approximation for few week long data segments. In the future, our analysis will be repeated using a more accurate realization of the WD distribution (e.g., [4]). Generally speaking, EMRIs are not detectable when their frequencies overlap with the WD confusion background, so the results in this paper should not change significantly. However, the location in frequency of the confusion limit will be important, i.e., which WD binaries are resolvable and which are not [5].

In the low-frequency regime, LISA can be regarded as consisting of two independent Michelson detectors which are denoted I and II [6]. We combine the responses from these two detectors to produce the power spectrum. In this first analysis, we add the spectrograms from the I and II responses before performing the search. An alternative algorithm might involve searching each spectrogram separately and looking for coincidences. Such refinements will be investigated in the future. For the combined spectrogram, we calculate the power ‘density’, $\rho(i, k)$, by computing the average power within a rectangular box centred at each point (i, k) ,

$$\rho(i, k) = \frac{1}{m} \sum_{a=-n/2}^{n/2} \sum_{b=-l/2}^{l/2} P(i+a, k+b), \quad (2)$$

where n, l are the lengths of the box edges in the time and frequency dimensions respectively, and $m = n \times l$ is the number of data points enclosed. To search for a possible signal, we vary the size of the box, and for each choice of n and l , (n_j, l_j) say, we search for any points at which $\rho(i, k)$ exceeds a threshold, ρ_j . A detection occurs when there is at least one point

with $\rho(i, k) > \rho_j$ for at least one box size. In the first algorithm, we use only box sizes $(n_j, l_j) = (2^{n_t}, 2^{n_f})$, for all possible integer choices of n_t and n_f . The thresholds, ρ_j , are set to ensure a small overall false alarm probability (FAP) for the search. This will be discussed in more detail later.

The optimal box size for a given source will be large enough to include much of the signal power but small enough to avoid too much noise contribution. The optimum will be source specific due to the wide variation in EMRI waveforms. The inspiral of a $0.6M_\odot$ white dwarf occurs much more slowly than that of a $10M_\odot$ black hole, so in the first case, the optimal box size is likely to be longer in the time dimension. GWs from an inspiral into a rapidly spinning black hole or from a highly eccentric inspiral orbit are characterized by many frequency harmonics, often close together. In that case, a box that is wider in the frequency dimension may perform well. In designing a search, a balance must therefore be struck between having sensitivity to a range of sources and increasing the reach of the search for a specific source.

In section 2 we discuss the statistics of this search in the absence of a signal. In section 3 we discuss how to assign thresholds for the different box sizes and in section 4 we illustrate the performance of the algorithm on a set of trial waveforms. We discuss some remaining issues and possible ways to develop time–frequency techniques in section 5.

2. Statistics of the search in the absence of a signal

2.1. Distribution of maximum power for one box size

If no signal is present, then the power $P(i, k)$ at each point in the unbinned plane is distributed as a χ^2 with four degrees of freedom. There are four degrees of freedom since there are two independent Michelson responses and the real and imaginary parts of n_k^i are independent for each. Over many noise realizations, the total binned power, $m\rho(i, k)$, at a specific point in the spectrogram will be distributed like a χ^2 with $4m$ degrees of freedom. The probability that $m\rho(i, k)$ exceeds a threshold ρ_j is therefore $1 - Q_{\chi_{4m}^2}(\rho_j)$ (where $Q_{\chi_X^2}$ is the cumulative distribution function for a χ^2 with X degrees of freedom). If q independent samples are drawn from such a distribution, the probability that the maximum, ρ_m , is less than a specified value, ρ_0 , is $P(\rho_m < \rho_0) = [Q_{\chi_{4m}^2}(\rho_0)]^q$. However, in the binned t–f plane, $m\rho(i, k)$ is not independent at nearby points since neighbouring bins overlap. There are at least N/m independent samples in the plane, but less than N (where N is the total number of pixels in the unbinned spectrogram). In [1], we used the estimate that there are $N/(m/4)$ independent points. In practice, the distribution of the maximum is not analytically tractable, but it can be computed numerically using Monte Carlo simulations.

From a Monte Carlo simulation of a quarter of a million noise realizations, we have computed the distribution of the maximum power in the absence of a signal for each box size used. These distributions are much more accurate than the approximations used in [1]. The distributions were computed assuming we had 3×2^{25} s (~ 3 years) of LISA data, sampled at a cadence of 8 s and divided this into 2^{20} s (~ 2 weeks) sections. This gives a time–frequency spectrogram with 96 points in time and 65 536 points in frequency and hence $7 \times 17 = 119$ possible box sizes of the form $n = 2^{n_t}, l = 2^{n_f}$. For each noise realization, we computed the time–frequency spectrogram of the LISA mission and then calculated the maximum power in each binned spectrogram. Figure 1 illustrates one typical distribution, with $n = 8$ and $l = 2048$. It also shows the approximate theoretical distribution described above, under the assumption that there are $q = N, N/m$ or $N/(m/4)$ independent samples in the t–f plane. As expected, the true distribution lies between the theoretical distributions for N and N/m independent points. Our first approximation of $N/(m/4)$ independent points




Figure 1. Distribution of maximum power for box size $n = 8, l = 2048$, as computed from Monte Carlo simulations. Also shown is the theoretical distribution described in the text, under the assumption that there are $N, N/(m/64), N/(m/4)$ or N/m independent points in the binned plane.

[1] is not particularly accurate. Assuming $N/(m/64)$ independent points gives a better approximation to the distribution in this case, although the shape is clearly different. This number is surprisingly large, indicating that boxes with significant overlap are still effectively independent. This suggests that the maximum is usually determined by a few unusually bright pixels, rather than sizable regions of high power. The effective number of independent boxes is box size specific, but always lies between N/m and N .

2.2. Overall search false alarm probability

Assigning a threshold, ρ_j , for a given box size amounts to fixing the false alarm probability for that box size, FAP_j (i.e., the probability that this threshold is exceeded by pure noise). If the binned distributions were independent, the overall false alarm probability for the search, FAP , would be $FAP = 1 - \prod_j (1 - FAP_j)$. But as this is not the case, we must make use of the Monte Carlo simulation results again. From the numerical distributions we can compute a threshold for each box size that sets the FAP_j to a specified value. We then use the Monte Carlo data and count the fraction of pure noise realizations in which the maximum power exceeds the threshold for at least one box size. This is the overall search FAP . Since the t - f plane encompasses the whole of the LISA data stream, the FAP is the fraction of LISA missions in which we would expect to have a false alarm. In figure 2, the search FAP is shown as a function of the (equal) FAP_j assigned to each box size. For small FAP_j , $FAP \sim 100 FAP_j$, indicating that effectively 100 of the 119 box sizes used are independent.

3. Choice of thresholds

At a fixed search FAP , a threshold choice corresponds to dividing up the FAP among the different box sizes. If we have no prior knowledge of the signal, the natural way to do this




Figure 2. Overall FAP of the search as a function of the FAP assigned to each box size. This assumes all box sizes are assigned equal FAP_j .

is to treat all the box sizes equally, i.e., set $FAP_j = \text{constant}$. This can be done accurately by determining the thresholds from the numerical distributions, as in the previous section. In [1] we set thresholds to give uniform FAP_j 's under the theoretical approximation that ρ_m^j was distributed as the maximum of $N/(m_j/4)$ independent samples from a $\chi_{4m_j}^2$ distribution. Using that approach, the true FAP_j 's are not equal, but we determined the true FAP of the search for each threshold choice using a small Monte Carlo simulation.

If we are searching for a particular source, the threshold choice can be tuned. To illustrate this, we use the ‘typical’ EMRI described in [1] (source ‘A’). This is the last three years of the inspiral of a $10M_\odot$ black hole into a 10^6M_\odot SMBH of spin $a = 0.8M$, with orbital eccentricity $e_0 = 0.4$ at the start of the observation and inclination $i_0 = 45^\circ$. The frequency spread, time spread and SNR of our target waveform in principle allow us to guess theoretically which box size will be optimal for detecting it. However, the presence of noise means that the source will not always be detected in the same box size and indeed it will not always be the same part of the inspiral that is detected (different box sizes are optimal for different stages of the inspiral). To illustrate this, we ran a Monte Carlo simulation of the search over several hundred noise realizations with source A present in the data, and recorded every time a threshold was exceeded, and for which box size this occurred. The thresholds were set to give equal FAP_j to all box sizes. When this procedure was repeated without the signal, the number of detections was roughly independent of the box size, as we would expect. In figure 3 we show the fraction of ‘detections’ (i.e., a threshold being exceeded) as a function of the box size label. It is clear that there are preferred box sizes for detecting this source, but more than one. For a given noise realization, the source might be detected (i.e., exceed the box threshold ρ_j) with more than one box size, but these box sizes do vary from realization to realization. Thus, a targeted search for source A should still use a number of box sizes.

In figure 4 we compare the performance of the algorithm at detecting source A under these various choices for the thresholds. We show the receiver operator characteristic (ROC)




Figure 3. Fraction of detections in each box size when the signal A is present and absent. The box sizes have been numbered sequentially—boxes numbered 1–17 have $n_t = 1, n_f = 1, 2, 4, \dots$, boxes numbered 18–34 have $n_t = 2, n_f = 1, 2, 4, \dots$ etc. Thresholds were assigned to give equal FAP_j to each box size.

curve, which plots the detection rate versus the overall search FAP. The detection rate was estimated by adding the signal waveform to a sequence of noise realizations, and running the search algorithm. We used ~ 6500 noise realizations, giving an accuracy in the detection rate of $\sim 0.02\%$.

The theoretical threshold assignment used in [1] gives performance that is comparable with the true ‘equal FAP_j ’ threshold assignment. A targeted search, which uses only the eleven box sizes which have detection rates greater than 0.03 in figure 3, and assigns equal FAP_j to each, performs significantly better than the blind search. However, if we use the set of box sizes tuned for source A to detect other sources, the performance can be worse. For example, when searching for the source ‘S’ (described below) at a distance of 1.2 Gpc and at a search FAP of 1%, the detection rate using thresholds tuned for source A is 25% compared to 50% in the untargeted search. The search can also be tuned for source S, and a final data analysis scheme might involve using a targeted search for one type of inspiral followed by a targeted search for another type and so on. However, using multiple tuned searches increases the overall false alarm probability and eventually will do no better than a blind search. Without a much better knowledge of the likely parameters of EMRI sources in the LISA data, it is probably best to treat all box sizes equally. In the next section, we will use thresholds that give equal FAP_j to all box sizes.

4. Algorithm performance

To assess the algorithm’s performance, we computed ROC curves for a sequence of inspirals at a number of distances. The inspiral waveforms were computed using the ‘numerical kludge’ approximation [7–9], with detector modulations added using the low frequency approximation




Figure 4. ROC of the search for detection of source A at 2 Gpc under different assignments of the thresholds. ‘Equal FAP, all bins’ is the threshold choice that assigns equal false alarm probability to all box sizes used. ‘Theoretical’ is the threshold choice used in [1], which assigns equal false alarm probability to all box sizes, assuming the maximum power is distributed as the maximum of $N/(m/4)$ independent points drawn from a χ_{4m}^2 distribution. ‘Equal FAP, targeted’ uses a reduced number of box sizes, tuned for this source, and assigns equal FAP_j to each. ‘Random’ indicates the performance of a random search, for which the detection rate and false alarm probability are equal.

to the LISA response [6]. The detector response can be treated more accurately [10–12] but the low frequency response is sufficient for this analysis. Our ‘typical’ source, ‘A’, has the parameters specified above and was placed at a sky position $\cos(\theta_S) = 0.5$, $\phi_S = 1$, with orientation chosen as $\cos(\theta_K) = -0.5$, $\phi_K = 4$. The angles θ_S and ϕ_S are the ecliptic colatitude and azimuth of the source on the sky, while the angles θ_K and ϕ_K are the ecliptic colatitude and azimuth to which the spin of the SMBH would point in the Solar System frame. The other trial waveforms were obtained by changing one or two of the inspiral parameters. The waveform parameters for the inspirals considered are listed in table 1. For each inspiral, the initial semi-latus rectum, p , of the orbit was chosen such that the plunge would occur a few days after the end of the LISA observation, i.e., in each case we are focusing on approximately the last three years of inspiral. Table 1 also lists the approximate signal-to-noise ratio that would be achieved in a fully coherent matched filtering search, using both Michelson data streams (I and II), if the source was at a distance of 1 Gpc.

For each trial waveform we performed a Monte Carlo simulation, with the source at distances of 0.8, 1.2, 1.4, 2 and 3 Gpc. Source ‘A’ was additionally placed at distances of 1, 1.75, 2.25 and 2.5 Gpc. In figure 5 we show the ROC for detection of source ‘A’ at each distance, in figure 6 we show the ROCs for detection of all the sources at a distance of 2 Gpc and in table 2, we summarize the detection rates for each source at each distance, if the thresholds are set to give an overall search false alarm probability of $\sim 1\%$. Since this is the false alarm rate per LISA mission, we could afford to increase this number in the final search, up to 10% or more. However, for the present we use 1% in order to compare the performance

Table 1. Parameters and signal-to-noise ratios at 1 Gpc for trial waveforms. Unspecified parameters are the same as source ‘A’, as given in the text.

Label	Parameters	Initial p/M	SNR
A	See text	10.3	155
B	$M = 3 \times 10^5 M_\odot$	18.25	119
C	$M = 3 \times 10^6 M_\odot$	6.5	110
D	$m = 0.6 M_\odot, M = 3 \times 10^5 M_\odot$	9.405	14.1
E	$m = 0.6 M_\odot$	5.83	21.0
F	$m = 0.6 M_\odot, M = 3 \times 10^6 M_\odot$	4.511	15.7
G	$m = 100 M_\odot$	17.78	382
H	$a = 0.95 M$	10.07	170
I	$a = 0.5 M$	10.74	132
J	$a = 0.1 M$	11.31	108
K	$e_0 = 0$	10.42	147
L	$e_0 = 0.1$	10.41	150
M	$e_0 = 0.25$	10.385	151
N	$e_0 = 0.7$	9.71	159
O	$\iota = 0$	9.925	223
P	$\iota = 30^\circ$	10.1	189
Q	$\iota = 60^\circ$	10.59	115
R	$\iota = 120^\circ$	12.126	57.7
S	$\iota = 150^\circ$	12.82	79.6
T	$\iota = 180^\circ$	13.11	87.0
ExtrinsA	$\cos(\theta_S) = 0.99$	10.3	117
ExtrinsB	$\cos(\theta_S) = 0.01$	10.3	162
ExtrinsC	$\cos(\theta_K) = 0.99$	10.3	118
ExtrinsD	$\cos(\theta_K) = 0.01$	10.3	146
ExtrinsE	$\phi_K = 0.01$	10.3	111
ExtrinsF	$\phi_K = 2.$	10.3	111

of this technique directly to that of the semi-coherent matched filtering search [2]. In the latter search, a matched filtering SNR of ~ 35 is required for detection when the overall search FAP is set at 1%. No ROC curves for the semi-coherent search are presently available, but the detection rate is estimated to be $\sim 50\%$ at threshold.

With an overall search false alarm probability of 1%, we see from figure 5 that source A is easily detectable out to 2 Gpc. The detection rate at greater distances decreases quite rapidly, although it is still $\sim 20\%$ at 2.25 Gpc. The absolute limit of the search for source A is approximately 3 Gpc, at which point we do no better than a random search. This compares quite favourably with the performance of the semi-coherent matched filtering search [2], which in principle can detect this source out to ~ 4.5 Gpc in a Euclidean universe. In fact, in [2], the reach of the search was limited at a redshift $z \approx 1$, or ~ 3.5 Gpc, by astrophysical considerations, since there is considerable uncertainty about the formation of EMRI sources at higher redshift.

Varying the parameters of the inspiral has a noticeable effect on the detectability. The white dwarf inspirals (sources D, E and F) are essentially not detectable (they would be detectable very close by, but we do not expect any such sources closer than ~ 0.5 Gpc). However, this is because these sources are weak and the SNR is so low that they would not even be detected by the semi-coherent technique. The intermediate mass black hole inspiral (source G) can be detected to great distances since it is very loud, but it is presently unclear whether any such sources are likely to exist [13, 14]. The inspirals into a lower (source B)




Figure 5. ROC performance of the search for detecting the typical EMRI, source ‘A’, at distances of 0.8, 1, 1.2, 1.4, 1.75, 2, 2.25, 2.5 and 3 Gpc (from uppermost downward). The lowermost solid line indicates the performance of a random search, as in figure 4. Note that the curves for 0.8–1.4 Gpc lie on top of the upper axis, with 100% detection rate at all the false alarm rates considered.

or higher (source C) mass SMBH are less detectable, the former becoming marginally detectable at 2 Gpc and the latter becoming marginally detectable at about 2.25 Gpc. This is a consequence of the shape of the LISA noise curve. Using the semi-coherent technique these sources are similarly less detectable than source A, although they can still be detected out to ~ 3.5 Gpc.

Varying the eccentricity (sources K–N) has a moderate effect on detectability—lower eccentricity sources can be detected somewhat further away. This is in contrast to the semi-coherent technique, for which eccentricity has little effect on detectability. Low eccentricity inspirals have power spread over fewer frequency components, and so the power is more likely to be concentrated in certain regions in the t – f plane, making it easier for this algorithm to detect them. Increasing the central SMBH spin (sources H–J) appears to increase the detectability. Inspirals into highly spinning black holes will be at higher frequencies (and thus over the floor of the instrumental noise) and will plunge closer to the central SMBH. Both of these properties increase the SNR of the source and therefore are likely to increase the detectability by either this t – f method or the semi-coherent method, although this has not been investigated for the latter technique. Increasing the inclination of the orbit with respect to the equatorial plane of the SMBH (sources O–T) also affects detectability, with lower inclination orbits being more detectable. The inclination is defined [7] such that an inclination of 0° is a prograde equatorial orbit, while an inclination of 180° is a retrograde equatorial orbit. The increased detectability at low inclination is again largely due to the fact that such orbits plunge closer to the SMBH and thus have higher SNR.

Finally, changing the extrinsic parameters (source position on the sky and inclination of the SMBH spin with respect to our line of sight) also affects detectability (sources




Figure 6. ROC performance of the search for detecting the trial EMRI events at a distance of 2 Gpc. Each plot shows the typical inspiral source 'A', and the ROC performance of a random search. The other trial sources have been divided among the four figures and have the parameters specified in table 1.

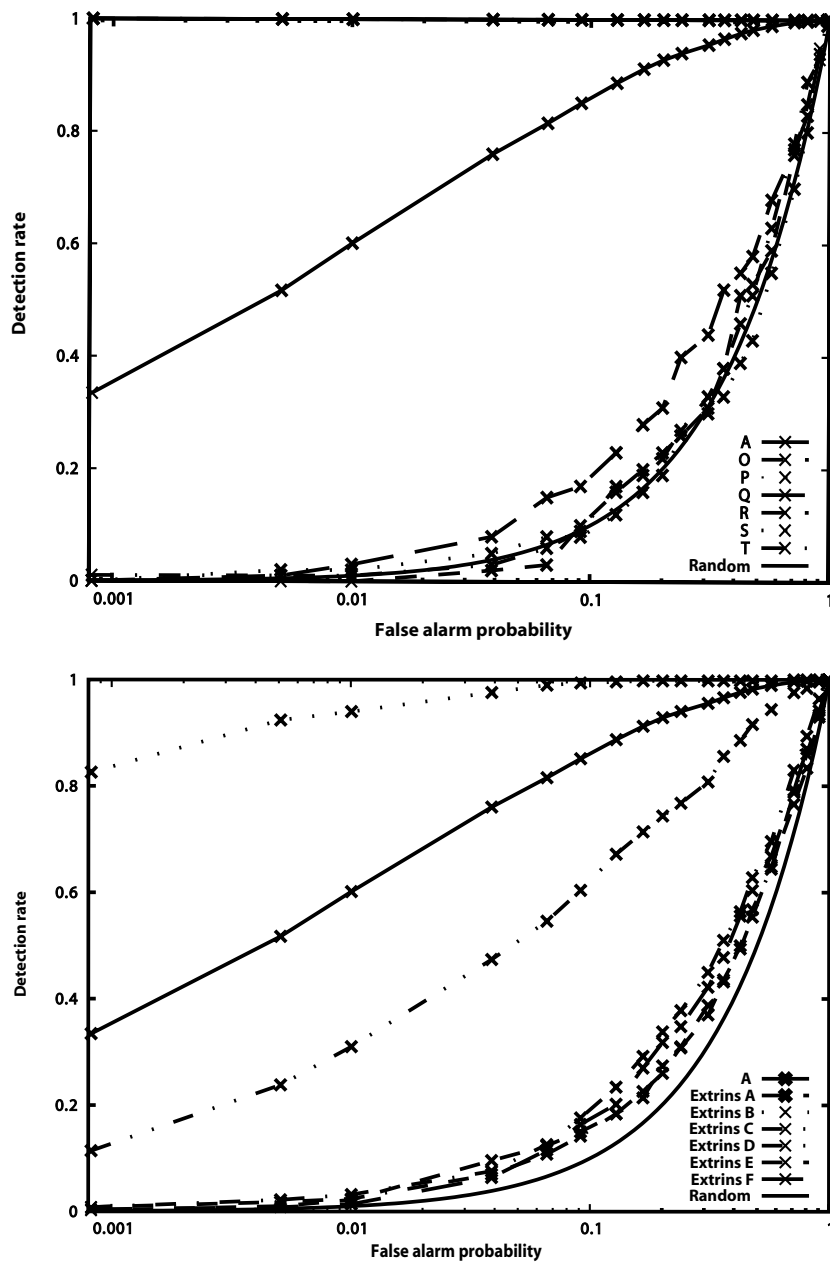


Figure 6. (Continued.)

ExtrinsA–ExtrinsF). By ensuring that the source is in a sky position to which LISA is very sensitive, and that the black hole spin is oriented favourably with respect to the detector, the detectability can be increased. Similarly, placing the source unfavourably can make detection marginal at 2 Gpc. However, the extrinsic parameters chosen for source A seem to represent an ‘average’ response performance.

Table 2. Detection rates for trial waveforms at various distances. Thresholds were set using the numerical probability distributions, and with an overall search false alarm probability of 1%.

	0.8 Gpc	1.2 Gpc	1.4 Gpc	2 Gpc	3 Gpc
A	1	1	1	0.60	0.02
B	1	1	0.93	0.04	0.01
C	1	1	1	0.10	0.02
D	0.00	0.00	0.02	0.01	0.01
E	0.01	0.00	0.00	0.02	0.00
F	0.03	0.02	0.02	0.01	0.02
G	1	1	1	1	1
H	1	1	1	0.96	0.01
I	1	1	1	0.17	0.00
J	1	1	0.85	0.02	0.01
K	1	1	1	1	0.51
L	1	1	1	1	0.29
M	1	1	1	1	0.07
N	1	1	0.99	0.22	0.00
O	1	1	1	1	0.63
P	1	1	1	1	0.10
Q	1	1	0.85	0.03	0.00
R	0.8	0.02	0.04	0.00	0.01
S	1	0.53	0.1	0.02	0.02
T	1	0.96	0.36	0.01	0.01
ExtrinsA	1	1	0.65	0.02	0.01
ExtrinsB	1	1	1	0.94	0.03
ExtrinsC	1	1	0.82	0.03	0.02
ExtrinsD	1	1	1	0.31	0.02
ExtrinsE	1	0.99	0.57	0.03	0.02
ExtrinsF	1	0.99	0.62	0.02	0.02

5. Conclusion

We have presented a thorough analysis of the time–frequency method to detect EMRIs with LISA that was originally proposed in [1]. We have accurately computed the statistics of this search and examined its performance on a wide range of possible EMRI signals. We find that this algorithm is able to detect many different EMRI events out to distances of 1–3 Gpc, depending on the source parameters. In an untargeted search, a typical source can be detected at 2 Gpc with a detection rate of 60% at a search false alarm probability of 1%. Lower eccentricity sources, which have less frequency spreading, can be detected as far away as 3 Gpc with detection rates of 50% at the same overall FAP. The reach of the search can be extended by increasing the allowed FAP or by using a targeted search. By comparison, the semi-coherent matched filtering algorithm [2] can reach ~ 4.5 Gpc for an overall FAP of 1%, but at a presently undetermined detection rate (perhaps $\sim 50\%$). Broadly speaking, this time frequency search has better than half the reach of the semi-coherent search [2], but at a tiny fraction of the computational cost. This is very encouraging, given the simplicity of the technique, and suggests that a method like this could be a valuable first step for detecting the loudest EMRI events in the LISA data. The approach is very similar to the ‘excess power’ technique used in LIGO [15], which was designed to search for bursting signals. Although we have concentrated on EMRIs, this technique will also be sensitive to other types of source, including bursts, galactic white dwarf binaries etc and so it may have many applications.

However, one significant issue that we have not yet considered is that of confusion. The LISA data will be dominated by a superposition of signals from many different astrophysical sources. So far, we have only examined the problem of detecting a single source in instrumental noise (plus an astrophysical foreground of white dwarf binaries). This algorithm should be able to make detections when there are multiple sources if the sources are reasonably well separated in time and frequency or of widely different signal-to-noise ratios, but it is not clear how it will fair on overlapping sources of comparable brightness. It is likely that the performance will be significantly impaired in such cases, due to the simplicity of the algorithm. Additionally, a shortcoming of this simple technique is that while it can tell us that a source is present, it does not provide much information about the properties or parameters of that source. In principle, a time–frequency analysis can provide at least the frequency and rate of frequency drift of an inspiral source, but this requires more sophisticated algorithms. One possibility would be to use boxes of different (i.e., non-rectangular) shapes, or to use a pattern recognition algorithm in the (binned or unbinned) t – f plane to search for tracks corresponding to inspirals. These more sophisticated techniques should also cope more readily with confusion. Finally, it will be instructive to test this and future algorithms using more accurate inspiral waveforms. The simple kludged waveforms used here are purely quadrupolar, and so their frequency structure is not as complex as true EMRI waveforms which might influence our results.

Despite these remaining issues, the performance of this algorithm is very promising and suggests that a time–frequency method of some kind is likely to be a useful first step in detecting EMRIs and other types of source in the LISA data. This provides a good incentive for developing and testing t – f algorithms in the ways suggested above and this work is now underway.

Acknowledgments

We thank Curt Cutler, Teviet Creighton, Leor Barack and Kip Thorne for critical discussions of this work. This work was supported in part by the Max Planck Institut fuer Gravitationsphysik (Albert Einstein Institut) (LW), by NASA grants NAG5-12834 and NAG5-10707 (JG) and by St. Catharine’s College, Cambridge (JG).

References

- [1] Wen L and Gair J R 2005 *Class. Quantum Grav.* **22** S445
- [2] Gair J R, Barack L, Creighton T, Cutler C, Larson S L, Phinney E S and Vallisneri M 2004 *Class. Quantum Grav.* **21** S1595
- [3] Barack L and Cutler C 2004 *Phys. Rev. D* **69** 082005
- [4] Nelemans G, Yungelson L R and Portegies Zwart S F 2001 *Astron. Astrophys.* **375** 890
- [5] Timpano S E, Rubbo L J and Cornish N J 2005 *Preprint gr-qc/0504071*
- [6] Cutler C 1998 *Phys. Rev. D* **57** 7089
- [7] Glampedakis K, Hughes S A and Kennefick D 2002 *Phys. Rev. D* **66** 064005
- [8] Creighton T, Gair J R, Hughes S A and Vallisneri M 2005 in preparation
- [9] Gair J R and Glampedakis K 2005 in preparation
- [10] Cornish N J and Rubbo L J 2003 *Phys. Rev. D* **67** 022001
Cornish N J and Rubbo L J 2003 *Phys. Rev. D* **67** 029905 (erratum)
- [11] Rubbo L J, Cornish N J and Poujade O 2004 *Phys. Rev. D* **69** 082003
- [12] Vallisneri M 2005 *Phys. Rev. D* **71** 022001
- [13] Madau P and Rees M J 2001 *Astrophys. J. Lett.* **551** L27
- [14] Miller M C 2002 *Astrophys. J.* **581** 438
- [15] Anderson W G, Brady P R, Creighton J D E and Flanagan E E 2001 *Phys. Rev. D* **63** 042003