Template Lattices for a Cross-Correlation Search for Gravitational Waves from Scorpius X-1

Katelyn J. Wagner [1], John T. Whelan [2,3], Jared K. Wofford [1,3], and Karl Wette [4,5]

E-mail: kjw4822@rit.edu

E-mail: john.whelan@astro.rit.edu

[1]Center for Computational Relativity and Gravitation and School of Physics and Astronomy, Rochester Institute of Technology, 84 Lomb Memorial Drive, Rochester, NY 14623, United States of America

[2]Center for Computational Relativity and Gravitation and School of Mathematical Sciences, Rochester Institute of Technology, 85 Lomb Memorial Drive, Rochester, NY 14623, United States of America

[3]Institute for Theoretical Physics, Goethe University Frankfurt, Max-von-Laue Str. 1, D-60438 Frankfurt am Main, Germany

[4]Centre for Gravitational Astrophysics, Australian National University, Canberra ACT 2601, Australia

[5]ARC Centre of Excellence for Gravitational Wave Discovery (OzGrav), Hawthorn VIC 3122, Australia

Abstract. We describe the application of the lattice covering problem to the placement of templates in a search for continuous gravitational waves from the low-mass X-Ray binary Scorpius X-1. Efficient placement of templates to cover the parameter space at a given maximum mismatch is an application of the sphere covering problem, for which an implementation is available in the `LatticeTiling` software library. In the case of Sco X-1, potential correlations, in both the prior uncertainty and the mismatch metric, between the orbital period and orbital phase, lead to complications in the efficient construction of the lattice. We define a shearing coordinate transformation which simultaneously minimizes both of these sources of correlation, and allows us to take advantage of the small prior orbital period uncertainty. The resulting lattices have a factor of about 3 fewer templates than the corresponding parameter space grids constructed by the prior straightforward method, allowing a more sensitive search at the same computing cost and maximum mismatch.

1. Introduction

Scorpius X-1 (Sco X-1) is a compact object in a binary system with a low-mass companion star. [1, 2] It is believed to be a rapidly spinning neutron star and a promising source of continuous gravitational waves [3]. The signal received by an observatory such as LIGO[4], Virgo[5] or KAGRA[6] depends on the parameters of the system, and a search for that signal loses sensitivity if the incorrect values are used for those parameters. Several of the parameters are uncertain, and one method to ensure that the signal is not missed is to perform the search at each point in a template bank covering the relevant parameter space. These include the projected semimajor axis $a_p = a \sin i$ of the neutron star's orbit,‡ the orbital period P_{orb}, and the time t_{asc} at which the neutron star crosses the ascending node as measured in the solar-system barycenter.

The loss of signal-to-noise ratio (SNR) associated with an incorrect choice of parameters is, in a generic Taylor expansion, a quadratic function of the parameter offsets. This allows us to write the fractional loss in SNR, also known as the mismatch, as a squared distance using a metric on parameter space. In general, this metric will vary over the parameter space (i.e., the associated geometry will have intrinsic curvature), but we can divide the parameter space into small enough pieces that the space is approximately flat, and the metric can be assumed to be constant. In that case, there exists a transformation to Euclidean coordinates. The problem of placing templates so that the mismatch of any point in parameter space from the nearest template is no more than some maximum mismatch μ is then equivalent to the problem of covering the corresponding Euclidean space with spheres of radius $\sqrt{\mu}$. The most efficient covering in $n \leq 5$ dimensions is the lattice family A_n^*, which includes the hexagonal lattice A_2^* and the body-centered cubic lattice A_3^*. For example, the density of lattice points for A_4^* is a factor of 2.8 lower than the corresponding hypercubic ($\mathbb{Z}^4$) lattice.

We use the `LatticeTiling` module in the LIGO Algorithms Library[7] (`lalsuite`) to investigate efficient lattice coverings for the parameter space of a search for Sco X-1 using advanced LIGO data. We show how the search can be made more efficient by: replacing a hypercubic grid with an A_n^* lattice; accounting for the elliptical boundaries associated with the correlated prior uncertainties between orbital period and orbital phase; defining a sheared coordinate change such that a particular combination of the orbital period and orbital phase is unresolved, and explicitly searching only in the other three dimensions of the parameter space. These improvements allow the search to be carried out using fewer computational resources. Alternatively, since the search method we use is tunable, with a trade-off between computational cost and sensitivity, the more efficient lattice allows a more sensitive search to be done at the same computing cost.

The plan of this paper is as follows: In section 2.1, we briefly summarize the cross-correlation search for continuous gravitational waves (GWs) as applied to Sco X-1. In section 2.2 we describe the existing method of template placement "by hand" in a

‡ The angle i is the inclination of the binary orbit with respect to the line of site, not to be confused with ι, which is in inclination of the neutron star spin to the line of site.

rectangular grid. In section 2.3 we describe the application of the sphere covering problem to generation of template lattices, which is implemented in the `LatticeTiling` module in the `lalsuite` software library[7], as described in [8]. In section 3 we consider the specific features of the parameter space for the Sco X-1 search which impact our search: section 3.1 describes the orbital priors, especially the relationship between orbital period and phase. In section 3.2 we consider the standard coordinates, where the time of ascension describing the orbital phase has been propagated in time to the epoch of the gravitational wave search, inducing prior correlations between the period and time of ascension. In section 3.3 we show how a shearing transformation can be used to define a modified period parameter whose prior uncertainty is independent of the uncertainty on the propagated time of ascension. In section 4 we construct a number of lattices and compare the numbers of templates and modelled computing costs. Finally section 6 contains conclusions and implications of this work.

2. Background

2.1. Cross-Correlation Search for Scorpius X-1

The model-based cross-correlation method [9] has been developed to search for continuous GWs, most notably from the low-mass X-ray binary Sco X-1 [10] and applied to mock data [11] as well as observational data from Advanced LIGO's first and second science runs [12, 13]. It is a semi-coherent method where the data are divided into short segments of duration $200\,\mathrm{s} \lesssim T_{\mathrm{sft}} \lesssim 2000\,\mathrm{s}$, which we call "SFTs" because we construct a Short Fourier Transform from each of them. A detection statistic is constructed including correlations between pairs of segments separated by a coherence time T_{max} or less.§ The sensitivity of the search scales with the number of included pairs; when T_{max} is much less than the total observation time, the detectable GW strain is proportional to $T_{\mathrm{max}}^{-1/4}$. Since the search is computationally limited and the computing cost increases with T_{max}, the search can be tuned to trade computational cost for sensitivity. This tuning can also be done across the parameter space, with different parameter space regions being assigned different T_{max} values. Typically, one uses more computing resources in regions of parameter space which are more likely to contain the signal, where the search is inherently more sensitive, and where it is inherently computationally cheaper.

The output of the search is a detection statistic, which is normalized to have unit variance. The value of ρ can then be seen as a SNR for the search. In the presence of a signal of intrinsic amplitude h_0, the expectation value $E[\rho] \propto h_0^2$, with the proportionality constant being a measure of the sensitivity of the search. Because the model-based statistic is constructed using signal parameters such as intrinsic frequency and parameters influencing the Doppler modulation of the signal, such as sky position

§ In this paper we describe the original "demod" implementation of the search. At low frequencies, the search can be made more efficient by using resampling to reimplement the loop over data and the search over frequencies, as described in [14], but the considerations for the template bank in the orbital parameters are similar.

and the binary orbit of the neutron star, the expected SNR in the presence of a signal will be reduced if the template model parameters differ from those of the signal. For parameters which are unknown or insufficiently constrained, the search is run repeatedly at different points in parameter space to try to find a point close to the true signal. If the parameter values for a search point are $\{\lambda_i\}$ and the corresponding true values of the signal are $\{\lambda_i^s\}$, we can define the *mismatch* μ as the fractional loss in expected SNR:

$$\mu = 1 - \frac{E\left[\rho\right]_{\{\lambda_i\}}}{E\left[\rho\right]_{\{\lambda_i^s\}}}. \tag{2.1}$$

A Taylor expansion in the n parameters $\{\lambda_i\}$ gives‖

$$\mu \approx \sum_{i=1}^{n}\sum_{j=1}^{n} g_{ij}(\lambda_i - \lambda_i^s)(\lambda_j - \lambda_j^s), \tag{2.2}$$

where the matrix with elements $\{g_{ij}\}$ acts as a metric on parameter space, thereafter referred to for brevity as "the metric $\{g_{ij}\}$". The general form of the metric for the cross-correlation search is [10]

$$g_{ij} \approx \frac{1}{2}\langle \Delta\Phi_{\alpha,i}\Delta\Phi_{\alpha,j}\rangle_\alpha \ , \tag{2.3}$$

where α represents a pair of SFTs, $\Delta\Phi_\alpha$ is the difference in modelled signal phase between the SFTs in the pair, $\langle\cdot\rangle_\alpha$ is an average over SFT pairs weighted by the antenna patterns and sensitivity of the detectors involved, and $_{,i} = \frac{\partial}{\partial\lambda i}$ is a partial derivative with respect to the parameter λ_i.

2.2. Simple Rectangular Template Placement

The cross-correlation analyses run to date examined a parameter space divided up into rectangular regions, small enough to assume a constant metric. Then, a set of discrete points is placed over the parameter space which lie on a rectangular grid with spacing $\delta\lambda_i$ in the λ_i direction, using what we refer to as the "by hand" method. The number of points used is

$$N_i = \left\lceil \frac{\lambda_i^{\max} - \lambda_i^{\min}}{\delta\lambda_i}, \right\rceil \tag{2.4}$$

where $\lceil\cdot\rceil$ indicates rounding to up to the next integer. The spacing $\delta\lambda_i$ is chosen to be

$$\delta\lambda_i = \sqrt{\frac{\mu_i}{g_{ii}}} \tag{2.5}$$

so that the mismatch between adjacent points¶ in the λ_i direction is $g_{ii}(\delta\lambda_i)^2 = \mu_i$. If the metric is approximately diagonal, $g_{ij} = g_{ii}\delta_{ij}$, then the point in the parameter space

‖ This assumes that the expected SNR $E\left[\rho\right]$ is a local maximum at the true signal point $\lambda_i = \lambda_i^s$. This is not quite true, as shown in Section IV.A of [10], but it is a good starting point.

¶ Note that for historical reasons, μ_i is defined as the mismatch between adjacent points in the grid, rather than the maximum mismatch between some point in the parameter space and the nearest grid point. This is the origin of the factor of $\frac{1}{4}$ appearing in (2.6).

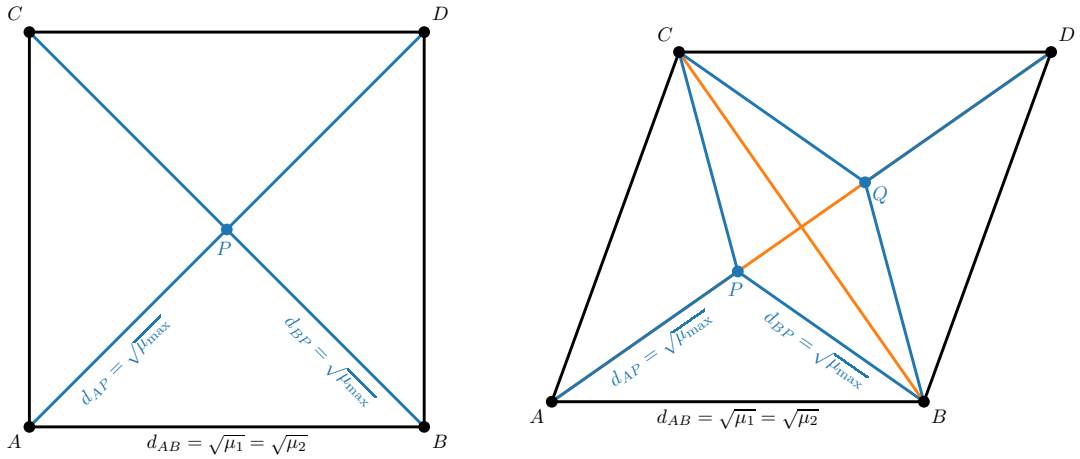

Figure 1. Illustration of maximum mismatch of the by-hand grid of section 2.2 when the metric is not diagonal. Specializing to the case of $n = 2$ and $\mu_1 = \mu_2$, we transform to Euclidean coordinates as described in section 2.3. Left: If the metric is diagonal, this is just a scaling which transforms the rectangle defined by four grid points $ABDC$ into a square with sides of length $\sqrt{\mu_1} = \sqrt{\mu_2}$. The point in parameter space farthest from any grid point is the center P of the square, with $d_{AP} = d_{BP} = d_{CP} = d_{DP} = \sqrt{\mu_1/2}$. Right: If the metric is not diagonal, this square becomes a rhombus $ABDC$. Defining AD to be the long diagonal, the point farthest from any vertex is not the center, but a point P on the long diagonal with $d_{AP} = d_{BP} = d_{CP}$. (There is an equivalent point Q on the other side of the center with $d_{DQ} = d_{BQ} = d_{CQ}$.) We see that APB (or equivalently APC or DQB or DQC) is an isoceles triangle with $d_{AP} = d_{BP} = \sqrt{\mu_{\max}}$ and $d_{AB} = \sqrt{\mu_1}$. In the case of a non-diagonal metric, $\angle APB$ is an obtuse angle, and $d_{AP} = d_{BP} < d_{AB}/\sqrt{2}$, so $\mu_{\max} > \mu_1/2 = (\mu_1 + \mu_2)/4$.

farthest (in the sense of the metric) from any grid point is $\frac{\delta\lambda_i}{2}$ away in the λ_i direction, and has a total mismatch of

$$\mu_{\max} = \sum_{i=1}^{n} g_{ii} \left(\frac{\delta\lambda_i}{2}\right)^2 = \frac{1}{4}\sum_{i=1}^{n} \mu_i \tag{2.6}$$

If the metric is not diagonal, the procedure described above will lead to a maximum mismatch greater than that given in (2.6), as illustrated in figure 1. This approach is conservative and can result in much larger template banks if the metric contains large correlations. The number of templates could be reduced by accounting for the metric correlations, which will be discussed later in section 3.3.

2.3. Covering Lattices

The general problem of choosing a set of template points with a prescribed maximum mismatch distance $\mu_{\max}$ between any point in the parameter space and the nearest template is an application of the *sphere covering problem* [15]. Since we treat the metric $\{g_{ij}\}$ as approximately constant, there is always a linear transformation of the

parameters $\{\lambda_i\}$ into Euclidean coordinates $\{x_i\}$; the mismatch between two points separated by parameter differences $\{\Delta\lambda_i\}$ is then

$$\sum_{i=1}^{n}\sum_{j=1}^{n} g_{ij}(\Delta\lambda_i)(\Delta\lambda_j) = \sum_{i=1}^{n}(\Delta x_i)^2 \ . \tag{2.7}$$

The template placement problem is then simplified to one of placing (hyper-)spheres of radius $\sqrt{\mu_{\max}}$ in the $\{x_i\}$ space so that every point of the region of interest is covered by at least one sphere. To efficiently cover the space, the overlap between spheres should be minimized. This is quantified using the normalized thickness or center density θ, which is the average number of templates per unit volume for the unit sphere.

A sphere covering based on a repeating pattern is known as a lattice. The number of templates required to cover the space will be at a minimum when the lattice has the smallest thickness θ. A perfect lattice has a thickness of 1.

The simplest lattice is the cubic lattice $\mathbb{Z}^n$, which has points equally spaced in each of the (Euclidean) coordinate directions. The "by hand" lattice of section 2.2 is an example of a $\mathbb{Z}^n$ lattice, if the metric $\{g_{ij}\}$ is diagonal and all of the mismatches μ_i are chosen to be equal. A more efficient lattice is A_n^*, which is a general analogue of the hexagonal lattice. For the sphere covering problem, the thinnest lattice is the A_n^* lattice, which in two dimensions has a hexagonal principal cell. The principal cell is the set of points closest to given point in a lattice, and the vertices are locations where covering spheres intersect. It has been shown for $n \leq 5$ that A_n^* is the most efficient covering lattice, i.e., has the smallest thickness θ [15], and for higher dimensions it is typically close to the most efficient covering [16]. Since a more efficient lattice allows the same volume of parameter space to be covered with fewer templates, it can reduce the necessary computing cost at a given sensitivity.[+]

Construction of $\mathbb{Z}^n$ and A_n^* lattices in physical coordinates $\{\lambda_i\}$ given a constant mismatch metric $\{g_{ij}\}$ is implemented by the `LatticeTiling` module in the `lalsuite` software library[7], as described in [8]. A particular challenge is ensuring that the area within the boundaries of a search region is completely covered, which sometimes requires retaining templates whose parameters lie outside the search region. We shall see that this can necessitate some care in choosing coordinates to take advantage of underresolved directions in parameter space.

3. Parameter Space for Sco X-1 Search

3.1. Observational Priors

The GW signal produced by a spinning neutron star is the system is nearly periodic in the neutron star's rest frame, and Doppler shifted as a result of the motion of the detector as the Earth rotates and moves in its orbit and, in the case of a low-mass X-ray binary such as Sco X-1, of the neutron star in its own orbit with its binary companion.

[+] But see [17].

For an accreting neutron star in approximate spin equilibrium, the frequency f_0 of the system can be approximated as constant.* The Doppler shift from detector motion is primarily affected by sky position, which for Sco X-1 is well enough known[19] that its uncertainty does not affect the search. The Doppler shift from the binary motion is affected by five orbital parameters: eccentricity, orientation, projected orbital speed, orbital period, and orbital phase[20]. The orbit of Sco X-1 is believed to be nearly circular[21], and the orientation parameters, inclination ι and polarization angle ψ, are averaged out through time-dependent antenna functions [9]. Thus the search needs to cover only three orbital parameters: projected speed, period, and phase.

The best constraints on these come from [22]. The constraint on orbital period P_{orb} is Gaussian, with a mean of $P_0 = 68023.86\,\mathrm{s}$ and a standard deviation of $\sigma_{P_{\mathrm{orb}}} = 0.043\,\mathrm{s}$. The orbital phase is described by time of ascension t_{asc}, which is the time at which the neutron star crosses the plane of the sky moving away from the observer (i.e., crosses the ascending node). The constraint on this is also Gaussian, with a mean of $t_{\mathrm{asc},0} = 974416624$ GPS (2010–Nov–21 23:16:49 UTC) and a standard deviation of $\sigma_{t_{\mathrm{asc}}} = 50\,\mathrm{s}$. These estimates are uncorrelated, as shown in the left panel of figure 2, but if we convert the time of ascension to a subsequent equivalent time $t'_{\mathrm{asc}} = t_{\mathrm{asc}} + n_{\mathrm{orb}} P_{\mathrm{orb}}$, a correlation is induced, as described in section 3.2 and shown in the right panel of figure 2. The constraints on the orbital velocity of the neutron star in [22] are described in terms of the amplitude of the component of velocity along the line of sight, known as K_1, and consist of constraints that $40\,\mathrm{km/s} \leq K_1 \leq 90\,\mathrm{km/s}$, but without a well-determined probability density between those limits. Searches for GWs from Sco X-1 typically use a uniform prior distribution on this parameter. The parameter used is also typically written as the line-of-sight component of the semimajor axis of the orbit, $a_p = \frac{K_1 P_{\mathrm{orb}}}{2\pi}$. Since the relative uncertainty on P_{orb} is much less than on K_1, one assumes a uniform prior on a_p for $1.44\,\mathrm{lt\text{-}s} \leq a_p \leq 3.25\,\mathrm{lt\text{-}s}$, where the units on a_p are given in light-seconds.

3.2. Standard Search Coordinates

The phase derivatives $\{\Delta\Phi_{\alpha,i}\}$ appearing in (2.3) are computed in [10] for the standard search coordinates $\{\lambda_i\} \equiv \{f_0, a_p, t'_{\mathrm{asc}}, P_{\mathrm{orb}}\}$. Here we introduce the t'_{asc} as the time of ascension at a given point in the propagated 2019 coordinates, indicated by the prime. For long searches which evenly sample the orbital phase of the binary, the non-negligible

* In practice, this equilibrium will be imperfect, leading to some "spin wandering", but the impact of deviations from equilibrium was shown in [10] to be limited when the coherence time is not too long, especially with the levels of spin wandering predicted by [18].

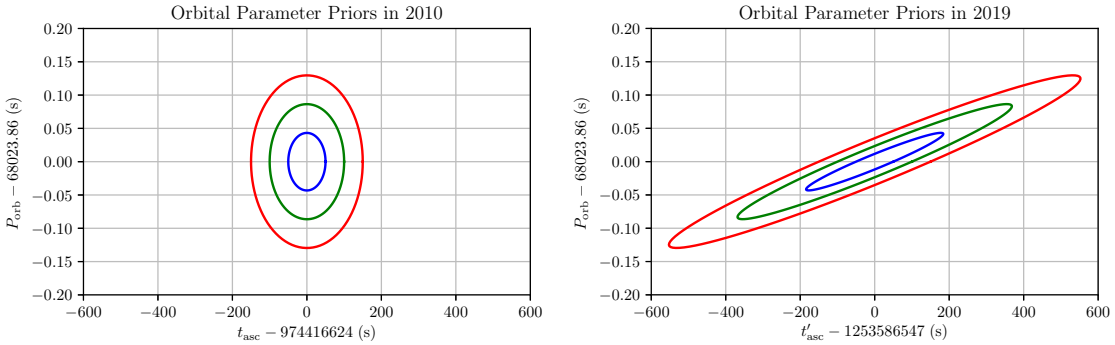

Figure 2. Orbital parameter constraints from [22]. If the time of ascension $t_{\rm asc}$ is quoted in 2010 (left panel) its uncertainty is uncorrelated with the orbital period $P_{\rm orb}$. If we propagate forward by $n_{\rm orb}$ orbits to determine an equivalent time of ascension $t'_{\rm asc} = t_{\rm asc} + n_{\rm orb}P_{\rm orb}$, this introduces correlations. In each case the level surfaces of the probability distribution are shown, for which χ^2, defined in (3.7) and (3.8) equals 1^1, 2^2 and 3^2. We refer to these as at normalized distances of 1σ, 2σ, and 3σ, and they correspond to cumulative probabilities of 39.3%, 86.5%, and 98.9%, respectively.

metric elements have the approximate form♯

$$g_{f_0 f_0} \approx 2\pi^2 \left\langle \Delta t_\alpha^2 \right\rangle_\alpha \approx \frac{2\pi^2}{3} T_{\rm max}^2, \tag{3.1a}$$

$$g_{a_p a_p} = 4\pi^2 f_0^2 \left\langle \sin^2 \frac{\pi \Delta t_\alpha}{P_{\rm orb}} \right\rangle_\alpha \approx 2\pi^2 f_0^2 \left(1 - \operatorname{sinc} \frac{2T_{\rm max}}{P_{\rm orb}}\right), \tag{3.1b}$$

$$\begin{pmatrix} g'_{t'_{\rm asc}t'_{\rm asc}} & g'_{t'_{\rm asc}P_{\rm orb}} \\ g'_{P_{\rm orb}t'_{\rm asc}} & g'_{P_{\rm orb}P_{\rm orb}} \end{pmatrix} \approx \begin{pmatrix} 1 & \frac{-(t'_{\rm asc} - \langle \bar{t}_\alpha \rangle_\alpha)}{P_{\rm orb}} \\ \frac{-(t'_{\rm asc} - \langle \bar{t}_\alpha \rangle_\alpha)}{P_{\rm orb}} & \frac{\langle (t'_{\rm asc} - \bar{t}_\alpha)^2 \rangle_\alpha}{P_{\rm orb}} \end{pmatrix} \frac{16\pi^4 f_0^2 a_p^2}{P_{\rm orb}^2} \left\langle \sin^2 \frac{\pi \Delta t_\alpha}{P_{\rm orb}} \right\rangle_\alpha, \tag{3.1c}$$

where Δt_α is the difference between the timestamps of the two SFTs in pair α, and $\bar{t}_\alpha$ is their mean. Note that the implementation in `lalsuite` [7] uses the exact metric elements, which include additional (generally small) off-diagonal elements.

If we define the midpoint of the run (according to the weighted average $\langle\cdot\rangle_\alpha$) to be $\mu_{\rm obs} = \langle \bar{t}_\alpha \rangle_\alpha$ and the variance as $\sigma_{\rm obs}^2 = \langle (\bar{t}_\alpha - \mu_{\rm obs})^2 \rangle_\alpha$ The metric elements on the $t'_{\rm asc}, P_{\rm orb}$ subspace become

$$g'_{t'_{\rm asc}t'_{\rm asc}} \approx \frac{16\pi^4 f_0^2 a_p^2}{P_{\rm orb}^2} \left\langle \sin^2 \frac{\pi \Delta t_\alpha}{P_{\rm orb}} \right\rangle_\alpha, \tag{3.2a}$$

$$g'_{t'_{\rm asc}P_{\rm orb}} \approx \left(\frac{-(t'_{\rm asc} - \mu_{\rm obs})}{P_{\rm orb}}\right) g'_{t'_{\rm asc}t'_{\rm asc}}, \tag{3.2b}$$

$$g'_{P_{\rm orb}P_{\rm orb}} \approx \left(\frac{(t'_{\rm asc} - \mu_{\rm obs})^2 + \sigma_{\rm obs}^2}{P_{\rm orb}^2}\right) g'_{t'_{\rm asc}t'_{\rm asc}}. \tag{3.2c}$$

♯ Note that the original formula for $g'_{t'_{\rm asc}P_{\rm orb}}$–(4.20h) in [10]–contains a sign error which has not been relevant previously because approximate form of $g'_{t'_{\rm asc}P_{\rm orb}}$ has only been used to set it to zero.

Note that, as shown in [10], if we ignore any data gaps and the noise and antenna pattern weighting, for an observing run of duration $T_{\rm obs}$ and coherence time $T_{\rm max}$, $\sigma^2_{\rm obs} \approx \frac{T^2_{\rm obs}}{12}$, and $\left\langle \sin^2 \frac{\pi \Delta t_\alpha}{P_{\rm orb}} \right\rangle_\alpha \approx \frac{1}{2}\left(1 - \mathrm{sinc}\, \frac{2T_{\rm max}}{P_{\rm orb}}\right)$.

If we choose $n_{\rm orb}$ so that

$$t'_{\rm asc,0} = t_{\rm asc,0} + n_{\rm orb} P_0 \tag{3.3}$$

is as close as possible to $\mu_{\rm obs}$, we can minimize the magnitude of

$$g'_{t'_{\rm asc} P_{\rm orb}} \approx \left(\frac{t_{\rm asc,0} + n_{\rm orb} P_0 - \mu_{\rm obs}}{P_0}\right) g'_{t'_{\rm asc} t'_{\rm asc}}. \tag{3.4}$$

This achieved by taking

$$n_{\rm orb} = \left\lfloor \frac{\mu_{\rm obs} - t_{\rm asc,0}}{P_0} \right\rceil \tag{3.5}$$

where $\lfloor \cdot \rceil$ indicates rounding to the nearest integer.

To give a concrete example, we consider the LIGO-Virgo O3 data run [23] which began on 2019–Apr–01 00:00:00 UTC (GPS 1238112018), continued until a commissioning break at 2019–Oct–01 00:00:00 UTC (GPS 1253923218), resumed on 2019–Nov–01 15:00:00 UTC (GPS 1256655618), and ended on 2020–Mar–27 17:00:00 UTC (GPS 1269363618). Neglecting variability of antenna patterns and noise spectra, as well as any data gaps other than the commissioning break, we find an average time of $\mu_{\rm obs}$ = GPS 1253589161 ≡2019–Sep–27 03:12:23 UTC. This translates into an optimal $n_{\rm orb} = 4104$, corresponding to $t'_{\rm asc,0}$ = GPS 1253586547 ≡2019–Sep–27 02:28:49 UTC.††

The joint prior on $t'_{\rm asc}$ and $P_{\rm orb}$ will remain a multivariate Gaussian, but now with a non-diagonal variance-covariance matrix. The marginal prior on $t'_{\rm asc}$ will be a Gaussian with mean $t'_{\rm asc,0}$ and variance

$$\sigma^2_{t'_{\rm asc}} = \sigma^2_{t_{\rm asc}} + n^2_{\rm orb}\sigma^2_{P_{\rm orb}}. \tag{3.6}$$

The joint prior can be illustrated by plotting level curves of the quantity

$$\chi^2 = \left(\frac{t_{\rm asc} - t_{\rm asc,0}}{\sigma_{t_{\rm asc}}}\right)^2 + \left(\frac{P_{\rm orb} - P_0}{\sigma_{P_{\rm orb}}}\right)^2, \tag{3.7}$$

whose prior distribution is a chi-squared with two degrees of freedom (figure 2, right panel). A bit of algebra shows that

$$\chi^2 = \frac{\sigma^2_{t'_{\rm asc}}}{\sigma^2_{t_{\rm asc}}}\left[\left(\frac{P_{\rm orb} - P_0}{\sigma_{P_{\rm orb}}}\right)^2 - \frac{2n_{\rm orb}\sigma_{P_{\rm orb}}}{\sigma_{t'_{\rm asc}}}\left(\frac{t'_{\rm asc} - t'_{\rm asc,0}}{\sigma_{t'_{\rm asc}}}\right)\left(\frac{P_{\rm orb} - P_0}{\sigma_{P_{\rm orb}}}\right) + \left(\frac{t'_{\rm asc} - t'_{\rm asc,0}}{\sigma_{t'_{\rm asc}}}\right)^2\right]. \tag{3.8}$$

†† The actual values for O3 including duty cycle, noise weighting and antenna patterns, and using the exact form of the metric, will be slightly different, but we will use the values above for illustration in this paper.

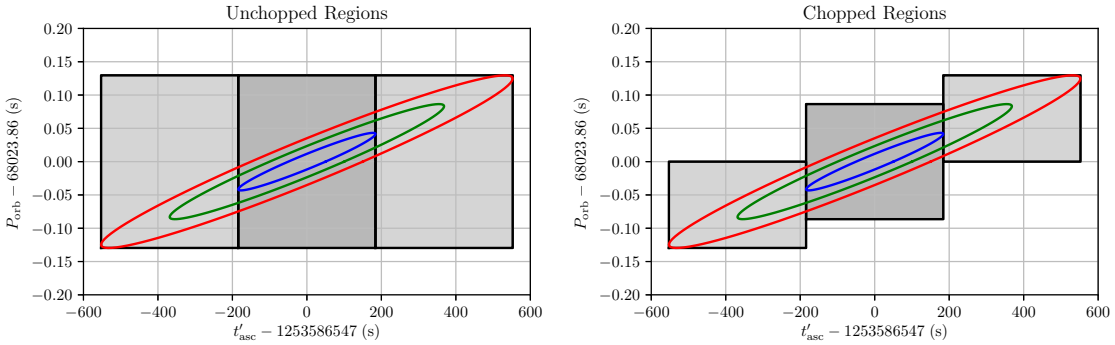

Figure 3. Left: rectangular search region boundaries cover the entire $t'_{\mathrm{asc}}, P_{\mathrm{orb}}$ region, as used in O1 CrossCorr [12]. The darker region bounded by one-sigma in the t_{asc} direction indicates a higher likelihood of finding the signal in that region of parameter space. Right: Rectangular regions more closely concentrated on the uncertainty ellipses than the O1 search regions. The areas in $t'_{\mathrm{asc}}, P_{\mathrm{orb}}$ have been "chopped" to eliminate extra parameter space area where a signal is not likely to be found.

In previous searches, rectangular boundaries have been used in all coordinate directions in the parameter space. For the O1 search, these regions covered out to 3σ of the marginal priors on t'_{asc} and P_{orb}, as shown in Figure 1 of [12]. If we use a similar approach in O3 (figure 3, left panel), the search regions cover a large area of $t'_{\mathrm{asc}}, P_{\mathrm{orb}}$ parameter space with negligible prior probability. Since the middle third of the t'_{asc} range is searched separately (at a higher coherence time T_{max}, since the prior probability density is higher there), a simple approach can reduce the over-coverage of the search region. The P_{orb} search range is different for each of the rectangular regions covering different ranges of t'_{asc}, discarding regions in which the prior $\chi^2 \gtrsim 3^2$. These "chopped" regions are shown in the right panel of figure 3. The chopped regions can be achieved without significant modification to the previously existing search code.

To further improve the efficiency of the parameter space coverage, we can define an elliptical boundary function which sets the range of P_{orb} continuously as a function of t'_{asc}. This function can be used in the `LatticeTiling` module to restrict template placement to those needed to cover the prior ellipse corresponding to $\chi^2 \leq k^2$ for a particular k:

$$\frac{P_{\mathrm{orb}} - P_0}{\sigma_{P_{\mathrm{orb}}}} \in \frac{n_{\mathrm{orb}}\sigma_{P_{\mathrm{orb}}}}{\sigma_{t'_{\mathrm{asc}}}} \left(\frac{t'_{\mathrm{asc}} - t'_{\mathrm{asc},0}}{\sigma_{t'_{\mathrm{asc}}}}\right) \pm \frac{\sigma_{t_{\mathrm{asc}}}}{\sigma_{t'_{\mathrm{asc}}}} \sqrt{k^2 - \left(\frac{t'_{\mathrm{asc}} - t'_{\mathrm{asc},0}}{\sigma_{t'_{\mathrm{asc}}}}\right)^2}. \tag{3.9}$$

This is used to define a search region, together with a constant boundary on t'_{asc}:

$$t'_{\mathrm{asc},0} - k\sigma_{t'_{\mathrm{asc}}} \leq t'_{\mathrm{asc,min}} \leq t'_{\mathrm{asc}} \leq t'_{\mathrm{asc,max}} \leq t'_{\mathrm{asc},0} + k\sigma_{t'_{\mathrm{asc}}}, \tag{3.10}$$

and illustrated in the left panel of figure 4. Note that we choose $k = 3.3$ rather than $k = 3$ as the boundary, since the former encloses 99.6% of the prior probability, while the latter would enclose only 98.9%.

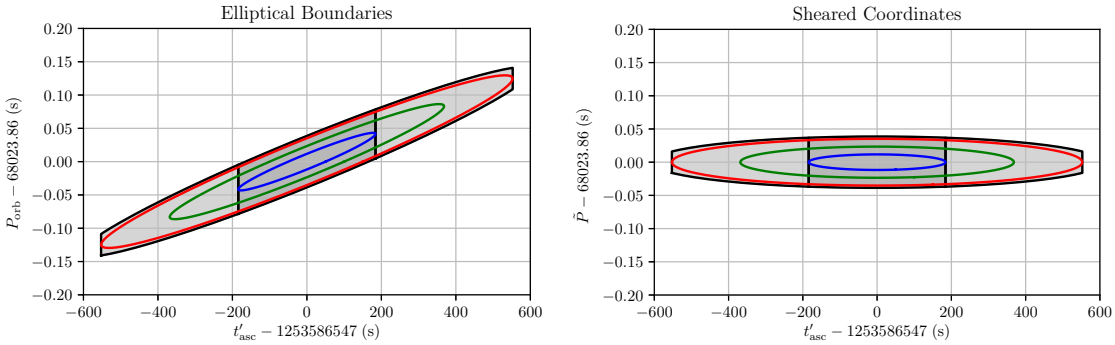

Figure 4. Left: Search region boundaries in $t'_{\mathrm{asc}}, P_{\mathrm{orb}}$ now define a boundary function, where the range of P_{orb} is computed as a function of t'_{asc}, defined in (3.9). We choose $k = 3.3$ for the boundary. For reference, the colored ellipses are level surfaces at 1σ, 2σ, and 3σ, i.e., $\chi^2 = 1^2, 2^2$, and 3^3, as defined in (3.8). The darker shaded region lies within $\pm 1\sigma$ of the marginal distribution on t'_{asc}, where the signal is more likely to be found, and the lighter-shaded region is between $\pm 3\sigma$ and $\pm 1\sigma$. Dividing up the search regions based on t'_{asc} rather than χ^2 is more efficient since t'_{asc} is always resolved, and P_{orb} may not be. Right: the same search regions in a "sheared" set of coordinates $t'_{\mathrm{asc}}, \tilde{P}$, where $\tilde{P}$ is a linear combination of P_{orb} and t'_{asc}, defined in (3.11), which aligns the constant-χ^2 ellipses with the coordinate axes.

3.3. Sheared Coordinates

The joint prior uncertainty in t'_{asc},P_{orb} space complicates the placement of lattice points neatly in coordinate directions. The fact that the semimajor axis of the uncertainty ellipses does not lie in a coordinate direction forces rows of lattice points calculated from a diagonal metric to be placed over a complicated area in parameter space. This difficulty in placing points is illustrated in the figures shown in section 4. A coordinate transformation can be performed that preserves the diagonal metric and shears the coordinates from $(t'_{\mathrm{asc}}, P_{\mathrm{orb}})$ to $(t'_{\mathrm{asc}}, \tilde{P})$, aligning the semimajor axis of the uncertainty ellipses with the coordinate directions, as shown in the right panel of figure 4. The lattice points are then chosen in a straightforward way, before a transformation is then performed back to the physical coordinates. In particular, this simplifies the question of whether multiple templates are necessary to cover the period direction. Looking at the right panel of figure 2 or the left panel of figure 4, we see that the marginal uncertainty in P_{orb} is considerably larger than the conditional uncertainty at a particular value of t'_{asc}. Changing coordinates to $\tilde{P}$, which is observationally uncorrelated with t'_{asc}, allows us to cover a range of period values corresponding to this smaller marginal uncertainty.

We can accomplish this coordinate transformation by subtracting from P_{orb} the centerline of the observational uncertainty ellipse and defining

$$\tilde{P} = P_{\mathrm{orb}} - \frac{n_{\mathrm{orb}}\sigma_{P_{\mathrm{orb}}}}{\sigma_{t'_{\mathrm{asc}}}} \left(\frac{t'_{\mathrm{asc}} - t'_{\mathrm{asc},0}}{\sigma_{t'_{\mathrm{asc}}}} \right) \sigma_{P_{\mathrm{orb}}}, \tag{3.11}$$

so that

$$\chi^2 = \left(\frac{t'_{\rm asc} - t'_{\rm asc,0}}{\sigma_{t'_{\rm asc}}}\right)^2 + \left(\frac{\tilde{P} - P_0}{\sigma_{\tilde{P}}}\right)^2, \tag{3.12}$$

and the priors on $t'_{\rm asc}$ and $\tilde{P}$ are once again independent Gaussians. Note that

$$\sigma_{\tilde{P}} = \left(\frac{\sigma_{t_{\rm asc}}}{\sigma_{t'_{\rm asc}}}\right)\sigma_{P_{\rm orb}}, \tag{3.13}$$

so the area of the uncertainty ellipse is the same in all three sets of coordinates: $(t_{\rm asc}, P_{\rm orb})$, $(t'_{\rm asc}, P_{\rm orb})$, and $(t'_{\rm asc}, \tilde{P})$.

This transformation affects the metric:

$$\tilde{g}_{t'_{\rm asc}t'_{\rm asc}} = g'_{t'_{\rm asc}t'_{\rm asc}} + 2\left(\frac{\partial P_{\rm orb}}{\partial t'_{\rm asc}}\right)_{\tilde{P}} g'_{t'_{\rm asc}P_{\rm orb}} + \left(\frac{\partial P_{\rm orb}}{\partial t'_{\rm asc}}\right)^2_{\tilde{P}} g'_{P_{\rm orb}P_{\rm orb}}, \tag{3.14a}$$

$$\tilde{g}_{t'_{\rm asc}\tilde{P}} = g'_{t'_{\rm asc}P_{\rm orb}} + \left(\frac{\partial P_{\rm orb}}{\partial t'_{\rm asc}}\right)_{\tilde{P}} g'_{P_{\rm orb}P_{\rm orb}}, \tag{3.14b}$$

$$\tilde{g}_{\tilde{P}\tilde{P}} = g'_{P_{\rm orb}P_{\rm orb}}, \tag{3.14c}$$

where

$$\left(\frac{\partial P_{\rm orb}}{\partial t'_{\rm asc}}\right)_{\tilde{P}} = n_{\rm orb}\left(\frac{\sigma_{P_{\rm orb}}}{\sigma_{t'_{\rm asc}}}\right)^2 = n_{\rm orb}\frac{\sigma^2_{P_{\rm orb}}}{\sigma^2_{t_{\rm asc}} + n^2_{\rm orb}\sigma^2_{P_{\rm orb}}}. \tag{3.15}$$

In order to make the metric as close to diagonal as possible in these coordinates, we should choose a different $n_{\rm orb}$ from that defined in (3.5). Instead we make (substituting (3.4), (3.15), and (3.2c) into (3.14b) and looking at the most likely point $t_{\rm asc} = t_{\rm asc,0}, P_{\rm orb} = P_0$)

$$\tilde{g}_{t'_{\rm asc}\tilde{P}} \approx \left(n_{\rm orb} - \frac{\mu_{\rm obs} - t_{\rm asc,0}}{P_0} + n_{\rm orb}\left(\frac{\sigma_{P_{\rm orb}}}{\sigma_{t'_{\rm asc}}}\right)^2 \frac{(t'_{\rm asc,0} - \mu_{\rm obs})^2 + \sigma^2_{\rm obs}}{P_0^2}\right) g'_{t'_{\rm asc}t'_{\rm asc}} \tag{3.16}$$

close to zero. If we set this to zero and solve algebraically for $n_{\rm orb}$, we get

$$n_{\rm orb} \approx \frac{\mu_{\rm obs} - t_{\rm asc,0}}{P_0}\left(1 + \left(\frac{\sigma_{P_{\rm orb}}}{\sigma_{t'_{\rm asc}}}\right)^2 \frac{(t'_{\rm asc,0} - \mu_{\rm obs})^2 + \sigma^2_{\rm obs}}{P_0^2}\right)^{-1}. \tag{3.17}$$

Since the definitions of $t'_{\rm asc,0}$ and $\sigma_{t'_{\rm asc}}$ depend on $n_{\rm orb}$ as well, we need to solve iteratively for the optimal $n_{\rm orb}$ to minimize the metric correlation in these sheared coordinates. This converges quickly, giving, for the reference values used in this paper, $n_{\rm orb} = 4108$, corresponding to $t'_{\rm asc,0}$ = GPS 1253858643 ≡2019–Sep–30 06:03:45 UTC.† With this choice, we have coordinates $t'_{\rm asc}$ and $\tilde{P}$ with no prior correlation and negligible correlation in the search metric.

† Again, the actual best value using the data with gaps, antenna patterns and variable noise level, as well as the exact metric, will be slightly different, but the relationship between the choices of $n_{\rm orb}$ optimized for sheared and unsheared coordinates is illustrative.

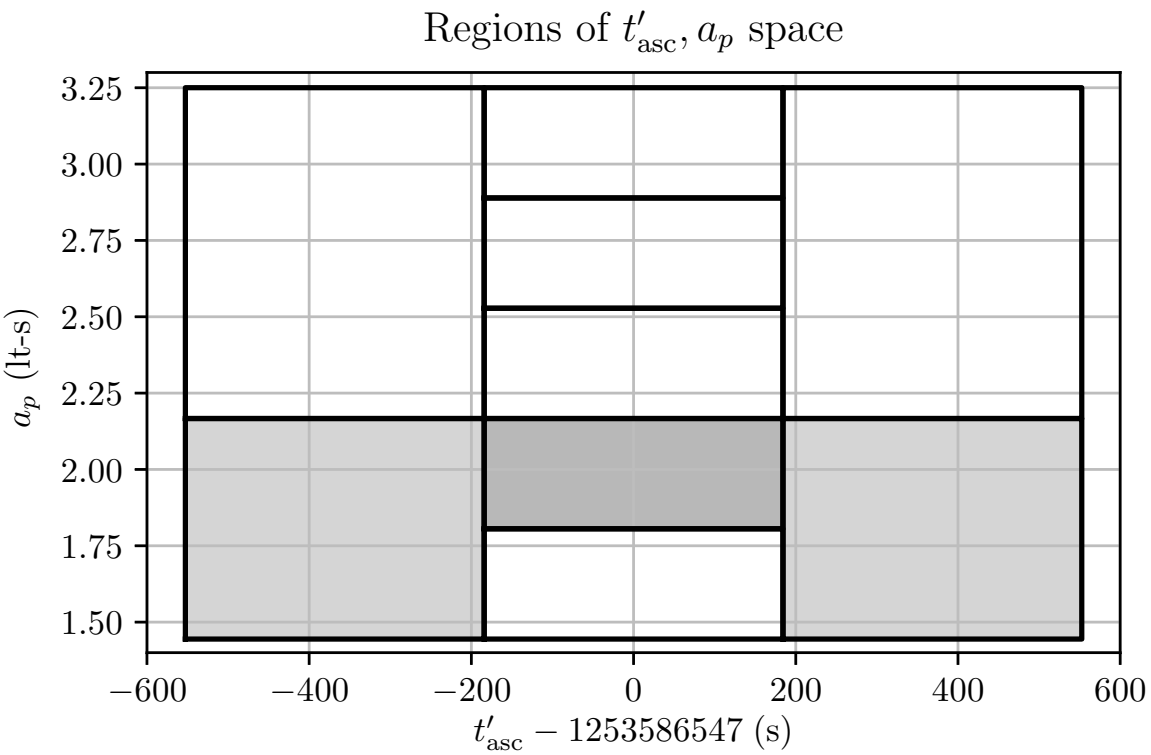

Figure 5. Cells of t'_{asc}, a_p space for sample lattice construction. Each rectangular cell has its own coherence time T_{max}, corresponding to a coherence time used in [12], and we construct a lattice in each of these cells. The range of orbital period values for each cell is a function of t'_{asc}, as illustrated in figure 3 or figure 4. We construct the lattice in all 9 of these regions, and include the template counts in the computing cost estimate. For the three shaded regions, we also include the templates in the corresponding $t'_{\mathrm{asc}}, P_{\mathrm{orb}}$ or $t'_{\mathrm{asc}}, \tilde{P}$ plot of the lattice.

4. Example Lattices and Results

To quantify the reduction in number of search templates and computing costs at a given mismatch, we construct sample lattices of each type for a variety of representative regions in parameter space. For each choice of coordinate system and lattice type, we construct 9×14 lattices, corresponding to the nine regions of orbital parameter space $(t'_{\mathrm{asc}}, P_{\mathrm{orb}}, a_p)$ or $(t'_{\mathrm{asc}}, \tilde{P}, a_p)$ shown in figure 5 and fourteen frequency bands beginning at 25 Hz and ending at 2000 Hz. Each of these regions has its own T_{max} value taken from the search in [12]. In that search, the frequency f_0 was split into ranges of width 0.05 Hz, and a search job covered that range of frequencies along with one of the orbital parameter space regions. Rather than constructing the full set of 9×39500 lattices covering all the bands from 25 Hz to 2000 Hz, we choose one 0.0005 Hz range from the middle of each band, construct the nine lattices (one for each orbital parameter cell) corresponding to that range, and scale up the number of templates by the number of such ranges in the band. Since the computing cost scales roughly with the number of

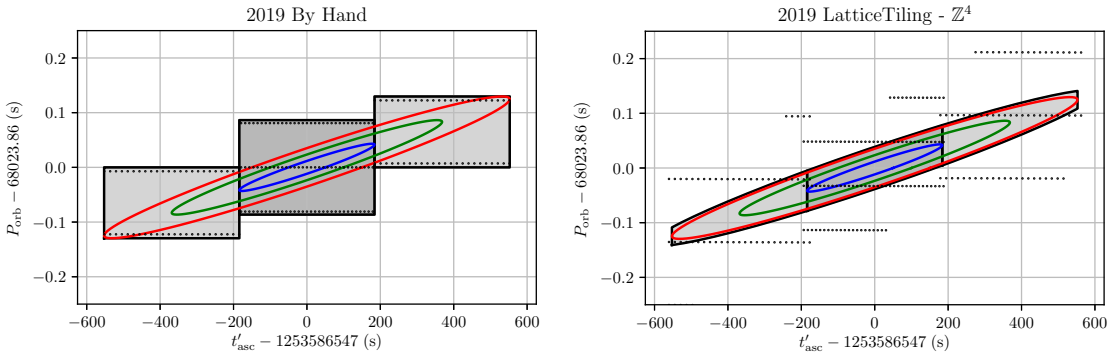

Figure 6. Left: Lattice setups that most closely resemble what was done in previous CrossCorr searches, but with the inclusion of "chopped" search regions in $t'_{\mathrm{asc}}, P_{\mathrm{orb}}$. The lattice points are placed by computing the spacing in each direction given the mismatch and the metric (section 2.2). Points are then placed to cover the uncertainty ellipses in each of the three rectangular search regions, where the darker region bounded by one-sigma in t'_{asc} represents the region where we are most likely to find a signal if it is present.The total number of templates for this setup is 1.060×10^{12}. Right: Implementing the elliptical boundary function and using `LatticeTiling` to place a cubic lattice changes where the points are placed. The total number of templates for this setup is 1.682×10^{12}.

templates times the number of SFT pairs, we approximate the computing cost for each band i and cell c as $N_{ic}^{\mathrm{pair}} N_{ic}^{\mathrm{tmplt}}$. We estimate the number of SFT pairs for data from N_{det} detectors as in [10] by

$$N_{ic}^{\mathrm{pair}} \approx N_{\mathrm{det}}^2 \frac{T_{\mathrm{obs}} T_{ic}^{\mathrm{max}}}{T_i^{\mathrm{SFT}}}, \tag{4.1}$$

where we show explicity that the SFT duration depends on the frequency band i while the coherence time depends on the frequency band i and orbital parameter space cell c. Note that this is an overestimate of the absolute number of pairs, because we computed the T_{obs} using the start and end times of the two parts of O3 rather than an actual set of data segments reflecting the true duty cycle. In addition to the total computing cost $\sum_i \sum_c N_{ic}^{\mathrm{pair}} N_{ic}^{\mathrm{tmplt}}$ for each lattice, we also plot the lattice points projected onto the $t'_{\mathrm{asc}}, P_{\mathrm{orb}}$ or $t'_{\mathrm{asc}}, \tilde{P}$ plane, limiting attention for the plots to the shaded cells in figure 5.

Figure 6 shows two implementations of cubic ($\mathbb{Z}^4$) lattices, both using the original by hand method described in section 2.2 and using the `LatticeTiling` module. The main difference between the two methods is in how they handle the boundaries of the elliptical search region. The by-hand method uses the chopped regions illustrated in the right panel of figure 3, while the `LatticeTiling` method uses the elliptical boundaries of figure 4. Note that while `LatticeTiling` uses a smaller region of parameter space, it actually requires more templates (a total over the whole parameter space of 1.682×10^{12} versus 1.060×10^{12} for the by-hand method) for the cubic ($\mathbb{Z}^4$) lattice because of its conservative approach to covering the boundaries. Ordinarily this would be a small effect, but since only two or three templates are required in the P_{orb} direction, it is

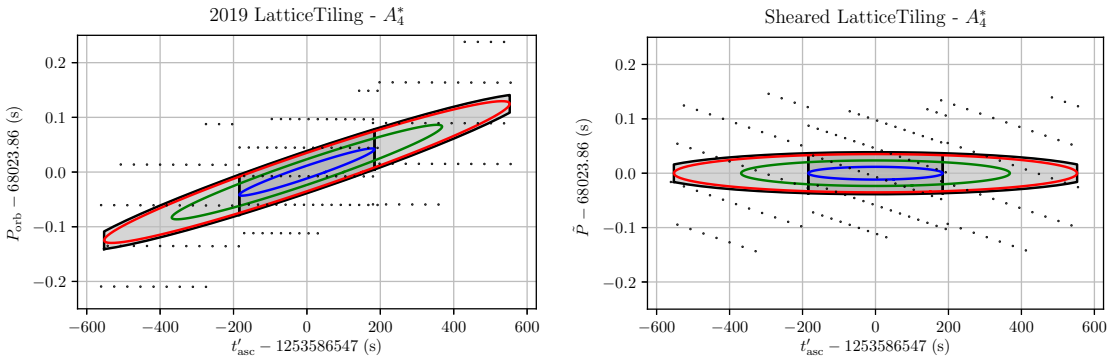

Figure 7. Left: Implementing the elliptical boundary function discussed in section 3.2 requires the use of `LatticeTiling` [8] for template placement. This setup shows a use of an A_4^* lattice with the elliptical boundary function and lattice template points placed by `LatticeTiling`. Note that the density of templates is increased in the one-sigma region. The total number of templates used here (across four-dimensional parameter space) is 6.006×10^{11}. Right: After performing the shearing transformation discussed in section 3.3, we use `LatticeTiling` to place templates in $t'_{\mathrm{asc}}, \tilde{P}$ space. This figure shows an A_4^* lattice over the sheared uncertainty ellipses using the elliptical boundary function. Note that the primary axes of the uncertainty ellipses are aligned with the coordinate axes in this area-preserving transformation, and that template density is again greater in the one-sigma region. The total number of templates here is 5.931×10^{11}.

significant in this case, which motivates the special handing of the P_{orb} coordinate which follows.

If we change the lattice from $\mathbb{Z}^4$ to A_4^*, we obtain the lattice shown in the left panel of figure 7. The use of a more efficient lattice has reduced the total number of templates to 6.006×10^{11}, but we can see from the figure that the templates extend well beyond the boundaries of the search region. In the right panel, we construct the lattice in the sheared coordinates $t'_{\mathrm{asc}}, \tilde{P}$ defined in section 3.3, which simplifies the boundaries of the search region, but produces lattices with comparable numbers of templates (5.931×10^{11} total). In these coordinates, the mismatch metric has a non-negligible off-diagonal component $\tilde{g}_{t'_{\mathrm{asc}}\tilde{P}}$, so the template lattice is constructed using a basis which looks "slanted" in these coordinates.

We can make the metric approximately diagonal, as described in section 3.3 by choosing a different value of n_{orb} derived from (3.17); for the example considered, this means changing n_{orb} from 4104 to 4108. The resulting lattice is shown in the left panel of figure 8. Note that the total number of templates is comparable to the other A_4^* lattices, a total 5.918×10^{11} across the whole parameter space. The fact that all of the A_4^* lattices have comparable numbers of templates indicates that the `LatticeTiling` module is behaving consistently, even when the coordinates being used have metric correlations or oddly-shaped boundaries. However, it is clearly not taking full advantage of the narrow range of plausible $\tilde{P}$ values. The underlying issue is that `LatticeTiling`, by the nature of its boundary-covering algorithm [8], uses a minimum of two templates in a coordinate

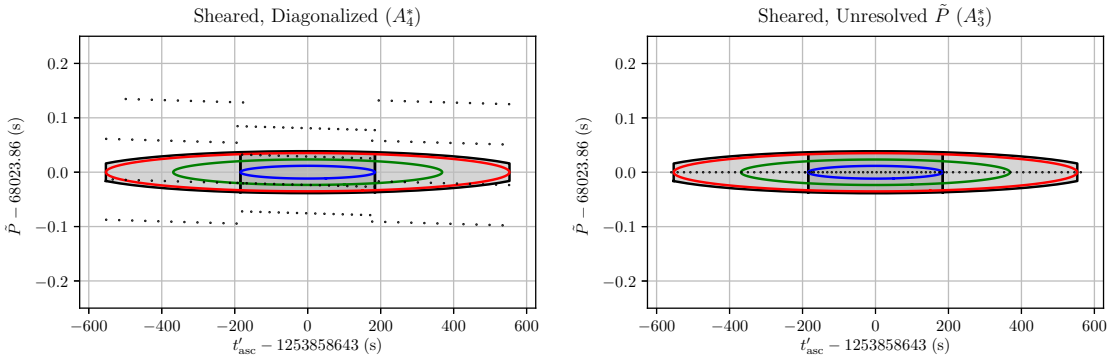

Figure 8. Left: The number of orbits used to propagate $t_{\rm asc}$ into 2019 coordinates was chosen based on what would diagonalize the standard/ 2019 metric. This is what introduced the slant in the template rows seen in figure 7. Choosing a different $n_{\rm orb}$ eliminates the slant, to produce this figure, showing a lattice covering using an A_4^* lattice, `LatticeTiling` to place the templates, and the sheared coordinates with a diagonalized metric to align the uncertainty ellipses with the primary axes of the parameter space. The total number of templates is 5.918×10^{11}. Right: Noticing that the spacing between template rows in $t'_{\rm asc}, \tilde{P}$ seemed to be larger than the cross-section of the uncertainty ellipse, we perform a calculation described in section 3.3 to determine whether the orbital period needs to be resolved in the sheared coordinates. After finding that it does not, for our search, we fix $\tilde{P} = P_0$, forcing `LatticeTiling` to place a single row of lattice templates along the centerline of the sheared uncertainty ellipse. Note that the template density is still greater in the one-sigma region. Here, the total number of templates is 3.867×10^{11}, our best result for template count and an improvement from the original setup by a factor of about 3.

direction, even if a single template would be sufficient to cover the space at the desired minimum mismatch.

The change to $t'_{\rm asc}, \tilde{P}$ coordinates, in which both the prior uncertainty and mismatch metric are approximately uncorrelated, allows us to take advantage of the small prior uncertainty in $\tilde{P}$. If we limit attention to lattices with all their templates on the hypersurface $\tilde{P} = P_0$, the mismatch between a signal with parameters $\{\lambda_i^s\}$ and a template point $\{\lambda_i\}$ will be

$$\mu = \tilde{g}_{\tilde{P}\tilde{P}}(\tilde{P}^s - P_0)^2 + 2\sum_\alpha g_{\alpha\tilde{P}}(\lambda_\alpha^s - \lambda_\alpha)(\tilde{P}^s - P_0) + \sum_\alpha \sum_\beta g_{\alpha\beta}(\lambda_\alpha^s - \lambda_\alpha)(\lambda_\alpha^s - \lambda_\alpha), \quad (4.2)$$

where $\{\lambda_\alpha\} = \{t'_{\rm asc}, f_0, a_p\}$ are the other three coordinates of the parameter space and $\tilde{g}_{\tilde{P}\tilde{P}}$ is the sheared metric element for orbital period. If we assume the metric is approximately diagonal, this becomes

$$\mu \approx \tilde{g}_{\tilde{P}\tilde{P}}(\tilde{P}^s - P_0)^2 + \mu^\parallel, \quad (4.3)$$

As shown in the Appendix, the general expression is

$$\mu \approx \frac{(\tilde{P}^s - P_0)^2}{\tilde{g}^{\tilde{P}\tilde{P}}} + \mu^\parallel \quad (4.4)$$

Table 1. Comparing Estimates of Raw Computing Cost: we display the chosen coordinates, the number of orbits needed to propogate $t'_{\rm asc}$ to obtain a diagonal metric, the type of search region boundary used, and the type of lattice structure. Then, we show the number of templates required to cover all of parameter space using a given lattice and estimate the computing cost by multiplying the number of lattice templates by the number of SFT pairs.

Coordinates	$n_{\rm orb}$	Boundary	Type	$\sum_{i,c} N_{ic}^{\rm tmplt}$	$\sum_{i,c} N_{ic}^{\rm pair} N_{ic}^{\rm tmplt}$
$t'_{\rm asc}, P_{\rm orb}, a_p, f_0$	4104	Chopped	$\mathbb{Z}^4$	1.060×10^{12}	1.434×10^{18}
$t'_{\rm asc}, P_{\rm orb}, a_p, f_0$	4104	Elliptical	$\mathbb{Z}^4$	1.682×10^{12}	2.085×10^{18}
$t'_{\rm asc}, P_{\rm orb}, a_p, f_0$	4104	Elliptical	A_4^*	6.006×10^{11}	7.439×10^{17}
$t'_{\rm asc}, \tilde{P}, a_p, f_0$	4104	Elliptical	A_4^*	5.931×10^{11}	7.352×10^{17}
$t'_{\rm asc}, \tilde{P}, a_p, f_0$	4108	Elliptical	A_4^*	5.918×10^{11}	7.316×10^{17}
$t'_{\rm asc}, a_p, f_0; \tilde{P} = P_0$	4108	Elliptical	A_3^*	3.867×10^{11}	4.928×10^{17}
Same with reallocated mismatch				3.431×10^{11}	4.483×10^{17}

Since the prior uncertainty ellipse with $\chi^2 \leq k^2$ (see (3.12) and figure 4) has $(\tilde{P} - P_0)^2 \leq k^2 \sigma_{\tilde{P}}^2$, we can obtain a lattice with $\mu < \mu_{\rm max}$ everywhere if we construct a three-dimensional lattice with

$$\mu_{\rm max}^{\parallel} \leq \mu_{\rm max} - \frac{k^2 \sigma_{\tilde{P}}^2}{\tilde{g}^{\tilde{P}\tilde{P}}} \tag{4.5}$$

A conservative approach is to allocate a mismatch of $\frac{\mu_{\rm max}}{4}$ to the $\tilde{P}$ direction and $\frac{3\mu_{\rm max}}{4}$ to the other three directions. Then we proceed as follows:

- If $\frac{k^2 \sigma_{\tilde{P}}^2}{\tilde{g}^{\tilde{P}\tilde{P}}} > \frac{\mu_{\rm max}}{4}$, we construct an A_4^* lattice covering the full four-dimensional parameter space as usual.
- If $\frac{k^2 \sigma_{\tilde{P}}^2}{\tilde{g}^{\tilde{P}\tilde{P}}} \leq \frac{\mu_{\rm max}}{4}$, we construct a three-dimensional A_3^* lattice with maximum mismatch $\mu_{\rm max}^{\parallel} = \frac{3\mu_{\rm max}}{4}$ and $\tilde{P} = P_0$ at all lattice points. (In `LatticeTiling` we accomplish this by setting the search region to have zero width in the $\tilde{P}$ direction.)

Following this approach produces the most efficient lattice, with 3.867×10^{11} total templates, illustrated in the right panel of figure 8. A slightly more agressive approach would be to "reallocate" any unused mismatch if $\frac{k^2 \sigma_{\tilde{P}}^2}{\tilde{g}^{\tilde{P}\tilde{P}}} < \frac{\mu_{\rm max}}{4}$, and set to the maxiumum mismatch of the A_3^* lattice to

$$\mu_{\rm max}^{\parallel} = \mu_{\rm max} - \frac{k^2 \sigma_{\tilde{P}}^2}{\tilde{g}^{\tilde{P}\tilde{P}}}. \; ; \tag{4.6}$$

This leads to a slightly smaller number of templates (3.431×10^{11}).

The properties of the different lattices are summarized in table 1. In addition to the total number of templates $\sum_{i,c} N_{ic}^{\rm tmplt}$ across all of the parameter space cells, we also show the sum $\sum_{i,c} N_{ic}^{\rm pair} N_{ic}^{\rm tmplt}$ which should roughly scale with the computing cost. Roughly speaking, replacing the by-hand cubic lattice with an A_n^* lattice reduces to overall computing cost by a factor of 2, while enforcing unresolved $\tilde{P}$ when possible

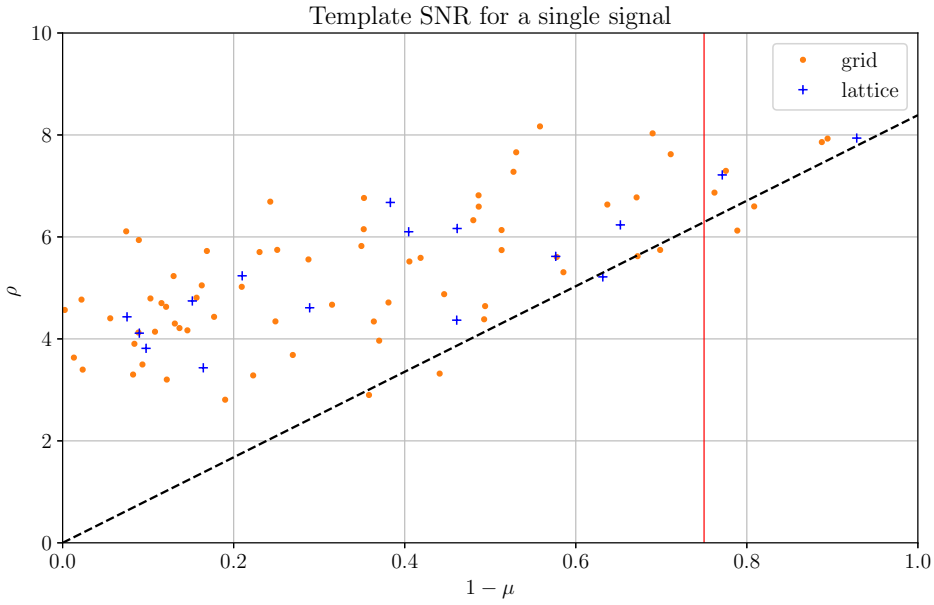

Figure 9. A comparison of the recovered values of SNR ρ at various parameter space points for a representative simulated signal. We compare the "by hand" hypercubic grid in the original coordinates $\{t'_{\mathrm{asc}}, P_{\mathrm{orb}}, a_p, f_0\}$ to an A_3^* lattice generated in the coordinates $t'_{\mathrm{asc}}, a_p, f_0$ with the sheared period coordinate $\tilde{P}$ set to P_0. We plot ρ versus the theoretical fraction of $E[\rho]$ from the metric mismatch μ; see (2.1). The dashed line is $\rho = \rho^{\mathrm{targ}}(1-\mu)$. We see that the points approximately follow that line for small mismatch, but the quadratic metric approximation breaks down farther away. Each template placement method is guaranteed to have at least one point with $\mu \leq \mu_{\mathrm{max}} = 0.25$, but since the A_3^* lattice has fewer and more efficiently placed points, its closest point to the true signal parameters is not likely to be as close as the closest point in the grid. Note that the highest ρ value in the grid or lattice need not be the closest to the true signal, due to noise fluctuations.

reduces the cost by a further factor of 1.5, for an overall improvement of a factor of 3 resulting from the enhancements described in this paper.

5. Recovery of Simulated Signals

To test that our improved lattice is as sensitive to signals with unknown parameters as the original grid search method, we generated data containing simulated signals and analyzed them with the lattices and grids considered in section 4. The data generated were Gaussian white noise with a power spectral density of $5.14 \times 10^{-24}\,\mathrm{Hz}^{1/2}$. In order to generate data with similar parameter-space metric components $\{g_{ij}\}$ to the O3 run, but with a smaller volume of data, we generated a set of 24-hour stretches of data separated in time by three weeks, over the entire duration of O3a and also of O3b. It was necessary to do this, rather than generating a single contiguous stretch of data, to obtain a similar σ_{obs}.

The frequencies of the signals were chosen between 300 and 400 Hz, and the

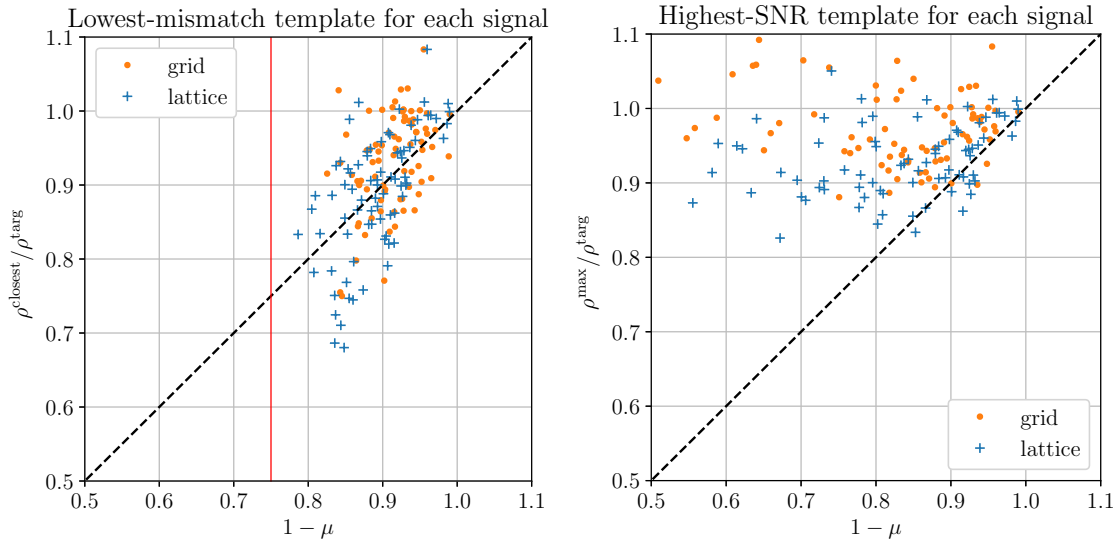

Figure 10. Recovered SNR ρ for each signal as a fraction of targeded SNR ρ^{targ} from searching with the true signal parameters. As in figure 9, we compare the by-hand hypercubic grid to the A_3^* lattice with $\tilde{P} = P_0$. At left is the SNR from the point in the grid or lattice with the smallest mismatch distance to the true signal parameters. At right is the maximum SNR in the lattice or grid, which we see from figure 9 need not be for the template closest to the true signal parameters. For reference, we plot the theoretical curve $\rho/\rho^{\text{targ}} = 1 - \mu$. The lattice does indeed recover the signals with a fractional SNR loss less than the designed maximum mismatch $\mu_{\text{max}} = 0.25$, although the grid, which requires more points to obtain the same maximum mismatch, generally finds points with slightly higher SNR because of this redundancy.

parameters a_p, t_{asc} and P_{orb} were randomly chosen according to the astrophysical priors described in section 3.1, i.e., $a_p \sim \text{Uniform}(1.44\,\text{lt-s}, 3.25\,\text{lt-s})$ and P_{orb} and t_{asc} from a bivariate normal distribution with means $P_0 = 68023.86$ and $t_{\text{asc},0} = 974416624$ GPS (2010–Nov–21 23:16:49 UTC), and standard deviations $\sigma_{P_{\text{orb}}} = 0.043\,\text{s}$ and $\sigma_{t_{\text{asc}}} = 50\,\text{s}$, and t'_{asc} generated by propagating t_{asc} to the middle of the data set, using the P_{orb} value generated for the signal. The initial phase of the GW signal was drawn from a uniform distribution, and the orientation angles of the neutron star spin were drawn from an isotropic distribution. The signal amplitude h_0 was generated to give an estimated expected SNR between 9.0 and 11.0. This amplitude depends on the randomly chosen value of the neutron star inclination, as detailed in [10].

The data were generated using `lalapps_Makefakedata_v5` from the `lalsuite` library [7]. For each signal a targeted search was run using a single template at the injected parameter values $\{\lambda_i^s\}$ to find the actual SNR ρ^{targ} of the signal in the simulated data. We then conducted searches using different template placement methods. These were each conducted in small parameter space regions containing, but not centered on, the point $\{\lambda_i^s\}$ at which the signal was injected. The regions typically covered a search range equivalent to 6 or 7 grid points in the f_0, a_p and t_{asc} directions, and the full prior range in P_{orb} out to $P_0 \pm 3\sigma_{P_{\text{orb}}}$, or the appropriate elliptical boundary. For simplicity

in comparisons, we use the same coherence time $T_{\max} = 2550\,\mathrm{s}$.

We performed the search using the template placement methods described in section 4. Here we report on the comparison between the original grid (or "by hand") method shown in e.g., the left panel of figure 6, and the A_3^* lattice in sheared coordinates with $\tilde{P}$ unresolved, as shown in the last line of table 1. In figure 9 we show, for a single representative signal, the SNR ρ at various points in the grid or lattice of each search. The search methods tend to recover the signal with comparable SNRs although the lattice has fewer "redundant" points close to the true signal parameter values.

For the full ensemble of 80 signals, we can look at the distribution of either the highest SNR ρ in each search (the "best" template), or the ρ of the template in each search with the lowest mismatch μ, which we refer to as the "closest" template. We show in figure 10 the fraction of the targeted-search SNR ρ^{targ} recovered in each search for the closest or best template in each of the two searches considered. We see that the grid method from previous searches, and the improved method with and A_3^* lattice and unresolved $\tilde{P}$ both recover signals with an SNR loss of less than the nominal maximum mismatch of $\mu_{\max} = 0.25$. The grid method generally obtains slightly higher SNRs due to the extra templates, which increase the chance of catching an upward fluctuation or a point significantly closer than the maximum mismatch would indicate. However, the lattice method achieves the specified maximum mismatch with fewer templates and less computing cost.

6. Conclusions

In this paper we have discussed changes to the lattice used in the template-based cross-correlation search for continuous gravitational waves from Scorpius X-1. We detailed the setup of our parameter space and explained how previous searches used lattices in the same parameter space. We then gave four major improvements to improve the lattice setup, using fewer templates for a given computing cost. We first showed that there is a reduction in template count by switching from a hypercubic lattice to an A_n^* lattice in section 2.2. Then, we defined an elliptical boundary function in section 3.2 to improve the shape of the search region in t'_{asc} and P_{orb} to be more focused on the section of parameter space within the prior ellipses. In section 3.3 we defined an area-preserving shearing transformation that aligned the axes of the prior ellipses with the coordinate axes. This simplifies the task of using `LatticeTiling` to place a horizontal row of templates in parameter space. Finally, we compared the cross-section of the prior ellipses in t'_{asc} and $\tilde{P}$ to determine whether $\tilde{P}$ needed to be resolved, and determined that it did not in section 3.3. This allowed us to use an A_3^* lattice, and reduced the template count by an overall factor of ~ 3 compared to the original grid search.

The improvements described in this paper are being incorporated into the LIGO-Virgo-KAGRA data analysis infrastructure [7] and may be used for GW searches in O3 and beyond. Note that the sheared period coordinate $\tilde{P}$ is likely to remain unresolved in future searches, even though the original orbital period P_{orb} would not be. For future

observations with observation epoch $\mu_{\rm obs}$ even further from $t_{\rm asc,0}$, the uncertainty $\sigma_{t'_{\rm asc}}$ in time of ascension will grow, and the uncertainty $\sigma_{\tilde{P}}$ in $\tilde{P}$ will shrink as shown in (3.13). This is unlikely to be offset by finer resolution due to increased coherence time (see (3.1c)) since the sinc function in (3.1b) saturates at $T_{\rm max} \sim 0.7 P_{\rm orb}$. This level was reached using resampling below 180 Hz in [13] but it is unlikely to be practical at higher frequencies.

The reduction in template count allows for a more sensitive search in two ways. First, since the search is computationally limited, reducing the number of templates needed at a given coherence time will allow us to use longer coherence times at the same computational cost. The quantitative benefit of this is somewhat limited, since the cost of searching a three-dimensional parameter space scales, for smaller coherence times, like $T_{\rm max}^4$, and the amplitude sensitivity on h_0 scales as $T_{\rm max}^{1/4}$, so a factor of 3 in computing time nominally translates to a 7% improvement in sensitivity. However, there are other benefits not captured in this number; for instance, the use of the sheared period coordinate and unresolved $\tilde{P}$ allows the parameter space to remain effectively three- rather than four-dimensional. Second, the reduction in template numbers reduces the trials factor on parameter space and therefore the expected loudest background outliers[24], enabling the use of a lower threshold for followups.

Acknowledgments

We wish to thank Chris Messenger, as well as the members of the LIGO Scientific Collaboration and Virgo Collaboration continuous waves group, for useful feedback. KJW, JTW, and JKW were supported by NSF grants PHY-1806824 and PHY-2110460. KW was supported by the Australian Research Council Centre of Excellence for Gravitational Wave Discovery (OzGrav) through project number CE170100004. This paper has been assigned LIGO Document Number LIGO-P2000502-v4.

References

[1] Fomalont E B, Geldzahler B J and Bradshaw C F 2001 Scorpius X-1: The Evolution and Nature of the Twin Compact Radio Lobes *Astrophys. J.* **558** 283–301 (*Preprint* arXiv:astro-ph/0104372)

[2] Steeghs D and Casares J 2002 The Mass Donor of Scorpius X-1 Revealed *Astrophys. J.* **568** 273–278

[3] Watts A L, Krishnan B, Bildsten L and Schutz B F 2008 Detecting gravitational wave emission from the known accreting neutron stars *Monthly Notices of the Royal Astronomical Society* **389** 839–868 ISSN 1365-2966 URL `http://dx.doi.org/10.1111/j.1365-2966.2008.13594.x`

[4] Aasi J *et al* (LIGO Scientific) 2015 Advanced LIGO *Class. Quant. Grav.* **32** 074001 (*Preprint* arXiv:1411.4547)

[5] Acernese F *et al* (VIRGO) 2015 Advanced Virgo: a second-generation interferometric gravitational wave detector *Class. Quant. Grav.* **32** 024001 (*Preprint* arXiv:1408.3978)

[6] Aso Y, Michimura Y, Somiya K, Ando M, Miyakawa O, Sekiguchi T, Tatsumi D and Yamamoto H (The KAGRA Collaboration) 2013 Interferometer design of the kagra gravitational wave detector *Phys. Rev. D* **88**(4) 043007 URL `https://link.aps.org/doi/10.1103/PhysRevD.88.043007`

[7] LIGO Scientific Collaboration 2018 LIGO Algorithm Library - LALSuite free software (GPL)

[8] Wette K 2014 Lattice template placement for coherent all-sky searches for gravitational-wave pulsars *Phys. Rev. D.* **90**(12) 122010
[9] Dhurandhar S, Krishnan B, Mukhopadhyay H and Whelan J T 2008 The cross-correlation search for periodic gravitational waves *Phys. Rev. D.* **77** 082001
[10] Whelan J T, Sundaresan S, Zhang Y and Peiris P 2015 Model-Based Cross-Correlation Search for Gravitational Waves from Scorpius X-1 *Phys. Rev. D.* **91** 102005
[11] Messenger C, Bulten H J, Crowder S G, Dergachev V, Galloway D K, Goetz E, Jonker R J G, Lasky P D, Meadors G D, Melatos A, Premachandra S, Riles K, Sammut L, Thrane E H, Whelan J T and Zhang Y 2015 Gravitational waves from Scorpius X-1: A comparison of search methods and prospects for detection with advanced detectors *Phys. Rev. D.* **92** 023006
[12] Abbott B P *et al* (LIGO Scientific Collaboration and Virgo Collaboration) 2017 Upper Limits on Gravitational Waves from Scorpius X-1 from a Model-based Cross-correlation Search in Advanced LIGO Data *The Astrophysical Journal* **847** 47
[13] Zhang Y, Papa M A, Krishnan B and Watts A L 2021 Search for Continuous Gravitational Waves from Scorpius X-1 in LIGO O2 Data *Astrophys. J. Lett.* **906** L14
[14] Meadors G D, Krishnan B, Papa M A r, Whelan J T and Zhang Y 2018 Resampling to accelerate cross-correlation searches for continuous gravitational waves from binary systems *Phys. Rev. D.* **97** 044017 (*Preprint* arXiv:1712.06515)
[15] Conway J and Sloane N 1998 *Sphere Packings, Lattices and Groups* Grundlehren der mathematischen Wissenschaften (Springer New York) ISBN 9780387985855
[16] Prix R 2007 Template-based searches for gravitational waves: Efficient lattice covering of flat parameter spaces *Class. Quant. Grav.* **24** S481–S490
[17] Allen B and Shoom A A 2021 Template banks based on $\mathbb{Z}^n$ and A_n^* lattices *arXiv e-prints* arXiv:2102.11631 (*Preprint* arXiv:2102.11631)
[18] Mukherjee A, Messenger C and Riles K 2018 Accretion-induced spin-wandering effects on the neutron star in Scorpius X-1: Implications for continuous gravitational wave searches *Phys. Rev. D.* **97** 043016 (*Preprint* arXiv:1710.06185)
[19] Skrutskie M F *et al* 2006 The Two Micron All Sky Survey (2MASS) *The Astronomical Journal* **131** 1163–1183
[20] Leaci P and Prix R 2015 Directed searches for continuous gravitational waves from binary systems: parameter-space metrics and optimal Scorpius X-1 sensitivity *Phys. Rev. D.* **91** 102003 (*Preprint* arXiv:1502.00914)
[21] Galloway D K, Premachandra S, Steeghs D, Marsh T, Casares J and Cornelisse R 2014 Precision Ephemerides for Gravitational-wave Searches. I. Sco X-1 *Astrophys. J.* **781** 14 (*Preprint* arXiv: 1311.6246)
[22] Wang L, Steeghs D, Galloway D K, Marsh T and Casares J 2018 Precision Ephemerides for Gravitational-wave Searches - III. Revised system parameters of Sco X-1 *Mon. Not. R. Astron. Soc.* **478** 5174–5183
[23] Abbott B P *et al* 2020 Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA *Living Reviews in Relativity* **23** URL `https://link.springer.com/article/10.1007/s41114-020-00026-9`
[24] Tenorio R, Modafferi L M, Keitel D and Sintes A M 2021 Empirically estimating the distribution of the loudest candidate from a gravitational-wave search *arXiv e-prints* arXiv:2111.12032 (*Preprint* arXiv:2111.12032)

Appendix A. Fixing the Sheared Period Coordinate

Consider how we handle the mismatch when $\tilde{P}$ is underresolved. Let the search region be contained within the range $P_0 - \Delta\tilde{P} \leq \tilde{P} \leq P_0 + \Delta\tilde{P}$ and let $\{\lambda_\alpha\} = \{f_0, a_p, t'_{\rm asc}\}$ be the other three search coordinates. If we construct a template lattice in $\{\lambda_\alpha\}$ with

a maximum mismatch $\mu^{\parallel}_{\rm max}$, we can ask what is the mismatch between a point in that lattice and a point on the $\tilde{P}$ boundary, with $\tilde{P} = P_0 + \Delta\tilde{P}$. If $\Delta\lambda_\alpha$ is separation from a grid point, the total mismatch will be

$$\mu = \tilde{g}_{\tilde{P}\tilde{P}}(\Delta\tilde{P})^2 + 2\sum_\alpha g_{\alpha\tilde{P}}(\Delta\lambda_\alpha)(\Delta\tilde{P}) + \underbrace{\sum_\alpha \sum_\beta g_{\alpha\beta}(\Delta\lambda_\alpha)(\Delta\lambda_\beta)}_{\mu^{\parallel}} . \tag{A.1}$$

If the metric is approximately diagonal, this becomes

$$\mu = \tilde{g}_{\tilde{P}\tilde{P}}(\Delta\tilde{P})^2 + \mu^{\parallel}. \tag{A.2}$$

One conservative approach is to say that as long as $\tilde{g}_{\tilde{P}\tilde{P}}(\Delta\tilde{P})^2 < \frac{\mu_{\rm max}}{4}$, we will set $\tilde{P}$ to P_0 and define a template lattice in the other three coordinates with $\mu^{\parallel}_{\rm max} = \frac{3}{4}\mu_{\rm max}$. More generally, we could choose

$$\mu^{\parallel}_{\rm max} = \mu_{\rm max} - \tilde{g}_{\tilde{P}\tilde{P}}(\Delta\tilde{P})^2 \tag{A.3}$$

which will work as long as $\tilde{g}_{\tilde{P}\tilde{P}}(\Delta\tilde{P})^2 < \mu_{\rm max}$.

In general, though, the metric might not be diagonal, and in particular $\tilde{g}_{t'_{\rm asc}\tilde{P}}$ might be non-negligible. To see how the mismatch for a point on the $\tilde{P}$ boundary changes, consider the case of a two-dimensional lattice in $t'_{\rm asc}$ and $\tilde{P}$, so that the mismatch is

$$\mu = \tilde{g}_{\tilde{P}\tilde{P}}(\Delta\tilde{P})^2 + 2\tilde{g}_{t'_{\rm asc}\tilde{P}}(\Delta t'_{\rm asc})(\Delta\tilde{P}) + \tilde{g}_{t'_{\rm asc}t'_{\rm asc}}(\Delta t'_{\rm asc})^2. \tag{A.4}$$

Suppose the spacing in the $t'_{\rm asc}$ direction is

$$\delta t'_{\rm asc} = 2\sqrt{\frac{\mu^{\parallel}}{\tilde{g}_{t'_{\rm asc}t'_{\rm asc}}}} . \tag{A.5}$$

Consider two adjacent lattice points separated by $\delta t'_{\rm asc}$, and a point in between them, which has $\Delta t'_{\rm asc} = t > 0$ from one point and $\Delta t'_{\rm asc} = t - \delta t'_{\rm asc} < 0$ from the second one. A point with this $t'_{\rm asc}$ value and $\tilde{P} = P_0 + \Delta\tilde{P}$ will have the maximum possible mismatch if it is the same mismatch distance away from the two nearest grid points. This means we're looking for the t which solves

$$\tilde{g}_{\tilde{P}\tilde{P}}(\Delta\tilde{P})^2 + 2\tilde{g}_{t'_{\rm asc}\tilde{P}}(\Delta\tilde{P})t + \tilde{g}_{t'_{\rm asc}t'_{\rm asc}}t^2 = \tilde{g}_{\tilde{P}\tilde{P}}(\Delta\tilde{P})^2 + 2\tilde{g}_{t'_{\rm asc}\tilde{P}}(\Delta\tilde{P})(t-\delta t'_{\rm asc}) + \tilde{g}_{t'_{\rm asc}t'_{\rm asc}}(t-\delta t'_{\rm asc})^2. \tag{A.6}$$

A bit of cancellation gives us

$$0 = -2\tilde{g}_{t'_{\rm asc}\tilde{P}}(\Delta\tilde{P})(\delta t'_{\rm asc}) - 2\tilde{g}_{t'_{\rm asc}t'_{\rm asc}}(\delta t'_{\rm asc})t + \tilde{g}_{t'_{\rm asc}t'_{\rm asc}}(\delta t'_{\rm asc})^2, \tag{A.7}$$

or

$$t = \frac{\delta t'_{\rm asc}}{2} - \frac{\tilde{g}_{t'_{\rm asc}\tilde{P}}}{\tilde{g}_{t'_{\rm asc}t'_{\rm asc}}}(\Delta\tilde{P}) . \tag{A.8}$$

As a quick sanity check, we see that this reduces to $t = \frac{\delta t'_{\text{asc}}}{2}$ when $\tilde{g}_{t'_{\text{asc}}\tilde{P}} = 0$, as we expect. Plugging this back into the mismatch equation gives

$$\begin{aligned}\mu_{\text{max}} &= \tilde{g}_{\tilde{P}\tilde{P}}(\Delta\tilde{P})^2 + 2\tilde{g}_{t'_{\text{asc}}\tilde{P}}(\Delta\tilde{P})\left(\frac{\delta t'_{\text{asc}}}{2} - \frac{\tilde{g}_{t'_{\text{asc}}\tilde{P}}}{\tilde{g}_{t'_{\text{asc}}t'_{\text{asc}}}}(\Delta\tilde{P})\right) + \tilde{g}_{t'_{\text{asc}}t'_{\text{asc}}}\left(\frac{\delta t'_{\text{asc}}}{2} - \frac{\tilde{g}_{t'_{\text{asc}}\tilde{P}}}{\tilde{g}_{t'_{\text{asc}}t'_{\text{asc}}}}(\Delta\tilde{P})\right)^2 \\ &= \tilde{g}_{t'_{\text{asc}}t'_{\text{asc}}}\left(\frac{\delta t'_{\text{asc}}}{2}\right)^2 + \frac{\tilde{g}_{\tilde{P}\tilde{P}}\tilde{g}_{t'_{\text{asc}}t'_{\text{asc}}} - \tilde{g}^2_{t'_{\text{asc}}\tilde{P}}}{\tilde{g}_{t'_{\text{asc}}t'_{\text{asc}}}}(\Delta\tilde{P})^2 = \mu^{\parallel}_{\text{max}} + \frac{(\Delta\tilde{P})^2}{\tilde{g}^{\tilde{P}\tilde{P}}}\end{aligned} \tag{A.9}$$

where $\tilde{g}^{ij}$ is the inverse matrix to $\tilde{g}_{ij}$.