Flat parameter-space metric for all-sky searches for gravitational-wave pulsars

Karl Wette[*] and Reinhard Prix[†]

*Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut), D-30167 Hannover,
Germany and Leibniz Universität Hannover, D-30167 Hannover, Germany*

All-sky, broadband, coherent searches for gravitational-wave pulsars are computationally limited. It is therefore important to make efficient use of available computational resources, notably by minimizing the number of templates used to cover the signal parameter space of sky position and frequency evolution. For searches over the sky, however, the required template density (determined by the parameter-space metric) is different at each sky position, which makes it difficult in practice to achieve an efficient covering. Previous work on this problem has found various choices of sky and frequency coordinates that render the parameter-space metric approximately constant but that are limited to coherent integration times of either less than a few days or greater than several months. These limitations restrict the sensitivity achievable by hierarchical all-sky searches and hinder the development of follow-up pipelines for interesting gravitational-wave pulsar candidates. We present a new flat parameter-space metric approximation and associated sky and frequency coordinates, which do not suffer from these limitations. Furthermore, the new metric is numerically well conditioned, which facilitates its practical use.

I. INTRODUCTION

Gravitational-wave pulsars are rapidly rotating neutron stars that are hypothesized to emit continuous, narrow-band, quasisinusoidal gravitational waves. Nonaxisymmetric distortions of the neutron star, unstable fluid oscillations such as r modes, and free precession due to misaligned symmetry and rotation axes have been proposed as possible emission mechanisms; see Refs. [1,2] for reviews. It remains uncertain, however, whether any of these mechanisms can generate gravitational waves strong enough to be detectable by large-scale ground-based interferometric detectors such as LIGO [3] or Virgo [4]. Several searches using data from the first generation of these detectors have been performed; see Refs. [5–7] for recent results. While energy-conservation-based upper limits on gravitational waves have been beaten for some individual sources [5,8,9], to date no gravitational-wave pulsars have been detected. Second-generation gravitational-wave interferometers such as Advanced LIGO [10], which are currently being constructed, may, however, be sufficiently sensitive to make a first detection [7,11,12].

The data analysis challenge of searching for gravitational-wave pulsar signals in long stretches of interferometer data is formidable. Aside from searches for gravitational waves from known radio and x-ray pulsars, which target a single gravitational-wave template assumed to be phase locked to the known electromagnetic signal [5,9], searches for undiscovered gravitational-wave pulsars must cover a vast parameter space of potential signals. For example, it is not feasible to perform a coherent search over the entire sky and a broad band of signal frequencies, despite the harnessing of $\sim 10^{21}$ floating-point operations of computing power through Einstein@Home [13], a distributed computing platform [7]. This has led to the development of numerous hierarchical pipelines [14–17], where several sensitive, computationally expensive coherent searches of different data segments are incoherently combined using a less sensitive, but computationally cheaper, semicoherent search. Recent progress has been made on the optimal combination of coherent and semicoherent searches [13,18] and on the accurate estimation of the overall search sensitivity [19].

A central issue in constructing a sensitive search for gravitational-wave pulsars is determining the bank of template signals to search over. The signal template of a gravitational-wave pulsar [20] is parametrized by its sky position, often given in terms of right ascension α and declination δ, and its frequency evolution, given most commonly by an initial frequency at some reference time, f_0, and a series of frequency time derivatives, or spindowns: $f_1 \equiv \dot{f}$, $f_2 \equiv \ddot{f}$, etc., up to as many as required. The signal parameters define a manifold $\mathcal{P}$ containing points $\vec{\lambda} = (\alpha, \delta, f_0, f_1, f_2, \ldots)$, each of which corresponds to a signal template; the parameters $\vec{\lambda}_0$ are then coordinates in $\mathcal{P}$. The search must select a finite subset of the members of $\mathcal{P}$, which in turn defines a finite bank of templates to search over. It is improbable, however, that any real signal present in the data will possess parameters precisely matching one of the chosen templates. At best, the real signal will be sufficiently close to one of the templates that it can be recovered with some loss in its signal-to-noise ratio.

An optimal template bank should contain a minimum number of templates, in order to reduce computational cost, with the constraint that any real signal will be recovered with some maximum acceptable loss in the signal-to-noise

---

[*]karl.wette@aei.mpg.de
[†]reinhard.prix@aei.mpg.de

ratio. To achieve these constraints, the parameter space is associated with a *metric*, or distance function, g [14,21,22]. Given two points $\vec{\lambda}_0$, $\vec{\lambda} \in \mathcal{P}$, the metric computes the *mismatch* μ, which gives the loss in the signal-to-noise ratio that would result, were $\vec{\lambda}_0$ a real signal and $\vec{\lambda}$ a nearby template. The template bank is then a finite subset of points $\{\vec{\lambda}_n\} \subset \mathcal{P}$, such that the mismatch between any point $\vec{\lambda}_0 \in \mathcal{P}$ and the "closest" template-bank member $\vec{\lambda} \in \{\vec{\lambda}_n\}$ is bounded by some prescribed maximum mismatch $\mu_{\max}$. If the metric is independent of $\vec{\lambda}$, i.e., it is explicitly *flat*, results from the theory of lattices can be used to place template points on a regular grid, such that the number of points required is minimized [23,24].

The most persistent obstacle to performing optimal template placement for gravitational-wave pulsar searches has been finding a set of coordinates $\vec{\lambda}$, with respect to which the metric is (approximately) constant. For searches targeting a particular point in the sky, where the search is only over the frequency evolution parameters $(f_0, f_1, f_2, \ldots)$, the metric does satisfy this property [25,26], and optimal template placement is used in a search targeting the supernova remnant Cassiopeia A [8]. For searches over the sky, however, the metric is not constant with respect to the angular coordinates (α, δ) [25,27]. An additional practical issue, noted in Ref. [25], is that the metric, when expressed in conventional coordinates, is numerically highly ill conditioned. This makes it very difficult to, for example, compute the transformations of the metric required to implement optimal template placement.

Several alternative sky coordinates and approximate phase models have been developed, with respect to which the metric is constant: the *linear phase models* of Refs. [28,29] and the *global correlation coordinates* of Refs. [17,30]. The principal limitation of these approaches is that, in a hierarchical pipeline, the time span of the data segments that can be coherently searched is restricted to less than a few days [17,29] or greater than several months [25,29]. The segment time span is a free parameter when designing a search for gravitational-wave pulsars; one would ideally choose it based solely on trade offs between sensitivity and computational cost, as detailed in Ref. [13]. These restrictions, however, prevent the sensitivity of an all-sky search from being improved by increasing the length of the coherently searched data segments beyond a few days, and it is not computationally feasible to perform an all-sky search with month-long coherent segments. They also hinder the development of follow-up pipelines [31], where one would like to perform more sensitive searches targeting a small number of interesting gravitational-wave pulsar candidates.

In this paper we present a new solution to these long-standing problems: an explicitly flat parameter-space metric approximation, and associated coordinates, without limitations on the coherent time span and where the metric is well
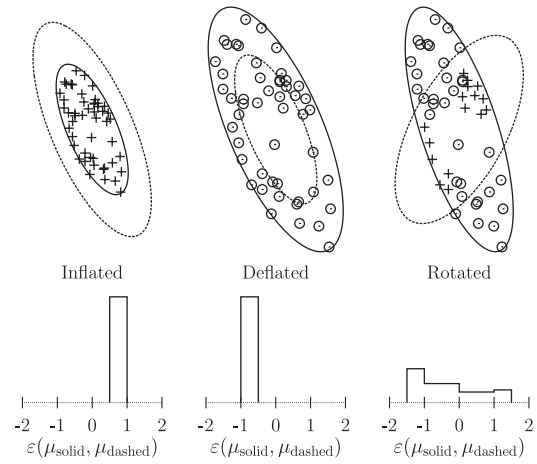

FIG. 1. Illustration of the behavior of the relative error $\varepsilon(\mu_{\text{solid}}, \mu_{\text{dashed}})$ between mismatches computed by a reference metric μ_{solid} and a transformed metric μ_{dashed}. Top: the reference and transformed metrics are plotted as solid and dashed ellipses. The transformed metric has been (left to right), inflated, deflated, and rotated relative to the reference metric. The circles and pluses represent 50 random points, sampled uniformly with respect to the reference metric; points where $\varepsilon(\mu_{\text{solid}}, \mu_{\text{dashed}}) \leq 0$ are plotted as circles, and points where $\varepsilon(\mu_{\text{solid}}, \mu_{\text{dashed}}) > 0$ are plotted as pluses. Bottom: histograms of $\varepsilon(\mu_{\text{solid}}, \mu_{\text{dashed}})$ of the 50 plotted points.

conditioned. Section II of this paper presents relevant background, and Sec. III examines prior research on the parameter-space metric. Section IV presents a new parameter-space metric approximation that is explicitly flat but that embeds $\mathcal{P}$ in a higher-dimensional space. Section V then demonstrates how to reduce the dimensionality of the new metric back to the dimensionality of $\mathcal{P}$, while remaining constant and improving its numerical conditionedness. Section VI discusses the potential uses of the new metric. Details of the numerical simulations presented throughout this paper are found in the Appendix.

When comparing different predicted and/or calculated mismatches, μ_a and μ_b, we compute their *relative error*, which we define following Ref. [25] to be

$$\varepsilon(\mu_a, \mu_b) = \frac{\mu_a - \mu_b}{0.5(\mu_a + \mu_b)}, \qquad \mu_a, \mu_b \geq 0. \quad (1)$$

This definition of relative error is bounded within the range $[-2, 2]$, even for large differences $|\mu_a - \mu_b| \gg 1$, while for $|\varepsilon(\mu_a, \mu_b)| \ll 1$, it approaches more common definitions, e.g., $(\mu_a - \mu_b)/\mu_b$. The behavior of the relative error is illustrated in Fig. 1.

II. BACKGROUND

This section presents background information relevant to this paper. We introduce the gravitational-wave pulsar signal model (Sec. II A), the concept of a parameter-space

metric (Sec. II B), and a useful approximation known as the *phase metric* (Sec. II C).

A. Signal model

A gravitational-wave pulsar signal $h(t, \mathcal{A}, \vec{\lambda})$, as seen in a detector, can be expressed as the sum of the products of four time-independent amplitudes, $\mathcal{A}^i$, and four time-dependent basis waveforms $h_i(t, \vec{\lambda})$ [20]:

$$h(t, \mathcal{A}, \vec{\lambda}) = \sum_{i=1}^{4} \mathcal{A}^i h_i(t, \vec{\lambda}), \qquad (2)$$

where t is the time at the detector. The $\mathcal{A}^i$ are functions of the gravitational-wave strain amplitude h_0, the pulsar's angles of inclination ι and polarization ψ, and the wave's initial phase ϕ_0 at a reference time t_0. The $h_i(t, \vec{\lambda})$ depend on the response functions $F_+(t, \vec{\lambda})$ and $F_\times(t, \vec{\lambda})$ of the detector and on the gravitational-wave phase evolution

$$\frac{\phi(\tau, \vec{\lambda})}{2\pi} = \sum_{s=0}^{s_{\max}} f^{(s)} \frac{(\tau - t_0)^{s+1}}{(s+1)!}, \qquad (3)$$

where τ is the arrival time of a wavefront at the Solar System barycenter (SSB), and $f^{(s)}(t_0) \equiv d^s f / d\tau^s |_{t_0}$ are the time derivatives of the signal frequency $f(\tau)$ at the SSB. The gravitational-wave phase at the detector is found by substituting

$$\tau(t, \vec{\lambda}) - t_0 = (t - t_0) + \frac{\vec{r}(t) \cdot \vec{n}}{c} + \Delta_{\text{relativistic}}, \qquad (4)$$

where t is the arrival time of the wavefront at the detector, $\vec{r}(t)$ is the detector position vector relative to the SSB, $\vec{n}$ is a unit vector pointing from the SSB to the pulsar's position in the sky, and $\Delta_{\text{relativistic}}$ represents the relativistic Einstein and Shapiro delays. The second term of Eq. (4) is also known as the Rømer delay.

The result of coherently matched filtering the signal model with detector data and maximizing over the unknown amplitudes $\mathcal{A}^i$ is known in this context as the $\mathcal{F}$ statistic [20,32]. For a signal with parameters $(\mathcal{A}^i, \vec{\lambda}_0)$, the $\mathcal{F}$ statistic in a template $\vec{\lambda}$ follows a noncentral χ^2 distribution with 4 degrees of freedom and a noncentrality parameter given by the squared signal-to-noise ratio (SNR) $\rho^2(\mathcal{A}, \vec{\lambda}_0; \vec{\lambda})$. For perfectly matched signals,

$$\rho^2(\mathcal{A}, \vec{\lambda}_0; \vec{\lambda}_0) = \frac{2}{S_h(f_0)} \int_{t_0 - T/2}^{t_0 + T/2} dt\, h^2(t, \mathcal{A}, \vec{\lambda}_0), \quad (5)$$

where $S_h(f_0)$ is the (single-sided) power spectral density (PSD) of the detector noise at the signal frequency f_0, and T is the time spanned by the coherently analyzed data. For simplicity, in this paper we will assume that the detector data is continuous (i.e., contains no gaps) and that the PSD in a sufficiently small frequency band surrounding a signal is constant in time and frequency. These limitations are readily addressed in a real implementation of the $\mathcal{F}$ statistic [33].

B. Metric

The mismatch μ between a signal with parameters $\vec{\lambda}_0$ and a nearby template with parameters $\vec{\lambda}$ is defined in terms of the squared SNR [14,25]:

$$\mu_0 = \frac{\rho^2(\mathcal{A}, \vec{\lambda}_0; \vec{\lambda}_0) - \rho^2(\mathcal{A}, \vec{\lambda}_0; \vec{\lambda})}{\rho^2(\mathcal{A}, \vec{\lambda}_0; \vec{\lambda}_0)}, \qquad (6)$$

where the mismatched SNR $\rho^2(\mathcal{A}, \vec{\lambda}_0; \vec{\lambda})$ is given by Eq. (28) of Ref. [25]. In the numerical simulations presented in this paper, μ_0 is calculated as follows: a gravitational-wave pulsar signal is generated with parameters $\vec{\lambda}_0$ and searched for using the $\mathcal{F}$ statistic at points $\vec{\lambda}_0$ and $\vec{\lambda}$, returning the values $\mathcal{F}(\vec{\lambda}_0)$ and $\mathcal{F}(\vec{\lambda})$, respectively. The mismatch is then calculated using Eq. (6) and the relation

$$2\mathcal{F}(\vec{\lambda}) = E[2\mathcal{F}(\vec{\lambda})] = 4 + \rho^2(\mathcal{A}, \vec{\lambda}_0; \vec{\lambda}). \qquad (7)$$

No simulated noise is added to the gravitational-wave pulsar signal, and thus $2\mathcal{F}(\vec{\lambda})$ is equal to its expectation value $E[2\mathcal{F}(\vec{\lambda})]$.

If the difference $\Delta\vec{\lambda} = \vec{\lambda} - \vec{\lambda}_0$ is small enough, $\rho^2(\mathcal{A}, \vec{\lambda}_0; \vec{\lambda})$ can be Taylor expanded with respect to $\rho^2(\mathcal{A}, \vec{\lambda}_0; \vec{\lambda}_0)$. The $\mathcal{F}$-statistic mismatch μ_0 is then approximated by

$$\mu_0 \approx \mu_{\mathbf{g}} \equiv \Delta\vec{\lambda} \cdot \mathbf{g}\Delta\vec{\lambda}. \qquad (8)$$

The metric mismatch $\mu_{\mathbf{g}}$ is calculated via the metric $\mathbf{g}$, for which the coefficients are

$$g(\lambda_i, \lambda_j) = \frac{-1}{2\rho^2(\mathcal{A}, \vec{\lambda}_0; \vec{\lambda}_0)} \frac{\partial \rho^2(\mathcal{A}, \vec{\lambda}_0; \vec{\lambda})}{\partial \lambda_i \partial \lambda_j} \bigg|_{\vec{\lambda} = \vec{\lambda}_0}. \qquad (9)$$

There are no terms proportional to the first derivatives of $\rho^2(\mathcal{A}, \vec{\lambda}_0; \vec{\lambda})$ with respect to $\vec{\lambda}$, since by definition $\rho^2(\mathcal{A}, \vec{\lambda}_0; \vec{\lambda})$ is a maximum at the signal location $\vec{\lambda}_0$.

The matrix $\mathbf{g}$ is positive definite by construction [25], and thus the region $\mu \leq \mu_{\max}$ forms an ellipsoid, centered on $\vec{\lambda}_0$, in the parameter space $\mathcal{P}$. If $\mathbf{g}$ is flat, each template point $\vec{\lambda}_n \in \mathcal{P}$ will be surrounded by an identical ellipsoid. We can then apply a global coordinate transformation to $\mathcal{P}$ that maps the ellipsoids to spheres, each with a template point at its center. The problem of template placement is now equivalent to the sphere-covering problem in lattice theory [23], and the solution that minimizes the number of template points is to place them at the vertices of a lattice that is known to achieve the best possible covering. The best choice of lattice depends on the dimensionality of $\mathcal{P}$; for example, in two dimensions it is the hexagonal lattice [34]. If $\mathbf{g}$ is not flat (or just nonconstant), however, other

methods of template placement, such as random or stochastic algorithms [35–37], must be employed.

C. Phase metric

For the $\mathcal{F}$ statistic, the $g(\lambda_i, \lambda_j)$ are complicated functions depending on the unknown amplitudes $\mathcal{A}$, as described in Ref. [25], which include both derivatives of the amplitude modulation $F_+(t, \vec{\lambda})$ and $F_\times(t, \vec{\lambda})$ of the signal, and derivatives of the phase modulation, given by $\phi(t, \vec{\lambda})$. If T is large compared to a day, however, the contribution of the more rapid ($\gtrsim 100/s$) phase modulation dominates that of the slower ($\lesssim 1/\text{day}$) amplitude modulation. In this limit, the metric reduces to a simplified form, known as the phase metric, for which the coefficients involve only derivatives of $\phi(t, \vec{\lambda})$:

$$g(\lambda_i, \lambda_j) = \left[\frac{\partial \phi(t, \vec{\lambda})}{\partial \lambda_i}, \frac{\partial \phi(t, \vec{\lambda})}{\partial \lambda_j} \right], \quad (10)$$

where we define the operators

$$[x(t), y(t)] = \langle x(t)y(t) \rangle - \langle x(t) \rangle \langle y(t) \rangle, \quad (11)$$

$$\langle x(t) \rangle = \frac{1}{T} \int_{t_0 - T/2}^{t_0 + T/2} dt\, x(t). \quad (12)$$

An equivalent expression for the metric was also obtained in Ref. [14] by instead assuming a simplified signal model where the amplitude motion is discarded. If the phase $\phi(t, \vec{\lambda})$ is linear in the coordinates $\vec{\lambda}$, then the $g(\lambda_i, \lambda_j)$ are independent of $\vec{\lambda}$, and **g** is therefore flat. Thus, the problem of finding a constant metric approximation is reduced to one of linearizing $\phi(t, \vec{\lambda})$ with respect to its coordinates.

To linearize $\phi(t, \vec{\lambda})$ with respect to the frequency and spindown coordinates $f^{(s)}$, we first substitute Eq. (4) into Eq. (3), neglecting the relativistic terms that are not important for template placement:

$$\frac{\phi(t, \vec{\lambda})}{2\pi} \approx \sum_{s=0}^{s_{\max}} \frac{f^{(s)}}{(s+1)!} \left[\Delta t + \frac{\vec{r}(t) \cdot \vec{n}}{c} \right]^{s+1}, \quad (13)$$

where $\Delta t = t - t_0$. We now expand the factor $[\dots]^{s+1}$ and retain only the first two leading-order terms in Δt. This approximation can be made because $\vec{r}(t) \cdot \vec{n}/c \lesssim 500$ sec (the approximate light travel time from the Sun to the Earth) while $\Delta t \sim T$, and so $\Delta t \gg \vec{r}(t) \cdot \vec{n}/c$ for $T \gtrsim$ days. The approximate Eq. (13) now reads

$$\frac{\phi(t, \vec{\lambda})}{2\pi} \approx \sum_{s=0}^{s_{\max}} f^{(s)} \frac{\Delta t^{s+1}}{(s+1)!} + \frac{\vec{r}(t) \cdot \vec{n}}{c} \sum_{s=0}^{s_{\max}} f^{(s)} \frac{\Delta t^s}{s!}. \quad (14)$$

The summation in the second right-hand-side term is precisely $f(t)$, the instantaneous frequency of the signal at time t. Using the same argument, we see that the derivatives of $\phi(t, \vec{\lambda})$ with respect to the $f^{(s)}$ [which will appear in

Eq. (10)] of the second term will be small relative to the first, i.e., $\Delta t^{s+1} \gg \Delta t^s \vec{r}(t) \cdot \vec{n}/c$. Hence, $f(t)$ may be approximated by some constant $f_{\max}$, usually chosen conservatively to be the maximum of $f(t)$ over T. The resulting approximate phase is now

$$\frac{\phi(t, \vec{\lambda})}{2\pi} \approx \sum_{s=0}^{s_{\max}} f^{(s)} \frac{\Delta t^{s+1}}{(s+1)!} + \frac{\vec{r}(t) \cdot \vec{n}}{c} f_{\max}. \quad (15)$$

III. PRIOR WORK

While it is straightforward to obtain a linear phase model $\phi(t, \vec{\lambda})$ with respect to the frequency and spindown coordinates $f^{(s)}$, as shown in Sec. II C, the same cannot be said of the sky coordinates, which enter Eq. (3) through the sky position vector $\vec{n}$. If, for example, we choose right ascension α and declination δ as sky coordinates, then the derivatives of $\phi(t, \vec{\lambda})$ themselves depend on the sky coordinates,

$$d\phi \propto (r_z \cos \delta - r_y \sin \alpha \sin \delta - r_x \cos \alpha \sin \delta) d\delta$$
$$+ (r_y \cos \alpha \cos \delta - r_x \sin \alpha \cos \delta) d\alpha + \cdots, \quad (16)$$

where $\vec{n} = (\cos \alpha \cos \delta, \sin \alpha \cos \delta, \sin \delta)$, and $\vec{r}(t) = (r_x, r_y, r_z)$ are expressed in equatorial coordinates (x, y, z). If, instead, two components of the vector $\vec{n} = (n_x, n_y, n_z)$ are chosen, e.g., n_x and n_y, the constraint $|\vec{n}| = 1$ requires that the third component is a function of the other two, i.e., $n_z = (1 - n_x^2 - n_y^2)^{1/2}$, and so the derivatives of $\phi(t, \vec{\lambda})$ still depend on the coordinates:

$$d\phi \propto \left(r_x - r_z \frac{n_x}{n_z} \right) dn_x + \left(r_y - r_z \frac{n_y}{n_z} \right) dn_y + \cdots. \quad (17)$$

This section presents two prior approaches to this problem: the linear phase models (Sec. III A) and the global correlation coordinates (Sec. III B). The new approach to this problem taken in this paper is presented in Sec. IV.

A. Linear phase models

The linear phase models of [28,29] express $\vec{n} = (n_X, n_Y, n_Z)$ in ecliptic coordinates (X, Y, Z) and adopt the X and Y components (n_X, n_Y) as sky coordinates. The restriction $|\vec{n}| = 1$ then requires $n_Z = \sqrt{1 - n_X^2 - n_Y^2}$. We write

$$\vec{r}(t) \cdot \vec{n} = r_X(t)n_X + r_Y(t)n_Y + r_Z(t)n_Z \quad (18)$$

$$= [r_{sX}(t) + r_{oX}(t)]n_X$$
$$+ [r_{sY}(t) + r_{oY}(t)]n_Y + r_{sZ}(t)n_Z, \quad (19)$$

where the detector position vector $\vec{r}(t) = \vec{r}_s(t) + \vec{r}_o(t)$ is decomposed into its diurnal and orbital components in ecliptic coordinates, $\vec{r}_s(t) = [r_{sX}(t), r_{sY}(t), r_{sZ}(t)]$ and $\vec{r}_o(t) = [r_{oX}(t), r_{oY}(t), 0]$, respectively. (This assumes a
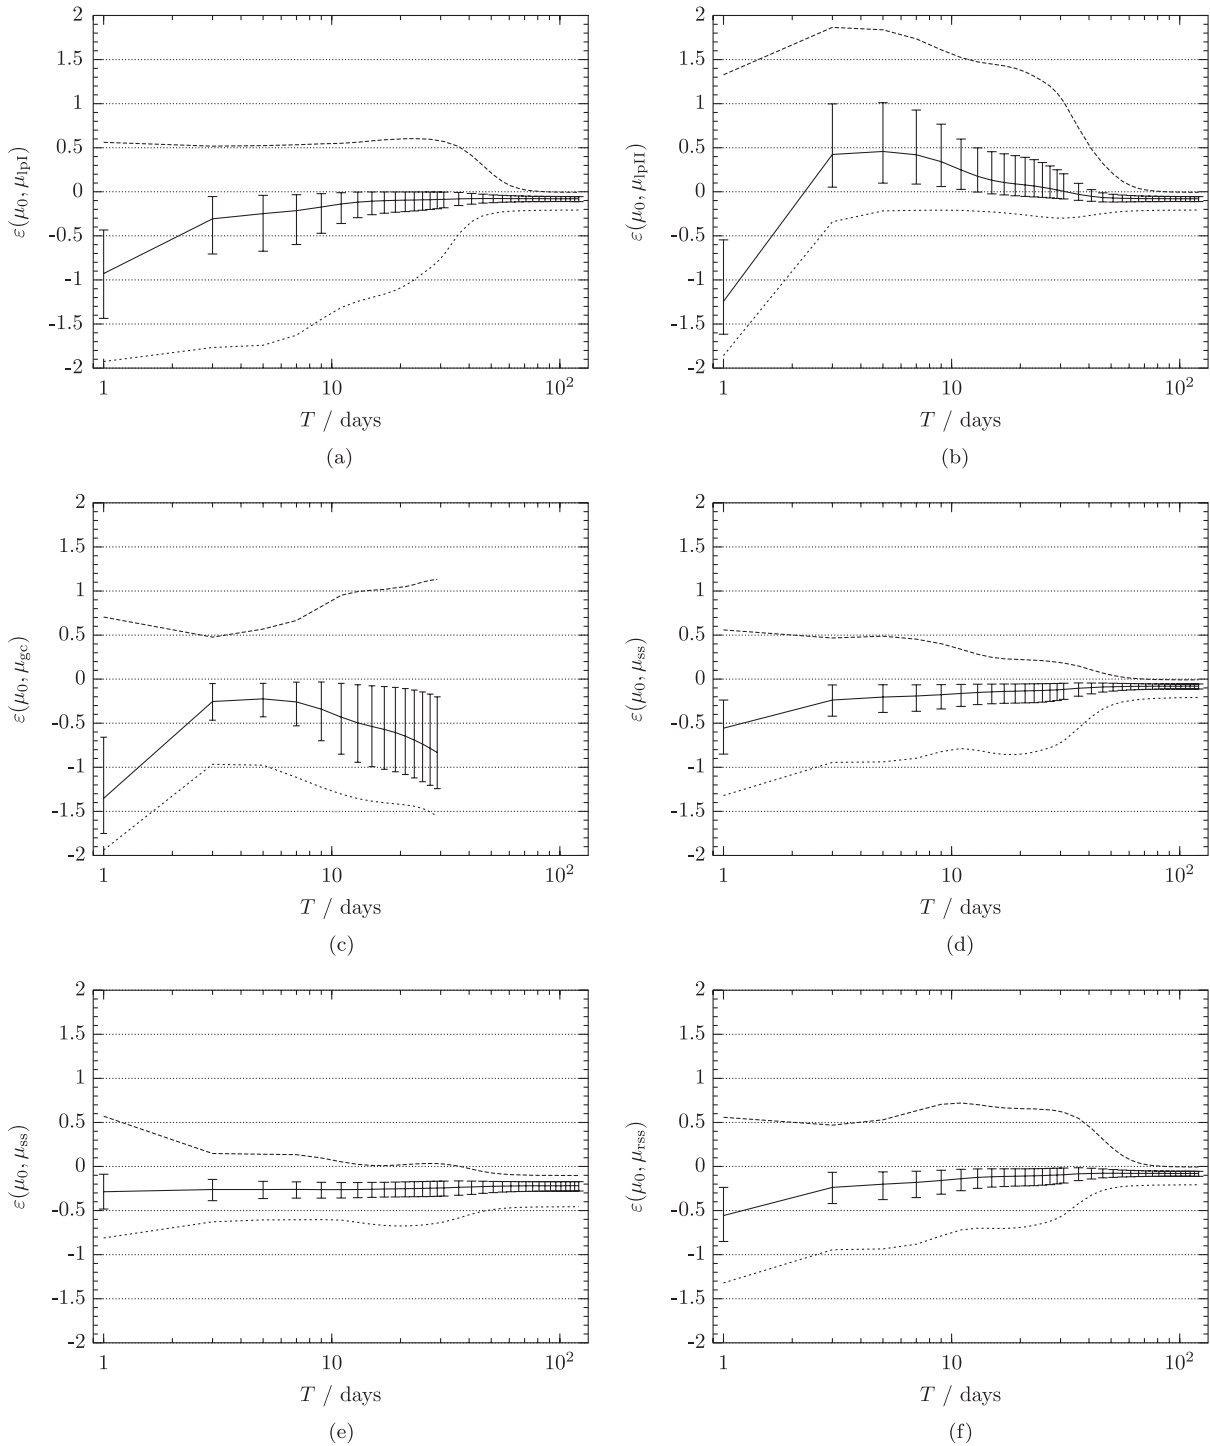
FIG. 2.   Relative errors as a function of T between mismatches calculated from the $\mathcal{F}$-statistic, μ_0, and predicted by: (a) the phase metric using linear phase model I, $\varepsilon(\mu_0, \mu_{\mathrm{lpI}})$, for $\mu_0 \leq 0.2$; (b) the phase metric using linear phase model II, $\varepsilon(\mu_0, \mu_{\mathrm{lpII}})$ for $\mu_0 \leq 0.2$; (c) the global correlation metric, $\varepsilon(\mu_0, \mu_{\mathrm{gc}})$ for $\mu_0 \leq 0.2$ and $T < 30$ days; (d), (e) the super-sky metric, $\varepsilon(\mu_0, \mu_{\mathrm{ss}})$ for $\mu_0 \leq 0.2$ and $0.2 \leq \mu_0 \leq 0.6$ respectively; and (f) the reduced supersky metric, $\varepsilon(\mu_0, \mu_{\mathrm{rss}})$ for $\mu_0 \leq 0.2$. Plotted are the median (solid line), the 25th–75th percentile range (error bars), and the 2.5th (dotted line) and 97.5th (dashed line) percentiles. Only first spindown is used.

planar Earth orbit, which is not exactly satisfied in reality, due to, e.g., the Earth–Moon interaction.) We see that $\vec{r}(t) \cdot \vec{n}$ is linear in n_X and n_Y only if $r_{\mathrm{s}Z}(t) = 0$, i.e., only if the diurnal motion of the detector in the ecliptic Z direction is neglected. Two linear phase models are presented in Ref. [28], which achieve this; in linear phase model I, $r_{\mathrm{s}Z}(t)$ is discarded, and in linear phase model II (also known as the *orbital metric* [25]), the entire diurnal motion

$\vec{r}_s(t)$ is discarded. Simulations investigating the accuracy of signal parameter estimation using the linear phase models were presented in Ref. [28].

Figures 2(a) and 2(b) show the relative errors $\varepsilon(\mu_0, \mu_{\rm lpI})$ and $\varepsilon(\mu_0, \mu_{\rm lpII})$ between mismatches predicted by the phase metric using the linear phase models I and II vs the mismatch μ_0 calculated from the $\mathcal{F}$ statistic, as given in Eq. (6). Linear phase model I [Fig. 2(a)] models the $\mathcal{F}$-statistic mismatch well, converging to $|\varepsilon(\mu_0, \mu_{\rm lpI})| \lesssim 0.1$ for $T \gtrsim 20$ days. The spread of errors can, however, be quite large: the 25–75th percentile range (error bars) is $\gtrsim 0.5$ for $T \lesssim 10$ days; the 2.5–97.5th percentile range (dotted to dashed lines) is $\gtrsim 1$ for $T \lesssim 35$ days. Linear phase model II [Fig. 2(b)] is also a reasonable approximation in the longer-T limit, consistent with similar simulations presented in Ref. [25]. It initially underestimates μ_0 [i.e., $\varepsilon(\mu_0, \mu_{\rm lpII}) > 0$], up to 0.5 at $T \sim 5$ days, before converging to an overestimate [$\varepsilon(\mu_0, \mu_{\rm lpII}) < 0$] of ~ 0.1 for $T \gtrsim 30$ days.

The relative errors between the linear phase models and the $\mathcal{F}$ statistic exhibit features common to most of the phase metrics examined in this paper. In general, the phase metrics tend to overestimate the $\mathcal{F}$-statistic mismatch [i.e., $\varepsilon(\mu_0, \mu_{...}) < 0$], which would lead to a conservative, overdense template bank. This is due to the neglect of higher-order terms in $\Delta\vec{\lambda}$ in the derivation of the metric [Eq. (8)], which leads to the metric-predicted mismatch typically being larger than the $\mathcal{F}$-statistic mismatch for the same $\Delta\vec{\lambda}$. This effect prevents the phase metric from exactly predicting the $\mathcal{F}$-statistic mismatch at long T. In addition, at $T \sim 1$ day, the diurnal amplitude modulation of the gravitational-wave pulsar signal, neglected in the phase metric approximation, changes the size and orientation of the full $\mathcal{F}$-statistic metric [25] relative to the phase metric. As illustrated in Fig. 1, this can lead to further overestimation of the $\mathcal{F}$-statistic mismatch by the phase metric.

B. Global correlation coordinates

The global correlation coordinates of Refs. [17,30] adopt the x and y coefficients of $\vec{n}$ in equatorial coordinates, (n_x, n_y), as sky coordinates. Decomposing the detector position vector into its diurnal and orbital components in equatorial coordinates, $\vec{r}_s(t) = [r_{sx}(t), r_{sy}(t), 0]$ and $\vec{r}_o(t)$, respectively, we have

$$\vec{r}(t) \cdot \vec{n} = r_{sx}(t)n_x + r_{sy}(t)n_y + \vec{r}_o(t) \cdot \vec{n}. \quad (20)$$

The global correlation coordinates then absorb $\vec{r}_o(t) \cdot \vec{n}$ into the frequency and spindown parameters $f^{(s)}$, by Taylor-expanding $\vec{r}_o(t)$ in time and defining new parameters $\nu^{(s)}(t)$ that are functions of $f^{(s)}$ and $\vec{n}$. Assuming that no second- or higher-order spindowns are required, i.e., $f^{(s)} = 0$ for $s > 1$, the global correlation coordinates $\nu(t)$ at time t are functions of the frequency $f(t) = f + \dot{f}\Delta t$ and spindown $\dot{f}$ at time t [30]:

$$\nu(t) = f(t) + \frac{f(t)\dot{\vec{r}}_o(t) + \dot{f}\vec{r}_o(t)}{c} \cdot \vec{n}, \quad (21a)$$

$$\dot{\nu}(t) = \dot{f} + \frac{f(t)\ddot{\vec{r}}_o(t) + 2\dot{f}\dot{\vec{r}}_o(t)}{c} \cdot \vec{n}. \quad (21b)$$

The phase is then linear in the coordinates n_x, n_y, $\nu(t)$, and $\dot{\nu}(t)$. The approximation of $\vec{r}_o(t)$ as a Taylor series limits the validity of these coordinates to $T \lesssim 2$–10 days, depending on the frequency searched [17]. A similar linearized phase model is also presented in Ref. [29]; they found that the adequacy of this model was limited to $T \lesssim 8$–14 days, depending on the search frequency, number
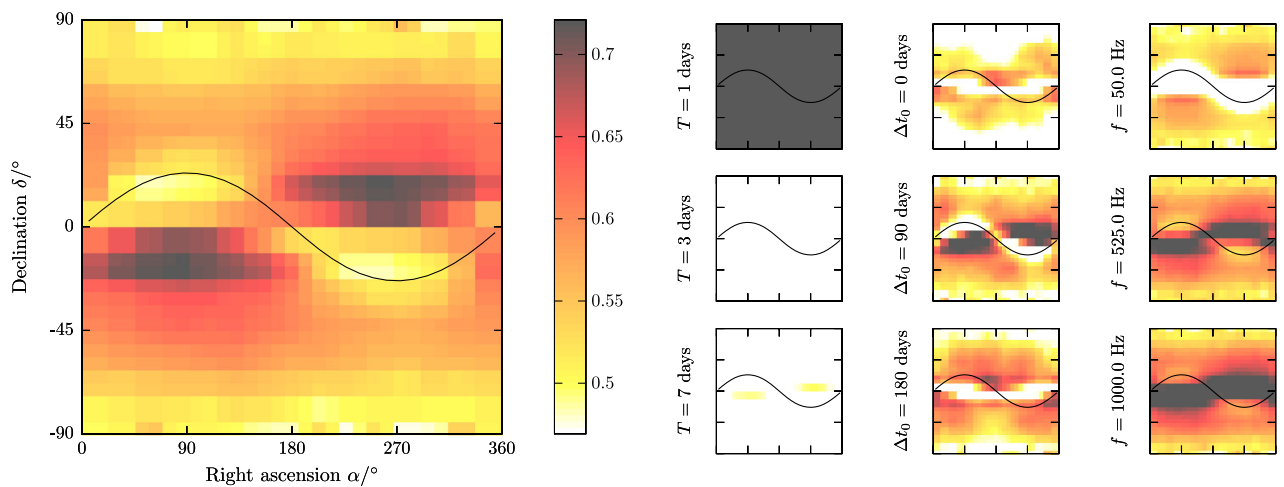

FIG. 3 (color online). Median magnitude of the relative error $|\varepsilon(\mu_0, \mu_{\rm gc})|$, as a function of α and δ, for $\mu_0 \leq 0.2$. The ecliptic equator is overplotted in black. Left: median $|\varepsilon(\mu_0, \mu_{\rm gc})|$ over time spans $1 \leq T \leq 29$ days and over all simulation reference times Δt_0 and maximum frequencies $f_{\rm max}$; see the Appendix. Right: median $|\varepsilon(\mu_0, \mu_{\rm gc})|$ at fixed values of T, Δt_0, and $f_{\rm max}$; axis ranges and color values are the same as for the left-hand-side plot. Only the first spindown is used.

of spindowns, and the parameter estimation accuracy required. An examination of the limitations of the global correlation method is presented in Ref. [38].

Figure 2(c) plots the relative error $\varepsilon(\mu_0, \mu_{gc})$ between mismatches predicted by the global correlation metric and calculated from the $\mathcal{F}$ statistic via Eq. (6), up to $T < 30$ days. The global correlation coordinates perform best when $3 \lesssim T \lesssim 7$ days; below ~ 3 days and above ~ 15 days, $|\varepsilon(\mu_0, \mu_{gc})| \gtrsim 0.5$, broadly consistent with the domain of validity found in Ref. [17]. Unlike other phase metrics examined in this paper, the global correlation coordinates perform worse at long T due to the breakdown of the Taylor expansion of the orbital motion. Figure 3 plots the median error magnitude $|\varepsilon(\mu_0, \mu_{gc})|$ as a function of sky position, over the full ranges of simulation parameters and at fixed values of T, Δt_0, and f_{max}. The smallest $|\varepsilon(\mu_0, \mu_{gc})|$ are along the ecliptic equator and at the poles; the largest are at the points $\alpha = 180° \pm 90°$, $\delta = \pm 20°$. The errors are independent of sky position at fixed T but become sky-position dependent when considering fixed Δt_0 and f_{max}.

IV. SUPERSKY METRIC

In this paper, to linearize the phase metric $\phi(t, \vec{\lambda})$ with respect to the sky coordinates $\vec{n}$, we simply relax the constraint that $|\vec{n}| = 1$ and instead consider each of the three components of $\vec{n}$ to be independent. It follows that the phase variation $d\phi(t, \vec{\lambda})$ is independent of sky position:

$$\mathrm{d}\phi \propto \vec{r}(t) \cdot \mathrm{d}\vec{n} + \cdots. \qquad (22)$$

This idea is the foundation of the parameter-space metric described in this paper. We refer to the phase metric expressed in the three sky coordinates $\vec{n} \in \mathbb{R}^3$ as the *supersky* metric $\mathbf{g}_{ss}$; it is given, e.g., in equatorial coordinates by

$$\mathbf{g}_{ss} = \begin{pmatrix} g_{n_x,n_x} & g_{n_x,n_y} & g_{n_x,n_z} & g_{n_x,f} & g_{n_x,\dot{f}} & \cdots \\ g_{n_x,n_y} & g_{n_y,n_y} & g_{n_y,n_z} & g_{n_y,f} & g_{n_y,\dot{f}} & \cdots \\ g_{n_x,n_z} & g_{n_y,n_z} & g_{n_z,n_z} & g_{n_z,f} & g_{n_z,\dot{f}} & \cdots \\ g_{n_x,f} & g_{n_y,f} & g_{n_z,f} & g_{f,f} & g_{f,\dot{f}} & \cdots \\ g_{n_x,\dot{f}} & g_{n_y,\dot{f}} & g_{n_z,\dot{f}} & g_{f,\dot{f}} & g_{\dot{f},\dot{f}} & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}, \qquad (23)$$

where the elements g_{λ_i, λ_j} are given by Eq. (10).

Geometrically, relaxing the constraint $|\vec{n}| = 1$ amounts to embedding the two-dimensional physical sky in three-dimensional space $\mathbb{R}^3$, as illustrated in Fig. 4. The physical sky metric is recovered by reimposing the restriction $|\vec{n}| = 1$, which is equivalent to finding the intersection of the sky sphere with a supersky metric ellipsoid centered on a point on the sky sphere. From Fig. 4 it is evident that, while the supersky metric ellipsoids have the same shape and orientation, regardless of their location, their intersections with the
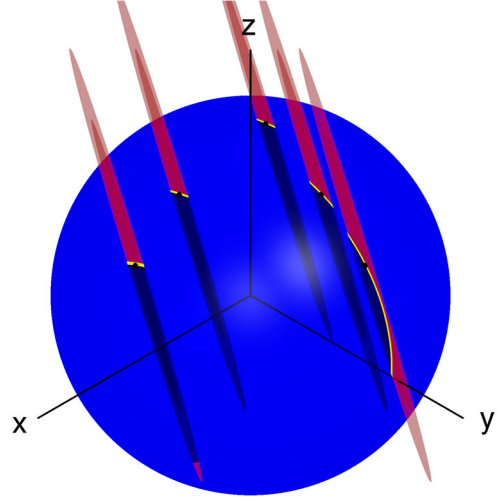

FIG. 4 (color online). Metric ellipsoids of the supersky metric $\mathbf{g}_{ss}$, at five example points, for $T = 4$ days and $\mu_{max} = 30$. Their intersections with the sky sphere $|\vec{n}| = 1$ reproduce the physical sky metric.

sky sphere produce shapes of differing sizes and orientations. This confirms that the metric is flat in the supersky $\vec{n} \in \mathbb{R}^3$ but not on the physical sky $|\vec{n}| = 1$.

Figures 2(d) and 2(e) plot the relative error $\varepsilon(\mu_0, \mu_{ss})$ between mismatch predicted by the supersky metric, and calculated from the $\mathcal{F}$ statistic, for $\mu_0 \leq 0.2$ and $0.2 \leq \mu_0 \leq 0.6$, respectively. For template placement, we are most interested in the median error and whether the supersky metric significantly underestimates the $\mathcal{F}$-statistic mismatch, which would lead to undercovering of parameter-space regions. For $\mu_0 \leq 0.2$, the median error is $|\varepsilon(\mu_0, \mu_{ss})| \lesssim 0.3$ for $T \gtrsim 2$ days and $|\varepsilon(\mu_0, \mu_{ss})| \lesssim 0.2$ for $T \gtrsim 7$ days; the 25–75th percentile range (error bars) is within ~ 0.5 for $T \gtrsim 1$ day and within ~ 0.2 for $T \gtrsim 40$ days. The supersky metric overestimates the $\mathcal{F}$-statistic mismatch on average [$\varepsilon(\mu_0, \mu_{ss}) < 0$], leading to slightly conservative template placement. Only for a small number of trials will the supersky metric significantly underestimate the $\mathcal{F}$-statistic mismatch [$\varepsilon(\mu_0, \mu_{ss}) > 0$], e.g., at $T \sim 3$ days, $\varepsilon(\mu_0, \mu_{ss}) > 0.5$ for 2.5% of trials (above dashed line). For $0.2 \leq \mu_0 \leq 0.6$, the median error magnitude is only slightly worse, $|\varepsilon(\mu_0, \mu_{ss})| \lesssim 0.25$.

Figure 5 plots the median error magnitude $|\varepsilon(\mu_0, \mu_{ss})|$ as a function of sky position, over the full ranges of simulation parameters, and at fixed values of T, Δt_0, and f_{max}. Generally, the $|\varepsilon(\mu_0, \mu_{ss})|$ are largest in two bands above and below the ecliptic equator (plotted in black) and smallest near the poles. The area of the sky where the error is large decreases as a function of T, increases as a function of $f_{max} \lesssim 500$ Hz, and is a more complicated function of the reference time of the data.

While testing the supersky metric by generating an actual template bank is beyond the scope of this paper, we believe the performance of the metric, demonstrated
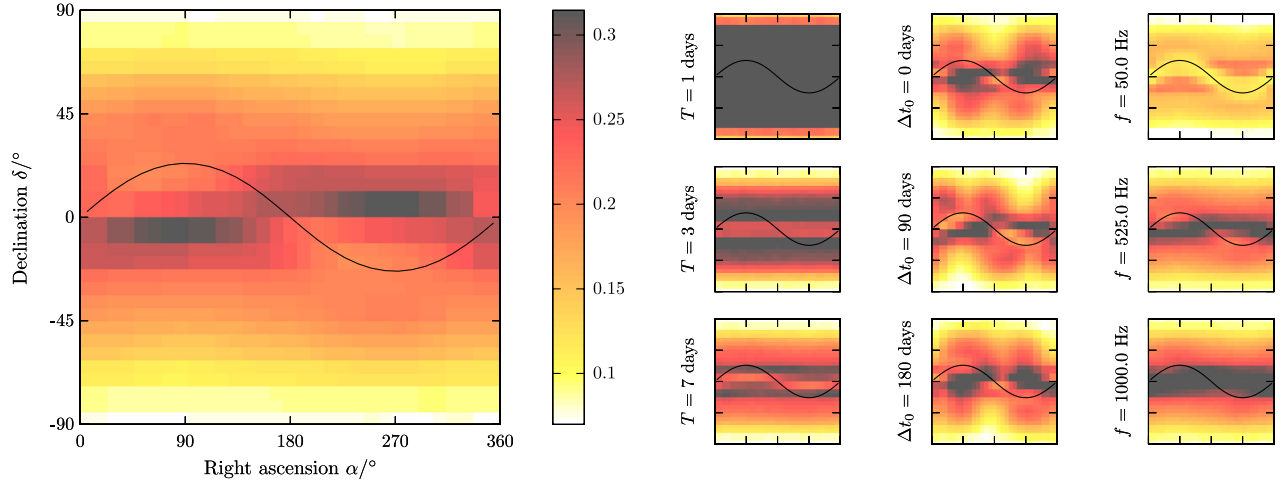
FIG. 5 (color online). Median magnitude of the relative error $|\varepsilon(\mu_0, \mu_{\mathrm{ss}})|$, as a function of α and δ, for $\mu_0 \lesssim 0.2$. The ecliptic equator is over-plotted in black. Left: median $|\varepsilon(\mu_0, \mu_{\mathrm{ss}})|$ over time-spans $1 \leq T \leq 29$ days, and all simulation reference times Δt_0 and maximum frequencies f_{max}; see the Appendix. Right: median $|\varepsilon(\mu_0, \mu_{\mathrm{ss}})|$ at fixed values of T, Δt_0, and f_{max}; axis ranges and color values are the same as for the left-hand-side plot. Only first spindown is used.

here, is sufficiently accurate for this task. That the phase metric, despite its approximations, still closely models the $\mathcal{F}$-statistic mismatch (even at large mismatches of 0.6) demonstrates the relative importance of maintaining phase coherence (as opposed to amplitude consistency) between a gravitational-wave pulsar signal and a search template. In Sec. IV B, we demonstrate the suitability of the coordinates $\vec{n}$ as sky coordinates, as opposed to the angular coordinates (α, δ).

A. Comparison to the $\mathcal{F}$-statistic metric

The results presented in the previous section may be compared to similar simulations performed in Ref. [25]. Figure 15 of that paper compares measured $\mathcal{F}$-statistic mismatches, μ_0, to mismatches predicted by the full $\mathcal{F}$-statistic metric, $\mu_{\mathcal{F}}$; while Figs. 4 and 5 of that paper compare $\mu_{\mathcal{F}}$ to mismatches predicted by the phase metric, which is essentially μ_{ss}. Taken together, these three figures show, for mismatches up to 0.5, relative errors $\varepsilon(\mu_0, \mu_{\mathcal{F}}) \lesssim 0.05$ and $\varepsilon(\mu_{\mathcal{F}}, \mu_{\mathrm{ss}}) \lesssim 0.4$ for $T \sim 1$ day and $\varepsilon(\mu_0, \mu_{\mathcal{F}}) \sim 0 \pm 0.01$ and $\varepsilon(\mu_{\mathcal{F}}, \mu_{\mathrm{ss}}) \sim 0 \pm 0.05$ for $T \gtrsim 7$ days.

While the simulations of Ref. [25] suggest better agreement between the $\mathcal{F}$ statistic and the phase/supersky metric than is shown in, e.g., Fig. 2(d), one must first account for a subtle difference between the two sets of simulations: the distributions of the randomly generated sky coordinate offsets $|\Delta \alpha|$ and $|\Delta \delta|$ from which the mismatches are calculated. The simulations in Ref. [25] drew offsets uniformly distributed in $|\Delta \alpha|$ and $|\Delta \delta|$ and furthermore excluded "large" sky offsets, quantified by

$$\Delta \hat{\Omega} \approx 10^{-4} f T \sqrt{(\Delta \alpha \cos \delta)^2 + (\Delta \delta)^2} \gtrsim 5. \quad (24)$$

In contrast, the simulations in this paper sample offsets uniformly in coordinates defined by the eigenvectors of

the supersky metric, i.e., the axes of the metric ellipsoids (see the Appendix for details). As seen in Fig. 4, the supersky metric ellipsoids are typically highly elongated along one semimajor axis. Giving equal weight to each metric semimajor axis when sampling offsets, as is done in this paper, leads to a greater number of large sky offsets than is achieved by uniform sampling in $|\Delta \alpha|$ and $|\Delta \delta|$ (see the discussion in Sec. IV E 2 of Ref. [25]). Figure 6
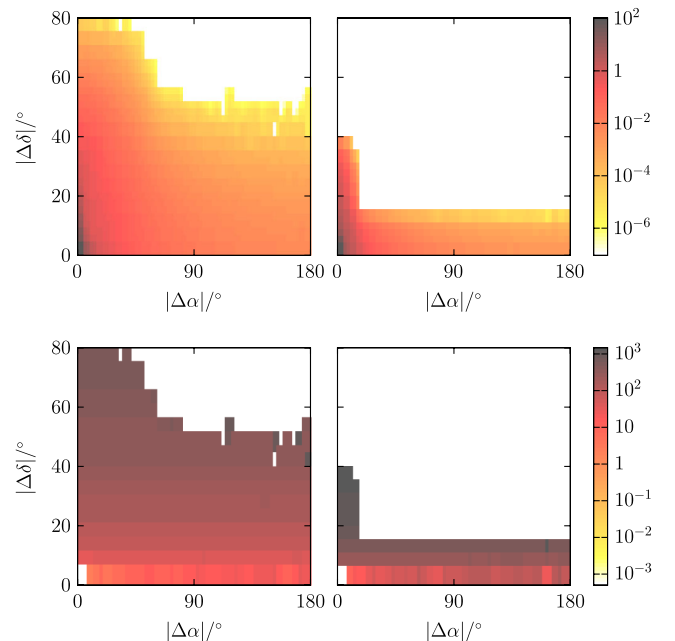

FIG. 6 (color online). Probability density of simulated coordinate offsets $|\Delta \alpha|$ and $|\Delta \delta|$ (top row) and the minimum of $\Delta \hat{\Omega}$ [Eq. (24)] at fixed $|\Delta \alpha|$ and $|\Delta \delta|$ (bottom row), over all simulation reference times and maximum frequencies and at $T = 1$ day (left column) and $T = 7$ days (right column).
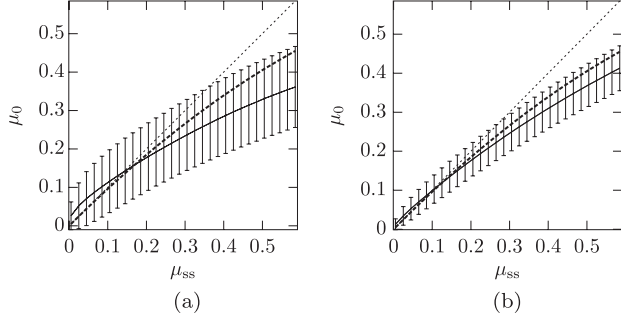
FIG. 7. Mean (solid line) and standard deviation (error bars) of μ_0 as a function of μ_{ss}, over all simulation reference times and maximum frequencies, and for (a) $T = 1$ day and (b) $T = 7$ days. The dotted lines plot $\mu_0 = \mu_{ss}$; the dashed lines plot $\mu_0 = \mu_{ss} - 0.38\mu_{ss}^2$.
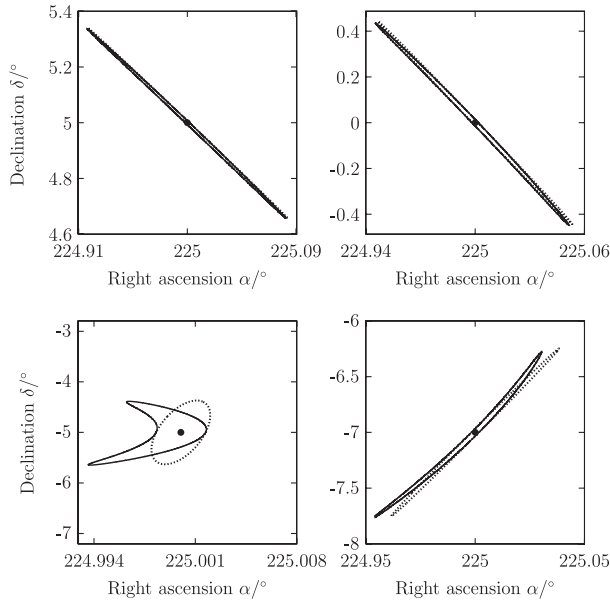

FIG. 8. True mismatch region $\mu \leq \mu_{max}$ of the sky metric (solid line) and the local metric ellipse in (α, δ) given by Eq. (25) (dotted line) at four example points (black dots) for $T = 2$ days and $\mu_{max} = 0.3$.

plots, at $T = 1$ and 7 days, the probability density of $|\Delta\alpha|$ and $|\Delta\delta|$, and the minimum sampled value of $\Delta\hat{\Omega}$ for simulations presented in this paper. The distributions of $|\Delta\alpha|$ and $|\Delta\delta|$ are far from uniform and contain few points that would satisfy the $\Delta\hat{\Omega} \lesssim 5$ cutoff used in Ref. [25]. We expect differences between the $\mathcal{F}$ statistic and the supersky metric to be magnified at larger coordinate offsets, and so it is unsurprising that the errors shown in, e.g., Figure 2(d) are larger than those found in Ref. [25].

Figure 7 plots the mean and standard deviation of μ_0 vs μ_{ss}, at $T = 1$ and 7 days. It is qualitatively similar to Fig. 10 of Ref. [25], which plots μ_0 vs $\mu_{\mathcal{F}}$ at $T = 0.5$ and 2.5 days. In both cases, higher-order terms in $\Delta\vec{\lambda}$, neglected
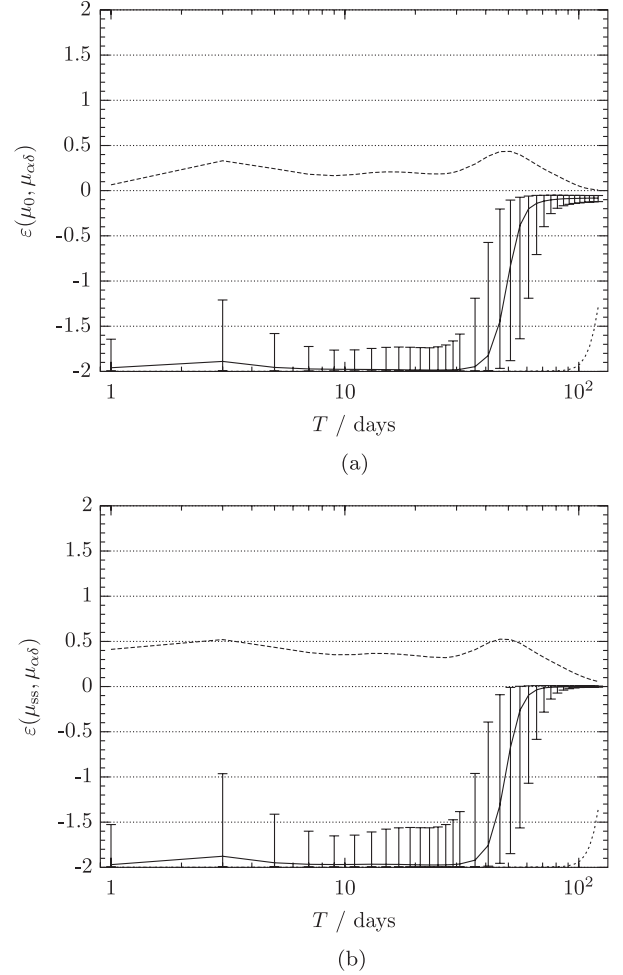

FIG. 9. Relative errors (a) $\varepsilon(\mu_0, \mu_{\alpha\delta})$ and (b) $\varepsilon(\mu_{ss}, \mu_{\alpha\delta})$ between mismatches predicted by the sky metric in α–δ coordinates, $\mu_{\alpha\delta}$, predicted by the supersky metric, μ_{ss}, and calculated from the $\mathcal{F}$ statistic, μ_0, as a function of T for $\mu_0 \lesssim 0.2$ and $\mu_{ss} \lesssim 0.2$, respectively. Plotted are the median (solid line), the 25–75th percentile range (error bars), and the 2.5th (dotted line) and 97.5th (dashed line) percentiles. Only the first spin-down is used.

in the derivation of Eq. (8), result in both $\mu_{\mathcal{F}}$ and μ_{ss} overestimating μ_0 at larger mismatches. The empirical fit $\mu_0 = \mu_{\mathcal{F}} - 0.38\mu_{\mathcal{F}}^2$ to the behavior of the full $\mathcal{F}$-statistic metric, found in Ref. [25], is not as good a fit to the behavior of μ_{ss}, suggesting that the approximations made in deriving the phase metric lead to further overestimation of μ_0.

B. Comparison to sky metric in α–δ coordinates

In contrast to the supersky coordinates $\vec{n}$, the angular coordinates (α, δ) are a poor choice of sky coordinates for the purpose of predicting mismatch. This is because, since the sky metric expressed in α–δ coordinates is itself a function of α and δ, the mismatch in these coordinates is generally calculated by evaluating the metric at a given point, e.g., (α_0, δ_0), and computing

$$\mu_{\alpha\delta} \approx \begin{pmatrix} \alpha - \alpha_0 \\ \delta - \delta_0 \\ \Delta f \\ \vdots \end{pmatrix} \cdot \mathbf{g}|_{\alpha=\alpha_0, \delta=\delta_0} \begin{pmatrix} \alpha - \alpha_0 \\ \delta - \delta_0 \\ \Delta f \\ \vdots \end{pmatrix}. \quad (25)$$

As can be seen in Fig. 4, this is often a very poor approximation, since the sky metric can change noticeably as a function of α and δ. To illustrate, Fig. 8 plots, for four example points, the true mismatch region of the sky metric, given by the intersection of the supersky metric ellipsoid and sky sphere, alongside the local metric ellipse, which is found from the intersection of the supersky metric ellipsoid with a plane tangent to the sky sphere at the chosen point. The agreement between the two changes from very good (top-left plot) to very poor (bottom-left plot) within a 10 deg change in δ.

Figure 9 plots the relative errors $\varepsilon(\mu_0, \mu_{\alpha\delta})$ and $\varepsilon(\mu_{ss}, \mu_{\alpha\delta})$ between mismatches predicted by the sky metric in α–δ coordinates, by the supersky metric, and calculated from the $\mathcal{F}$ statistic. While the sky metric in α–δ coordinates begins to perform better at longer $T \gtrsim 60$ days, it performs extremely poorly at shorter T. This is a consequence of the poor approximation in Eq. (25), which, as shown in Fig. 8, can cause the sky region covered by the local metric ellipsoid at (α_0, δ_0) to be very different from the true mismatch region. This difference in covered sky regions tends to show up as an overestimation of, e.g., μ_0 by $\mu_{\alpha\delta}$ [$\varepsilon(\mu_0, \mu_{\alpha\delta}) < 0$], as illustrated in Fig. 1.

V. REDUCED SUPERSKY METRIC

The fundamental idea behind the supersky metric, presented in Sec. IV, is also its main drawback when it comes to template placement: it embeds the two-dimensional physical sky in three-dimensional space. This means that we cannot simply fill the three-dimensional supersky space with templates, since only a small fraction of them will correspond to physical sky positions (i.e., satisfy $|\vec{n}| = 1$). Instead, we must find a way to reduce the dimensionality of the supersky metric to two dimensions, while preserving flatness, and without making assumptions that introduce significant errors between the supersky and $\mathcal{F}$-statistic mismatches.

The approach taken in this paper is to derive a new set of sky coordinates (n_a, n_b, n_c) and frequency and spindown coordinates $(\nu, \dot{\nu}, \ldots)$, such that the supersky metric in these coordinates, $\mathbf{g}_{ss}^{\dagger}$, is (nearly) diagonal:

$$\mathbf{g}_{ss}^{\dagger} = \begin{pmatrix} g_{n_a,n_a}^{\dagger} & 0 & 0 & 0 & 0 & \cdots \\ 0 & g_{n_b,n_b}^{\dagger} & 0 & 0 & 0 & \cdots \\ 0 & 0 & g_{n_c,n_c}^{\dagger} & 0 & 0 & \cdots \\ 0 & 0 & 0 & g_{\nu,\nu}^{\dagger} & g_{\nu,\dot{\nu}}^{\dagger} & \cdots \\ 0 & 0 & 0 & g_{\nu,\dot{\nu}}^{\dagger} & g_{\dot{\nu},\dot{\nu}}^{\dagger} & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}. \quad (26)$$
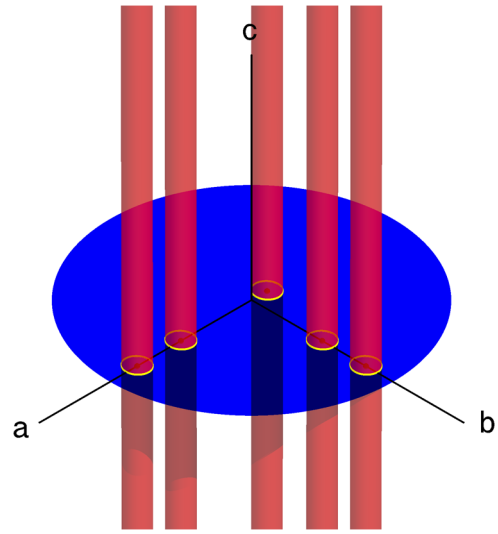

FIG. 10 (color online). Metric ellipsoids of the supersky metric, in aligned coordinates (n_a, n_b, n_c) [see Sec. V C], for $T = 4$ days and $\mu_{\max} = 5$ at five example points. By removing the c dimension, we produce a projected metric on the two-dimensional a–b plane.

Suppose that the new sky coordinates are chosen such that the inequalities $g_{n_a,n_a}^{\dagger} \geq g_{n_b,n_b}^{\dagger} \geq g_{n_c,n_c}^{\dagger}$ always hold. The metric ellipsoids of the supersky metric in these coordinates, shown in Fig. 10, are longest along the c axis; equivalently, the mismatch is most insensitive to changes in the coordinates n_c. Dropping the c dimension, therefore, will introduce the smallest possible error in calculating mismatches, relative to using the full three-dimensional supersky metric. This is geometrically equivalent to projecting the supersky metric onto the two-dimensional a–b plane, as seen in Fig. 10. The *reduced supersky* metric, $\mathbf{g}_{rss}$, in the coordinates $(n_a, n_b, \nu, \dot{\nu}, \ldots)$ reduces the sky dimensionality to two dimensions while remaining constant. Its derivation is presented in the remainder of this section.

A. Diagonalization and condition numbers

A straightforward approach to diagonalizing the supersky metric would be to decompose it as $\mathbf{g}_{ss} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^{T}$, where $\mathbf{\Lambda}$ is a diagonal matrix of eigenvalues and $\mathbf{Q}$ is an orthogonal matrix for which the columns are the corresponding eigenvectors; its transpose, $\mathbf{Q}^{T}$, gives the linear transform from the original to the new supersky coordinates. This approach is complicated, however, by a practical difficulty: the numerical stability of the matrix $\mathbf{g}_{ss}$. A useful measure of numerical stability of a matrix is its *condition number*, which for a real symmetric matrix, $\mathbf{g}_{ss}$, is given by the absolute ratio of the largest to the smallest of its eigenvalues, e.g., Ref. [39]. Generally, when the condition number is of the same order as the numerical precision of the matrix components, computations using the matrix become unreliable [25]. When computed in Systeme International units (i.e., f in Hz, $\dot{f}$ in
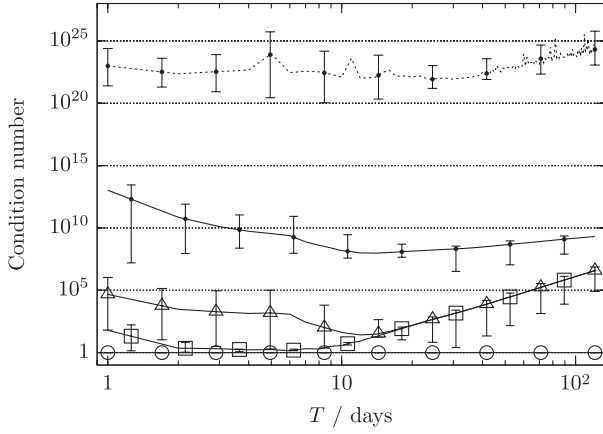
FIG. 11. Condition numbers of the supersky metric, in Systeme International units, $\mathbf{g}_{ss}$ (points, dotted line), and after diagonal rescaling, $\bar{\mathbf{g}}_{ss}$ (points, solid line); the diagonal-rescaled metric $\bar{\mathbf{g}}'_{ss}$, derived in Sec. V B (triangles); and the diagonal-rescaled metrics $\bar{\mathbf{g}}^{\dagger\dagger}_{ss}$ (squares) and $\bar{\mathbf{g}}^{\dagger}_{ss}$ (circles), derived in Sec. V C. Plotted as functions of T are the mean (lines) and the minimum-to-maximum range (error bars) of condition numbers over all simulation reference times Δt_0; see the Appendix. Only the first spindown is used.

Hz/s, etc.), the condition number of $\mathbf{g}_{ss}$, plotted in Fig. 11, is $\gtrsim 10^{20}$, much larger than the ~ 2 in 10^{16} numerical precision of double-precision floating-point computer arithmetic. Figure 12 shows the effect of the metric's ill-conditionedness, by plotting the relative error $|\lambda_{ss} - \lambda^{\dagger}_{ss}|/\lambda^{\dagger}_{ss}$ between the smallest eigenvalue, λ_{ss}, computed from the untransformed metric $\mathbf{g}_{ss}$ and the smallest eigenvalue $\lambda^{\dagger}_{ss}$ computed from the diagonalized metric $\mathbf{g}^{\dagger}_{ss}$ [Eq. (26)], obtained in Sec. V D. When the metric includes only the first spindown, relative errors of $\gtrsim 1$ are possible
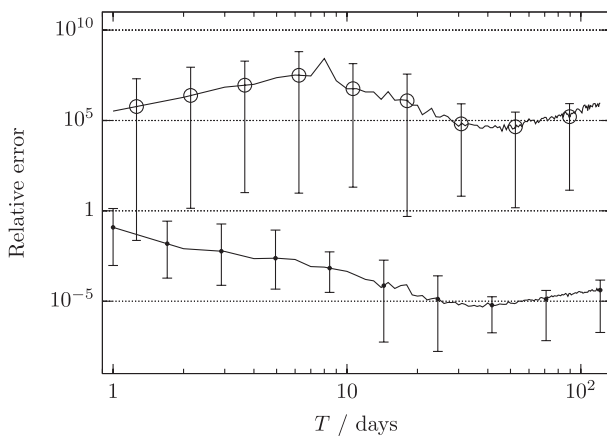

FIG. 12. Relative error in computing the smallest eigenvalue of the supersky metric in Systeme International units, including the first spindown (points) and a second spindown (circles). Plotted as functions of T are the mean (lines) and the minimum-to-maximum range (error bars) of errors over all simulation reference times Δt_0; see the Appendix.

for $T \lesssim 2$ days; when a second spindown is included, the errors increase by several orders of magnitude.

One simple method of reducing the condition number of a matrix, $\mathbf{g}$, is the following diagonal rescaling:

$$\bar{g}(\lambda_i, \lambda_j) = \frac{g(\lambda_i, \lambda_j)}{\sqrt{g(\lambda_i, \lambda_i)g(\lambda_j, \lambda_j)}}, \qquad (27)$$

where the $\bar{g}(\lambda_i, \lambda_j)$ are elements of a rescaled matrix, $\bar{\mathbf{g}}$. This particular diagonal rescaling reduces the condition number to within a factor n of the smallest condition number achievable by any diagonal rescaling, n being the dimensionality of the matrix [39]. Even after applying Eq. (27) to the supersky metric, resulting in $\bar{\mathbf{g}}_{ss}$, its condition number is still of the order $\sim 10^8$–10^{15} (Fig. 11).

We therefore apply a series of transformations to the supersky metric, described in the following sections, which are designed to both diagonalize the metric and further reduce its condition number.

B. Modeling the orbital motion

The first transformation takes advantage of a well-known property of the gravitational-wave pulsar signal model: that for coherent time spans T much shorter than 1 year, the change in signal phase due to Earth's orbital motion closely resembles a change in phase due to the frequency evolution of the pulsar. This is because, for $\Delta t \ll T_o = 1$ year, and assuming a circular orbit about the Sun, the orbital motion of Earth is modeled by terms involving $\sin 2\pi \Delta t/T_o$ and $\cos 2\pi \Delta t/T_o$; for small $\Delta t/T_o$, these can be Taylor expanded as a series of terms in $(\Delta t/T_o)^s$, which are then absorbed into the frequency and spindown terms in the signal phase [Eq. (15)]. This property of the signal model is the basis of the global correlation method [17,30].

Here, we hypothesize that the similarity between the orbital motion and frequency evolution terms in the signal model leads to a linear relation between the corresponding components in the supersky metric and hence to an ill-conditioned matrix $\mathbf{g}_{ss}$. We introduce an intermediate set of coordinates $(\vec{n}, f^{(s)\prime})$ that take advantage of this relation and result in a metric, $\mathbf{g}'_{ss}$, with a greatly reduced condition number. The derivation of these coordinates consists of three steps: splitting the diurnal and orbital motion of the detector into separate sky coordinates (Sec. V B 1), performing a least-squares fit to the orbital motion using the frequency and spindown coordinates (Sec. V B 2), and recombining the diurnal and orbital motion to recover the supersky coordinates (Sec. V B 3).

1. Splitting the diurnal and orbital motions

Recall that the supersky coordinates $\vec{n}$ enter the signal phase $\phi(t, \vec{\lambda})$ through the expression $\vec{r}(t) \cdot \vec{n}$ [Eq. (15)]. We split $\vec{r}(t) = \vec{r}_s(t) + \vec{r}_o(t)$ into its diurnal and orbital components, $\vec{r}_s(t)$ and $\vec{r}_o(t)$, respectively:

$$\vec{r}(t) \cdot \vec{n} = \vec{r}_s(t) \cdot \vec{n} + \vec{r}_o(t) \cdot \vec{n}. \tag{28}$$

We now relax the constraint that the $\vec{n}$ that multiplies $r_s(t)$ is the same $\vec{n}$ that multiplies $r_o(t)$. Instead, we introduce new sky position vectors, $\vec{n}_s$ and $\vec{n}_o$, and write

$$\vec{r}(t) \cdot \vec{n} = \vec{r}_s(t) \cdot \vec{n}_s + \vec{r}_o(t) \cdot \vec{n}_o, \tag{29}$$

where $\vec{n}_s$ and $\vec{n}_o$ are now treated as independent sets of coordinates. If we now write $\vec{r}_s(t) = [r_{sx}(t), r_{sy}(t), r_{sz}]$ and $\vec{n}_s = (n_{sx}, n_{sy}, n_{sz})$ in equatorial coordinates and $\vec{r}_o(t) = [r_{oX}(t), r_{oY}(t), r_{oZ}(t)]$ and $\vec{n}_o = (n_{oX}, n_{oY}, n_{oZ})$ in ecliptic coordinates, we have

$$\vec{r}(t) \cdot \vec{n} = r_{sx}(t)n_{sx} + r_{sy}(t)n_{sy} + r_{sz}n_{sz} + r_{oX}(t)n_{oX}$$
$$+ r_{oY}(t)n_{oY} + r_{oZ}(t)n_{oZ}. \tag{30}$$

Since r_{sz} is a constant (i.e., there is no motion of the detector with respect to the Earth in the equatorial z direction), $g(n_{sz}, \lambda_i) = 0$ for any coordinate λ_i, and so we ignore the term $r_{sz}n_{sz}$. (Since, however, Earth's orbit does include motion in the ecliptic Z direction, due to, e.g., its interaction with the Moon, we do not neglect the term $r_{oZ}(t)n_{sZ}$.) We are left with five independent sky coordinates,

$$\vec{n}_{ess} = (n_{sx}, n_{sy}, n_{oX}, n_{oY}, n_{oZ}), \tag{31}$$

and we write the metric in these *expanded supersky* coordinates as $\mathbf{g}_{ess}$. Just as the supersky coordinates $\vec{n}$ embedded the two-dimensional physical sky in three dimensions, the expanded supersky coordinates embed the three-dimensional supersky in five dimensions. Reimposing the constraint $\vec{n}_s = \vec{n}_o$ recovers the supersky metric $\mathbf{g}_{ss}$.

2. Least-squares linear fit to the orbital motion

We apply the diagonal rescaling of Eq. (27) to the expanded supersky metric:

$$\bar{g}_{ess}(\lambda_i, \lambda_j) = \frac{g_{ess}(\lambda_i, \lambda_j)}{\sqrt{g_{ess}(\lambda_i, \lambda_i)g_{ess}(\lambda_j, \lambda_j)}}, \tag{32}$$

where $\lambda_i, \lambda_j \in \vec{\lambda} = (\vec{n}_{ess}, f^{(s)})$. We now hypothesize an (approximately) linear relationship between the rows of $\bar{\mathbf{g}}_{ess}$ corresponding to orbital motion in the ecliptic X and Y directions and the rows of $\bar{\mathbf{g}}_{ess}$ corresponding to frequency evolution; i.e., we write

$$\bar{g}_{ess}(\lambda_i, n_{o\sigma}) = \sum_{s=0}^{s_{max}} \bar{g}_{ess}(\lambda_i, f^{(s)})C(f^{(s)}, n_{o\sigma})$$
$$+ \delta\bar{g}_{ess}(\lambda_i, n_{o\sigma}), \tag{33}$$

where $\sigma \in \{X, Y\}$, $C(f^{(s)}, n_{o\sigma})$ are the components of a $(1 + s_{max}) \times 2$ matrix of fitting coefficients $\mathbf{C}$, and $\delta\bar{g}_{ess}(\lambda_i, n_{o\sigma})$ are components of a $(6 + s_{max}) \times 2$ matrix of residuals $\boldsymbol{\delta}\bar{\mathbf{g}}_{ess}$. Equations (33) are overdetermined since

they represent $(6 + s_{max}) \times 2$ equations, one for each element $\bar{g}_{ess}(\lambda_i, n_{o\sigma})$, for only $(1 + s_{max}) \times 2$ unknowns $C(f^{(s)}, n_{o\sigma})$. Using linear least squares for each $\sigma \in \{X, Y\}$, i.e., minimizing the objective function $\sum_i[\delta\bar{g}_{ess}(\lambda_i, n_{o\sigma})]^2$ over the vector $C(f^{(s)}, n_{o\sigma})$ at fixed σ, yields the fitting coefficients

$$C(f^{(s)}, n_{o\sigma}) = \sum_{s'=0}^{s_{max}} A^{-1}(f^{(s)}, f^{(s')})B(f^{(s')}, n_{o\sigma}), \tag{34}$$

where the $(1 + s_{max}) \times 2$ matrix $\mathbf{B}$ has components

$$B(f^{(s)}, n_{o\sigma}) = \sum_i \bar{g}_{ess}(f^{(s)}, \lambda_i)\bar{g}_{ess}(\lambda_i, n_{o\sigma}), \tag{35}$$

and $A^{-1}(f^{(s)}, f^{(s')})$ are the elements of the inverse of the symmetric $(1 + s_{max}) \times (1 + s_{max})$ matrix $\mathbf{A}$, with components

$$A(f^{(s)}, f^{(s')}) = \sum_i \bar{g}_{ess}(f^{(s)}, \lambda_i)\bar{g}_{ess}(\lambda_i, f^{(s')}). \tag{36}$$

Next, we perform a coordinate transformation that results in a metric, $\bar{\mathbf{g}}''_{ess}$. In this metric, the linear fits to $\bar{g}_{ess}(\lambda_i, n_{o\sigma})$ are subtracted, leaving only the residuals:

$$\bar{g}''_{ess}(\lambda_i, n_{o\sigma}) = \delta\bar{g}_{ess}(\lambda_i, n_{o\sigma}). \tag{37}$$

This is achieved by the coordinate transform

$$f^{(s)''} = f^{(s)} + \sum_{\sigma \in \{X,Y\}} n_{o\sigma}C(f^{(s)}, n_{o\sigma})\sqrt{\frac{g_{ess}(n_{o\sigma}, n_{o\sigma})}{g_{ess}(f^{(s)}, f^{(s)})}}, \tag{38}$$

with the sky coordinates $\vec{n}_{ess}$ remaining unchanged. The new frequency and spindown coordinates $f^{(s)''}$ are now linear functions of sky position, similar to the global correlation coordinates [c.f. Eqs. (21)].

3. Recombining the diurnal and orbital motions

Finally, we reimpose the constraint $\vec{n}_s = \vec{n}_o$ and reverse the diagonal rescaling given by Eq. (32). This gives the final product of this section, the metric $\mathbf{g}'_{ss}$. The coordinates of this metric are $(\vec{n}, f^{(s)'})$, where $\vec{n}$ is the three-dimensional sky position vector, and

$$f^{(s)'} = f^{(s)} + \vec{\Gamma}^s \cdot \vec{n}. \tag{39}$$

The components of the vectors $\vec{\Gamma}^s$ are given in ecliptic coordinates by

$$\Gamma_X^s = C(f^{(s)}, n_{oX})\sqrt{\frac{g_{ess}(n_{oX}, n_{oX})}{g_{ess}(f^{(s)}, f^{(s)})}}, \tag{40a}$$

$$\Gamma_Y^s = C(f^{(s)}, n_{oY})\sqrt{\frac{g_{ess}(n_{oY}, n_{oY})}{g_{ess}(f^{(s)}, f^{(s)})}}, \tag{40b}$$

$$\Gamma_Z^s = 0 \tag{40c}$$

and in equatorial coordinates by

$$\Gamma^s_x = \Gamma^s_X, \tag{41a}$$
$$\Gamma^s_y = \Gamma^s_Y \cos \epsilon, \tag{41b}$$
$$\Gamma^s_z = \Gamma^s_Y \sin \epsilon, \tag{41c}$$

where ϵ is Earth's inclination angle with respect to the ecliptic Z direction.

In summary, the coordinate transformation presented in this section consists of the following steps:

 (i) Compute the phase metric $\mathbf{g}_{ess}$, using Eq. (10), in the expanded supersky coordinates of Eq. (31). The signal phase is found by substituting Eq. (30) in Eq. (15).
 (ii) Rescale the metric using Eq. (32), which gives the metric $\bar{\mathbf{g}}_{ess}$.
 (iii) Compute the matrices $\mathbf{A}$, using Eq. (36), and $\mathbf{B}$, using Eq. (35).
 (iv) Compute the matrix $\mathbf{C}$ from Eq. (34).
 (v) Compute the vectors $\vec{\Gamma}^s$ from Eqs. (40) and (41).
 (vi) Apply the inverse of the coordinate transformation given by Eq. (39) to the supersky metric $\mathbf{g}_{ss}$, which gives the metric $\mathbf{g}'_{ss}$.

The condition number of this metric after diagonal rescaling, $\bar{\mathbf{g}}'_{ss}$, is plotted in Fig. 11; it is reduced to ~ 1–10^6 and is ~ 3–8 orders of magnitude less than the condition number of the untransformed metric, $\bar{\mathbf{g}}_{ss}$.

Figure 13 plots the metric ellipsoids of $\mathbf{g}'_{ss}$, with $\mu_{max} = 5$. The areas of intersections of the ellipsoids of $\mathbf{g}'_{ss}$ with the sky sphere $|\vec{n}| = 1$ are larger (and more circular) than for the ellipsoids of $\mathbf{g}_{ss}$, plotted in Fig. 4 with $\mu_{max} = 30$, which implies a coarser sky metric induced by $\mathbf{g}'_{ss}$ than by the untransformed metric, $\mathbf{g}_{ss}$. This can be intuitively

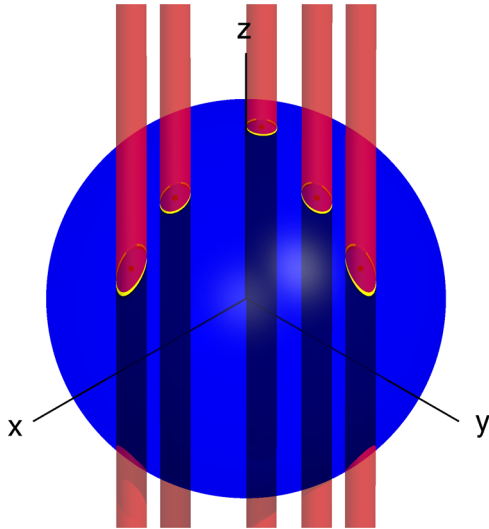

FIG. 13 (color online). Metric ellipsoids of the transformed supersky metric $\mathbf{g}'_{ss}$, in equatorial coordinates $\vec{n}$, for $T = 4$ days and $\mu_{max} = 5$, at five example points and their intersections with the sky sphere $|\vec{n}| = 1$.

understood as follows: given the constraint $\mu \leq \mu_{max}$, the resolution in the supersky coordinates, e.g., $|\Delta n_x|$, are inversely proportional to the corresponding supersky metric elements; e.g., $g_{ss}(n_x, n_x)(\Delta n_x)^2 = \mu \leq \mu_{max}$ implies $|\Delta n_x| \propto g_{ss}(n_x, n_x)^{-1/2}$. The magnitude of $g_{ss}(n_x, n_x)$ is proportional to the magnitude of the detector position vector, $|\vec{r}(t)|^2$, which is dominated by the magnitude of its orbital component, $|\vec{r}_o(t)|^2 \gg |\vec{r}_s(t)|^2$. Hence, $|\Delta n_x| \propto 1/|\vec{r}_o(t)|$, and the coordinate resolution is largely determined by the orbital motion. If, however, the fitting of the orbital motion described in Sec. V B 2 is effective, the magnitude of the transformed supersky metric elements, e.g., $g'_{ss}(n_x, n_x)$, will instead be proportional to magnitude of the unfitted diurnal motion, $|\vec{r}_s(t)|^2$, and hence the coordinate resolution $|\Delta n_x| \propto 1/|\vec{r}_s(t)|$ will be coarser by a factor $\sim |\vec{r}_o(t)|/|\vec{r}_s(t)|$.

C. Diagonalizing the metric

Now that we have a transformed supersky metric, $\mathbf{g}'_{ss}$, with a greatly reduced condition number, we can reliably apply two further transformations to diagonalize it and arrive at the metric $\mathbf{g}^\dagger_{ss}$ in the form shown in Eq. (26).

The first transformation removes the elements of $\mathbf{g}'_{ss}$ that couple the sky coordinates $\vec{n}$ to the frequency and spin-down coordinates $f^{(s)\prime}$. We write the matrix $\mathbf{g}'_{ss}$ as a block matrix of the form

$$\mathbf{g}'_{ss} = \begin{pmatrix} \mathbf{g}'_{nn} & \mathbf{g}'_{nf} \\ \mathbf{g}'^{T}_{nf} & \mathbf{g}'_{ff} \end{pmatrix}, \tag{42}$$

where $\mathbf{g}'_{nn}$, $\mathbf{g}'_{nf}$, $\mathbf{g}'_{ff}$ are matrices with elements $g'_{ss}(n_i, n_j)$, $g'_{ss}(n_i, f^{(s)\prime})$, and $g'_{ss}(f^{(s)\prime}, f^{(s')\prime})$, respectively. We then require a coordinate transformation that sets $\mathbf{g}'_{nf} = \mathbf{0}$. This condition is satisfied by the coordinates $(\vec{n}, \nu^{(s)})$, where

$$\nu^{(s)} = f^{(s)} + \vec{\Delta}^s \cdot \vec{n}, \tag{43}$$

$$\vec{\Delta}^s = \vec{\Gamma}^s + \mathbf{g}'^{-1}_{ff} \mathbf{g}'^{T}_{nf}. \tag{44}$$

The vectors Δ^s introduce additional shifts in frequency and spindown that are linear in the sky position $\vec{n}$. The metric in these coordinates is given by

$$\mathbf{g}^{\dagger\dagger}_{ss} = \begin{pmatrix} \mathbf{g}^{\dagger\dagger}_{nn} & \mathbf{0} \\ \mathbf{0} & \mathbf{g}'_{ff} \end{pmatrix}, \tag{45}$$

where

$$\mathbf{g}^{\dagger\dagger}_{nn} = \mathbf{g}'_{nn} - \mathbf{g}'_{nf} \mathbf{g}'^{-1}_{ff} \mathbf{g}'^{T}_{nf}. \tag{46}$$

The right-hand side of Eq. (46) is the Schur complement of the block form of $\mathbf{g}'_{ss}$ given in Eq. (42). Figure 11 plots the condition number of the diagonal-rescaled $\bar{\mathbf{g}}^{\dagger\dagger}_{ss}$, which for $T \lesssim 10$ days is reduced, relative to $\bar{\mathbf{g}}'_{ss}$, by a few orders of magnitude.

The second transformation diagonalizes the sky coordinate block $\mathbf{g}_{nn}^{\dagger\dagger}$ of the matrix $\mathbf{g}_{ss}^{\dagger\dagger}$. We introduce *aligned supersky* coordinates (n_a, n_b, n_c), such that the metric takes the form [c.f. Eq. (26)]

$$
\mathbf{g}_{ss}^{\dagger} = \begin{pmatrix} g_{n_a,n_a}^{\dagger} & 0 & 0 & \\ 0 & g_{n_b,n_b}^{\dagger} & 0 & \mathbf{0} \\ 0 & 0 & g_{n_c,n_c}^{\dagger} & \\ & \mathbf{0} & & \mathbf{g}_{ff}' \end{pmatrix}. \quad (47)
$$

This is achieved by eigendecomposing $\mathbf{g}_{nn}^{\dagger\dagger}$ as

$$
\mathbf{g}_{nn}^{\dagger\dagger} = \mathbf{Q}^{\dagger}\mathbf{\Lambda}^{\dagger}\mathbf{Q}^{\dagger\mathrm{T}}, \quad (48)
$$

where the elements of the diagonal matrix $\mathbf{\Lambda}^{\dagger}$ are the eigenvalues $g_{n_a,n_a}^{\dagger} \geq g_{n_b,n_b}^{\dagger} \geq g_{n_c,n_c}^{\dagger}$, and the columns of $\mathbf{Q}$ are the corresponding eigenvectors $\vec{Q}_{n_a}$, $\vec{Q}_{n_b}$, and $\vec{Q}_{n_c}$. The aligned supersky coordinates are then defined by

$$
n_a = \vec{Q}_{n_a} \cdot \vec{n}, \qquad n_b = \vec{Q}_{n_b} \cdot \vec{n}, \qquad n_c = \vec{Q}_{n_c} \cdot \vec{n}. \quad (49)
$$

In summary, we have found the diagonal metric $\mathbf{g}_{ss}^{\dagger}$ of Eq. (26), and associated coordinates $(n_a, n_b, n_c, \nu^{(s)})$, by performing the following steps:

(i) Compute the metric $\mathbf{g}_{ss}'$ and vectors $\vec{\Gamma}^s$, by following the procedure described in Sec. V B.

(ii) Partition $\mathbf{g}_{ss}'$ into a block matrix, as given by Eq. (42).

(iii) Compute the matrix $\mathbf{g}_{nn}^{\dagger\dagger}$, given by Eq. (46).

(iv) Eigendecompose $\mathbf{g}_{nn}^{\dagger\dagger}$ into eigenvectors $\vec{Q}_{n_a}$, $\vec{Q}_{n_b}$, $\vec{Q}_{n_c}$ and eigenvalues $g_{n_a,n_a}^{\dagger}$, $g_{n_b,n_b}^{\dagger}$, $g_{n_c,n_c}^{\dagger}$, following Eq. (48).

(v) Compute the vectors $\vec{\Delta}^s$, given by Eq. (44).

The metric $\mathbf{g}_{ss}^{\dagger}$ is then given by Eq. (47), the sky coordinates (n_a, n_b, n_c) by Eqs. (49), and the frequency and spindown coordinates $\nu^{(s)}$ by Eq. (43). Figure 11 plots the condition number of $\mathbf{g}_{ss}^{\dagger}$, which after diagonal rescaling is reduced to ~ 1.

It should be noted that the transformation from $\mathbf{g}_{ss}$ to $\mathbf{g}_{ss}^{\dagger}$ is invertible, and no approximations or assumptions are made in its derivation. Thus, the mismatch predicted by $\mathbf{g}_{ss}^{\dagger}$ will be identical to that predicted by $\mathbf{g}_{ss}$.

D. Reducing the sky dimensionality

As described at the beginning of Sec. V, the reduced supersky metric, $\mathbf{g}_{rss}$, takes the diagonalized supersky metric $\mathbf{g}_{ss}^{\dagger}$, derived in Secs. V B and V C, and removes the dimension corresponding to n_c. This is equivalent to projecting the metric onto the a–b plane, as illustrated in Fig. 10. As a consequence, the mismatch predicted by the reduced supersky metric, μ_{rss}, will always be smaller than (or equal to) that predicted by the untransformed metric, μ_{ss}, the difference being $\mu_{ss} - \mu_{rss} = g_{n_c,n_c}^{\dagger}(\Delta n_c)^2 \geq 0$, where Δn_c is a coordinate offset in n_c. Since $g_{n_c,n_c}^{\dagger}$ is, by

construction, the smallest of the elements of the sky-sky block of $\mathbf{g}_{ss}^{\dagger}$, this represents the smallest error that can be achieved by projecting the sky metric from three to two dimensions.

Figures 2(f) and 14 plot the relative errors, $\varepsilon(\mu_0, \mu_{rss})$ and $\varepsilon(\mu_{ss}, \mu_{rss})$, respectively, between mismatches predicted by the reduced and untransformed supersky metrics and calculated from the $\mathcal{F}$ statistic. The performance of the reduced supersky metric, relative to the $\mathcal{F}$ statistic [Fig. 2(f)], is similar to that of the untransformed supersky metric [Fig. 2(d)]. The one noticeable difference is that there are more trials where the reduced supersky metric significantly underestimates the $\mathcal{F}$-statistic mismatch [$\varepsilon(\mu_0, \mu_{rss}) > 0$]; for 2.5% of trials (above dashed line), $\varepsilon(\mu_0, \mu_{rss}) \gtrsim 0.5$ for $T \lesssim 40$ days. The origin of this difference is evident in Fig. 14; while the median error $\varepsilon(\mu_{ss}, \mu_{rss}) \sim 0$, and the 25–75th percentile range is $\lesssim 0.1$, 2.5% of trials have $\varepsilon(\mu_{ss}, \mu_{rss}) > 0.5$ for $10 \lesssim T \lesssim 45$ days.

Figure 15 plots the median error magnitude $|\varepsilon(\mu_{ss}, \mu_{rss})|$ as a function of sky position, over the full ranges of simulation parameters, and at fixed values of T, Δt_0, and f_{max}. The largest $|\varepsilon(\mu_{ss}, \mu_{rss})|$ and the source of the underestimation of μ_{ss} by μ_{rss} observed in Fig. 14 occur along the ecliptic equator, when averaged over T and also at $T = 41$ days. Note, however, that at $T = 7$ days the error is largest along the *equatorial* equator $\delta = 0$; at $T = 23$ days, the error region is transitioning between equatorial and ecliptic equators. The largest error is also a function of the reference time: compare $\Delta t_0 = 90$ days with $\Delta t_0 = 0$ and 180 days. This suggests that the source of the underestimation can occur either along the equatorial or ecliptic equators, depending on T, and is also a function of Δt_0.
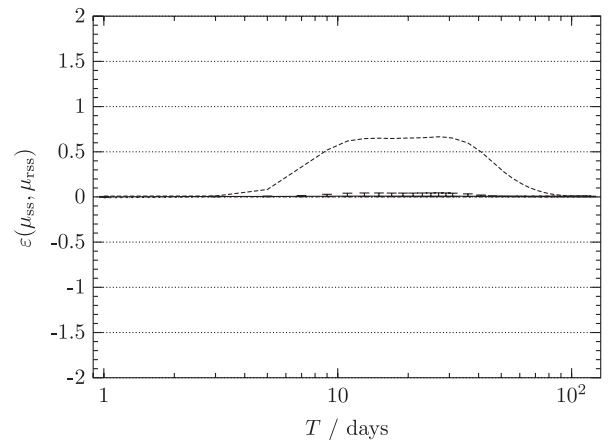

FIG. 14. Relative errors $\varepsilon(\mu_{ss}, \mu_{rss})$ between mismatches predicted by the supersky metric, μ_{ss}, and reduced supersky metric, μ_{rss}, as a function of T, for $\mu_{ss} \leq 0.2$. Plotted are the median (solid line), the 25–75th percentile range (error bars), and the 2.5th (dotted line) and 97.5th (dashed line) percentiles. Only the first spindown is used.
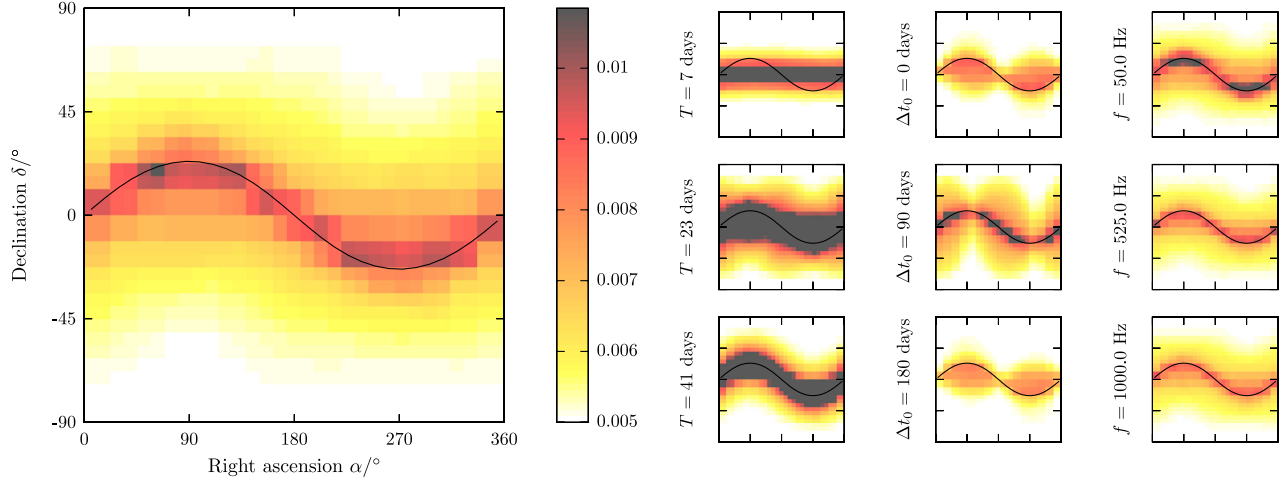
FIG. 15 (color online).   Median magnitude of the relative error $|\varepsilon(\mu_{ss}, \mu_{rss})|$, as a function of α and δ, for $\mu_{ss} \lesssim 0.2$. The ecliptic equator is overplotted in black. Left: median $|\varepsilon(\mu_{ss}, \mu_{rss})|$ over time spans $1 \leq T \leq 121$ days and over all simulation reference times Δt_0 and maximum frequencies f_{max}; see the Appendix. Right: median $|\varepsilon(\mu_{ss}, \mu_{rss})|$ at fixed values of T, Δt_0, and f_{max}; axis ranges and color values are the same as for the left-hand-side plot. Only the first spindown is used.

Figure 16 plots two quantities as functions of T and Δt_0. The first, plotted in Fig. 16(a), is the ratio

$$R = \frac{g^{\dagger}_{n_c, n_c}}{g^{\dagger}_{n_b, n_b}} \qquad (50)$$

of eigenvalues of the diagonalized supersky metric $\mathbf{g}^{\dagger}_{ss}$. This quantity is a proxy for the error introduced by dropping the term $g^{\dagger}_{n_c, n_c}(\Delta n_c)^2$ from the reduced supersky mismatch, relative to the size of the next-largest term in the sky-sky mismatch, $g^{\dagger}_{n_b, n_b}(\Delta n_b)^2$. We see that R increases over the period $10 \lesssim T \lesssim 45$, broadly consistent with the rise and fall of the 97.5th percentile line plotted in Fig. 14. It is also strongly a function of reference time, being largest at $\Delta t_0 \sim$ 75 and 275 days, similar to the median $|\varepsilon(\mu_{ss}, \mu_{rss})|$ plotted in Fig. 15. Note that $\Delta t_0 \sim 75$ and 275 days roughly coincide with the vernal and autumnal equinoxes [40], coordinated

universal time (UTC) $\sim$2007-03-21 ($\Delta t_0 \sim 79$ days) and UTC $\sim$2007-09-23 ($\Delta t_0 \sim 265$ days), respectively.

The second quantity, plotted in Fig. 16(b), is the angle

$$\beta = \frac{\cos^{-1}|\vec{z} \cdot \vec{Q}_{n_c}|}{\epsilon} \qquad (51)$$

between the equatorial z axis $\vec{z}$ and the aligned c axis $\vec{Q}_{n_c}$, as a fraction of Earth's inclination angle ϵ. When $\beta \sim 0$, at $T \lesssim 10$ days, $\vec{z} \cdot \vec{Q}_{n_c} \sim 1$, and the a–b plane is aligned with the equatorial x–y plane. Likewise, when $\beta \sim 1$, at $T \gtrsim 45$ days, $\vec{z} \cdot \vec{Q}_{n_c} \sim \cos \epsilon$, and the a–b plane is aligned with the ecliptic X–Y plane. In the period $10 \lesssim T \lesssim$ 45 days, the aligned supersky coordinates (n_a, n_b, n_c) transition from "equatoriallike" to "eclipticlike" coordinates, at a rate dependent on Δt_0; note that the Δt_0 where this transition occurs most rapidly correlate with the Δt_0 where R is largest.

From the results presented in this section, we may deduce the behavior of the reduced supersky metric as a function of T. When $T \lesssim 10$ days ($\lesssim 3\%$ of 1 year), the fitting performed in Sec. V B removes the orbital motion of Earth from the signal phase, leaving only the diurnal motion of the detector. This motion is in a plane parallel to the equatorial equator, and hence the a–b plane is aligned with the equatorial x–y plane, and the aligned supersky coordinates resemble equatorial coordinates [Fig. 16(b)]. The diagonalized supersky metric ellipsoids are highly elongated perpendicular to the plane of motion, i.e., in the equatorial z direction [Fig. 13], and hence the ratio of eigenvalues R is small [Fig. 16(a)].

As T increases to between 3% and 12% of 1 year ($10 \lesssim T \lesssim 45$ days), the change in signal phase due to Earth's orbital motion no longer closely resembles a change in phase due to frequency evolution, essentially because it
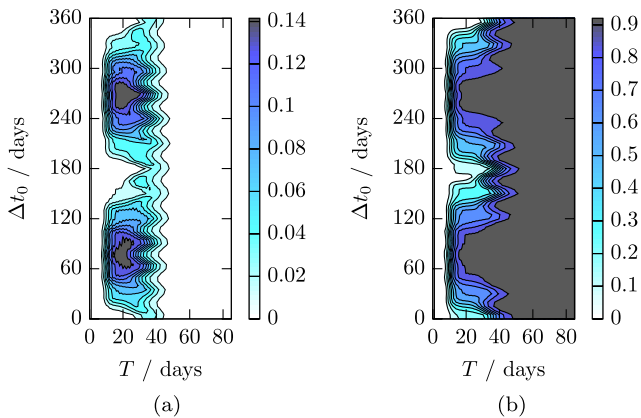

FIG. 16 (color online).   (a) Ratio of diagonalized supersky eigenvalues R and (b) orientation of the a–b plane β, as functions of T and Δt_0. Only the first spindown is used.

can no longer be modeled as circular motion, which can then be Taylor expanded (see Sec. V B). Consequentially, the modeling of the orbital motion becomes less effective, and the residuals of the fit, $\delta\bar{g}_{ess}(\lambda_i, n_{o\sigma})$ [Eq. (33)], become larger. The motion of the detector is therefore a combination of diurnal motion in the equatorial x–y plane and the residual, unfitted orbital motion in the ecliptic X–Y plane. The a–b plane is oriented between the equatorial x–y and ecliptic X–Y planes [Fig. 16(b)]; its exact position is determined by the relative contributions of the diurnal and residual orbital motions. Since the detector motion is no longer two-dimensional, the diagonalized supersky metric ellipsoids are no longer highly elongated perpendicular to the a–b plane, and hence R increases [Fig. 16(a)]. Because Earth's orbit is elliptical, the effectiveness of the orbital motion modeling depends on the reference time: at the equinoxes, where the Earth-Sun distance is changing the most rapidly, Earth's orbit is most poorly approximated by a circular motion, and hence the orbital motion modeling is least effective [Fig. 16(a)].

Eventually, for long enough T, the orbital motion can no longer be fitted, and the residuals become equal to the full orbital motion, i.e., $\delta\bar{g}_{ess}(\lambda_i, n_{o\sigma}) \sim \bar{g}_{ess}(\lambda_i, n_{o\sigma})$ [Eq. (33)]. The detector motion is then dominated by the full orbital motion, which is much larger than the smaller diurnal motion (as noted in Sec. V B 2), and the a–b plane is aligned with the plane of Earth's orbit, i.e., the ecliptic X–Y plane. The diagonalized supersky metric ellipsoids return to being highly elongated perpendicular to the plane of motion, i.e., the ecliptic Z direction [Fig. 13], and the ratio R is again small [Fig. 16(a)].

E. Second spindown

The results presented so far in this paper, with the exception of Fig. 12, have considered parameter-space metrics that include only the first spindown, $\dot{f}$, in the gravitational-wave pulsar signal model. In this section, we briefly consider the effect of adding the second spindown, $\ddot{f}$, to the reduced supersky metric derived in Sec. V. The second spindown is important for some potential gravitational-wave pulsars, such as young neutron stars in supernova remnants [26] and at the Galactic center [41].

Figure 17 plots the relative errors $\varepsilon(\mu_0, \mu_{rss})$ and $\varepsilon(\mu_{ss}, \mu_{rss})$ between mismatches predicted by the reduced and untransformed supersky metrics, respectively, and calculated from the $\mathcal{F}$ statistic, where both the first and second spindowns are included in the metrics. The median errors are similar to those plotted in Figs. 2(f) and 14, which include only the first spindown. In Fig. 17(a), however, we observe a larger number of trials outside of the 2.5–97.5th percentile range (dotted to dashed lines) at longer T, relative to Fig. 2(f). In Fig. 17(b), the rise and fall of the 97.5th percentile line spans $30 \lesssim T \lesssim 70$ days, whereas in Fig. 14 it spans $10 \lesssim T \lesssim 45$ days.

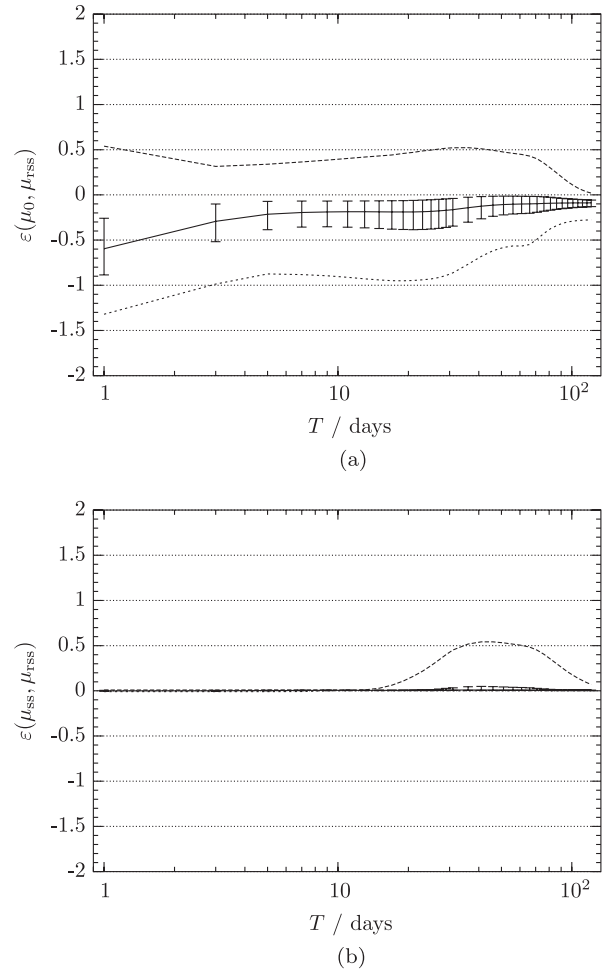FIG. 17. Relative errors (a) $\varepsilon(\mu_0, \mu_{rss})$ and (b) $\varepsilon(\mu_{ss}, \mu_{rss})$ between mismatches calculated from the $\mathcal{F}$ statistic, μ_0, and predicted by the supersky metric, μ_{ss}, and reduced supersky metric, μ_{rss}, as a function of T, for $\mu_0 \leq 0.2$ and $\mu_{ss} \leq 0.2$, respectively. Plotted are the median (solid line), the 25–75th percentile range (error bars), and the 2.5th (dotted line) and 97.5th (dashed line) percentiles. Both the first and second spindowns are used.
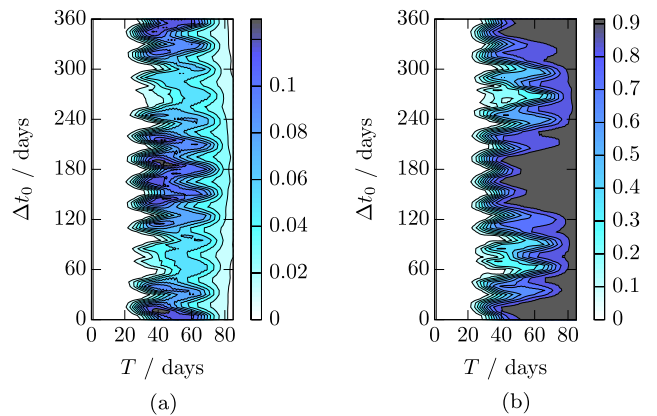

FIG. 18 (color online). (a) Ratio of diagonalized supersky eigenvalues R and (b) orientation of a–b plane β as functions of T and Δt_0. Both the first and second spindowns are used.

Figure 18 plots R [Eq. (50)] and β [Eq. (51)] as functions of T and Δt_0. Whereas in Fig. 16(a) the ratio R is largest at the equinoxes, in Fig. 18(a) it is largest at the solstices. Comparing Fig. 18(a) to Fig. 16(b), we see that, with the second spindown included, the orbital motion fitting is effective for longer T and that the transition of the aligned supersky coordinates from equatoriallike to eclipticlike occurs over the later period $20 \lesssim T \lesssim 75$ days, consistent with the rise and fall of the 97.5th percentile line in Fig. 17(b). In short, including the second spindown changes the behavior of the reduced supersky metric as a function of T.

VI. DISCUSSION

In this paper we have presented a new explicitly flat metric approximation, the reduced supersky metric $\mathbf{g}_{\mathrm{rss}}$, with associated coordinates $(n_a, n_b, \nu, \dot{\nu}, \ldots)$, for performing all-sky coherent searches for gravitational-wave pulsars. Unlike previous work, the reduced supersky metric places no limitation on the time span T that can be coherently analyzed. In addition, compared to previous metrics, the reduced supersky metric is well conditioned.

Figure 19 compares the median absolute relative errors in the predicted mismatch of the supersky, reduced supersky, linear phase model I, and global correlation metrics. The supersky metrics perform better at predicting the $\mathcal{F}$-statistic mismatch than both the linear phase model I and global correlation metrics, for $T \lesssim 10$ days, and are similar to linear phase model I for larger T. Additionally, as shown in Fig. 20, the reduced supersky metric has the practical advantage of a much smaller condition number (after diagonal rescaling) than both the supersky and linear phase model I metrics.

Future work will focus on extending this method to the semicoherent metric, which is required for placing templates in a hierarchical search, where coherently analyzed data segments are incoherently combined. In principle, by
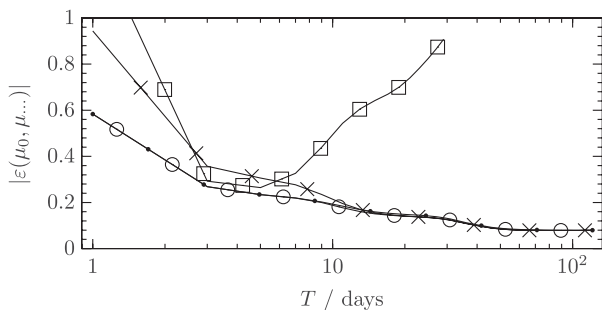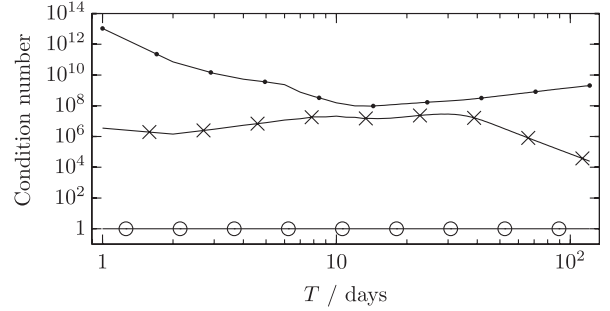

FIG. 20. Mean condition numbers, as functions of T, of the diagonally rescaled supersky (points), reduced supersky (circles), and linear phase model I (crosses) metrics, over all simulation reference times Δt_0; see the Appendix. Only the first spindown is used.

allowing T to be increased, the reduced supersky metric can help improve the sensitivity of all-sky hierarchical searches.

Another interesting application of the reduced supersky metric would be in follow-up pipelines for interesting gravitational-wave pulsar candidates [31]. Starting with a list of candidates from an initial all-sky search, the follow up refines the candidates using a series of semicoherent searches with increasing coherence time T (thereby increasing sensitivity). Finally the follow up concludes by performing a fully coherent search of the remaining candidates. It would be most convenient for such a follow-up pipeline to be able to make use of a parameter-space metric that remains valid for all T, without having to switch and/or interpolate between different metric approximants. Therefore, the reduced supersky metric could assist in performing more refining searches, with a smoother increase in T, thereby allowing more candidates to be followed up and increasing the chances of detection.

ACKNOWLEDGMENTS

APPENDIX: NUMERICAL SIMULATIONS

This appendix details the numerical simulations presented in this paper.

The relative error comparisons presented in Figs. 2, 3, 5, 9, 15, 17, and 19 were produced by Monte Carlo simulations, as follows. For each trial, pairs of gravitational-wave signal parameters, given in the reduced supersky coordinates of Sec. V by $\vec{\lambda}'_1 = (n_{a1}, n_{b1}, \nu_1^{(s)})$ and $\vec{\lambda}'_2 = (n_{a2}, n_{b2}, \nu_2^{(s)})$, were randomly chosen, such that the mismatch μ_{rss} between them, computed using the reduced

FIG. 19. Median absolute relative errors, as functions of T, between mismatches calculated from the $\mathcal{F}$ statistic and predicted by the supersky metric, $|\varepsilon(\mu_0, \mu_{\mathrm{ss}})|$ (points); the reduced supersky metric, $|\varepsilon(\mu_0, \mu_{\mathrm{rss}})|$ (circles); the linear phase model I metric, $|\varepsilon(\mu_0, \mu_{\mathrm{lpI}})|$ (crosses); and the global correlation metric, $|\varepsilon(\mu_0, \mu_{\mathrm{gc}})|$ (squares), for $\mu_0 \leq 0.2$. Only the first spindown is used.

supersky metric $\mathbf{g}_{\mathrm{rss}}$, was uniformly distributed within $0 \le \mu_{\mathrm{rss}} \le 0.6$. To achieve this, a random point $\vec{p}$ was chosen uniformly within the $(3 + s_{\max})$-dimensional unit sphere and used to compute $\vec{\lambda}'_2 - \vec{\lambda}'_1 = \sqrt{0.6}\mathbf{G}_{\mathrm{rss}}^{-1}\vec{p}$, where $\mathbf{G}_{\mathrm{rss}}$ is the Cholesky decomposition of $\mathbf{g}_{\mathrm{rss}}$ (i.e., $\mathbf{G}_{\mathrm{rss}}^{\mathrm{T}}\mathbf{G}_{\mathrm{rss}} = \mathbf{g}_{\mathrm{rss}}$). The reason for using the reduced supersky metric here was that a well-conditioned metric is required to compute the Cholesky decomposition, which in turn is needed for sampling uniformly with respect to the metric. Once the offset $\vec{\lambda}'_2 - \vec{\lambda}'_1$ was determined in this way, we could chose $\vec{\lambda}'_1$ uniformly in (n_{a1}, n_{b1}) over the unit disc and $\nu_1^{(s)}$ uniformly in frequency and spin-down(s). We then computed $n_{c1} = \sqrt{1 - n_{a1}^2 - n_{b1}^2}$ and transformed the coordinates $(n_{a1}, n_{b1}, n_{c1}, \nu_1^{(s)})$ back to the untransformed supersky coordinates $\vec{\lambda}_1 = (\vec{n}_1, f_1^{(s)})$ and similarly for n_{c2} and $\vec{\lambda}_2$. These coordinates were then used to compute the untransformed supersky mismatch μ_{ss} and, with additional coordinate transformations, the linear phase model mismatches μ_{lpI} and μ_{lpII}, the global correlation mismatch μ_{gc}, and the mismatch in α–δ coordinates $\mu_{\alpha\delta}$. The global correlation metric was computed from the equations in Ref. [17]; all other metrics were computed numerically using the software package

LALPULSAR [42], using the standard Jet Propulsion Laboratory ephemerides for Earth's orbital motion. Finally, the $\mathcal{F}$-statistic mismatch μ_0 was calculated as described in Sec. II B, using the implementation of the $\mathcal{F}$ statistic in LALPULSAR. Template banks are usually generated with mismatches of the order $\mu_{\max} \sim 0.2$–0.3 [e.g. [7,8]]; in this paper we generally restricted the sampled mismatches to $\mu_0 \le 0.2$.

The above procedure was parametrized by the coherent time span T and reference time $\Delta t_0 = t_0 -$ UTC 2007-01-01 00:03:06 of the simulated data and the maximum frequency $f_{\max}$ of the simulated signals. We performed the simulations at 34 values of T, from 1 to 31 days in steps of 2 days and from 36 to 121 days in steps of 5 days; 25 values of Δt_0, from 0 to 360 days in steps of 15 days; and 5 values of $f_{\max} \in \{50, 287.5, 525, 762.5, 1000\}$ Hz. In total, $\sim 10^8$ trials were performed. When plotted, however, the trials were filtered by restrictions on the mismatches, e.g., $\mu_0 \le 0.2$ in Fig. 2(a). The data from the simulations were also used to generate Figs. 6 and 7.

Figures 11, 12, 16, 18, and 20 were created by computing the various super sky metrics $\mathbf{g}_{\mathrm{ss}}$, $\mathbf{g}_{\mathrm{rss}}$, etc., using LALPULSAR, at: 121 values of T, from 1 to 121 days in steps of 1 day; 73 values of Δt_0, from 0 to 360 days in steps of 5 days; and at $f_{\max} = 1000$ Hz.

[1] R. Prix, *Neutron Stars and Pulsars*, edited by W. Becker, Astrophysics and Space Science Library Vol. 357 (Springer-Verlag, Berlin, 2009), p. 651.

[2] B. Sathyaprakash and B. F. Schutz, Living Rev. Relativity **12**, 2 (2009).

[3] B. P. Abbott *et al.* (LIGO Scientific Collaboration), Rep. Prog. Phys. **72**, 076901 (2009).

[4] T. Accadia *et al.* (Virgo Collaboration), JINST **7**, P03012 (2012).

[5] B. P. Abbott *et al.* (LIGO Scientific Collaboration), Astrophys. J. **713**, 671 (2010).

[6] J. Abadie *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Phys. Rev. D **85**, 022001 (2012).

[7] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Phys. Rev. D **87**, 042001 (2013).

[8] J. Abadie *et al.* (LIGO Scientific Collaboration), Astrophys. J. **722**, 1504 (2010).

[9] J. Abadie *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Astrophys. J. **737**, 93 (2011).

[10] G. M. Harry (LIGO Scientific Collaboration), Classical Quantum Gravity **27**, 084006 (2010).

[11] B. Abbott *et al.* (LIGO Scientific Collaboration), Astrophys. J. **683**, L45 (2008).

[12] B. Knispel and B. Allen, Phys. Rev. D **78**, 044031 (2008).

[13] R. Prix and M. Shaltev, Phys. Rev. D **85**, 084010 (2012).

[14] P. R. Brady, T. Creighton, C. Cutler, and B. F. Schutz, Phys. Rev. D **57**, 2101 (1998).

[15] B. Krishnan, A. M. Sintes, M. A. Papa, B. F. Schutz, S. Frasca, and C. Palomba, Phys. Rev. D **70**, 082001 (2004).

[16] V. Dergachev, Classical Quantum Gravity **27**, 205017 (2010).

[17] H. J. Pletsch, Phys. Rev. D **82**, 042002 (2010).

[18] C. Cutler, I. Gholami, and B. Krishnan, Phys. Rev. D **72**, 042004 (2005).

[19] K. Wette, Phys. Rev. D **85**, 042003 (2012).

[20] P. Jaranowski, A. Królak, and B. F. Schutz, Phys. Rev. D **58**, 063001 (1998).

[21] R. Balasubramanian, B. S. Sathyaprakash, and S. V. Dhurandhar, Phys. Rev. D **53**, 3033 (1996).

[22] B. J. Owen, Phys. Rev. D **53**, 6749 (1996).

[23] R. Prix, Classical Quantum Gravity **24**, S481 (2007).

[24] K. W. Wette, Ph.D. thesis, The Australian National University, 2009.

[25] R. Prix, Phys. Rev. D **75**, 023004 (2007).

[26] K. Wette *et al.*, Classical Quantum Gravity **25**, 235011 (2008).

[27] D. Whitbeck, Ph.D. thesis, The Pennsylvania State University, 2006.

[28] P. Jaranowski and A. Królak, Phys. Rev. D **59**, 063003 (1999).

[29] P. Astone, K. M. Borkowski, P. Jaranowski, and A. Królak, Phys. Rev. D **65**, 042003 (2002).

[30] H. J. Pletsch and B. Allen, Phys. Rev. Lett. **103**, 181102 (2009).

[31] M. Shaltev and R. Prix, Phys. Rev. D **87**, 084057 (2013).

[32] C. Cutler and B. F. Schutz, Phys. Rev. D **72**, 063006 (2005).

[33] R. Prix, Technical Report No. LIGO-T0900149-v3, https://dcc.ligo.org/LIGO-T0900149-v3/public.

[34] J. H. Conway and N. J. A. Sloane, *Sphere Packings, Lattices and Groups*, Grundlehren der Mathematischen Wissenshaften Vol. 290 (Springer-Verlag, Berlin, 1988).

[35] C. Messenger, R. Prix, and M. A. Papa, Phys. Rev. D **79**, 104017 (2009).

[36] I. W. Harry, B. Allen, and B. S. Sathyaprakash, Phys. Rev. D **80**, 104014 (2009).

[37] G. M. Manca and M. Vallisneri, Phys. Rev. D **81**, 024004 (2010).

[38] G. M. Manca and R. Prix (unpublished).

[39] N. J. Higham, *Accuracy and Stability of Numerical Algorithms* (SIAM, Philadelphia, 2002), 2nd ed.

[40] United States Naval Observatory, "Earth's Seasons: Equinoxes, Solstices, Perihelion, and Aphelion, 2000-2020," http://aa.usno.navy.mil/data/docs/EarthSeasons.php.

[41] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Phys. Rev. D **88**, 102002 (2013).

[42] Available from https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html.