LIGO-P070030-00-Z

Template-based searches for gravitational waves: efficient lattice covering of flat parameter spaces

Reinhard Prix

Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut), Am Mühlenberg 1, 14476 Golm, Germany

E-mail: `reinhard.prix@aei.mpg.de`

Abstract. We show that the construction of optimal template banks for matched-filtering searches is an example of the *sphere covering problem*. For parameter spaces with constant-coefficient metrics a (near-) optimal template bank is achieved by the A_n^* lattice, which is the best lattice-covering in dimensions $n \le 5$, and is close to the best covering known for dimensions $n \le 16$. Generally this provides a *substantially* more efficient covering than the simpler hyper-cubic lattice. We present an algorithm for generating lattice template banks for constant-coefficient metrics and we illustrate its implementation by generating A_n^* template banks in $n = 2, 3, 4$ dimensions.

1. Introduction

The detection of gravitational waves (GWs) in the noisy data of a detector ideally requires the knowledge of the signal waveform, in order to coherently correlate the data with the expected signal by *matched filtering*. Depending on the type of astrophysical sources considered, however, one typically only knows a parametrized family of possible waveforms (or approximations thereof). The unknown parameters of these waveforms could be, for example, the frequency and sky-position of spinning neutron stars, or the masses and spins of inspiraling binary systems. Parameter spaces of such wide-parameter searches typically have between 1 and $\lesssim 10$ dimensions, depending on computational constraints and on how much astrophysical information is available to restrict the search space a-priori.

One can obviously search only a finite subset of points in this parameter-space, and this subset constitutes the "template bank" or search grid. The templates must *cover* the parameter space, i.e. they must be placed densely enough that no signal in this space can lose more than a certain fraction of its power (called *mismatch*) at the closest template. However, coherently correlating the data with every template is computationally expensive and increases the expected number of statistical false-alarm candidates. An *optimal* template bank therefore contains the smallest number of templates such that the worst-case mismatch does not exceed a given limit.

It was realized early on that a geometric approach is very useful to construct a template bank, in particular the introduction of a parameter-space *metric* [1, 2] based on the mismatch. This provides a natural measure of distance in parameter space and allows one to "correctly" place templates, in the sense that the spacing is not too wide and the maximal mismatch will not be exceeded. Less attention, however, was devoted to the problem of *optimally* placing templates once the metric is known, and often a simple – but highly suboptimal – hyper-cubic template grid was used or the problem was confused with *sphere packing* [2, 3, 4]. We will see in the following that finding an optimal template bank is an instance of the *sphere covering* problem, which in a sense is the "opposite" of the sphere packing problem. The covering problem is highly non-trivial (as is the packing problem), which is illustrated by the fact that even in Euclidean space the solution is only known in $n = 2$ dimensions, a partial solution (restricted to lattices) is known in $n \leq 5$ dimensions, while an optimal solution for higher dimensions is unknown (cf. [5, 6]). The main motivation of the present work is to develop a method for constructing efficient template banks by using the results about Euclidean sphere covering.

2. Template-based searches and parameter-space metric

A wide class of searches for GWs can be characterized as *template based*, in the sense that one searches for signals belonging to a family of waveforms $s(t;\boldsymbol{\lambda})$, which depend on a vector of parameters $\{\boldsymbol{\lambda}\}^i = \lambda^i$. The strain $x(t)$ measured by a detector contains (usually dominating) noise $n(t)$ in addition to possible weak GW signals $s(t;\boldsymbol{\lambda}_{\rm s})$, i.e. $x(t) = n(t) + s(t;\boldsymbol{\lambda}_{\rm s})$. One has to construct a *detection statistic*, $\mathcal{F}(\boldsymbol{\lambda};x)$ say, which is a scalar observable characterizing the probability of a signal with parameters $\boldsymbol{\lambda}$ being present in the data $x(t)$. Due to the random noise fluctuations $n(t)$ in the data, the detection statistic is a random variable, and generally (assuming $\mathcal{F}$ is unbiased) its expectation value $\overline{\mathcal{F}}(\boldsymbol{\lambda};\boldsymbol{\lambda}_{\rm s}) \equiv E[\mathcal{F}(\boldsymbol{\lambda};x)]$ has a (local) maximum at the location of a signal $\boldsymbol{\lambda} = \boldsymbol{\lambda}_{\rm s}$, i.e.

$$\left.\frac{\partial\overline{\mathcal{F}}(\boldsymbol{\lambda};\boldsymbol{\lambda}_{\rm s})}{\partial\boldsymbol{\lambda}}\right|_{\boldsymbol{\lambda}=\boldsymbol{\lambda}_{\rm s}} = 0\,. \tag{1}$$

Taylor-expanding the expected detection-statistic $\overline{\mathcal{F}}$ in small offsets $\Delta\boldsymbol{\lambda} = \boldsymbol{\lambda} - \boldsymbol{\lambda}_{\rm s}$ around the signal location $\boldsymbol{\lambda}_{\rm s}$ therefore reads as

$$\overline{\mathcal{F}}(\boldsymbol{\lambda};\boldsymbol{\lambda}_{\rm s}) = \overline{\mathcal{F}}(\boldsymbol{\lambda}_{\rm s};\boldsymbol{\lambda}_{\rm s}) + \frac{1}{2}\left.\frac{\partial^2\overline{\mathcal{F}}(\boldsymbol{\lambda};\boldsymbol{\lambda}_{\rm s})}{\partial\lambda^i\,\partial\lambda^j}\right|_{\boldsymbol{\lambda}_{\rm s}} \Delta\lambda^i\,\Delta\lambda^j + \mathcal{O}(\Delta\lambda^3)\,, \tag{2}$$

where the matrix of second derivatives of $\overline{\mathcal{F}}$ is negative definite. Here and in the following we use automatic summation over repeated parameter indices $i, j, \ldots$. The perfect-match detection statistic $\mathcal{F}(\boldsymbol{\lambda}_{\rm s};\boldsymbol{\lambda}_{\rm s})$ is typically used to define the optimal signal-to-noise ratio (SNR). We can now introduce the *mismatch* m, which characterizes the fractional loss of SNR at a parameter-space point $\boldsymbol{\lambda}$, with respect to the signal location $\boldsymbol{\lambda}_{\rm s}$, namely

$$m(\boldsymbol{\lambda};\boldsymbol{\lambda}_{\rm s}) \equiv \frac{\overline{\mathcal{F}}(\boldsymbol{\lambda}_{\rm s};\boldsymbol{\lambda}_{\rm s}) - \overline{\mathcal{F}}(\boldsymbol{\lambda};\boldsymbol{\lambda}_{\rm s})}{\overline{\mathcal{F}}(\boldsymbol{\lambda}_{\rm s};\boldsymbol{\lambda}_{\rm s})}\,, \tag{3}$$

Inserting the local expansion (2), we find

$$m(\boldsymbol{\lambda};\boldsymbol{\lambda}_{\rm s}) = g_{ij}(\boldsymbol{\lambda}_{\rm s})\,\Delta\lambda^i\Delta\lambda^j + \mathcal{O}(\Delta\lambda^3)\,, \tag{4}$$

where we defined the positive-definite *metric tensor* $g_{ij} \equiv -\frac{1}{2}\partial_i\partial_j\overline{\mathcal{F}}$, and $\partial_i \equiv \partial/\partial\lambda^i$. Searching a parameter space $\mathbb{P}(\lambda^i,\, g_{ij})$, we need to compute the detection statistic $\mathcal{F}(x;\boldsymbol{\lambda}_\xi)$ for a discrete set of templates $\boldsymbol{\lambda}_\xi \in \mathbb{P}$. Generally one can distinguish two different approaches to this problem: one is a random sampling of $\mathbb{P}$ using Markov-chain Monte-Carlo (MCMC) algorithms (e.g. see [7, 8]), and the other consists of constructing a template bank $\mathbb{T} \equiv \{\boldsymbol{\lambda}_\xi\} \subset \mathbb{P}$ that *covers* the whole of $\mathbb{P}$, that is, no point $\boldsymbol{\lambda} \in \mathbb{P}$ exceeds a given maximal mismatch $m_{\max}$ to its closest template $\boldsymbol{\lambda}_\xi \in \mathbb{T}$, i.e.

$$\max_{\boldsymbol{\lambda}\in\mathbb{P}} \min_{\boldsymbol{\lambda}_\xi\in\mathbb{T}} m(\boldsymbol{\lambda};\boldsymbol{\lambda}_\xi) \le m_{\max}\,. \tag{5}$$

Here we focus on the construction of *optimal* template banks, namely those satisfying (5) with the smallest possible number of templates $\boldsymbol{\lambda}_\xi$. Using the local metric approximation (4), each template $\boldsymbol{\lambda}_\xi$ *covers* a region of parameter-space

$$B_\xi = \left\{\boldsymbol{\lambda} \in \mathbb{P} : g_{ij}(\boldsymbol{\lambda}_\xi)\, \Delta\lambda^i\, \Delta\lambda^j \le m_{\max}\,, \quad \Delta\boldsymbol{\lambda} \equiv \boldsymbol{\lambda} - \boldsymbol{\lambda}_\xi\right\}\,, \tag{6}$$

which is a *sphere* of radius

$$R = \sqrt{m_{\max}}\,, \tag{7}$$

in the metric space $\mathbb{P}(\lambda^i,\, g_{ij})$. We can therefore reformulate the definition of an optimal template bank as the set of (overlapping) spheres of covering radius R which *cover* the whole of $\mathbb{P}$ in the sense of (5) with the smallest number of spheres. This is known as the *sphere covering problem* [5], not to be confused with the – somewhat dual – *sphere packing problem*, which seeks to pack the largest number of non-overlapping "hard" spheres into a given volume. An optimal covering consists of an arrangement minimizing the density of overlapping spheres, while an optimal packing maximizes the density of non-overlapping spheres.

3. The Euclidean sphere covering problem

In this section we summarize the current status of the sphere covering problem as far as relevant for the construction of optimal template banks. There has been impressive progress in the study of the covering problem in recent years, e.g. see [5] for a general overview and [6] for a more recent update. Unfortunately, all of these studies are restricted to Euclidean spaces $\mathbb{E}^n$, while the metric parameter spaces of GW searches are often curved. In the following we therefore make the assumption that $\mathbb{P}(\lambda^i,\, g_{ij})$ can be treated as *approximately* flat, or at least broken up into smaller pieces that can individually be treated as nearly flat. If the curvature of the metric is too strong, i.e. if the curvature radius is comparable to the covering radius, it will be difficult to make use of the Euclidean covering problem, and a different approach such as a stochastic template bank or an MCMC sampling might be more fruitful.

We further assume that we have found a coordinate system of $\mathbb{P}$ such that the metric components are (approximately) constant, i.e. $g_{ij}(\boldsymbol{\lambda}) \approx \mathrm{const}_{ij}$, and for simplicity of notation we assume in this section (without loss of generality) that we have chosen coordinates x^i where the constant-coefficient metric is Cartesian, i.e. $\mathbb{P} = \mathbb{E}^n(x^i,\, \delta_{ij})$.

A covering can consist of any arrangement of covering spheres, but currently all best coverings known are *lattices*, and we therefore restrict the discussion to *lattice coverings*, where the centers of the spheres (i.e. templates) form a lattice.

3.1. Basics on lattices

An n-dimensional lattice Λ can be defined as a discrete set of points $\boldsymbol{\nu}_\xi$ (forming an additive group) generated by

$$\boldsymbol{\nu}_\xi = \xi^i\, \boldsymbol{l}_{(i)}\,, \quad \text{with} \quad \xi^i \in \mathbb{Z}\,, \tag{8}$$

where $\{\boldsymbol{l}_{(i)}\}_{i=1}^{n}$ is called a *basis* of the lattice. Note that it is sometimes convenient to express the n basis vectors in a higher-dimensional Euclidean space, i.e. generally we can have $\boldsymbol{l}_{(i)} \in \mathbb{E}^m$ with $m \geq n$. When writing $\mathbb{E}^n$ in the following we refer to the subspace of $\mathbb{E}^m$ containing the n-dimensional lattice Λ. The $m \times n$ matrix $M^a{}_i \equiv l^a_{(i)}$ is called a *generator matrix* of the lattice, with the columns of M holding the components of the n lattice basis vectors, so we can also write the lattice Λ as

$$\Lambda = \{\boldsymbol{\nu}_\xi : \boldsymbol{\nu}_\xi = M\,\boldsymbol{\xi}\,,\ \ \boldsymbol{\xi} \in \mathbb{Z}^n\}\ . \tag{9}$$

The $n \times n$ matrix $A \equiv M^{\mathrm{T}} M$ is called the *Gram matrix* (where T denotes the transpose), which is symmetric and positive definite, and

$$A_{ij} = \boldsymbol{l}_{(i)} \cdot \boldsymbol{l}_{(j)} = \delta_{ab}\, l^a_{(i)}\, l^b_{(j)}, \tag{10}$$

i.e. its coefficients are the mutual scalar products of lattice basis vectors. Each choice

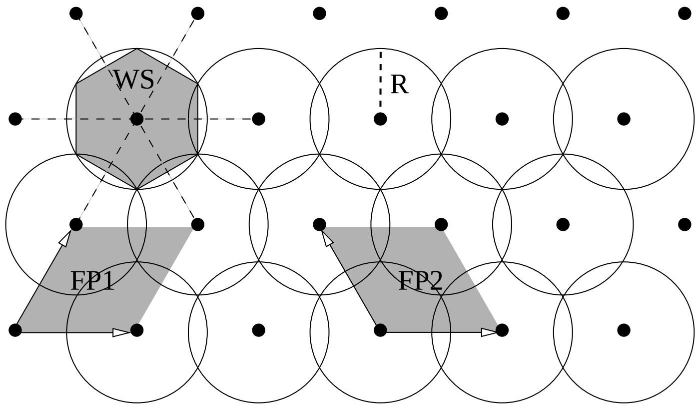

Figure 1. Hexagonal lattice (A_2^*) illustrating a 2-dimensional lattice covering. The shaded areas are different choices of fundamental regions for the lattice. FP1 and FP2 are two fundamental polytopes (11) associated with different choices of lattice basis vectors, WS is the Wigner-Seitz cell (13), and R is the covering radius.

of lattice basis vectors $\{\boldsymbol{l}_{(i)}\}$ defines a corresponding fundamental parallelotope FP, namely

$$\mathrm{FP}\left(\{\boldsymbol{l}_{(i)}\}\right) \equiv \left\{\boldsymbol{x} \in \mathbb{E}^n : \boldsymbol{x} = \theta^i\, \boldsymbol{l}_{(i)}\,,\ \ 0 \leq \theta^i < 1\right\}\,, \tag{11}$$

which is illustrated in figure 1. The FP is an example of a *fundamental region* for the lattice, i.e. a building block containing exactly one lattice point, which fills the whole space $\mathbb{E}^n$ when repeated. There are many different choices of basis and fundamental regions for the same lattice Λ, but they all have the same volume vol(Λ), given by

$$\mathrm{vol}(\Lambda) = \sqrt{\det A}\,, \tag{12}$$

and in the case where M is a square matrix we also have $\mathrm{vol}(\Lambda) = \det M$. One special choice of fundamental region is the *nearest-neighbor region*, often referred to as *Dirichlet-Voronoi* cell by Mathematicians, and more commonly known as Wigner-Seitz cell or Brillouin zone by Physicists, which is defined as

$$\mathrm{WS}(\Lambda) \equiv \{\boldsymbol{x} \in \mathbb{E}^n : \|\boldsymbol{x}\| \leq \|\boldsymbol{x} - \boldsymbol{\nu}_\xi\|,\ \ \boldsymbol{\nu}_\xi \in \Lambda\}\ , \tag{13}$$

where $\|\boldsymbol{x}\| = \sqrt{\boldsymbol{x}\cdot\boldsymbol{x}}$ is the standard Euclidean norm in $\mathbb{E}^n$. The vertices of the Wigner-Seitz cell are by construction local maxima of the distance function of points in $\mathbb{E}^n$ from the nearest grid point. The maximum distance of any points in $\mathbb{E}^n$ to the nearest point of the lattice is called the *covering radius* R, which corresponds to the *circumradius* of WS, as seen in figure 1.

Two lattices Λ_1 and Λ_2 with generator matrices M_1 and M_2 are *equivalent* if they can be transformed into one another by a rotation, reflection and change of scale, namely if the generator matrices satisfy

$$M_2 = c\, B\, M_1\, U\,, \tag{14}$$

where $c \in \mathbb{R}$ is a scale-factor, U is integer-valued $\det U = \pm 1$, which accounts for different choices of basis vectors, and B is a real orthogonal matrix, i.e. $B^{\mathrm{T}}\, B = \mathbb{I}$. The associated Gram matrices are therefore related by

$$A_2 = c^2\, U^{\mathrm{T}}\, A_1\, U\,, \tag{15}$$

and the fundamental volumes (12) of the two lattices are

$$\mathrm{vol}(\Lambda_2) = c^n\, \mathrm{vol}(\Lambda_1)\,. \tag{16}$$

Let us consider the example of a 2-dimensional hexagonal lattice, such as shown in figure 1. An obvious generator matrix is

$$M_1 = \begin{pmatrix} 1 & 1/2 \\ 0 & \sqrt{3}/2 \end{pmatrix}, \tag{17}$$

corresponding to FP1 in figure 1. However, sometimes it is more convenient to work with a generator matrix of the form

$$M_2 = \begin{pmatrix} 1 & 0 \\ -1 & 1 \\ 0 & -1 \end{pmatrix}, \tag{18}$$

which has simpler coefficients, but uses a 3-dimensional representation of the 2-dimensional lattice with all lattice points lying in the plane $x + y + z = 0$. One can verify that these two representations are equivalent in the sense of (14) with

$$c = \sqrt{2}\,, \quad U = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad B = \begin{pmatrix} 1/\sqrt{2} & -1/\sqrt{6} \\ -1/\sqrt{2} & -1/\sqrt{6} \\ 0 & \sqrt{2/3} \end{pmatrix}. \tag{19}$$

Such a higher-dimensional representation of the generator matrix will be useful later for the description of the n-dimensional A_n^* lattice.

3.2. Known results on optimal sphere covering

A sphere covering is characterized by its *thickness* Θ (sometimes also referred to as the *covering density*), which measures the fractional amount of overlap between the covering spheres, or the average number of spheres covering any point in $\mathbb{E}^n$. This can be expressed as the ratio of the volume of one covering sphere to the volume of the fundamental region of the lattice, i.e.

$$\Theta \equiv \frac{V_n\, R^n}{\mathrm{vol}(\Lambda)} \geq 1\,, \tag{20}$$

where R is the covering radius and V_n is the volume of the unit-sphere in n dimensions, namely $V_n = \pi^{n/2}/\Gamma(n/2 + 1)$. We also use the *normalized thickness* or *center density* θ, defined as

$$\theta \equiv \frac{\Theta}{V_n}\,, \tag{21}$$

which measures the number of centers (i.e. templates) per unit volume in the case of $R = 1$. Note that under a lattice transformation (14) the covering radius R obviously scales as $R_2 = c\,R_1$, and we therefore see from (16) that the thickness (20) and (21) is an invariant property of a lattice, i.e. $\theta_2 = \theta_1$. The covering problem consists of finding the covering with the lowest center density θ.

Kershner showed in 1939 (see [5]) that in $n = 2$ dimensions the most economical arrangement of circles covering the plane is the hexagonal lattice, shown in figure 1, which is equivalent to an A_2^* lattice. In dimensions $n = 3, 4, 5$ only the best *lattice covering* is known, and is given by A_n^* in all three cases. In three dimensions A_3^* is the well-known *body-centered-cubic* (bcc) lattice. Note that the best *packing* in $n = 2$ is also achieved by the hexagonal lattice, but for $n = 3$ the face-centered cubic (fcc) lattice is a denser packing than bcc. In higher dimensions the best lattice covering is still unknown. In the first edition (1988) of Conway&Sloane's book [5], the best lattice covering known in all dimensions up to $n \leq 23$ was the A_n^* lattice. Since then, however, this "record" has been broken in most dimensions $5 < n \leq 23$, e.g. see table 2 in [6], and see [9] for an up-to-date online version of the best covering lattices currently known. The A_n^* lattice has a center density of

$$\theta(A_n^*) = \sqrt{n+1}\left\{\frac{n(n+2)}{12(n+1)}\right\}^{n/2} . \qquad (22)$$

Consider on the other hand the hyper-cubic grid $\mathbb{Z}^n$: the Wigner-Seitz cell is a unit hypercube, so $\mathrm{vol}(\mathbb{Z}^n) = 1$, and the covering radius is the diagonal $R = \sqrt{n}/2$, therefore the center density (21) is found as

$$\theta(\mathbb{Z}^n) = 2^{-n}\, n^{n/2} , \qquad (23)$$

which is dramatically worse than A_n^* in higher dimensions, as can be seen from the thickness ratio

$$\kappa(n) \equiv \frac{\theta(\mathbb{Z}^n)}{\theta(A_n^*)} = \frac{3^{n/2}}{\sqrt{n+1}}\left(\frac{n+1}{n+2}\right)^{n/2} \overset{n\to\infty}{\sim} \frac{3^{n/2}}{\sqrt{ne}} . \qquad (24)$$

Table 1. Thickness ratio $\kappa(n) = \theta(\mathbb{Z}^n)/A_n^*$, and $\gamma(n) = \theta(\mathrm{best})/\theta(A_n^*)$ for dimensions $n \leq 17$.

n	2	3	4	5	6	7	8	9
$\kappa(n)$	1.3	1.9	2.8	4.3	6.8	10.9	17.7	28.9
$\gamma(n)$	1.0	1.0	1.0	1.0	0.97	0.95	0.86	0.97

n	10	11	12	13	14	15	16	17
$\kappa(n)$	47.4	78.2	130	216	359	601	1007	1692
$\gamma(n)$	0.98	0.88	0.99	0.86	0.82	0.86	1.0	0.68

There is a theoretical *lower* limit on the thickness of any covering, the Coxeter-Few-Rogers (CFR) bound τ_n (see [5]), namely $\theta_n \geq \tau_n/V_n$, where asymptotically $\tau_n \sim n/(e\sqrt{e})$ for $n \to \infty$. Figure 2 shows the normalized thickness θ as a function of dimension n for the A_n^* and hyper-cubic $\mathbb{Z}^n$ lattices, as well as the CFR limit and the best covering known. In table 1 we see that in dimensions $n > 5$ where A_n^* has been superseded as the best covering, the relative improvement $\gamma(n) \equiv \theta(\mathrm{best})/\theta(A_n^*)$ in thickness is typically quite small. The current "record holders" $\theta(\mathrm{best})$ are taken from

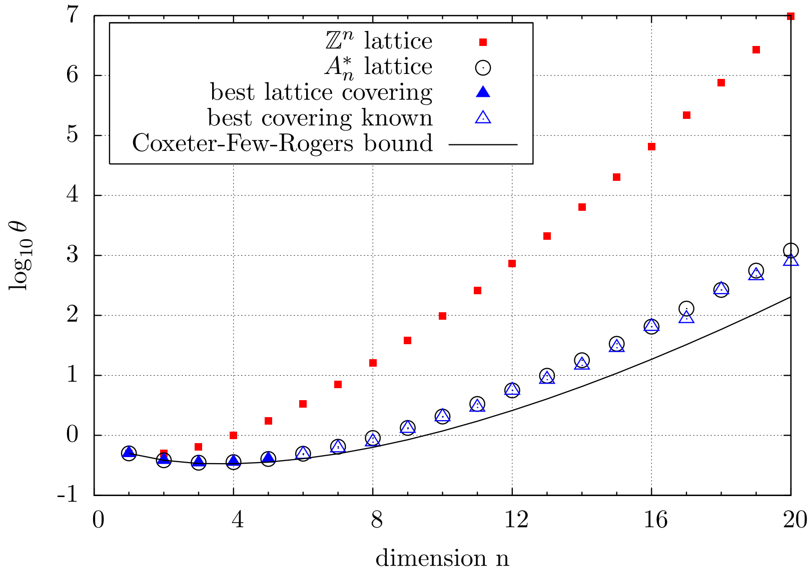

Figure 2. Normalized covering thickness θ as function of dimension n. Shown are the hyper-cubic lattice ($\mathbb{Z}^n$), the A_n^* lattice, the theoretical lower limit (CFR), the best lattice coverings in $1 \leq n \leq 5$, and the best covering *known* in $5 < n \leq 24$.

[9] as of 22 March 2007. In particular, for $n \leq 16$ the improvement $\gamma(n)$ is typically less than 18%, while the advantage $\kappa(n)$ of A_n^* compared to the hyper-cubic grid $\mathbb{Z}^n$ grows large very rapidly, as seen in table 1 and figure 2. For practical simplicity we therefore propose A_n^* as the covering lattice of choice for parameter spaces of dimension up to $n \leq 16$.

4. Lattice covering of template spaces

4.1. Template counting

We now return to template spaces $\mathbb{P}(\lambda^i, g_{ij})$ with constant-coefficient metrics g_{ij}, which only differs from the Cartesian case considered in the previous section by a simple coordinate-transformation. An infinitesimal parameter-space region $d^n\lambda$ has a volume dV measured by the metric, i.e. $dV = \sqrt{g}\, d^n\lambda$, where $g \equiv \det g_{ij}$. This can be integrated to yield the volume V of a finite region of parameter space as

$$V = \int_{\mathbb{P}} dV = \sqrt{g} \int_{\mathbb{P}} d^n\lambda \,, \tag{25}$$

where in the second expression we used the fact that g_{ij} is a constant-coefficient metric. The number of templates dN_p in dV is simply given by the inverse lattice volume, i.e.

$$dN_p = \frac{dV}{\mathrm{vol}(\Lambda)} \,. \tag{26}$$

With the relation (7) between covering radius and maximal mismatch, and using (20) and (21), we obtain

$$dN_p = \theta\, m_{\mathrm{max}}^{-n/2}\, dV \quad \Longrightarrow \quad N_p = \theta\, m_{\mathrm{max}}^{-n/2}\, \sqrt{g} \int_{\mathbb{P}} d^n\lambda \,, \tag{27}$$

which generalizes previous template counting expressions [2, 3, 4] to arbitrary lattices.

4.2. Practical implementation of lattice covering

In this section we present a practical algorithm for generating a lattice covering of given maximal mismatch $m_{\max}$ in a parameter space with constant-coefficient metric g_{ij}. The approach described here works for any lattice generator M, but in practice (cf. section 3.2) we will be most interested in the A_n^* lattice. The generator for A_n^* can be expressed (cf. [5]) as the $(n+1)\times n$ matrix,

$$M^a{}_j = \begin{pmatrix} 1 & 1 & \dots & 1 & \frac{-n}{n+1} \\ -1 & 0 & \dots & 0 & \frac{1}{n+1} \\ 0 & -1 & \dots & 0 & \frac{1}{n+1} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & -1 & \frac{1}{n+1} \\ 0 & 0 & \dots & 0 & \frac{1}{n+1} \end{pmatrix}, \tag{28}$$

which is a $(n+1)$-dimensional representation of the n-dimensional lattice, with the columns of M holding the n lattice basis vectors $\boldsymbol{l}_{(j)}$ expressed in $\mathbb{E}^{n+1}$, i.e. $M^a{}_j = l^a_{(j)}$, where we use index conventions $i, j = 1, \dots, n$ and $a, b = 1, \dots, n+1$. The volume of the fundamental region and the covering radius of this lattice are

$$\mathrm{vol}(A_n^*) = \frac{1}{\sqrt{n+1}}\,, \quad \text{and} \quad R = \sqrt{\frac{n(n+2)}{12(n+1)}}\,, \tag{29}$$

which together with (20) results in the (normalized) thickness $\theta(A_n^*)$ given in (22). In order to generate such a lattice in a parameter space $\mathbb{P}(\lambda^i, g_{ij})$, we need to express the generator $M^a{}_j$ in the λ^i coordinates, $\widetilde{M}^i{}_j$ say, such that the lattice of templates is generated by

$$\lambda^i_\xi = \widetilde{M}^i{}_j\,\xi^j\,, \quad \text{with} \quad \xi^j \in \mathbb{Z}\,. \tag{30}$$

This coordinate transformation is achieved in several steps:

(i) Reduce the $(n+1)\times n$ matrix $M^a{}_j$ to a full rank generator, $\widehat{M}^i{}_j$ say, by expressing the lattice basis vectors in a Euclidean basis spanning the n-dimensional subspace $\mathbb{E}^n$ of the lattice: this is achieved by a simple Gram-Schmidt procedure with respect to the Cartesian metric δ_{ab} using the $\{l^a_{(j)}\}$ to generate an orthonormal basis $\{e^a_{(j)}\}$ satisfying

$$\delta_{ab}\, e^a_{(i)}\, e^b_{(j)} = \delta_{ij}\,. \tag{31}$$

The full-rank generator $\widehat{M}^i{}_j$ is obtained from the projections of the lattice vectors $\{l^a_{(i)}\}$ in this orthonormal basis, namely

$$\widehat{M}^i{}_j = \widehat{l}^i_{(j)} = l^a_{(j)}\, e^b_{(i)}\, \delta_{ab} = e_{(i)a}\, M^a{}_j\,. \tag{32}$$

(ii) Translate the full-rank generator $\widehat{M}^i{}_j$ from Cartesian coordinates into the coordinate system λ^i with metric g_{ij}. For this we use a Gram-Schmidt orthonormalization with respect to the metric g_{ij}, with the lattice vectors $\{\widehat{l}^j_{(i)}\}$ as input to find an orthonormal basis $\{d^j_{(i)}\}$ satisfying

$$g_{ij}\, d^i_{(l)}\, d^j_{(k)} = \delta_{lk}\,. \tag{33}$$

This representation of an orthonormal basis in coordinates λ^i allows one to simply translate the lattice vectors in these coordinates as

$$\widetilde{l}^i_{(j)} = \widehat{l}^k_{(j)}\, d^i_{(k)} = d^i_{(k)}\, \widehat{M}^k{}_j\,. \tag{34}$$

(iii) Scale the generator to the desired covering radius $R = \sqrt{m_{\max}}$, and using (29) we obtain

$$\widetilde{M}^i{}_j = \sqrt{m_{\max}}\,\sqrt{\frac{12(n+1)}{n(n+2)}}\;\widetilde{l}^i_{(j)}\,, \tag{35}$$

which is a generator in the sense of (30) for an A_n^* template lattice with maximal mismatch $m_{\max}$.

This algorithm has been implemented in `XLALFindCoveringGenerator()` in LAL [10], and some tests of this code are presented in the next section.

4.3. Test of the implementation

In order to illustrate and test the implementation of the algorithm described in the previous section, we generate an A_n^* lattice in dimensions $n = 2, 3, 4$ respectively, with a maximal mismatch of $m_{\max} = 0.04$, i.e. a covering radius of $R = 0.2$. For generality we use non-Cartesian metrics $g_{ij} \neq \delta_{ij}$, as illustrated in the left panel of figure 3. We picked 100,000 points $\boldsymbol{\lambda} \in \mathbb{P}(\lambda^i, g_{ij})$ at random and compute their mismatch m

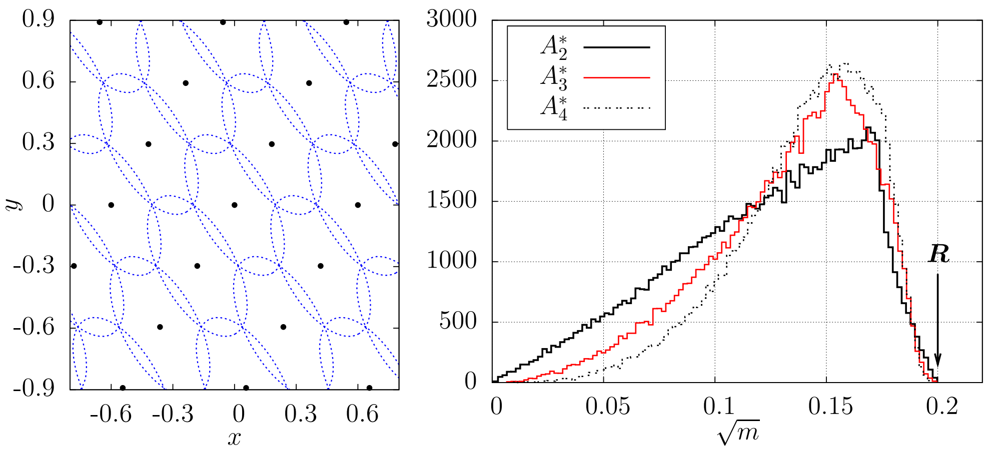

Figure 3. *Left panel:* Hexagonal (A_2^*) lattice covering with covering radius $R = 0.2$, using coordinates $\{x, y\}$ with the metric $g_{ij} = [1,\ 0.4;\ 0.4,\ 0.5]$. *Right panel:* Histogram of measured distances $\sqrt{m}$ in a Monte-Carlo sampling of 100,000 points from an A_n^* covering in $n = 2, 3, 4$ dimensions, using non-Cartesian metrics g_{ij}. The nominal covering radius in all three cases is $R = \sqrt{m_{\max}} = 0.2$.

(using the metric) to the nearest template $\boldsymbol{\lambda}_\xi$, which is a way of *measuring* the maximal mismatch of a template bank. The distribution of measured mismatch-distances $\sqrt{m}$ is presented in the right panel of figure 3, which shows that the mismatches are bounded by $m_{\max}$ as required by (5). From the total number of templates N_p in the parameter-space $\Delta\boldsymbol{\lambda}^n$ covered, we can measure the (normalized) thickness θ of the template bank via (27), namely

$$\theta = \frac{R^n}{\sqrt{g}}\,\frac{N_p}{\Delta\boldsymbol{\lambda}^n}\,, \tag{36}$$

which is found to agree within 0.2% with the theoretical values (22) in all three cases $n = 2, 3, 4$. The generated template banks in this example had of the order of $N_p \sim \mathcal{O}(10^4)$ templates, and this error can be attributed to boundary effects.

5. Discussion

The applicability of the lattice covering algorithm presented here is restricted to *explicitly flat* parameter spaces, which limits its usefulness to cases where we can find a coordinate system in which the parameter-space metric is (at least) *approximately* constant. Another difficulty stems from the fact that sometimes (such as in continuous-wave searches), even though the parameter-space metric can be approximated as flat, its coefficient matrix turns out to be highly ill-conditioned [11], which results in the lattice-construction algorithm to fail due to numerical difficulties. One therefore needs to *analytically* "factor out" this near-degeneracy of the metric before the lattice-covering procedure can be safely applied. More work is also required to deal with non-trivial parameter-space boundaries, in particular non-convex ones, which complicates the n-dimensional filling algorithm. The typical size of template banks might be as large as $N_p \sim \mathcal{O}(10^9)$ templates, therefore we need to be able to fill the parameter space "template per template", without keeping the whole lattice in memory at once.

Acknowledgments

I am grateful for helpful discussions with Bruce Allen, Andrzej Królak and Curt Cutler. This work was supported by the Max-Planck-Society (`AEI-2007-016`). This paper has been assigned LIGO Document Number `LIGO-P070030-00-Z`.

References

[1] R. Balasubramanian, B. S. Sathyaprakash, and S. V. Dhurandhar. Gravitational waves from coalescing binaries: Detection strategies and Monte Carlo estimation of parameters. *Phys. Rev. D.*, 53:3033, 1996.
[2] B. J. Owen. Search templates for gravitational waves from inspiraling binaries: Choice of template spacing. *Phys. Rev. D.*, 53:6749, 1996.
[3] B. J. Owen and B. S. Sathyaprakash. Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement. *Phys. Rev. D.*, 60(2):022002, 1999.
[4] P. R. Brady, T. Creighton, C. Cutler, and B. F. Schutz. Searching for periodic sources with LIGO. *Phys. Rev. D.*, 57:2101, 1998.
[5] J. H. Conway and N. J. A. Sloane. *Sphere packings, lattices and groups.* A Series of Comprehensive Studies in Mathematics. Springer, 1999.
[6] Achill Schürmann and Frank Vallentin. Computational approaches to lattice packing and covering problems. *Discrete and Computational Geometry*, 35:73, 2006.
[7] N. Christensen, R. J. Dupuis, G. Woan, and R. Meyer. Metropolis-Hastings algorithm for extracting periodic gravitational wave signals from laser interferometric detector data. *Phys. Rev. D.*, 70(2):022001, 2004.
[8] N. J. Cornish and J. Crowder. LISA data analysis using Markov chain Monte Carlo methods. *Phys. Rev. D.*, 72(4):043005, 2005.
[9] Achill Schürmann and Frank Vallentin. Geometry of lattices and algorithms. `http://www.math.uni-magdeburg.de/lattice_geometry/`, 2007. (cited March 22, 2007).
[10] LIGO Scientific Collaboration. LAL/LALApps: FreeSoftware (GPL) tools for data-analysis. `http://www.lsc-group.phys.uwm.edu/daswg/`.
[11] Reinhard Prix. Search for continuous gravitational waves: metric of the multi-detector F-statistic. *Phys. Rev. D.*, 75:023004, 2007. (preprint gr-qc/0606088).