

Random template banks and relaxed lattice coverings

C. Messenger,* R. Prix, and M. A. Papa

Albert-Einstein-Institut

(Dated: Tue Sep 30 17:14:05 2008 +0200)

(commitID: 4301865fd72d5bca56a3cc57f80b6c3af5e0c117)

(LIGO-P080090-01-Z)

Template-based searches for gravitational waves are often limited by the computational cost associated with searching large parameter spaces. The study of *efficient* template banks, in the sense of using the smallest number of templates, is therefore of great practical interest. The “traditional” approach to template-bank construction requires every point in parameter space to be *covered* by at least one template, which rapidly becomes inefficient at higher dimensions. Here we study an alternative approach, where any point in parameter space is covered only with a given probability $\eta < 1$. We find that by giving up complete coverage in this way, large reductions in the number of templates are possible, especially at higher dimensions. The prime examples studied here are “random template banks”, in which templates are placed randomly with uniform probability over the parameter space. In addition to its obvious simplicity, this method turns out to be surprisingly efficient. We analyze the statistical properties of such random template banks, and compare their efficiency to traditional lattice coverings. We further study “relaxed” lattice coverings (using \mathbb{Z}_n and A_n^* lattices), which similarly cover any signal location only with probability η . The relaxed A_n^* lattice is found to yield the most efficient template banks at low dimensions ($n \lesssim 10$), while random template banks increasingly outperform any other method at higher dimensions.

I. INTRODUCTION

Matched filtering is the optimal linear detection method for signals in additive stationary noise. This method is widely used in searches for gravitational waves (GWs) in the data of ground based detectors (LIGO [1, 2, 3], GEO [4, 5], VIRGO [6], TAMA [7]) as well as being among those proposed for future space based detectors [8]. The majority of signals that are being searched for have predictable waveforms but unknown waveform parameters. Matched filtering consists of processing the data with multiple waveforms (templates) each corresponding to a different set of waveform parameters. The different templates are spaced based on a metric on the space of signal parameters. The metric defines a distance measure directly related to the loss in the matched filter signal-to-noise ratio for a given template and signal [9, 10, 11, 12]. Using the metric, a template bank, usually in the form of a lattice, can be placed on the space such that the loss between any putative signal and at least a single template in the bank is less than a predefined maximum value.

Template placement for gravitational wave data analysis has proven to be a complicated and involved procedure even in the relatively low-dimensional spaces already searched [13, 14, 15, 16, 17, 18]. Here we discuss the possibility of adopting a seemingly far less complicated template placement method whereby we *randomly* position templates within our search space rather than placing them on a lattice.

It has recently been shown [19] that constructing op-

timal template banks can be interpreted as an instance of the mathematical *sphere covering problem*, and that results from this field of research can usefully be applied to template banks. For instance, the hyper-cubic \mathbb{Z}_n lattice covering is known to become extremely inefficient at higher dimensions compared to other lattices, in particular the A_n^* lattice, which provides a highly efficient covering for dimensions up to $n < 24$.

In practice, however, constructing lattice-based template banks often turns out to be problematic, due to the difficulties associated with adapting lattice coverings to curved parameter spaces and performing coordinate transformations to avoid non-constant metric components. Furthermore, even the best lattice covering becomes increasingly inefficient at higher dimensions ($n \gtrsim 10$, say), which makes this approach increasingly unsuitable for problems involving high-dimensional parameter-spaces.

A radically different approach to template-bank construction consists in relaxing the strict requirement of *complete* coverage for a given mismatch, and instead require coverage only with a certain confidence $\eta < 1$. This is a natural step for searches that employ statistical detection techniques which always involve a finite false-dismissal probability. Connected to this idea are new types of template banks, commonly referred to as “stochastic”, which have recently been studied and applied by various groups [20, 21, 22, 23]. Stochastic template banks are constructed by randomly placing templates on the parameter space, accompanied by a “pruning” step in which “superfluous” templates, which are deemed to lie too close to each other, are removed.

Here we study an even simpler approach, which we refer to as “random template banks”, in order to distinguish it from stochastic banks. This method consists

*Electronic address: chris.messenger@aei.mpg.de

of placing *the right number of templates* randomly, with probability density dependent on the metric determinant and without any additional pruning steps. Apart from the practical advantage of relative simplicity, this allows one to analyze the properties of such random template banks analytically and in great detail. For example, we can explicitly determine the number of templates $N_{\mathcal{R}}$ required to achieve any desired level of coverage confidence η . This paper presents the first detailed study of the properties of such template banks, and an explicit comparison of their efficiency to traditional full-coverage lattice template banks.

Despite their simplicity, random template banks are found to achieve astonishing levels of efficiency compared to traditional template banks, especially at higher dimensions. They outperform even the highly efficient A_n lattice in dimensions above $n \sim 6 - 7$ for covering confidences η in the range of 90% – 95%.

As a by-product of this study, we also analyze the properties of “relaxed lattice” coverings, which share a fundamental feature with random template banks: for any signal location, the nominal covering mismatch is guaranteed only with probability $\eta < 1$. This results in a coarser lattice and therefore a reduction of the number of templates. We find that these relaxed lattices generally result in the most efficient template banks at dimensions up to $n \sim 11$, where random template banks start to dominate.

The plan of this paper is as follows: first we review the general template-bank problem and traditional lattice-based template banks in Sec. II. In Sec. III we present a detailed analysis of random template banks: We calculate their template densities, and compare them to traditional lattice coverings, and we investigate some of the relevant statistical properties of random template banks. In Sec. IV we describe a modification to traditional lattice template banks, termed relaxed lattice covering. Sec. V provides a summary and discussion of the results.

II. “TRADITIONAL” TEMPLATE-BANK CONSTRUCTION

In this section we briefly review some fundamental concepts used in constructing template banks, namely the parameter-space metric [9, 10, 24] and lattice coverings [11, 12, 19]. One key feature of traditional template banks is that they require *complete* coverage of the parameter space, i.e. no point in parameter space is allowed to be further away from its closest template than a given maximal mismatch.

Consider an n -dimensional parameter space \mathcal{S}_n , with coordinates $\{\lambda^i\}_{i=1}^n$. Each point λ describes a set of parameters of a signal model, which we assume to be an accurate description of the true signal family $s(t; \lambda)$. Assume we measured data $x(t)$ containing a signal $s(t; \lambda_s)$ in addition to Gaussian additive noise $n(t)$, i.e. $x(t) = n(t) + s(t; \lambda_s)$. Typically one constructs a detection statis-

tic of the data, $X(\lambda; x)$, say, which is a scalar representing the probability that a signal with parameters λ is present in the data $x(t)$. Due to the random noise fluctuations $n(t)$, X is a random variable, but with the property that its expectation value $\bar{X}(\lambda; \lambda_s) \equiv E[X(\lambda; x)]$ has a maximum at the true location of the signal $\lambda = \lambda_s$. We can define a notion of *mismatch*, or squared length of a parameter offset $\Delta\lambda = \lambda - \lambda_s$, as the relative loss in the expected detection statistic due to this offset, i.e.

$$m(\Delta\lambda; \lambda_s) = 1 - \frac{\bar{X}(\lambda; \lambda_s)}{\bar{X}(\lambda_s; \lambda_s)} = g_{ij}(\lambda_s) \Delta\lambda^i \Delta\lambda^j + \dots, \quad (1)$$

where we use automatic summation over repeated indices i, j and the metric tensor g_{ij} is defined via Taylor expansion of the mismatch m in the small offset $\Delta\lambda$. Using this definition of the metric, the proper volume of the parameter space \mathcal{S}_n can now be expressed as

$$V_{\mathcal{S}_n} = \int_{\mathcal{S}_n} dV, \quad \text{with } dV \equiv \sqrt{g} d^n\lambda. \quad (2)$$

where $g \equiv \det g_{ij}$ is the determinant of the metric g_{ij} . A common approach to the problem of parameter-space covering is to use a lattice of templates. Template-based searches are often computationally expensive due to the large number of templates required, therefore much effort has gone into identifying the most efficient covering, namely the lattice that requires the fewest templates to achieve complete coverage of the parameter space [19]. Note that constructing lattice template banks in curved parameter spaces is highly impractical, and most of the following results implicitly assume that the parameter-space metric g_{ij} is *flat*, i.e. we can find coordinates in which g_{ij} is constant (i.e. independent of parameter-space location λ).

A parameter-space point λ is considered to be “covered” by a template $\lambda_{(k)}$ if its squared distance to the template is smaller than the given nominal mismatch m_* , i.e.

$$g_{ij} \Delta\lambda_{(k)}^i \Delta\lambda_{(k)}^j < m_*, \quad \text{with } \Delta\lambda_{(k)} \equiv \lambda - \lambda_{(k)}. \quad (3)$$

This is equivalent to saying that λ lies within the n -dimensional sphere of radius $R = \sqrt{m_*}$ centered on the template $\lambda_{(k)}$. The construction of efficient (complete) template banks is therefore an instance of the *sphere covering problem*, which asks for the sphere arrangement requiring the smallest number of overlapping spheres to completely cover an n -dimensional (Euclidean) space [19, 25].

A key quantity used in assessing the efficiency of a given sphere covering is its *thickness*. The thickness Θ is defined [25] as *the average number of n -dimensional spheres (templates) covering any point in the parameter space*. For a complete coverage, the thickness therefore satisfies by definition $\Theta \geq 1$ (where in practice equality can only be attained for $n = 1$, in higher dimensions there will always be some overlap between spheres). For

a lattice covering, the thickness can be conveniently expressed as

$$\Theta = \frac{V_n m_*^{n/2}}{V_\Lambda(m_*)}, \quad (4)$$

where V_n is the volume enclosed by an n -dimensional unit-sphere¹, namely

$$V_n = \frac{\pi^{n/2}}{\Gamma(n/2 + 1)}, \quad (5)$$

and V_Λ is the volume of a fundamental region of the lattice Λ , with covering radius $R = \sqrt{m_*}$. Note that under a linear rescaling c , lengths change like $R' = cR$, mismatches like $m' = c^2 m$, and lattice volumes like $V'_\Lambda = c^n V_\Lambda$. Therefore we see from Eq. (4) that the thickness Θ is a scale-invariant property, characterizing the geometric *structure* of a covering. In particular Θ is independent of mismatch m_* . A special instance of a fundamental lattice region is the *Voronoi cell* (also known as the *Wigner-Seitz cell*), which is the set of points that are closer to a given template than to any other template. Let us also define at this point the volume covered by a single template as

$$V_T \equiv V_n m_*^{n/2}. \quad (6)$$

In the following it will also be useful to introduce the *normalized thickness*,

$$\theta \equiv \frac{\Theta}{V_n}, \quad (7)$$

which corresponds to the number of templates per unit volume in the case of $m_* = 1$. Like the thickness Θ , this is a scale-invariant property of a covering, independent of mismatch m_* . As shown in [19], the total number of templates N of a covering can be expressed in terms of the normalized thickness as

$$N = \theta m_*^{-n/2} V_{\mathcal{S}_n}, \quad (8)$$

which shows that the total number of templates is proportional to the normalized thickness.

In the following we will focus on two lattices, namely the \mathbb{Z}_n (hyper-cubic) and the A_n^* lattice, known respectively for their simplicity and covering efficiency. The normalized thickness is known analytically for both lattices, namely

$$\theta_{\mathbb{Z}_n} = \frac{n^{n/2}}{2^n}, \quad (9)$$

$$\theta_{A_n^*} = \sqrt{n+1} \left[\frac{n(n+2)}{12(n+1)} \right]^{n/2}. \quad (10)$$

¹ We are using the geometers convention with respect to the definition of the n -dimensional sphere (or n -sphere) where the 1-sphere represents two points on a line, the 2-sphere is a circle, etc.

In the following we will mostly use the normalized thickness for comparing different covering strategies, since it is proportional to the total number of templates (Eq. (8)), which is the quantity we wish to minimize in order to reduce the computational cost of searching a parameter space.

III. RANDOM TEMPLATE BANKS

We now investigate the properties of a new type of template bank, which we call the “random template bank”, which consists of $N_{\mathcal{R}}$ templates placed randomly with uniform probability distribution (per proper volume) over the parameter space \mathcal{S}_n . Note that contrary to the lattice template-banks discussed in the previous section, nothing in the following requires the parameter-space metric $g_{ij}(\lambda)$ to be constant or flat. The only practical assumption we will make for simplicity is that the metric curvature radius has to be large compared to the covering radius of one template, so we can neglect metric curvature in the expression for the volume of a single template given by Eq. (6). Some practical issues arising from non-constant metrics will be discussed in Sec. III E.

A. Number of required random templates $N_{\mathcal{R}}$

Let us select a point $\lambda_s \in \mathcal{S}_n$ which we assume to be the location of a signal. We assume that the covering sphere with radius $R = \sqrt{m_*}$ centered on λ_s does not intersect the boundary of the parameter-space \mathcal{S}_n , which allows us to neglect boundary effects in the following discussion. Now consider a single randomly placed template with uniform probability distribution (per proper volume). What is the probability that this template does *not* cover (“miss”) within m_* the point λ_s ? This is equivalent to the probability that the template does not fall within the covering sphere centered on λ_s . The probability of falling within this volume is $V_T/V_{\mathcal{S}_n}$, and so the answer is simply

$$P(\text{miss}(m_*) | \mathcal{S}_n, N_{\mathcal{R}} = 1) = 1 - \frac{V_n m_*^{n/2}}{V_{\mathcal{S}_n}}. \quad (11)$$

If we were to place $N_{\mathcal{R}}$ templates randomly in this way, the probability that *none* of the templates cover the point λ_s is therefore

$$P(\text{miss}(m_*) | \mathcal{S}_n, N_{\mathcal{R}}) = \left(1 - \frac{V_n m_*^{n/2}}{V_{\mathcal{S}_n}} \right)^{N_{\mathcal{R}}}, \quad (12)$$

since in this construction each template location is independent of all previously placed templates. It follows that the probability that the point λ_s is covered (“hit”) by *at least one* template is

$$P(\text{hit}(m_*) | \mathcal{S}_n, N_{\mathcal{R}}) = 1 - \left(1 - \frac{V_n m_*^{n/2}}{V_{\mathcal{S}_n}} \right)^{N_{\mathcal{R}}}. \quad (13)$$

This expression shows that the probability of an (unknown) signal location $\lambda_s \in \mathcal{S}_n$ being covered by this random template bank is always < 1 and we would require $N_{\mathcal{R}} \rightarrow \infty$ templates to achieve certain complete coverage. However, we can relax the requirement on certainty of coverage and instead ask how many randomly placed templates do we need in order to obtain a probability η that an (unknown) signal location λ_s would be covered. This is simply given by the solution to $P(\text{hit}(m_*) | \mathcal{S}_n, N_{\mathcal{R}}) = \eta$, which yields

$$N_{\mathcal{R}}(\eta, m_*, \mathcal{S}_n) = \frac{\ln(1 - \eta)}{\ln\left(1 - m_*^{n/2} V_n/V_{\mathcal{S}_n}\right)}. \quad (14)$$

In practice we will mostly be interested in “large” parameter spaces, in the sense that the parameter-space volume $V_{\mathcal{S}_n}$ is very large compared to the volume V_T of one template, and we can therefore Taylor-expand Eq. (14) in the small quantity $m_*^{n/2} V_n/V_{\mathcal{S}_n} \ll 1$, which yields

$$N_{\mathcal{R}}(\eta, m_*, \mathcal{S}_n) \approx \frac{1}{V_n} \ln\left(\frac{1}{1 - \eta}\right) m_*^{-n/2} V_{\mathcal{S}_n}, \quad (15)$$

where the neglected higher-order terms in the expansion correspond to corrections to $N_{\mathcal{R}}$ of order $\mathcal{O}(\ln(1 - \eta)/2)$. We see that for a given parameter-space volume $V_{\mathcal{S}_n}$, the two parameters m_* and η completely determine the number $N_{\mathcal{R}}$ of random templates we need to place randomly on the parameter-space \mathcal{S}_n . We therefore introduce the notation ${}^{\eta}\mathcal{R}_n(m_*)$ to denote an n -dimensional random template bank with nominal mismatch m_* and covering confidence η . An illustrative example of a 2-dimensional random template bank ${}^{0.9}\mathcal{R}_2(0.1)$ is shown in Fig. 1 for a nominal mismatch of $m_* = 0.1$, covering confidence $\eta = 0.9$, and a parameter-space volume $V_{\mathcal{S}_n}$ requiring $N_{\mathcal{R}} = 100$ templates. By comparing Eq. (15) to the general expression in Eq. (8), we can directly read off the normalized thickness $\theta_{\mathcal{R}}(\eta)$ of a random template bank, namely

$$\theta_{\mathcal{R}}(\eta) = \frac{1}{V_n} \ln\left(\frac{1}{1 - \eta}\right), \quad (16)$$

and from Eq. (7) we obtain the corresponding thickness $\Theta_{\mathcal{R}}(\eta)$ as

$$\Theta_{\mathcal{R}}(\eta) = \ln\left(\frac{1}{1 - \eta}\right). \quad (17)$$

This expression reveals a very special property of random template banks compared to any lattice covering (e.g. see Eq. (9) and Eq. (10)), namely the thickness $\Theta_{\mathcal{R}}$ only depends on the covering confidence η , and is *independent* of the dimension n .

Figure 2 shows a comparison between the normalized thickness θ of the hyper-cubic (\mathbb{Z}_n) lattice covering, the A_n^* lattice covering and random template banks ${}^{\eta}\mathcal{R}_n$ with different covering confidences η . We see that random template banks beat the efficiency of the A_n^* lattice

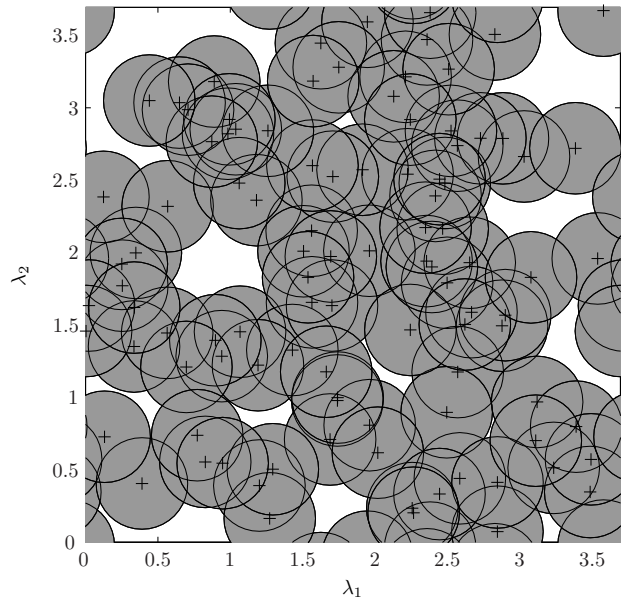


FIG. 1: Example realization of a random template bank ${}^{0.9}\mathcal{R}_2(0.1)$ in $n = 2$ dimensions (using periodic boundary conditions), with nominal mismatch $m_* = 0.1$ and covering confidence $\eta = 0.9$. The parameter-space volume $V_{\mathcal{S}_n}$ was chosen such that the resulting number of templates is $N_{\mathcal{R}} = 100$. Template locations are indicated by crosses and template boundaries (corresponding to the covering radius $R = \sqrt{m_*}$) are shown as black circles. The covered and uncovered volumes are shaded grey and white respectively.

(which is the best, or close to the best lattice covering currently known for dimensions up to $n < 24$ [19]) at sufficiently high dimension n , namely in $n = 10$ for ${}^{0.99}\mathcal{R}_n$, $n = 7$ for ${}^{0.95}\mathcal{R}_n$, and $n = 6$ for ${}^{0.90}\mathcal{R}_n$. Even more surprisingly, random template banks beat the theoretical lower bound (“Coxeter-Few-Rogers bound” [25]) on the thickness of *any* covering! This is possible because random template banks do not provide a *covering* in the strict sense, as they leave some fraction of parameter space uncovered.

B. Distribution of mismatches

Taking the derivative of Eq. (13) with respect to the nominal mismatch m_* , we obtain the probability density function (pdf) for signal mismatches m in a random template bank of given $N_{\mathcal{R}}$ and $V_{\mathcal{S}_n}$, namely

$$\text{pdf}(m|N_{\mathcal{R}}, \mathcal{S}_n) = \frac{n N_{\mathcal{R}} V_n m^{n/2-1}}{2V_{\mathcal{S}_n}} \left(1 - \frac{V_n m^{n/2}}{V_{\mathcal{S}_n}}\right)^{N_{\mathcal{R}}-1}, \quad (18)$$

which describes the probability of finding a template within the mismatch interval $[m, m + dm]$ of some location λ_s . This expression is somewhat inconvenient, however, as it depends on “extensive” parameter-space

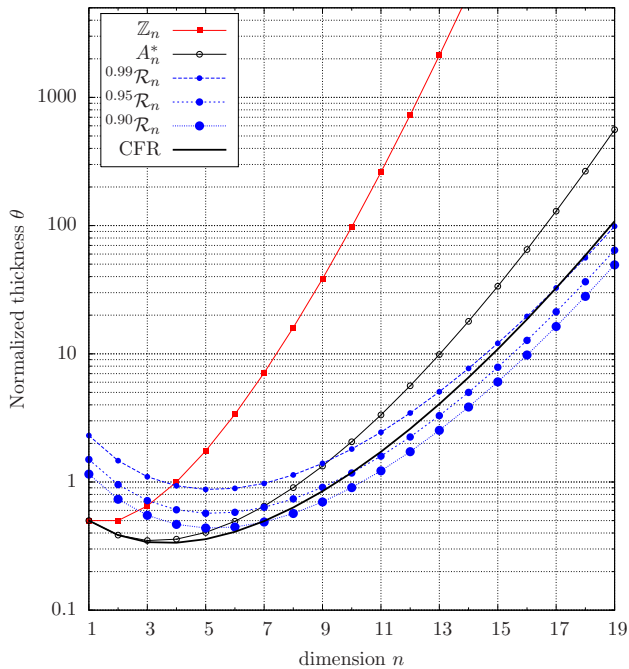


FIG. 2: Normalized thickness θ as a function of dimension n for hyper-cubic (\mathbb{Z}_n) and A_n^* lattice covering, and for random template banks ${}^\eta\mathcal{R}_n$ with different choices of covering confidence $\eta = 0.99, 0.95, 0.90$. Also plotted is the Coxeter-Few-Rogers bound (CFR), which is a theoretical lower bound on the thickness of any strict ($\eta = 1$) covering.

properties, namely the total number of templates $N_{\mathcal{R}}$ and the parameter-space volume V_{S_n} .

In order to rewrite this purely in terms of the “intensive” parameters m_* and η , let us first note that for given nominal mismatch m_* and covering confidence η , we can rewrite the required number of templates given by Eq. (15) as

$$N_{\mathcal{R}} = \Theta_{\mathcal{R}}(\eta) m_*^{-n/2} \mathcal{N}_0, \quad \text{with } \mathcal{N}_0 \equiv \frac{V_{S_n}}{V_n}, \quad (19)$$

where we defined \mathcal{N}_0 as the parameter-space volume V_{S_n} measured in units of the unit-sphere volume V_n . Using this expression, we can write the probability of a point being covered within mismatch m in a random template bank ${}^\eta\mathcal{R}(m_*)$, using Eq. (13), as

$$P(\text{hit}(m) | {}^\eta\mathcal{R}_n(m_*)) = 1 - \left(1 - \frac{m^{n/2}}{\mathcal{N}_0}\right)^{\mathcal{N}_0 \Theta_{\mathcal{R}} m_*^{-n/2}} \quad (20)$$

and using the assumption of a large parameter-space volume V_{S_n} , namely $\mathcal{N}_0 \gg 1$, we can express this as

$$P(\text{hit}(m) | {}^\eta\mathcal{R}_n(m_*)) \approx 1 - e^{-\Theta_{\mathcal{R}} \tilde{m}^{n/2}}, \quad (21)$$

where we defined \tilde{m} as the mismatch m measured in units

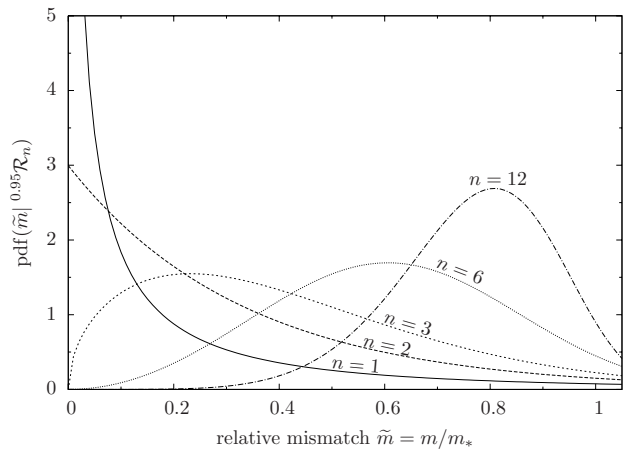


FIG. 3: Probability distribution Eq. (23) of relative mismatches $\tilde{m} = m/m_*$ in random template banks ${}^{0.95}\mathcal{R}_n$ of dimension $n = 1, 2, 3, 6, 12$.

of the nominal mismatch m_* , i.e.

$$\tilde{m} \equiv \frac{m}{m_*}. \quad (22)$$

Note that we obviously find $P(\text{hit}(m_*) | {}^\eta\mathcal{R}_n(m_*)) = \eta$. Contrary to Eq. (20), this expression for $P(\text{hit})$ only depends on the parameters \tilde{m} and η . The interpretation of Eq. (21) is: given a random template bank ${}^\eta\mathcal{R}_n(m_*)$ constructed for nominal mismatch m_* and covering confidence η , what is the probability of a signal location being covered with a mismatch of at most m . By differentiating this with respect to the relative mismatch \tilde{m} we obtain the probability density function of a template falling within the mismatch interval $[\tilde{m}, \tilde{m} + d\tilde{m}]$ of a signal, namely

$$\begin{aligned} \text{pdf}(\tilde{m} | {}^\eta\mathcal{R}_n) &\equiv \frac{d}{d\tilde{m}} P(\text{hit}(m) | {}^\eta\mathcal{R}_n(m_*)) \\ &= \frac{n}{2} \Theta_{\mathcal{R}} \tilde{m}^{n/2-1} e^{-\Theta_{\mathcal{R}} \tilde{m}^{n/2}}. \end{aligned} \quad (23)$$

Figure 3 shows a plot of the mismatch pdf for random template banks ${}^{0.95}\mathcal{R}_n$ in different dimensions n . We can see in this plot that at higher dimensions the bulk of the mismatch probability shifts to the higher end of relative mismatches \tilde{m} . This shows that random template banks become more and more efficient as the number of dimensions grows. The reason for this is that a bank that yields mismatch distributions peaked at lower values of the relative mismatch must be more densely populated than a bank with mismatch distribution peaked at higher values of the relative mismatch. In order to see this it is useful to stress that we construct template banks with the only constraint that the maximum mismatch be lower than a given fixed value, m_* (with probability η): if that constraint is satisfied the bank is acceptable. We can construct different banks that satisfy this constraint,

with different template densities. In a bank that is very densely populated the average distance to the nearest template is smaller than in a bank with lower template density. Hence in a more densely populated bank the distribution of mismatches peaks at a lower value of the relative mismatch compared to a lower-density bank. Among the two banks, the lower-density one has fewer templates, hence it is more efficient.

C. Spatial parameter-space coverage

In our discussion of random template banks so far we have focused on the probability η that an (unknown) signal location λ_s is covered by a template. If we were to construct a number of random template banks, this “confidence” therefore describes the fraction of template banks in which λ_s would be covered. A somewhat related, yet different, question is: given a realization of a random template bank ${}^n\mathcal{R}_n(m_*)$, what is the fraction \mathcal{C} of parameter space that is *actually* covered? This *spatial coverage* \mathcal{C} , as a property of an individual random-template-bank realization, is a random variable, and in the following we will analyze some of its relevant statistical properties.

For a given random-template-bank realization, we define the function $f(\lambda)$ on the parameter-space \mathcal{S}_n as

$$f(\lambda) = \begin{cases} 1 & \text{if } \lambda \text{ is covered} \\ 0 & \text{otherwise,} \end{cases} \quad (24)$$

describing whether a point λ is covered (“hit”) or not (“miss”) within mismatch m_* . Using this we can express the spatial coverage fraction \mathcal{C} as

$$\mathcal{C} = \frac{1}{V_{\mathcal{S}_n}} \int_{\mathcal{S}_n} f(\lambda) dV, \quad (25)$$

where $dV = \sqrt{g} d^n \lambda$ is the volume element associated with $d^n \lambda$. Given any point $\lambda \in \mathcal{S}_n$ the expectation value of $f(\lambda)$ over an ensemble of template banks is given by

$$\begin{aligned} E[f(\lambda)] &= \sum_{j=0}^1 j P(f(\lambda) = j) \\ &= P(\text{hit}(m_*) | {}^n\mathcal{R}_n(m_*)) \\ &= \eta, \end{aligned} \quad (26)$$

where we used the fact that the probability of any point λ being covered in ${}^n\mathcal{R}_n(m_*)$ is by construction given by the covering confidence η . Using this we can express the expectation value of the spatial coverage \mathcal{C} as

$$E[\mathcal{C}] = \frac{1}{V_{\mathcal{S}_n}} \int_{\mathcal{S}_n} E[f(\lambda)] dV = \eta, \quad (27)$$

showing that the expected spatial coverage \mathcal{C} is equal to the covering confidence η . Note that this result holds true despite the obvious existence of correlations in $f(\lambda)$ between neighboring points.

However, the question of the *variance* in spatial coverage \mathcal{C} over an ensemble of random template banks is substantially complicated by these spatial correlations, as we need to evaluate the following integral

$$\text{Var}[\mathcal{C}] = \frac{1}{V_{\mathcal{S}_n}^2} \left\{ \int_{\mathcal{S}_n} dV \int_{\mathcal{S}_n} dV' E[f(\lambda) f(\lambda')] \right\} - \eta^2. \quad (28)$$

Since we have no good handle on the spatial correlations in $f(\lambda)$ for a given realization of ${}^n\mathcal{R}_n$, we try to find some reasonable approximations. First we approximate the above integral as a discrete sum over finite volume elements ΔV . In addition we assume that $f(\lambda)$ in each volume element ΔV is completely uncorrelated with $f(\lambda')$ in any other volume element, and that $f(\lambda)$ is perfectly correlated within each volume element. Applying these approximations yields the following estimate for the variance:

$$\text{Var}[\mathcal{C}] \approx \frac{\Delta V}{V_{\mathcal{S}_n}} \eta(1 - \eta). \quad (29)$$

We have found that a reasonably good semi-empirical “guess” for the number N_I of independent “uncorrelated” volume elements in \mathcal{S}_n seems to be

$$N_I \sim 2n N_{\mathcal{R}}, \quad \text{and} \quad \Delta V = \frac{V_{\mathcal{S}_n}}{N_I}. \quad (30)$$

Note that we do currently not have a good theoretical understanding of this number, however, it seems to yield reasonably good quantitative agreement with the Monte-Carlo results on the coverage, as well as the worst-case mismatch discussed in the next section. We therefore present the analytic estimate as a useful indicator of the general behavior and trends at higher dimensions. Substituting this into Eq. (29), we obtain the following rough estimate for the coverage variance

$$\text{Var}[\mathcal{C}] \approx \frac{\eta(1 - \eta)}{2n N_{\mathcal{R}}}. \quad (31)$$

A key feature to notice is the inverse proportionality of the variance to the number of templates $N_{\mathcal{R}}$ and the parameter-space dimension n . In Fig. 4 we show example distributions of the spatial coverage \mathcal{C} , obtained through numerical simulations, and compare them to Gaussian distributions of mean η and variance given by Eq. (31). Each simulation used 2000 random-template-bank realizations (parameter space volumes were computed as a function of $N_{\mathcal{R}}$ and η and edge effects removed using periodic boundary conditions). A Monte-Carlo integration using 5×10^5 points was then performed on each realization to compute the spatial coverage \mathcal{C} . These simulations become very computationally intensive as the dimensionality of the space increases and hence we limited our test scenarios to $n \leq 6$ and $N_{\mathcal{R}} = 10^4$. The Gaussian distributions using mean $E[\mathcal{C}]$ and variance $\text{Var}[\mathcal{C}]$ seem to agree quite well with the results from the Monte-Carlo simulations for the range of dimensions considered.

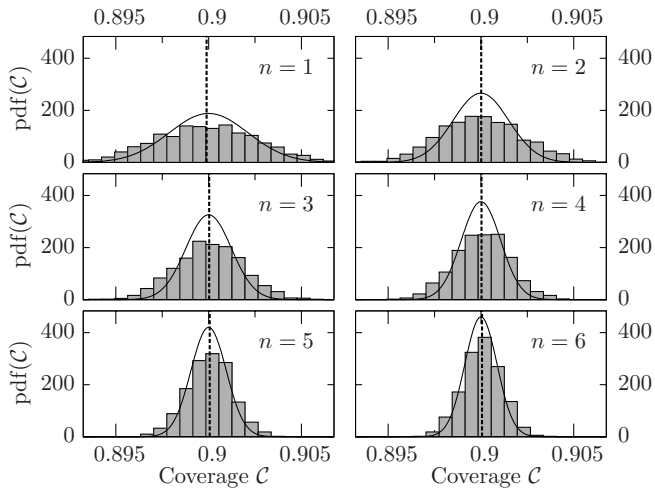


FIG. 4: Results of Monte-Carlo simulations of spatial coverage \mathcal{C} for realizations of ${}^{0.9}\mathcal{R}_n$ random template banks in $n = 1, \dots, 6$ dimensions, using $N_{\mathcal{R}} = 10^4$ templates. The histograms show the distribution of measured coverages \mathcal{C} , the solid black curves show a Gaussian distribution of mean η and variance given by Eq. (31), and the dashed black lines indicate the mean values of the measured distributions.

D. Expected worst-case mismatch

A related question to the spatial coverage \mathcal{C} is the worst-case mismatch m_w found in an individual realization of ${}^{\eta}\mathcal{R}_n(m_*)$, i.e. the “deepest hole” in the template bank. As seen in the previous section, a fraction $1 - \mathcal{C}$ of the parameter space \mathcal{S}_n will have mismatch $m > m_*$. Within this uncovered fraction there will be a worst-case location with the largest value m_w of mismatch. Similar to \mathcal{C} , the worst-case mismatch m_w is a random variable dependent on a given realization of ${}^{\eta}\mathcal{R}_n(m_*)$, and in this section we analyze the statistical properties of m_w over an ensemble of random-template-bank realizations.

The geometrical definition of the minimal mismatch $m(\lambda)$ in any parameter-space point λ for a given random template bank can be expressed as

$$m(\lambda) = \min_{k=1}^{N_{\mathcal{R}}} \left\{ g_{ij} \Delta\lambda_{(k)}^i \Delta\lambda_{(k)}^j \right\}, \quad (32)$$

where k indexes the $N_{\mathcal{R}}$ random templates $\lambda_{(k)}$ and $\Delta\lambda_{(k)} = \lambda - \lambda_{(k)}$. The worst-case mismatch m_w on the space \mathcal{S}_n is therefore

$$m_w \equiv \max_{\mathcal{S}_n} \{m(\lambda)\}. \quad (33)$$

Note that this is in fact the definition of the “covering mismatch” (or covering radius, squared) of a set of templates [25], and for lattices and other complete ($\eta = 1$) coverings, this will be equal to the nominal mismatch, i.e. $m_w = m_*$. In the case of random template banks, however, we control the nominal mismatch m_* , but the

worst-case mismatch could take on arbitrarily large values $m_w > m_*$. One very interesting approach to answering this question would be to use the fact that all “local” worst-case mismatches occur at Voronoi cell vertices [26] (being the so-called “holes” of the covering). For a given number of random templates (or “Voronoi seeds”) the expectation value of the number of vertices is known [27]. This would be a good start to defining a finite set of points for which to draw random mismatch values. The problem with this approach, however, is that the vertex density grows exponentially with dimension, and vertex mismatch values therefore will become highly correlated due to their relative “closeness”. In addition, while we know the mismatch distribution (see Eq. (23)) at randomly selected points λ , the mismatch distribution on Voronoi *vertices* (being very special points) is different, and a complicated function of $\text{pdf}(\tilde{m} | {}^{\eta}\mathcal{R}_n)$.

We therefore employ a more direct and crude approach, namely using the semi-empirical “guess” Eq. (30) for the number N_I of statistically “independent” locations in \mathcal{S}_n , together with the known mismatch distribution of Eq. (23). A linear rescaling of the whole template space will affect all mismatches equally, and therefore it will be more useful in the following to consider the relative worst-case mismatch $\tilde{m}_w \equiv m_w/m_*$, i.e. m_w measured in units of the nominal mismatch m_* .

The probability that the *largest* (relative) mismatch \tilde{m} of N_I independent trials falls within the interval $[\tilde{m}_w, \tilde{m}_w + d\tilde{m}]$ can be expressed as

$$\text{pdf}(\tilde{m}_w | \eta, N_I) d\tilde{m} = \binom{N_I}{1} p_1 p_0^{N_I-1}, \quad (34)$$

where p_1 is the probability of the mismatch in a single trial falling within $[\tilde{m}_w, \tilde{m}_w + d\tilde{m}]$, i.e.

$$p_1 \equiv \text{pdf}(\tilde{m}_w | {}^{\eta}\mathcal{R}_n) d\tilde{m}, \quad (35)$$

and p_0 is the probability of the mismatch in a single trial falling within $[0, \tilde{m}_w]$, i.e.

$$p_0 \equiv P(\text{hit}(m_w) | {}^{\eta}\mathcal{R}_n(m_*)). \quad (36)$$

Using Eq. (23) this can be rewritten as

$$\begin{aligned} \text{pdf}(\tilde{m}_w | {}^{\eta}\mathcal{R}_n, N_I) &= \frac{d}{d\tilde{m}} \Big|_{\tilde{m}_w} P(\text{hit}(\tilde{m}, \eta)^{N_I}) \\ &= \frac{n}{2} \Theta_{\mathcal{R}} N_I \tilde{m}_w^{n/2-1} e^{-\Theta_{\mathcal{R}} \tilde{m}_w^{n/2}} \left(1 - e^{-\Theta_{\mathcal{R}} \tilde{m}_w^{n/2}}\right)^{N_I-1}. \end{aligned} \quad (37)$$

Using this expression we can easily compute quantiles of the worst-case mismatch distribution, namely from

$$\begin{aligned} P(\tilde{m}_w \leq \tilde{m} | \eta, N_I) &= \int_0^{\tilde{m}} \text{pdf}(\tilde{m}_w | \eta, N_I) d\tilde{m}_w \\ &= \left(1 - e^{-\Theta_{\mathcal{R}} \tilde{m}^{n/2}}\right)^{N_I}, \end{aligned} \quad (38)$$

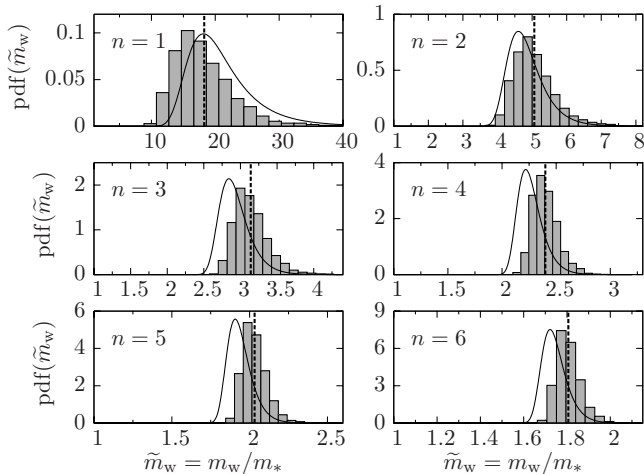


FIG. 5: Results of Monte-Carlo simulations of (relative) worst-case mismatch \tilde{m}_w for realizations of $^{0.9}\mathcal{R}_n$ random template banks in $n = 1, \dots, 6$ dimensions, using $N_{\mathcal{R}} = 10^4$ templates. The histograms show the distribution of measured (relative) worst-case mismatches \tilde{m}_w , the solid black curves show the estimated distribution Eq. (37) and the dashed black lines indicate the mean values of the measured distributions.

so for example we obtain the median worst-case mismatch as

$$\tilde{m}_w^{50\%} = \left[-\frac{1}{\Theta_{\mathcal{R}}} \ln \left(1 - 2^{-1/N_I} \right) \right]^{2/n}. \quad (39)$$

Using the semi-empirical guess, given in Eq. (30), for the number N_I of “independent” parameter-space points in \mathcal{S}_n , we obtain quantitative estimates for the distribution of the worst-case mismatch. Figure 5 shows a few example worst-case mismatch distributions, comparing the analytical estimate of Eq. (37) to the results of Monte-Carlo simulations. These simulations were run with 2000 random-template-bank realizations of a $^{0.9}\mathcal{R}_n$ bank containing $N_{\mathcal{R}} = 10^4$ templates each. The predicted distributions differ slightly from the simulations, but this should be expected given our crude estimation of N_I . With the exception of $n = 1$, the analytical expression tends to slightly underestimate the worst-case mismatch values. Nevertheless, our rough estimate seems to capture the overall trend to smaller values of $E[\tilde{m}_w]$ and $\text{Var}[\tilde{m}_w]$ with increasing dimension n . This behavior is also illustrated in Fig. 6, where we have plotted the analytic expectation value and variance (computed from Eq. (37)) of \tilde{m}_w as a function of dimension n for various covering confidences η and numbers of templates $N_{\mathcal{R}}$. We see that in higher dimensions the expectation value of the relative worst-case mismatch asymptotically approaches unity. We also see that as the number of templates $N_{\mathcal{R}}$ increases, corresponding to a larger number of mismatch “trials”, we obtain an increase in the expected value of \tilde{m}_w . In addition, we see that the standard deviation $\sigma[\tilde{m}_w] \equiv \sqrt{\text{Var}[\tilde{m}_w]}$ decreases both with increasing

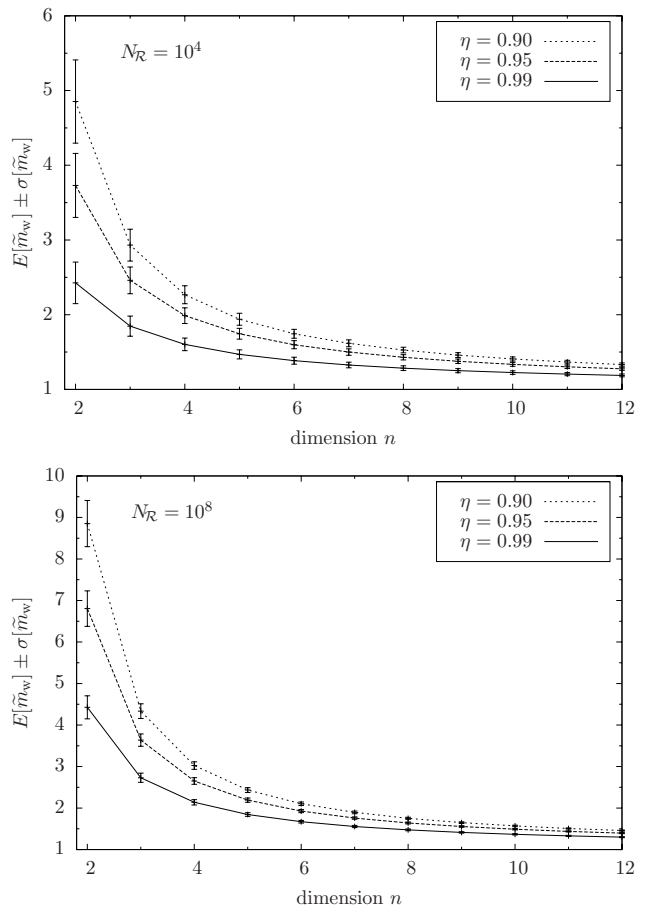


FIG. 6: Statistical properties of the worst-case mismatch \tilde{m}_w in random template banks as function of dimension n , covering confidence η and number of templates $N_{\mathcal{R}}$. This plot shows the expectation value $E[\tilde{m}_w]$ and standard deviation $\sigma[\tilde{m}_w]$ derived from the theoretical estimate Eq. (37). The upper panel shows the case for $N_{\mathcal{R}} = 10^4$ templates, while the lower panel is for $N_{\mathcal{R}} = 10^8$.

number of template $N_{\mathcal{R}}$ as well as with higher dimension n , so that worst-case mismatches become both smaller and more tightly constrained in these limits. For example, as shown in Fig. 5, for $n = 6$ our Monte-Carlo simulations give the mean value of \tilde{m}_w as 1.81 with a standard deviation of 0.06 for a random template bank containing $N_{\mathcal{R}} = 10^4$ templates with a covering confidence $\eta = 0.9$. In using such a template bank one would therefore expect with $> 99\%$ (3σ) confidence that the largest mismatch is smaller than $1.99 m_*$. We should also note that despite our rather crude estimate for the number of independent parameter space locations N_I our model is able to estimate the mean value of \tilde{m}_w to within $\approx 10\%$ of the values obtained from our simulations. This is in fact one of the key results from this investigation, that although random template banks by construction only provide incomplete coverage, $\mathcal{C} < 1$, the actual worst-case mismatch m_w in the uncovered regions is not expected to

be substantially larger than the nominal mismatch m_* , especially at higher dimensions.

E. Practical issues in curved parameter-spaces

As noted in the beginning of the section, the random template bank results apply without modification to curved parameter spaces and non-constant metrics. In order to construct such a template bank in practice, however, we need to generate a uniform random sampling *in proper volume*, which for non-constant metrics can be non-trivial. In order to see how to take account of non-constant metric components $g_{ij}(\lambda)$ in the random template placement, we note that Eq. (15) specifies a constant uniform probability density $\rho_{\mathcal{R}}$ of templates, namely

$$\rho_{\mathcal{R}} = \frac{N_{\mathcal{R}}}{V_{S_n}} = \frac{dN_{\mathcal{R}}}{dV} = \frac{1}{\sqrt{g}} \frac{dN_{\mathcal{R}}}{d^n \lambda}. \quad (40)$$

In other words, the non-constant template “pseudo-density” $\hat{\rho}_{\mathcal{R}}(\lambda)$ in *coordinate-space* satisfies

$$\hat{\rho}_{\mathcal{R}}(\lambda) \equiv \frac{dN_{\mathcal{R}}}{d^n \lambda} = \sqrt{g(\lambda)} \rho_{\mathcal{R}} = \sqrt{g(\lambda)} \theta_{\mathcal{R}} m_*^{-n/2}, \quad (41)$$

which specifies the required random sampling density in coordinate space.

There are various sophisticated and efficient methods for sampling from non-uniform distributions, Markov-Chain-Monte-Carlo (MCMC) methods, importance resampling, and rejection sampling are a few examples. In each of these sampling methods it is sufficient to know the density $\hat{\rho}_{\mathcal{R}}(\lambda)$ (and hence the metric determinant) only up to some normalising constant factor. However, one must know $N_{\mathcal{R}}$, the total number of random templates to draw, which requires accurate knowledge of the proper volume, defined by Eq. 2, and therefore ultimately one must also have accurate knowledge of the metric determinant.

The simplest method, applicable for slowly varying template densities is to decompose the parameter space \mathcal{S}_n into smaller patches $\mathcal{S}_n^{(j)}$, which are small enough so they can be approximated by a constant metric, and sampling it uniformly with template density given by Eq. (41) evaluated at the centre of each patch.

IV. RELAXED LATTICES

In the previous section we saw that random template banks ${}^n\mathcal{R}_n$ will outperform *any* covering at sufficiently high parameter-space dimension. One of the key features of random template banks, however, is that they do not actually provide a strict covering: any point is only covered with probability $\eta < 1$. This allows random template banks to beat even the theoretical (Coxeter-Few-Rogers) lower bound on the thickness of coverings.

In higher dimensions it seems to get extremely expensive (in terms of number of templates) to cover the “last few percent” of a parameter space. Relaxing the requirement of complete coverage therefore allows enormous gains in efficiency.

We can now apply this insight to lattice coverings, by relaxing the strict “minimax” prescription and instead requiring a mismatch m_* only with probability $\eta < 1$. This allows us to use a larger maximal mismatch $m_{\max} > m_*$ for the n -dimensional covering lattice $\Lambda_n(m_{\max})$, thereby reducing the required number of templates. The relation between these quantities is given by

$$\eta = \int_0^{m_*} \text{pdf}(m | \Lambda_n(m_{\max})) dm, \quad (42)$$

where $\text{pdf}(m | \Lambda_n(m_{\max}))$ is the probability distribution of mismatches m for sampled points within a lattice $\Lambda_n(m_{\max})$. A uniform linear rescaling of the whole template space will affect all mismatches equally, and therefore it will be more useful to introduce the relative mismatch $\hat{m} \equiv m/m_{\max}$, which is invariant under re-scalings. Note that contrary to Sec. III, here we use a relative mismatch defined with respect to m_{\max} of the lattice, while the nominal mismatch m_* is a-priori unknown and will be determined from Eq. (42). In Sec. III a relative mismatch $\tilde{m} \equiv m/m_*$ was used, because there was no strict maximal mismatch in this case and we directly prescribed the nominal mismatch m_* . The probability distribution of \hat{m} is $\text{pdf}(\hat{m} | \Lambda_n) = m_{\max} \text{pdf}(m | \Lambda_n(m_{\max}))$, such that $\int_0^1 \text{pdf}(\hat{m} | \Lambda_n) d\hat{m} = 1$. We can therefore restate Eq. (42) in the more useful form

$$\eta = \int_0^{\hat{m}_*} \text{pdf}(\hat{m} | \Lambda_n) d\hat{m}, \quad (43)$$

where we defined the “effective” relative mismatch $\hat{m}_* \equiv m_*/m_{\max} < 1$, which is determined by the lattice mismatch distribution $\text{pdf}(\hat{m} | \Lambda_n)$ and the covering confidence η . We define the linear *relaxation factor* r for the relaxed lattice ${}^n\Lambda_n$ as

$$r({}^n\Lambda_n) \equiv \sqrt{\frac{m_{\max}}{m_*}} = \hat{m}_*^{-1/2} > 1. \quad (44)$$

For a given nominal mismatch m_* and covering confidence η , the relaxation factor determines the maximal covering mismatch m_{\max} of the lattice as

$$m_{\max} = m_* r^2({}^n\Lambda_n), \quad (45)$$

which defines the “relaxed lattice” as ${}^n\Lambda(m_*) \equiv \Lambda(m_{\max})$. The normalized thickness $\theta_{\Lambda_n}(\eta)$ of a relaxed lattice ${}^n\Lambda_n$ is thereby reduced to

$$\theta_{\Lambda_n}(\eta) = \frac{\theta_{\Lambda_n}}{r^n({}^n\Lambda_n)}, \quad (46)$$

with respect to the thickness $\theta_{\Lambda_n} = \theta_{\Lambda_n}(\eta = 1)$ of a traditional covering lattice. Note that contrary to random

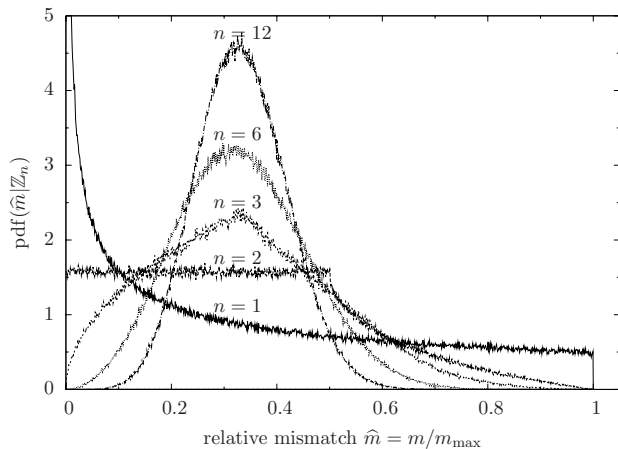


FIG. 7: Results of Monte-Carlo simulations (each using 10^6 points) of the distribution of relative mismatches \hat{m} in hyper-cubic (\mathbb{Z}_n) lattices in dimensions $n = 1, 2, 3, 6, 12$.

template banks, the spatial coverage fraction \mathcal{C} of relaxed lattices is not a random variable and corresponds exactly to the covering confidence, i.e. $\mathcal{C} = \eta$ and $\text{Var}[\mathcal{C}] = 0$. Furthermore, the worst-case mismatch is also exactly known, namely $m_w = m_{\max}$.

As seen in Eq. (43) and Eq. (44), the relaxation factor $r({}^\eta\Lambda_n)$ is determined from the mismatch distribution $\text{pdf}(\hat{m}|\Lambda_n)$. The distribution of mismatches depends strongly on the type of lattice Λ_n and the dimension n . Unfortunately, these mismatch distributions are generally not known analytically, and we need to resort to Monte-Carlo simulations to determine them.

In order to sample the lattice mismatch distribution, we uniformly pick points in parameter space, then find the closest lattice template and determine its relative mismatch $\hat{m} = m/m_{\max}$. For finding the closest template, we use an elegant and efficient method described in [28], which is based on the “fast quantizing” algorithms available for many well-known lattices (see [29] and [25] Chapter 20). Here we focus on two lattices only: the simple but inefficient hyper-cubic lattice \mathbb{Z}_n , and the highly efficient covering lattice A_n^* . Some resulting sampled mismatch distributions are shown in Figs. 7 and 8 for different dimensions n . The Monte-Carlo simulations used 10^6 sampling points for each lattice Λ_n (except for A_n^* in $n = 18, 19$, where 10^5 points were used), and the mismatches were binned into 1000 mismatch bins. The “jitter” seen in Figs. 7 and 8 illustrates the intrinsic sampling fluctuation in these simulations.

The resulting relaxation factors $r({}^\eta\Lambda_n)$ for covering confidences $\eta = 0.99, 0.95, 0.90$ obtained via Eq. (44) are shown in Fig. 9. The errors on the relaxation factors were determined using a Jackknife estimator (see [28], using 100 subsets) and are found to be below 0.04% in all cases. We see in Fig. 9 that the hypercubic (\mathbb{Z}_n) lattice can be relaxed substantially more than the A_n^* lattice, which is also apparent from the pdfs in Fig. 7 and 8: the

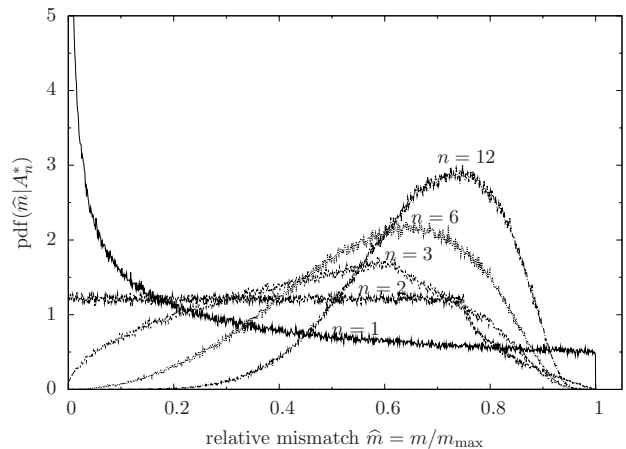


FIG. 8: Results of Monte-Carlo simulations (each using 10^6 points) of the distribution of relative mismatches \hat{m} in A_n^* lattices in dimensions $n = 1, 2, 3, 6, 12$.

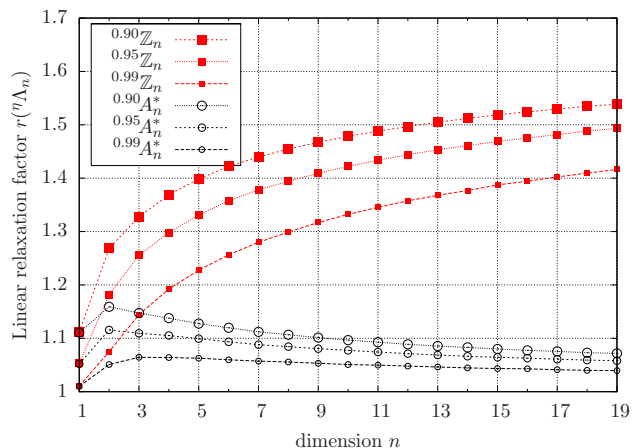


FIG. 9: Linear relaxation factors $r({}^\eta\Lambda)$ for \mathbb{Z}_n and A_n^* lattices and covering confidence $\eta = 0.90, 0.95, 0.99$.

mismatch distribution of \mathbb{Z}_n is much more “wasteful”, as it increasingly concentrates around $\hat{m} = 1/3$. The A_n^* lattice on the other hand, which is a highly efficient covering lattice, has the bulk of mismatches concentrated closer to the maximal mismatch $\hat{m} = 1$.

In Fig. 10 and Table I we show the resulting covering thickness $\theta_{\Lambda_n}(\eta)$ of the relaxed ${}^\eta\mathbb{Z}_n$ and ${}^\eta A_n^*$ lattices in comparison to random template banks ${}^\eta\mathcal{R}_n$. We see that while relaxed lattices are substantially more efficient than traditional complete-coverage lattices, at higher dimensions the random template banks eventually still outperform them. At low dimensions, $n \lesssim 10$ say, the relaxed ${}^\eta A_n^*$ lattice provides the most efficient covering we have found so far. However, having studied only relaxed \mathbb{Z}_n and A_n^* lattices so far, it is conceivable that other lattices, while not necessarily very good covering lattices, could provide an even better relaxed lattice than ${}^\eta A_n^*$. More work is required to study this possibility.

n	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
$^{1.0}\mathbb{Z}_n$	0.50	0.50	0.65	1.0	1.7	3.4	7.1	16	38	98	261	729	2e3	6e3	2e4	7e4	2e5	8e5	3e6
$^{1.0}A_n^*$	0.50	0.38	0.35	0.36	0.40	0.49	0.65	0.90	1.3	2.1	3.3	5.6	9.9	18	34	65	130	266	559
$^{0.99}\mathbb{Z}_n$	0.49	0.43	0.43	0.49	0.63	0.86	1.3	2.0	3.2	5.5	10.	19	36	73	149	322	699	2e3	4e3
$^{0.99}A_n^*$	0.49	0.35	0.29	0.28	0.30	0.35	0.44	0.59	0.83	1.3	2.0	3.2	5.5	9.7	18	33	65	131	268
$^{0.99}\mathcal{R}_n$	2.3	1.5	1.1	0.93	0.87	0.89	0.97	1.1	1.4	1.8	2.4	3.4	5.1	7.7	12	20	33	56	99
$^{0.95}\mathbb{Z}_n$	0.48	0.36	0.33	0.35	0.42	0.54	0.75	1.1	1.8	2.9	5.0	8.9	16	32	63	130	274	590	1e3
$^{0.95}A_n^*$	0.48	0.31	0.26	0.24	0.25	0.29	0.36	0.47	0.66	0.98	1.5	2.5	4.2	7.3	13	25	47	95	192
$^{0.95}\mathcal{R}_n$	1.5	0.95	0.72	0.61	0.57	0.58	0.63	0.74	0.91	1.2	1.6	2.2	3.3	5.0	7.9	13	21	36	64
$^{0.90}\mathbb{Z}_n$	0.45	0.31	0.28	0.28	0.33	0.41	0.55	0.80	1.2	2.0	3.3	5.8	10	20	38	77	160	339	748
$^{0.90}A_n^*$	0.45	0.29	0.23	0.21	0.22	0.25	0.31	0.40	0.56	0.82	1.3	2.0	3.4	5.9	11	20	38	75	150
$^{0.90}\mathcal{R}_n$	1.2	0.73	0.55	0.47	0.44	0.45	0.49	0.57	0.70	0.90	1.2	1.7	2.5	3.8	6.0	9.8	16	28	49

TABLE I: Normalized thickness θ in dimensions $n \leq 19$ for traditional lattice covering ($^{1.0}\mathbb{Z}_n$, $^{1.0}A_n^*$), relaxed lattice covering ($^{\eta}\mathbb{Z}_n$, $^{\eta}A_n^*$) and random template banks ($^{\eta}\mathcal{R}_n$) with covering confidences $\eta = 0.99, 0.95, 0.90$ respectively. Boldface indicates the lowest thickness at given covering confidence η and dimension n .

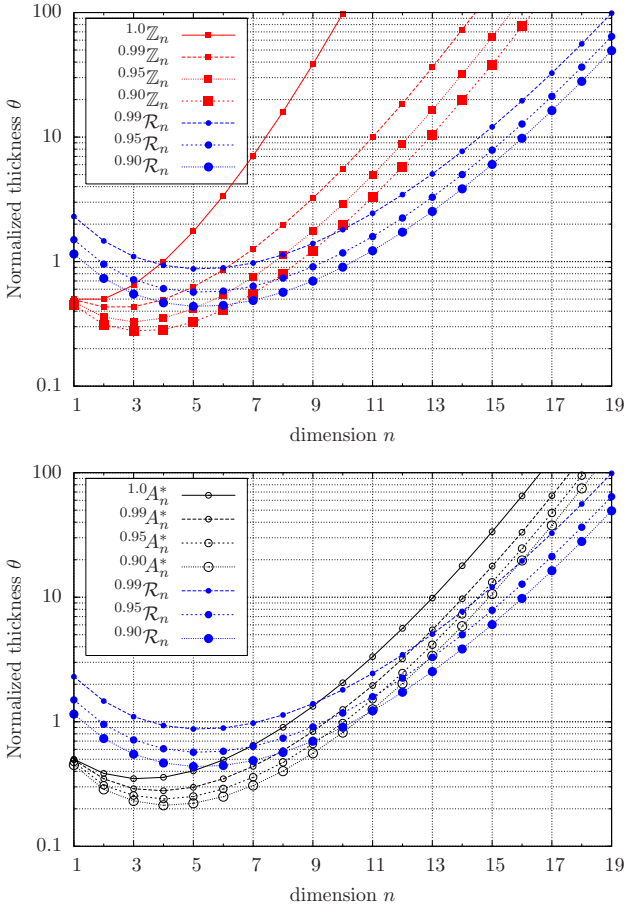


FIG. 10: Normalized thickness θ as function of dimension n for strict lattice covering Λ_n , relaxed lattices $^{\eta}\Lambda_n$ and random template banks $^{\eta}\mathcal{R}_n$ for values $\eta = 0.99, 0.95, 0.90$ of covering confidence. The upper panel is for the \mathbb{Z}_n lattice, while the lower panel shows the case of the A_n^* lattice.

V. DISCUSSION

Our results show that giving up deterministic certainty of coverage of a parameter space can result in large gains in efficiency, by substantially reducing the number of required templates. The prime example of such a *relaxed covering* is the *random template bank* construction, which we defined as templates placed randomly with uniform probability distribution (per proper volume) over the parameter space. Such a random template bank $^{\eta}\mathcal{R}_n(m_*)$ covers any signal location (excluding boundary effects) within mismatch m_* with probability $\eta < 1$. We have found that the template density of these random template banks can be significantly lower than that of even the most efficient (complete) covering, and that this advantage increases in higher parameter-space dimensions.

The exclusion of boundary effects in the random template bank analysis is valid for situations in which the projected length scale of the templates in each of the search dimensions is \ll the width of the parameter space in the corresponding dimension. We should note that boundary effects have also been excluded in the calculation for the lattice coverings (traditional and relaxed) and in practice boundary effects will be equally problematic for all template bank strategies.

Other studies have recently started to investigate a somewhat different random template placement strategy, referred to as “stochastic” template banks [20, 21, 22, 23]. The key difference of these methods is that they involve a “pruning stage” in the random template placement, which is aimed to remove templates that are too close to each other. Albeit not yet completely quantified, these methods could potentially produce even more efficient template banks for equivalent coverage. The disadvantage, however, is that the template placement procedure is more complicated and more computationally intensive,

requiring the comparison of each template with every other. Stochastic template banks might therefore be less suitable for higher dimensions and larger numbers of templates. In addition, it seems unclear how much efficiency can be gained by such a pruning step, especially in higher dimensions.

Relaxed coverings do not provide complete coverage of the template parameter space. Applying such a scheme therefore affects the final results as an additional uncertainty. For gravitational-wave searches such uncertainty can be well maintained at a level comparable to other uncertainties of the problem – typically of the order of a few percent – hence not significantly affecting the overall degree of confidence of the result.

In computationally limited searches, relaxed template banks allow significant reductions in computational cost. Through the re-investment of this saved computational cost, this can yield an increase in sensitivity and breadth of the search.

We have investigated some of the relevant statistical properties of random template banks, in particular the *spatial* parameter-space coverage fraction \mathcal{C} and the worst-case mismatch m_w expected in individual template-bank realizations. We have performed Monte-Carlo simulations to determine the statistical distribution of these quantities, and we have found a rough analytical estimate, which shows reasonably good agreement with the numerical results. These results show that the variance of the spatial coverage fraction \mathcal{C} is inversely proportional to both the dimensionality of the space and the number of random templates. More importantly, the worst-case mismatch $m_w > m_*$ is found to be typically of the same order of magnitude as the prescribed nominal mismatch m_* , and is rapidly approaching m_* with increasing dimension n . At $n = 4$, for example, the largest mismatch for $N_{\mathcal{R}} = 10^8$ templates and $\eta = 0.9$ has an expected value of $\sim 3m_*$ with a standard deviation of $0.09m_*$ (see Fig. 6). At $n = 12$, the distribution peaks at even lower values, with an expectation value of $\sim 1.5m_*$

and a low standard deviation of $0.014m_*$. Even though the coverage at the nominal mismatch m_* only holds in a statistical sense defined by the confidence η , in practice, at dimensions greater than $n = 11$, the worst-case loss is within a factor of 1.5 of the nominal value, with very high confidence.

Inspired by these results, we have also investigated the properties of “relaxed” lattice coverings, which follow the analogous prescription of a nominal covering mismatch m_* , achieved with probability $\eta < 1$ for any signal location. This leads to lattices with larger maximal mismatch $m_{\max} > m_*$, thereby reducing the required number of templates. We have analyzed the properties of relaxed \mathbb{Z}_n and A_n^* lattices, and we have found that the relaxed ${}^{\eta}A_n^*$ lattice provides the most efficient covering found so far for $n \lesssim 10$, for covering confidences $\eta \geq 0.90$. In higher dimensions, however, random template banks outperform any other method considered so far, including relaxed lattice coverings.

Possibly the greatest advantage of random template banks for “real-world” applications, however, is their practical simplicity. Constructing lattice template banks is notoriously difficult, especially regarding the handling of curved parameter spaces and non-constant metrics. Random template banks, on the other hand, are nearly trivial to construct, even in spaces with non-constant metrics: one only needs to adjust the spatial probability density according to the metric determinant.

Acknowledgments

We would like to thank Bruce Allen, Stephen Fairhurst, Graham Woan, Christian Roever, Karl Wette, Matthew Pitkin, John Veitch, Holger Pletsch, and the LIGO Scientific Collaboration continuous waves working group, for many useful discussions.

-
- [1] The LIGO Scientific Collaboration: B. Abbott et al., Nuclear Instruments and Methods in Physics Research A **517**, 154 (2004).
 - [2] B. C. Barish and R. Weiss, Physics Today **52**, 44 (1999).
 - [3] A. Abramovici, W. E. Althouse, R. W. P. Drever, Y. Gursel, S. Kawamura, F. J. Raab, D. Shoemaker, L. Sievers, R. E. Spero, and K. S. Thorne, Science **256**, 325 (1992).
 - [4] B. Willke, P. Aufmuth, C. Aulbert, S. Babak, R. Balasubramanian, B. W. Barr, S. Berukoff, S. Bose, G. Cagnoli, M. M. Casey, et al., Classical and Quantum Gravity **19**, 1377 (2002).
 - [5] S. Gossler et al., Classical and Quantum Gravity **19**, 1835 (2002).
 - [6] F. Acernese et al., Classical and Quantum Gravity **23**, 635 (2006).
 - [7] R. Takahashi and the TAMA Collaboration, Classical and Quantum Gravity **21**, S403 (2004).
 - [8] P.-L. Bender, K. Danzmann, and the LISA Study Team, Max-Planck-Institut für Quantenoptik Report p. 233 (1998), URL <http://www.srl.caltech.edu/lisa/documents/PrePhaseA.pdf>.
 - [9] R. Balasubramanian, B. S. Sathyaprakash, and S. V. Dhurandhar, Phys. Rev. D **53**, 3033 (1996).
 - [10] B. J. Owen, Phys. Rev. D **53**, 6749 (1996).
 - [11] B. J. Owen and B. S. Sathyaprakash, Phys. Rev. D **60**, 022002 (1999).
 - [12] P. R. Brady, T. Creighton, C. Cutler, and B. F. Schutz, Phys. Rev. D **57**, 2101 (1998).
 - [13] The LIGO Scientific Collaboration: B. Abbott et al., Phys. Rev. D **76**, 082001 (2007).
 - [14] The LIGO Scientific Collaboration: B. Abbott et al., *The Einstein@Home search for periodic gravitational waves in LIGO S4 data* (2008), accepted for publication in Phys.

- ical Review D (preprint arXiv:0804.1747v1 [gr-qc]).
- [15] The LIGO Scientific Collaboration: B. Abbott et al., Physical Review D (Particles, Fields, Gravitation, and Cosmology) **72**, 102004 (pages 22) (2005).
 - [16] The LIGO Scientific Collaboration: B. Abbott et al., Physical Review D (Particles, Fields, Gravitation, and Cosmology) **77**, 062002 (pages 13) (2008).
 - [17] The LIGO Scientific Collaboration: B. Abbott et al., Physical Review D (Particles, Fields, Gravitation, and Cosmology) **78**, 102004 (pages 19) (2008).
 - [18] The LIGO Scientific Collaboration: B. Abbott et al., The Astrophysical Journal **681**, 1419 (2008).
 - [19] R. Prix, Classical and Quantum Gravity **24**, 481 (2007), arXiv:0707.0428.
 - [20] S. Babak, *Building a stochastic template bank for detecting massive black hole binaries* (2008), (preprint arXiv:0801.4070v1 [gr-qc]).
 - [21] I. W. Harry, S. Fairhurst, B. S. Sathyaprakash, and B. Allen (2008), in preparation.
 - [22] I. W. Harry, S. Fairhurst, and B. S. Sathyaprakash, *A hierarchical search for gravitational waves from supermassive black hole binary mergers* (2008), (preprint arXiv:0804.3274v1 [gr-qc]).
 - [23] C. van den Broeck, I. Harry, B. S. Sathyaprakash, H. Tagoshi, and H. Takahashi (2008), in preparation.
 - [24] R. Prix, Phys. Rev. D **75**, 023004 (2007).
 - [25] J. H. Conway and N. J. A. Sloane, *Sphere Packings, Lattices, and Groups* (New York: Springer-Verlag, 1993), second ed.
 - [26] A. Okabe, B. Boots, K. Sugihara, and S. N. Chiu, *Spatial tessellations: Concepts and applications of Voronoi diagrams*, Probability and Statistics (Wiley, NYC, 2000), 2nd ed., 671 pages.
 - [27] A. Brakke, K, unpublished (1986), URL <http://www.susqu.edu/brakke/papers/voronoi.htm>.
 - [28] J. H. Conway and N. J. A. Sloane, SIAM Journal on Algebraic and Discrete Methods **5**, 294 (1984).
 - [29] J. Conway and N. Sloane, Information Theory, IEEE Transactions on **28**, 227 (1982), ISSN 0018-9448.