

OPTIMIZED BLIND GAMMA-RAY PULSAR SEARCHES AT FIXED COMPUTING BUDGET

HOLGER J. PLETSCH^{1,2} AND COLIN J. CLARK^{1,2}

¹ Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut), D-30167 Hannover, Germany

² Institut für Gravitationsphysik, Leibniz Universität Hannover, D-30167 Hannover, Germany; holger.pletsch@aei.mpg.de

Received 2014 June 6; accepted 2014 August 28; published 2014 October 15

ABSTRACT

The sensitivity of blind gamma-ray pulsar searches in multiple years worth of photon data, as from the *Fermi* LAT, is primarily limited by the finite computational resources available. Addressing this “needle in a haystack” problem, here we present methods for optimizing blind searches to achieve the highest sensitivity at fixed computing cost. For both coherent and semicoherent methods, we consider their statistical properties and study their search sensitivity under computational constraints. The results validate a multistage strategy, where the first stage scans the entire parameter space using an efficient semicoherent method and promising candidates are then refined through a fully coherent analysis. We also find that for the first stage of a blind search incoherent harmonic summing of powers is not worthwhile at fixed computing cost for typical gamma-ray pulsars. Further enhancing sensitivity, we present efficiency-improved interpolation techniques for the semicoherent search stage. Via realistic simulations we demonstrate that overall these optimizations can significantly lower the minimum detectable pulsed fraction by almost 50% at the same computational expense.

Key words: gamma rays: general – methods: data analysis – methods: statistical – pulsars: general

Online-only material: color figures

1. INTRODUCTION

The *Fermi* Large Area Telescope (LAT; Atwood et al. 2009) has an unprecedented sensitivity to detect the periodic gamma-ray emission from spinning neutron stars. Owing to the LAT, the number of detected gamma-ray pulsars has vastly increased from a handful to about 150 (for a recent review, see, e.g., Caraveo 2014), making these objects a dominant Galactic source class at GeV energies.

So far, the largest fraction of LAT-detected gamma-ray pulsars has been uncovered indirectly (Abdo et al. 2013). In this approach, pulsar ephemerides known from previous radio observations are used to assign rotational phases to the gamma-ray photons, which are then tested for pulsations. Dedicated radio searches at positions of unidentified gamma-ray sources in the *Fermi*-LAT Second Source Catalog (2FGL; Nolan et al. 2012) have been particularly successful in discovering many new radio pulsars, and have provided ephemerides for subsequent gamma-ray phase-folding (e.g., Ransom et al. 2011; Guillemot et al. 2012; Abdo et al. 2013).

The direct detection of new gamma-ray pulsars, which are not known beforehand from other wavelengths, requires blind searches for periodicity in the sparse gamma-ray photon data (e.g., Chandler et al. 2001). With the *Fermi*-LAT, for the first time such blind searches have been successful (Abdo et al. 2009). Notably, many of the gamma-ray pulsars found this way have so far remained undetected at radio wavelengths (Abdo et al. 2013), implying that blind searches are the only way to access this pulsar population. Currently, hundreds of *Fermi*-LAT sources still remain unidentified, but feature pulsar-like properties (Ackermann et al. 2012; Lee et al. 2012) and thus likely harbor undiscovered pulsars.

The key problem in blind searches for gamma-ray pulsars is the enormous computational demand involved, which is what limits the search sensitivity. Since the relevant pulsar parameters are unknown in advance, one has to search a dense grid covering a multidimensional parameter space. The number

of search grid points increases rapidly with longer observation times. For observations spanning multiple years, “brute-force” (most sensitive but most expensive) methods, which involve fully coherently tracking the pulsar rotational phase over the entire observational data time span, are unfeasible. Therefore, the efficiency of blind-search methods is crucial because optimal strategies are those that provide the best search sensitivity at a fixed computing cost. This is the main theme of this work.

The problem is generally best addressed by a multistage search scheme (e.g., Meinshausen et al. 2009). This also applies to blind searches for gravitational-wave pulsars, i.e., spinning neutron stars emitting periodic gravitational waves (Brady et al. 1998; Brady & Creighton 2000; Cutler et al. 2005; Prix & Shaltev 2012). The basic idea is that in a first stage, the entire search parameter space is scanned but employing a much lower resolution, and therefore at a much lower computing cost, which can most efficiently discard unpromising regions. This reduction in parameter resolution is accomplished by semicoherent methods, in which only time intervals of data much shorter than one year are coherently analyzed and whose results are then incoherently summed over multiple years. In subsequent stages, only small promising regions (i.e., pulsar candidates) are followed up with higher resolution at higher computational expense by using longer coherent integration times.

One semicoherent method appropriate for the first search stage in gamma-ray pulsar searches is the seminal “time differencing technique” by Atwood et al. (2006, hereafter A06). It can basically be seen as the application of the classic Blackman–Tukey method (Blackman & Tukey 1958) to gamma-ray data. To search along the f -dimension (estimating the power spectrum), A06 calculated the discrete Fourier transform (DFT) of the autocorrelation function between photon arrival times up to a maximum lag. This significantly improved the efficiency over earlier methods (e.g., Brady & Creighton 2000; Chandler et al. 2001, summing power of many DFTs from subintervals) because the autocorrelation function can be computed at

negligible cost thanks to the sparsity of the photon arrival times. The success of the A06 method has been spectacularly demonstrated by the blind-search discovery of 24 gamma-ray pulsars (Abdo et al. 2009; Saz Parkinson et al. 2010) within the first *Fermi* mission year.

Using further improved methods, in part originally developed for blind searches for gravitational-wave pulsars (Pletsch & Allen 2009; Pletsch 2010), analyzing about three years of LAT data revealed 10 new gamma-ray pulsars (Pletsch et al. 2012a, 2012c). Crucial methodological improvements included the use of an analytic metric on parameter space to construct the grid over both sky position and frequency derivative. This allowed pulsars to be found that are much farther from the LAT catalog sky position than was possible previously. In addition, a photon weighting scheme (first studied by Kerr 2011) was used for both photon selection and for the search computations to ensure near-optimal detection significance. For enlarged computational resources, we have recently moved this ongoing search effort onto the volunteer computing system *Einstein@Home*.³ So far, this has resulted in the discovery of another four young pulsars (Pletsch et al. 2013). Here, we give a more detailed description of the strategies and methods exploited in these searches and consider related questions one might face with when setting up a blind search. Could a fully coherent blind search using a subset of data perhaps be more sensitive than a semicoherent search using all of the data? Is harmonic summing worthwhile under computational constraints? What is the optimal search-grid point density to balance sensitivity versus computing effort? In addressing such questions, we present the technical framework to optimize the sensitivity of blind pulsar searches in gamma-ray data at fixed computing cost. Moreover, we present further important methodological advances to improve the overall blind-search efficiency.

The paper is organized as follows. In Section 2, we describe the statistical detection of pulsations in general. In Section 3, we discuss the statistical properties of coherent blind searches and study their computational cost scalings using the parameter-space metric. We also investigate the efficiency of harmonic summing for different pulse profiles. In Section 4, we describe the statistical properties of a semicoherent blind-search method and compare the respective computing demand using the semicoherent metric. Section 5 presents a collection of technical improvements for the implementation of the semicoherent search stage, including efficient interpolation methods and automated candidate follow-up procedures. We demonstrate the superiority from combining these advances through realistic simulations in Section 6. Finally, conclusions follow in Section 7.

2. STATISTICAL DETECTION OF PULSATIONS

In blind pulsar searches, the pulse profile (the periodic light curve) and the exact parameters describing the rotational evolution of the neutron star are unknown in advance. As Bickel et al. (2008) have pointed out, unless the pulse profile shape is precisely known, there is no universally optimal statistical test because any most powerful test for one template profile will not be most powerful against another. Any test can only be most sensitive to a finite-dimensional class of targets. Thus, for computational feasibility of a blind search, an efficient (potentially suboptimal) template pulse profile to test against should attain only modest reduction in detection sensitivity

compared to an optimal template. The construction of such a test can be guided by the profiles of known gamma-ray pulsars, which we will consider below.

For isolated pulsars, the search parameters describing the rotational phase of the neutron star is at least four-dimensional, consisting of frequency f , spindown rate \dot{f} , and sky position with right ascension α and declination δ . To the LAT-registered arrival times t_{LAT} , sky-position (α, δ)-dependent corrections (“barycentric corrections”) are applied in order to obtain the photon arrival times t at the solar system barycenter (SSB). Then, the rotational phase $\Phi(t)$ is described by

$$\Phi(t) = \phi_0 + 2\pi f(t - t_0) + 2\pi \dot{f} \frac{(t - t_0)^2}{2}, \quad (1)$$

where f and \dot{f} are defined at reference time t_0 , when the phase equals the constant ϕ_0 .

Apart from the arrival time, for each of N detected gamma-ray photons, indexed by j , the LAT also records the photon’s reconstructed energy and direction. From these, a weight, w_j , can be computed measuring the probability that it has originated from the target source (Bickel et al. 2008; Kerr 2011). Using these probability weights efficiently avoids testing different hard selection cuts on energy and direction (implying binary weights), providing near optimal pulsation detection sensitivity (Kerr 2011; Pletsch et al. 2012a).

The observed gamma-ray pulse profile $F(\Phi)$, the flux as a function of Φ , can be written as

$$F(\Phi) \propto \frac{1-p}{2\pi} + p F_s(\Phi), \quad (2)$$

where p is the pulsed fraction that is estimated by the number of pulsed gamma-ray photons divided by the total number of photons. $F_s(\Phi)$ represents the pulse profile (undisturbed by background) and is a probability density function on $[0, 2\pi]$, which can be expressed as a Fourier series

$$F_s(\Phi) = \frac{1}{2\pi} \left(1 + \sum_{n \neq 0} \alpha_n e^{in\Phi} \right), \quad (3)$$

with the complex Fourier coefficients α_n , defined at harmonic order n as

$$\alpha_n = \int_0^{2\pi} F_s(\Phi) e^{-in\Phi} d\Phi. \quad (4)$$

Hence, the total flux $F(\Phi)$ can be rewritten as

$$F(\Phi) \propto 1 + p \sum_{n \neq 0} \alpha_n e^{in\Phi}. \quad (5)$$

If $F_s(\Phi)$ is an exact sinusoidal pulse profile, then from Equation (4) it follows that $|\alpha_1| = 1/2$ and all other coefficients vanish, $|\alpha_{n>1}| = 0$. As another example, if the pulse profile $F_s(\Phi)$ is a Dirac delta function, i.e., the narrowest possible profile, then all coefficients are equal, $|\alpha_n| = 1$, implying equal Fourier power at all harmonic orders.

In general, the null hypothesis is given by $p = 0$, meaning that all phases are uniformly distributed (i.e., no pulsations). From the likelihood for photon arrival times, Bickel et al. (2008) derived a score test statistic Q_M for $p > 0$,

$$Q_M = \frac{1}{K^2} \sum_{n=1}^M |\alpha_n|^2 |A_n|^2, \quad (6)$$

³ <http://einstein.phys.uwm.edu/>

where we defined the normalization constant K (different from Bickel et al. 2008) as

$$K^2 = \frac{1}{2M} \sum_{n=1}^M |\alpha_n|^2, \quad (7)$$

and A_n is given by

$$A_n = \frac{1}{\kappa} \sum_{j=1}^N w_j e^{-in\phi(t_j)}, \quad (8)$$

with the time-dependent part of the phase $\phi(t) = \Phi(t) - \phi_0$ and the normalization constant κ defined as

$$\kappa^2 = \frac{1}{2} \sum_{j=1}^N w_j^2. \quad (9)$$

Thus, we denote by \mathcal{P}_n the coherent Fourier power at the n th harmonic,

$$\mathcal{P}_n = |A_n|^2 = \frac{1}{\kappa^2} \left| \sum_{j=1}^N w_j e^{-in\phi(t_j)} \right|^2. \quad (10)$$

Appealing to the Central Limit Theorem (since $N \gg 1$ in all practical cases), the normalization choice of Equation (9) has the convenient property that the coefficients $\Re(A_n)$ and $\Im(A_n)$ become independent Gaussian random variables with zero mean and unit variance under the null hypothesis. Therefore, to good approximation, each \mathcal{P}_n is χ^2 -distributed with two degrees of freedom, as will be discussed below. Thus, Q_M is the weighted sum of coherent Fourier powers,

$$Q_M = \sum_{n=1}^M \frac{|\alpha_n|^2}{K^2} \mathcal{P}_n. \quad (11)$$

Therefore, as noted by Bickel et al. (2008), the test statistic Q_M is invariant under phase shifts (i.e., independent of reference phase ϕ_0) and only depends on the amplitudes of the Fourier coefficients α_n , but not on their phases. Moreover, Beran (1969) showed earlier that if the pulse profile is known a priori, a test statistic following from Q_M for binary weights is locally most powerful for testing uniformity of a circular distribution, assuming unknown and weak (small p) signal strength.

3. COHERENT TEST STATISTICS

In what follows, we examine the sensitivity of coherent blind searches at fixed computational cost, taking into account the statistical properties and sensitivity scalings in terms of relevant quantities. For simplicity, during the remainder of this section here we assume hard photon selection cuts, i.e., binary weights only, $w_j \in \{0, 1\}$, such that \mathcal{P}_n reduces to

$$\mathcal{P}_n = \frac{2}{N} \left| \sum_{j=1}^N e^{-in\phi(t_j)} \right|^2. \quad (12)$$

However, the main conclusions obtained will also have applicability when arbitrary (i.e., non-binary) weights are used.

3.1. Statistical Properties

Under the null hypothesis $p = 0$ and assuming $N \gg 1$, the coherent power \mathcal{P}_n as of Equation (12) follows a central χ^2 distribution with two degrees of freedom (see Appendix A), whose the first two moments are

$$E_0[\mathcal{P}_n] = 2, \quad \text{Var}_0[\mathcal{P}_n] = 4. \quad (13)$$

Suppose the photon data contain a pulsed signal, $p > 0$, whose pulse profile can be expressed in terms of complex Fourier coefficients, γ_n as in Equation (4). In this case, we show in Appendix A that for moderately strong pulsed signals the distribution of \mathcal{P}_n can be well approximated by a noncentral χ^2 distribution (Groth 1975; Guidorzi 2011) with two degrees of freedom. Thus, in the perfect-match case (the pulsar parameters f , \dot{f} , and sky position are precisely known), the first two moments are approximately given by

$$E_p[\mathcal{P}_n] \approx 2 + 2p^2 N |\gamma_n|^2, \quad (14a)$$

$$\text{Var}_p[\mathcal{P}_n] \approx 4 + 8p^2 N |\gamma_n|^2, \quad (14b)$$

where pN photons are assumed to be ‘‘pulsed’’ and accordingly $(1-p)N$ photons are ‘‘non-pulsed’’ (i.e., background). Thus, the second summand in Equation (14a) represents the noncentrality parameter.⁴ We can also identify the amplitude signal-to-noise ratio (S/N) at the n th harmonic, $\theta_{\mathcal{P}_n}$, as

$$\theta_{\mathcal{P}_n}^2 = \frac{E_p[\mathcal{P}_n] - E_0[\mathcal{P}_n]}{\sqrt{\text{Var}_0[\mathcal{P}_n]}} \approx p^2 N |\gamma_n|^2. \quad (15)$$

Therefore, by comparison with Equation (14a), the noncentrality parameter is just $2\theta_{\mathcal{P}_n}^2$.

A similar calculation for Q_M , based on the above relations, shows that if $p = 0$,

$$E_0[Q_M] = 2M, \quad \text{Var}_0[Q_M] = \frac{4}{K^4} \sum_{n=1}^M |\alpha_n|^4, \quad (16)$$

and for $p > 0$, one obtains

$$E_p[Q_M] \approx 2M + \frac{2p^2 N}{K^2} \sum_{n=1}^M |\alpha_n|^2 |\gamma_n|^2. \quad (17)$$

Thus, the amplitude S/N θ_{Q_M} for the test statistic Q_M can be expressed as

$$\theta_{Q_M}^2 \approx \frac{p^2 N \sum_{n=1}^M |\alpha_n|^2 |\gamma_n|^2}{\sqrt{\sum_{n=1}^M |\alpha_n|^4}}. \quad (18)$$

A similar expression has been derived by Bickel et al. (2008), who used this parameter as an approximate measure of the sensitivity of the test statistic Q_M since the larger the S/N θ_{Q_M} , the higher the probability of detection. However, it is only an approximate sensitivity measure because any meaningful sensitivity comparison must be done at a fixed probability of false alarm as will be described below. Equation (18) also shows that the S/N is maximized if $|\alpha_n|^2 \propto |\gamma_n|^2$, i.e., when the template pulse profile α_n perfectly matches the γ_n , representing

⁴ A random variable X following a noncentral χ^2 -distribution with two degrees of freedom and noncentrality parameter λ has expectation value $2 + \lambda$.

the signal pulse profile. However, as Bickel et al. (2008) correctly note, practical blind searches can only test for a finite-dimensional class of template pulse profiles.

A particularly simple template profile for a given value of M is

$$|\alpha_n| = \begin{cases} 1, & n \leq M \\ 0, & n > M \end{cases} \quad (19)$$

With this choice, Q_M measures the coherent Fourier power summed over the first M harmonics, which we therefore refer to as incoherent harmonic summing. The resulting statistic is also known as Z_M^2 (Buccheri et al. 1983),

$$Z_M^2 = \sum_{n=1}^M \mathcal{P}_n. \quad (20)$$

Maximizing Z_M^2 over different values of M as $H = \max_{1 \leq M \leq 20} (Z_M^2 - 4M + 4)$ also recovers the widely used H -test by de Jager et al. (1989).

The template of Equation (19) has the additional benefit that the statistical distribution of Z_M^2 is known analytically. Therefore, we use this to obtain realistic sensitivity scalings for such coherent test statistics. Since \mathcal{P}_n is χ_2^2 -distributed,⁵ it follows that Z_M^2 is distributed as χ_{2M}^2 . Thus, one obtains

$$E_0 [Z_M^2] = 2M, \quad \text{Var}_0 [Z_M^2] = 4M, \quad (21)$$

and

$$E_p [Z_M^2] \approx 2M + 2\theta_M^2 \sqrt{M}. \quad (22)$$

Correspondingly, the S/N θ_M is written as

$$\theta_M^2 = \frac{1}{\sqrt{M}} \sum_{n=1}^M \theta_{\mathcal{P}_n}^2 = \frac{p^2 N}{\sqrt{M}} \sum_{n=1}^M |\gamma_n|^2. \quad (23)$$

In the Neyman–Pearson sense, we define search sensitivity from the lowest threshold pulsed fraction required to achieve a certain detection probability P_{DET}^* for a given number of photons N and at given false alarm probability P_{FA}^* . For Z_M^2 the false alarm probability is computed as

$$P_{\text{FA}}(Z_{M,\text{th}}^2) = \int_{Z_{M,\text{th}}^2}^{\infty} \chi_{2M}^2(Z_M^2; 0) dZ_M^2, \quad (24)$$

where $\chi_k^2(X; \lambda)$ denotes the probability density function for the χ_k^2 -distributed variable X with noncentrality parameter λ . The probability of detection for a noncentrality parameter of $2\theta_M^2 \sqrt{M}$ is

$$P_{\text{DET}}(Z_{M,\text{th}}^2, 2\theta_M^2 \sqrt{M}) = \int_{Z_{M,\text{th}}^2}^{\infty} \chi_{2M}^2(Z_M^2; 2\theta_M^2 \sqrt{M}) dZ_M^2. \quad (25)$$

The minimum detectable pulsed-fraction threshold for summing coherent power from M harmonics, $p_{\text{coh},M}$, is obtained by first inverting Equation (24) to get the threshold test-statistic value $Z_{M,\text{th}}^2(P_{\text{FA}}^*)$, which in a second step is substituted in Equation (25) to numerically find the required threshold S/N:

$$\theta_M^* = \theta_M(P_{\text{FA}}^*, P_{\text{DET}}^*). \quad (26)$$

⁵ We use the notation χ_k^2 to indicate a χ^2 distribution with k degrees of freedom.

Finally, Equation (23) can be used to convert the threshold S/N θ_M^* into $p_{\text{coh},M}$, which defines the coherent search sensitivity as

$$p_{\text{coh},M}^{-1} = \frac{\sqrt{N}}{M^{1/4} \theta_M^*} \left[\sum_{n=1}^M |\gamma_n|^2 \right]^{1/2}. \quad (27)$$

Assuming the overall photon count rate, $\mu = N/T_{\text{coh},1}$, is constant throughout the entire coherent integration time, $T_{\text{coh},1}$ then the search sensitivity increases with the well-known square-root scaling of $T_{\text{coh},1}$,

$$p_{\text{coh},M}^{-1} = \frac{\sqrt{\mu T_{\text{coh},1}}}{M^{1/4} \theta_M^*} \left[\sum_{n=1}^M |\gamma_n|^2 \right]^{1/2}. \quad (28)$$

Thus, we have obtained an expression for the search sensitivity, separating the two effects of photon count rate (or integration time) and pulse profile shape. Regarding the latter effect, Equation (28) reveals that the sensitivity only improves with including higher harmonics (i.e., increasing M) if the pulse profile shape is such that $(\sum_{n=1}^M |\gamma_n|^2)^{1/2}$ increases more quickly than the “statistical penalty” factor $M^{1/4} \theta_M^*$. While this is true for the narrowest possible pulse profile (a Dirac delta function), we show below that the same does not hold in general for typical gamma-ray pulsar profiles.

3.2. Effects of Pulse Profile on Sensitivity

From Equation (28) in the previous section, we have seen how the sensitivity for pulsation detection depends on the shape of the pulse profile, represented by the Fourier coefficients γ_n . Therefore, it is instructive to examine the change in sensitivity as a function of the number of harmonics M for some exemplary profiles. Thus, we consider the following ratio:

$$\frac{p_{\text{coh},M}^{-1}}{p_{\text{coh},1}^{-1}} = \frac{\theta_1^*}{M^{1/4} \theta_M^*} \frac{1}{|\gamma_1|} \left[\sum_{n=1}^M |\gamma_n|^2 \right]^{1/2}, \quad (29)$$

which compares in the statistical sense the search sensitivity of including M harmonics, compared to using the fundamental only (in the absence of any computational constraints).

In the ideal case, where all harmonics have equal power $|\gamma_n|^2 = 1$, the pulse profile is a Dirac delta function as described above. In this case, $(\sum_{n=1}^M |\gamma_n|^2)^{1/2} = M^{1/2}$, and the sensitivity is a monotonically increasing function of M at fixed detection probability, P_{DET}^* , and fixed false alarm probability, P_{FA}^* . To illustrate this, consider the following example, assuming that $P_{\text{FA}}^* = 1\%$ and $P_{\text{DET}}^* = 90\%$. Then, to good approximation, the corresponding S/N threshold θ_M^* can be described by

$$\theta_M^* \approx \left(3.715 + \frac{4.987}{\sqrt{M}} \right)^{1/2}. \quad (30)$$

Hence, with increasing M , the threshold S/N θ_M^* decreases and becomes constant in the limit of large M , in which case the statistical penalty factor $(M^{1/4} \theta_M^*)$ becomes $\propto M^{1/4}$. Since this scaling is slower than the pulse profile factor $(\sum_{n=1}^M |\gamma_n|^2)^{1/2} = M^{1/2}$ in this case, the sensitivity is monotonically increasing with M . This is also shown in Figure 1, using the exact values for θ_M^* that we calculated numerically.

To obtain a more realistic signal pulse-profile model, we considered those of the known gamma-ray pulsars. We carried




Figure 1. Sensitivity as a function of the number of harmonics M included in absence of computational constraints and for three different pulse-profile models. In each panel, we fixed the detection probability $P_{\text{DET}}^* = 90\%$ and the four curves correspond to different values of false alarm probability P_{FA}^* as shown by the legend. The upper panel is for a Dirac delta function pulse profile (implying equal Fourier power at all harmonics). The middle panel is for a typical pulse profile, obtained from the known gamma-ray pulsars by averaging those profiles that are mostly single-peaked (i.e., the γ_n values shown in the bottom left panel in Figure 2). The bottom panel is also for a realistic pulse profile, obtained from the known gamma-ray pulsars by averaging those profiles that are mostly two-peaked (i.e., the γ_n values shown in the bottom right panel in Figure 2). Since for these profiles the Fourier power $|\gamma_2|^2$ is highest at the second harmonic ($n = 2$), in this plot the vertical axis shows the sensitivity compared to a blind search which would report the highest detection significance at the second harmonic (i.e., “misidentify” the fundamental).

(A color version of this figure is available in the online journal.)

out a harmonic analysis of the pulse profile shapes of the 117 known gamma-ray pulsars listed in the second *Fermi* LAT pulsar catalog (Abdo et al. 2013) and computed their Fourier coefficients, γ_n . These are shown in Figure 2 (top panel) and illustrate that for most of the known gamma-ray pulsars the largest fraction of Fourier power is typically in a single harmonic that is either the first (mostly single-peaked profiles) or the second (mostly two-peaked profiles). Therefore, before computing an average profile (by averaging the $|\gamma_n|$), it makes sense to divide the pulsars into these two groups (based on whether or not $|\gamma_1| > |\gamma_2|$). These results, separately for each group, are displayed in the two bottom panels of Figure 2.

We use the resulting two sets of coefficients γ_n to calculate the sensitivity scaling with M from Equation (28) as also shown in Figure 1. Notice that for the typical pulse profiles, in contrast to the Dirac delta pulse-profile, when summing more than a certain number of harmonics, the sensitivity starts to decrease (at fixed P_{DET}^* and P_{FA}^*). This is because the Fourier powers $|\gamma_n|^2$ at the higher harmonics become vanishingly small and thus effectively only contribute “noise” when summed (i.e., the statistical penalty factor cannot be overcome anymore).

These results also illustrate the success of the H -test for targeted pulsation searches in gamma-ray data with known pulsar ephemerides, because this test maximizes the Fourier power sums over the first 20 harmonics. Maximizing only over fewer harmonics could likely already be sufficient (or even be more sensitive due to the reduced trials factor) in most cases, as suggested by Figure 1. Besides, further improvements over the H -test could also be achieved by employing one or more template profiles α_n that are more representative of the typical gamma-ray profile (than the delta function) to compute the Q_M




Figure 2. Harmonic analysis of pulse profiles of the 117 gamma-ray pulsars in the second *Fermi* LAT pulsar catalog (Abdo et al. 2013). Top panel: Fourier power $|\gamma_n|^2$ (color-coded) at the n th harmonic (vertical axis) for each of the 117 pulsars (horizontal axis). Bottom left panel: Fourier power $|\gamma_n|^2$ at the n th harmonic averaged over the 76 out of the 117 pulsars, whose power at the fundamental is highest (mostly single-peaked profiles). Bottom right panel: Fourier power $|\gamma_n|^2$ at the n th harmonic averaged over 41 out of the 117 pulsars, whose power at the second harmonic is highest (mostly two-peaked profiles).

(A color version of this figure is available in the online journal.)

test statistic. Using the average profile from the known pulsars from above for this seems the simplest first step. While also conducting a principal component analysis appears worthwhile, we defer a detailed study of this to future work.

So far, we have not considered the computational costs involved, which is only justifiable for computationally inexpensive targeted searches. In contrast, blind searches are limited by computational power. Therefore, in the following section, we will revisit the efficiency of harmonic summing under the constraint of a fixed computational cost.

3.3. Grid-point Counting for Coherent Search

In blind searches, the pulsar’s rotational and positional parameters are unknown a priori. Therefore, one has to construct a grid in the multidimensional search parameter space that is explicitly searched, i.e., the test statistic is to be computed at each grid point. Therefore the question arises: What is the most efficient scheme for constructing the search grid? If grid points are placed too far apart potential pulsar signals might be missed. On the other hand, it is highly inefficient to place grid points too closely together, because of redundancy resulting from strongly correlated nearby grid points. The problem of constructing efficient search grids has been intensively studied

in the context of gravitational-wave searches (see, e.g., Brady et al. 1998; Brady & Creighton 2000; Prix 2007; Pletsch & Allen 2009; Pletsch 2010), and we employ some of these concepts here.

The key element is a distance metric on the search space (Balasubramanian et al. 1996; Owen 1996). The metric provides an analytic geometric tool measuring the expected fractional loss in squared S/N for any given pulsar-signal location at a nearby grid point.

Let the vector \mathbf{u}_{sig} collect the actual pulsar signal parameters. In a blind search for isolated pulsars, this vector is at least four-dimensional, $\mathbf{u}_{\text{sig}} = (f_{\text{sig}}, \dot{f}_{\text{sig}}, \alpha_{\text{sig}}, \delta_{\text{sig}})$. For simplicity, we begin by considering the metric at the fundamental harmonic ($n = 1$). As will be shown below, it is subsequently straightforward to generalize the results to higher harmonic orders. Following Equation (15), let $\theta_{\mathcal{P}_1}(\mathbf{u}_{\text{sig}})$ denote the S/N for the perfect-match case, i.e., at the signal parameter-space location. In a blind search, the signal parameters generally will not coincide with a grid point \mathbf{u} , but will typically have some offset,

$$\Delta\mathbf{u} = \mathbf{u} - \mathbf{u}_{\text{sig}}. \quad (31)$$

These offsets lead to a (time-dependent) residual phase $\phi(t; \mathbf{u}) - \phi(t; \mathbf{u}_{\text{sig}})$ and therefore a fractional loss in squared S/N results, which is commonly referred to as mismatch,

$$m(\Delta\mathbf{u}) = 1 - \frac{\theta_{\mathcal{P}_1}^2(\mathbf{u})}{\theta_{\mathcal{P}_1}^2(\mathbf{u}_{\text{sig}})} = 1 - \frac{\theta_{\mathcal{P}_1}^2(\mathbf{u}_{\text{sig}} + \Delta\mathbf{u})}{\theta_{\mathcal{P}_1}^2(\mathbf{u}_{\text{sig}})}. \quad (32)$$

The metric is obtained from a Taylor expansion of the mismatch to second order in the offsets $\Delta\mathbf{u}$ at the signal location \mathbf{u}_{sig} ,

$$m(\Delta\mathbf{u}) \approx \sum_{k,\ell} G_{k\ell} \Delta u^k \Delta u^\ell + \mathcal{O}(\Delta u^3), \quad (33)$$

This equation defines a positive definite metric tensor G with components $G_{k\ell}$, where k and ℓ label the tensor indices. In Appendix B, we derive explicit expressions for the coherent metric for a simplified phase model that is appropriate for the purpose of grid construction. We also find that the resulting metric tensor G is diagonal, which greatly simplifies the grid construction. The results of this derivation will therefore be used in what follows.

As noted by Prix & Shaltev (2012), the probability distribution of signal mismatches in a given search grid constructed with a certain maximal mismatch m depends on the structure and dimensionality of the search parameter space. The corresponding average mismatch in each dimension, ξm , will generally be smaller by a characteristic geometric factor $\xi \in (0, 1)$, depending on the actual search-grid construction. For example, for hyper-cubical lattices, ξ is known to be $\xi = 1/3$. In order to construct a hyper-cubical grid in which the maximum mismatch due to an offset in each parameter is m , then the grid point spacing in each parameter should be

$$\Delta u^k = 2\sqrt{\frac{m}{G_{kk}}}. \quad (34)$$

Denote by \mathcal{U} the four-dimensional parameter space, spanned by \mathbf{u} , which is to be searched. Thus, when searching for pulsars with spin frequencies in the range $[0, f_{\text{max}}]$, with spin-down rates in the range $[\dot{f}_{\text{max}}, 0]$, and whose sky location is confined by the LAT to a region of area A_{sky} , the proper volume \mathcal{U} can be written as

$$\mathcal{U} = f_{\text{max}} | \dot{f}_{\text{max}} | A_{\text{sky}}. \quad (35)$$




Figure 3. Scaling of the determinant of the coherent metric G as function of the coherent integration time $T_{\text{coh},1}$ (black solid curve). The red dot-dashed curve shows the model of the coherent metric determinant from the approximation of Equation (39) used to estimate the computing cost scaling.

(A color version of this figure is available in the online journal.)

In principle, the metric coefficients (and hence also the grid point spacings) can vary throughout the parameter space. Indeed, for the metrics considered in this work, the grid point spacing in the sky dimensions depends on the spin frequency of the pulsar. In order to avoid having to construct a separate sky grid for each search frequency value, we adopt the conservative approach of using the highest frequency searched f_{max} for the sky grid construction. The metric (and hence also the grid point spacing) becomes uniform throughout \mathcal{U} . The total number of search-grid points $\mathcal{N}_{\text{coh},1}$ for a coherent blind search over \mathcal{U} is therefore simply the product of the number of grid points in each dimension.

$$\mathcal{N}_{\text{coh},1} = \mathcal{U} \prod_k \frac{1}{\Delta u^k} = \frac{1}{16} \mathcal{U} m^{-2} \sqrt{\det G}, \quad (36)$$

as G is found to be diagonal. In Appendix B, we derive that

$$\sqrt{\det G} = \frac{\pi^4}{\sqrt{135}} T_{\text{coh},1}^3 f^2 r_E^2 \Psi(T_{\text{coh},1}), \quad (37)$$

where we defined⁶

$$\Psi^2(T_{\text{coh},1}) = [1 + \text{sinc}(\Omega_E T_{\text{coh},1}/\pi) - 2 \text{sinc}^2(\Omega_E T_{\text{coh},1}/2\pi)] \times [1 - \text{sinc}(\Omega_E T_{\text{coh},1}/\pi)], \quad (38)$$

and where we have denoted the Earth's orbital angular frequency as $\Omega_E = 2\pi/1 \text{ yr}$ and the light travel-time from the Earth to the SSB as $r_E = 1 \text{ AU}/c \sim 500 \text{ s}$.

To analytically study the scaling of $\mathcal{N}_{\text{coh},1}$ as a function of $T_{\text{coh},1}$, the function $\Psi(T_{\text{coh},1})$ can be well approximated by

$$\Psi(T_{\text{coh},1}) \approx \begin{cases} \frac{\Omega_E^3 T_{\text{coh},1}^3}{12\sqrt{15}}, & T_{\text{coh},1} < 0.572 \text{ yr} \\ 1, & T_{\text{coh},1} \geq 0.572 \text{ yr} \end{cases}. \quad (39)$$

The validity of this approximation is illustrated in Figure 3. Hence, the total number of grid points required in a coherent search is

$$\mathcal{N}_{\text{coh},1} = \frac{\pi^4}{48\sqrt{15}} \left(\frac{\Omega_E^3}{12\sqrt{15}} \right)^{(a-3)/3} r_E^2 m^{-2} f_{\text{max}}^2 T_{\text{coh},1}^a \mathcal{U}, \quad (40)$$

⁶ We use the definition $\text{sinc}(x) = \sin(\pi x)/(\pi x)$ throughout this article.

where

$$a \approx \begin{cases} 6, & T_{\text{coh},1} < 0.572\text{yr} \\ 3, & T_{\text{coh},1} \geq 0.572\text{yr} \end{cases}. \quad (41)$$

Equation (40) tells us that for coherent integration times much shorter than half a year the sky metric components also still scale with $T_{\text{coh},1}$, such that $\mathcal{N}_{\text{coh},1}$ increases approximately as $T_{\text{coh},1}^6$. After half a year of coherent integration the sky metric components quickly approach the resolution saturation as the maximum baseline (1 AU) is reached, and thereafter become approximately independent of $T_{\text{coh},1}$. Therefore, $\mathcal{N}_{\text{coh},1}$ scales only as $T_{\text{coh},1}^3$ in this regime.

3.4. Coherent Search Sensitivity at Fixed Computing Cost

For computational efficiency, we use the fast Fourier transform (FFT) algorithm (Frigo & Johnson 2005) to scan the f dimension. There are two steps involved in calculating an FFT, each with an associated computational cost. First, it is necessary to construct a discrete time series by interpolating (e.g., by binning) the photon arrival times into equidistant samples. The cost of this step is proportional to the number of photon arrival times which must be interpolated. Second, the discrete time series must be transformed into a discretely sampled frequency spectrum using the FFT algorithm. For a maximum frequency of f_{max} and a coherent integration time of $T_{\text{coh},1}$, there are $f_{\text{max}}T_{\text{coh},1}$ frequency samples, and the computational cost of calculating the FFT is proportional to $f_{\text{max}}T_{\text{coh},1} \log_2(f_{\text{max}}T_{\text{coh},1})$. We assume that the cost of calculating the FFT is much larger than the cost of creating the discrete time series. Compared to the cost of computing \mathcal{P}_1 explicitly for N photon times at $f_{\text{max}}T_{\text{coh},1}$ frequencies, which is proportional to $Nf_{\text{max}}T_{\text{coh},1}$, it is clear that the FFT method offers more efficiency provided $N \gg \log_2(f_{\text{max}}T_{\text{coh},1})$.

The spacing of frequency samples output by the FFT is $1/T_{\text{coh},1}$. According to the metric (see Equation (B11a)), this implies a worst-case mismatch due to frequency offsets of $m = G_{ff}/(4T_{\text{coh},1}^2) = \pi^2/12 = 0.82$, which obviously also leads to a high average of mismatches. However, as we will discuss in Section 5.2, it is possible to reduce this mismatch at almost no extra computational cost by interpolating the frequency spectrum. In the following derivations, we therefore separate the total mismatch m_{tot} into two components: a constant mismatch due to the frequency spacing, m_f determined by the interpolation method used, which has a negligible effect on the overall computing cost; and the mismatch due to offsets in the remaining parameters, m , which can be freely varied to construct an optimal grid.

For every grid point in $\{\dot{f}, \alpha, \delta\}$, an FFT must be computed, and hence the overall computation time for the search is simply the cost of calculating one FFT multiplied by the number of FFTs that must be computed. The total cost, $C_{\text{coh},1}$ (measured in units of time), is

$$C_{\text{coh},1} = K_{\text{FFT}} f_{\text{max}} T_{\text{coh},1} \log_2(f_{\text{max}} T_{\text{coh},1}) \frac{\mathcal{N}_{\text{coh},1}}{\mathcal{N}_f}, \quad (42)$$

where K_{FFT} is an implementation and computing hardware dependent constant, and where \mathcal{N}_f is the number of frequency samples that would be calculated using a grid with an arbitrary maximum mismatch per dimension of m ,

$$\mathcal{N}_f = \frac{f_{\text{max}}}{2} \sqrt{\frac{G_{ff}}{m}} = \frac{\pi}{2\sqrt{3}m} f_{\text{max}} T_{\text{coh},1}. \quad (43)$$

The total computational cost is therefore

$$C_{\text{coh},1} = K_{\text{coh},a} m^{-3/2} T_{\text{coh},1}^a \log_2(T_{\text{coh},1} f_{\text{max}}), \quad (44)$$

where the constant $K_{\text{coh},a}$ depends on a ,

$$K_{\text{coh},a} = K_{\text{FFT}} \frac{\pi^3 r_E^2 f_{\text{max}}^2 \mathcal{U}}{24\sqrt{5}} \left(\frac{\Omega_E^3}{12\sqrt{15}} \right)^{(a-3)/3}. \quad (45)$$

For a search grid constructed with maximum mismatch $m_{\text{tot}} = m_f + 3m$, the search sensitivity will scale with the average mismatch $\langle m_{\text{tot}} \rangle = \langle m_f \rangle + 3\langle m \rangle$ as $\sqrt{1 - \langle m_{\text{tot}} \rangle}$ (Prix & Shaltev 2012). Thus, from Equation (28) it follows that the search sensitivity without harmonic summing scales as

$$p_{\text{coh},1}^{-1} = \frac{\sqrt{(1 - \langle m_{\text{tot}} \rangle) \mu T_{\text{coh},1}}}{\theta_1^*} |\gamma_1|. \quad (46)$$

For a computing cost $C_{\text{coh},1}$, Equation (44) can be used to obtain (numerically) the maximum $T_{\text{coh},1}$. Substituting this value of $T_{\text{coh},1}$ in Equation (46) finally yields the search sensitivity at the given computational cost.

3.5. Efficiency of Harmonic Summing at Fixed Computing Cost

Based on the results of the previous sections, we now investigate the efficiency of incoherent harmonic summing under computational cost constraints. More precisely, we address the question of whether it is more efficient in blind searches to sum M harmonics, or to instead use a longer coherent integration time without harmonic summing at the same computing cost.

Thus, we consider the test statistic Z_M^2 , which incoherently sums Fourier powers \mathcal{P}_n from M higher harmonics. In Appendix C, we derive the parameter space metric for the Z_M^2 statistic, denoted by \tilde{G} , and find that $\sqrt{\det \tilde{G}} = r^4 \sqrt{\det G}$, where r represents a refinement factor due to harmonic summing, and G is the metric tensor for \mathcal{P}_1 of Equation (37). Therefore, to ensure equal sensitivity throughout the original parameter space⁷ the required number of grid points increases by the factor of r^4 compared to using \mathcal{P}_1 only. The value of $r \geq 1$ depends on the pulse profile γ_n . For a sinusoidal pulse profile ($|\gamma_1| = 1/2$ and $|\gamma_{n>1}| = 0$), obviously, $r = 1$ (i.e., no refinement), and for a Dirac delta function ($|\gamma_n| = 1$), one finds $r \sim M$, as derived in Equation (C6). In principle, one could construct a grid with $r^4 \mathcal{N}_{\text{coh},1}$ points, and calculate and sum M values of \mathcal{P}_n at each point, leading to the cost of a harmonic summing search being simply Mr^4 times greater than that of a coherent search at the fundamental frequency with the same coherent integration time.

In practice, to utilize the efficiency of the FFT, it would be necessary to construct a sub-optimal grid in which the range in f and \dot{f} is extended by a factor of M , and the coherent powers summed appropriately over harmonics. The sky-grid in this case may still be constructed using the refinement factor r , leading to the computing cost being $M^2 r^2$ times $C_{\text{coh},1}$ at the same coherent integration time. While this method may quickly become infeasible due to the amount of memory required, we use this only as a theoretically efficient method to compare to an equally costly search using only the fundamental harmonic power.

⁷ This constraint is imposed to eliminate any detection bias in favor of pulsars with low frequencies and frequency derivatives, allowing for estimates of the true astrophysical pulsar populations.

Here, we assume that the small extra cost of actually summing the \mathcal{P}_n is negligible.⁸ The computational expense for incoherent harmonic summing, $C_{\text{coh},M}$, using the Z_M^2 statistic for a coherent integration time $T_{\text{coh},M}$ becomes

$$C_{\text{coh},M} = K_{\text{coh},a} m^{-3/2} T_{\text{coh},M}^a M^2 r^2 \log_2(T_{\text{coh},M} f_{\text{max}} M). \quad (47)$$

From Equation (27) above, we found that the search sensitivity of incoherent harmonic summing is given by

$$p_{\text{coh},M}^{-1} = \frac{\sqrt{(1 - \langle m_{\text{tot}} \rangle) \mu T_{\text{coh},M}}}{M^{1/4} \theta_M^*} \left[\sum_{n=1}^M |\gamma_n|^2 \right]^{1/2}. \quad (48)$$

Hence, to compare the search sensitivities $p_{\text{coh},1}^{-1}$ and $p_{\text{coh},M}^{-1}$ at fixed computing cost, in principle, the following steps are required. First, for a given computing cost $C_{\text{coh},1}$, Equations (44) and (46) provide the corresponding coherence time $T_{\text{coh},1}$ and sensitivity $p_{\text{coh},1}^{-1}$, respectively. Second, by equating $C_{\text{coh},1} = C_{\text{coh},M}$, Equation (47) then can be solved (numerically) for $T_{\text{coh},M}$, which finally is used to obtain the sensitivity $p_{\text{coh},M}^{-1}$ from Equation (48). It should be noted that in comparing $p_{\text{coh},1}^{-1}$ and $p_{\text{coh},M}^{-1}$ the same values of P_{FA}^* and P_{DET}^* must be assumed. We here also assume the same mismatch m in either case, because as shown in Appendix E, the optimal mismatch at fixed computing cost is independent of coherent integration time, number of harmonics summed, and computing power available. Notably, a similar result has been found previously by Prix & Shalvet (2012) in the context of gravitational-wave pulsar searches.

In the following, we describe an analytical approximation to the numerical approach above which we show to be sufficiently accurate for typical search setups. This approximation is based on ignoring the slowly varying \log_2 factors in Equations (44) and (47), such that

$$C_{\text{coh},M} \sim K_{\text{coh},a} m^{-3/2} T_{\text{coh},M}^a M^2 r^2. \quad (49)$$

Then from $C_{\text{coh},1} = C_{\text{coh},M}$, it immediately follows that $T_{\text{coh},M}$ must be shorter by the factor $(M^2 r^2)^{1/a}$,

$$T_{\text{coh},M} = T_{\text{coh},1} (M^2 r^2)^{-1/a}. \quad (50)$$

We show in Appendix D that the $T_{\text{coh},M}$ obtained from this approximation slightly overestimates the sensitivity $p_{\text{coh},M}^{-1}$, while being accurate to within less than about 1% for typical search setups. Using Equation (50) to substitute $T_{\text{coh},M}$ in Equation (48), one obtains for the ratio of search sensitivities,

$$\frac{p_{\text{coh},1}^{-1}}{p_{\text{coh},M}^{-1}} = \frac{M^{1/4+1/a} r^{1/a} \theta_M^*}{\theta_1^*} |\gamma_1| \left[\sum_{n=1}^M |\gamma_n|^2 \right]^{-1/2}, \quad (51)$$

which, remarkably, is independent of $T_{\text{coh},1}$ and $T_{\text{coh},M}$. This sensitivity ratio $p_{\text{coh},1}^{-1}/p_{\text{coh},M}^{-1}$ of Equation (51) is shown in Figure 4 and is found to be greater than unity for typical gamma-ray pulsars. Only for unrealistically narrow pulse profiles (i.e., a Dirac delta function), the sensitivity ratio can remain close to or slightly below unity. It also should be pointed out that we obtained these results despite the generous assumptions in favor

⁸ Note that this makes the computing cost estimate generous in favor of the harmonic summing approach in this comparison.

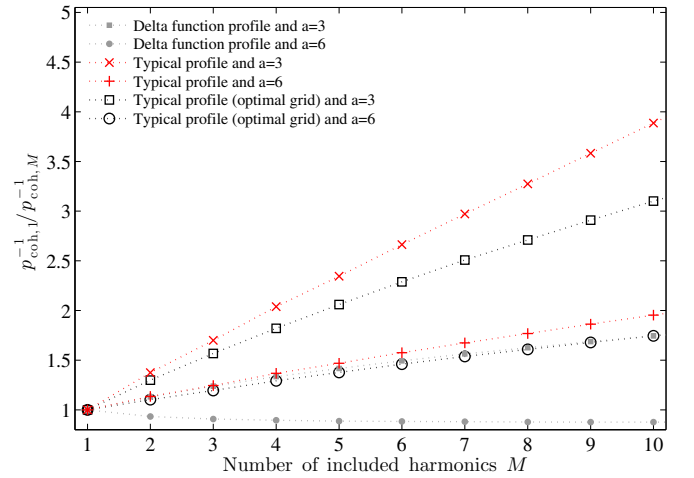


Figure 4. Ratio of search sensitivities $p_{\text{coh},1}^{-1}$ (without harmonic summing) and $p_{\text{coh},M}^{-1}$ (summing power from M harmonics) at fixed computational cost. The Z_M^2 test statistic used here assumes a delta function pulse profile, so optimizing the search grid for this profile leads to the curves shown by the filled squares and circles. The red crosses and pluses are for the same statistic and grid, but where the signal pulse profile is a more typical one (derived from averaging those of the known gamma-ray pulsars of Figure 1). The open squares and circles are for the same statistic, but using the same typical signal pulse profile and a grid that is also optimized for that same pulse profile. For each case, the results of two different scalings of the computing cost with $T_{\text{coh},M}^a$ are shown, corresponding to $a = 6$ and $a = 3$ (see the text for details). While all points shown are for $P_{\text{DET}}^* = 90\%$ and $P_{\text{FA}}^* = 1\%$, they remain qualitatively similar for lower P_{FA}^* values, too.

(A color version of this figure is available in the online journal.)

of the harmonic summing approach. First, we ignored the extra costs of summing the M power values. Second, we neglected the possible extra trials when one would maximize the test statistics over different M . Third, the analytical approximation of Equation (50) overestimates the true $T_{\text{coh},M}$ (and hence the sensitivity $p_{\text{coh},M}^{-1}$) as we show by numerical evaluation in Figure 11.

Hence, the basic moral is clear. For blind searches for isolated gamma-ray pulsars, whose sensitivity is limited by computing power rather than the amount of available data, a more sensitive search strategy is to employ a longer coherence time instead of using incoherent harmonic summing at the same computational cost.

4. SEMICOHERENT TEST STATISTICS

The key property of the semicoherent test statistics is that only pairs of photon arrival times (t_j, t_k) whose separation $\tau_{jk} = t_j - t_k$, also called lag, is at most T (which is much shorter than T_{obs}) are combined coherently, otherwise incoherently. Hence, we refer to T as the coherence window size and denote by R the ratio of total observational data time span T_{obs} of the semicoherent search and T ,

$$R = T_{\text{obs}}/T. \quad (52)$$

Compared to fully coherent methods, this semicoherent approach drastically reduces the computing cost since fewer search grid points are required (due to the lower parameter-space resolution as will be described in Section 4.2) at the expense of reduced search sensitivity. In Section 4.3, we argue that this trade-off is a profitable one, because at fixed given computing cost the overall search sensitivity of the semicoherent searches outperform fully coherent searches restricted to data spans shorter than T_{obs} by the computational constraints.

To derive a semicoherent test statistic, notice the (unnormalized) coherent Fourier power from Equation (10) for the fundamental frequency (first harmonic) can also be written in the following form:

$$\mathcal{P}_1 \propto \left| \sum_{j=1}^N w_j e^{-i\phi(t_j)} \right|^2 = \sum_{j,k=1}^N w_j w_k e^{-i[\phi(t_j)-\phi(t_k)]}. \quad (53)$$

Thus, the semicoherent statistic S_1 is formed by multiplying the terms in the above double sum with a real lag window $\hat{W}_T(\tau_{jk})$, such that

$$S_1 = \sum_{j,k=1}^N w_j w_k e^{-i[\phi(t_j)-\phi(t_k)]} \hat{W}_T(\tau_{jk}), \quad (54)$$

where the lag window has an effective size T ,

$$\int_{-\infty}^{\infty} \hat{W}_T(\tau) d\tau = T, \quad (55)$$

and thus must fall off rapidly outside the interval $[-T/2, T/2]$. Blackman & Tukey (1958) were the first to consider power spectral estimators of the form of S_1 , which can be seen as the Fourier transform of the lag-windowed covariance sequence (Stoica & Moses 2005). The semicoherent statistic S_1 is just a more general version of the classic Blackman–Tukey method (Blackman & Tukey 1958) in spectral analysis, e.g., if the phase model was simply $\phi(t_j) = 2\pi f t_j$ only. Hence, S_1 can also be seen as a local spectral average of \mathcal{P}_1 values over neighboring frequencies weighted according to the frequency response of \hat{W}_T (Stoica & Moses 2005).

As outlined in Pletsch et al. (2012a), for special forms of the lag window, S_1 can also be obtained by summing time-windowed coherent power from overlapping subsets of data. This implies a lag window that must be always positive semidefinite, because it is formed by the convolution of the time window with itself in this case (Stoica & Moses 2005), whereas the more general form as of Equation (54) in principle can have arbitrary lag windows.

In general, the choice of lag-window function $\hat{W}_T(\tau)$ has an impact on the sensitivity of the statistic S_1 . In tests with simulated LAT data, for the purpose of pulsation detection, we found that the best sensitivity is provided by the simple rectangular lag window,

$$\hat{W}_T^{\text{rect}}(\tau) = \begin{cases} 1, & |\tau| \leq T/2 \\ 0, & \text{otherwise} \end{cases}, \quad (56)$$

which also allows for an efficient implementation as will be described in more detail in Section 5. The usage of the rectangular lag window could also be motivated from the following viewpoint. Considering the significant sparseness of the LAT data, typically all pairs of photon times fall at different lags (for any practical sampling time, see Section 5.1). Therefore, one could argue that optimally (for minimum variance) all lags (i.e., all photon pairs) should be weighted equally when forming S_1 , which is exactly what $\hat{W}_T^{\text{rect}}(\tau)$ implements. Thus, in the remainder of this manuscript we will keep using the rectangular lag window $\hat{W}_T^{\text{rect}}(\tau)$ to calculate S_1 .

4.1. Statistical Properties

To examine the statistical properties of the semicoherent statistic, S_1 , it is useful to rewrite Equation (54) as

$$S_1 = \sum_{j=1}^N w_j^2 + 2 \sum_{j=1}^N \sum_{k=j+1}^N w_j w_k \cos[\phi(t_j) - \phi(t_k)] \hat{W}_T^{\text{rect}}(\tau_{jk}). \quad (57)$$

Under the null hypothesis, $p = 0$ and assuming $N \gg 1$, we show in Appendix F that S_1 follows a normal distribution, whose first two moments of the noise distribution of S_1 are

$$E_0[S_1] = \sum_{j=1}^N w_j^2, \quad (58)$$

$$\text{Var}_0[S_1] = 2 \sum_{j=1}^N \sum_{k=j+1}^N w_j^2 w_k^2 [\hat{W}_T^{\text{rect}}(\tau_{jk})]^2, \quad (59)$$

Now consider that the photon data contain a pulsed signal (i.e., $p > 0$) with a pulse profile defined by Fourier coefficients γ_n . Then, the expectation value of S_1 is obtained as

$$E_p[S_1] \approx E_0[S_1] + 2 E_p \left[\sum_{j=1}^N \sum_{k=j+1}^N w_j w_k \times \cos(\phi(t_j) - \phi(t_k)) \hat{W}_T^{\text{rect}}(\tau_{jk}) \right]. \quad (60)$$

Thus, for S_1 we can identify the amplitude S/N θ_{S_1} as

$$\begin{aligned} \theta_{S_1}^2 &= \frac{E_p[S_1] - E_0[S_1]}{\sqrt{\text{Var}_0[S_1]}} \\ &= \frac{\sqrt{2} E_p [\sum_{j=1}^N \sum_{k=j+1}^N w_j w_k \cos(\phi(t_j) - \phi(t_k)) \hat{W}_T^{\text{rect}}(\tau_{jk})]}{\sqrt{\sum_{j=1}^N \sum_{k=j+1}^N w_j^2 w_k^2 [\hat{W}_T^{\text{rect}}(\tau_{jk})]^2}}. \end{aligned} \quad (61)$$

To extract the scalings of the semicoherent S/N θ_{S_1} in terms of the relevant search parameters, we assume hard photon-selection cuts, i.e., binary photon weights, for the remainder of this section. Then Equation (57) reduces to

$$S_1 = N + 2 \sum_{j=1}^N \sum_{k=j+1}^N \cos[\phi(t_j) - \phi(t_k)] \hat{W}_T^{\text{rect}}(\tau_{jk}). \quad (62)$$

In this case, as derived in Appendix F, the first two moments of the noise distribution are

$$E_0[S_1] = N, \quad \text{Var}_0[S_1] \approx N^2 R^{-1}. \quad (63)$$

We show in Appendix F that for moderately strong signals, the first two moments of the distribution of S_1 are approximately given by

$$E_p[S_1] \approx N + p^2 N^2 |\gamma_1|^2 R^{-1}, \quad (64a)$$

$$\text{Var}_p[S_1] \approx N^2 R^{-1} (1 + 2p^2 N |\gamma_1|^2 R^{-1}), \quad (64b)$$

and the squared S/N of Equation (61) becomes

$$\theta_{S_1}^2 \approx p^2 N R^{-1/2} |\gamma_1|^2. \quad (65)$$

As shown in Appendix F, the probability density function of S_1 can be approximated by a normal distribution with the above expectation values and variances. The sensitivity of a semicoherent search is the lowest threshold pulsed fraction p for a given number of photons N and at given false alarm probability P_{FA}^* to achieve a certain detection probability P_{DET}^* . For a threshold $S_{1,\text{th}}$ the false alarm probability is computed as

$$P_{\text{FA}}(S_{1,\text{th}}) \approx \int_{S_{1,\text{th}}}^{\infty} \mathcal{N}\{S_1; E_0[S_1], \text{Var}_0[S_1]\} dS_1 \approx \frac{1}{2} \text{erfc}\left(\frac{S_{1,\text{th}} - E_0[S_1]}{\sqrt{2 \text{Var}_0[S_1]}}\right), \quad (66)$$

where, in this context, $\mathcal{N}\{X; \mu, \sigma^2\}$ denotes a normal distribution with mean μ and variance σ^2 , and should not be confused with the number of grid-points, $\mathcal{N}_{\text{coh},1}$. We compute the probability of detection using $\text{Var}_p[S_1] \approx \text{Var}_0[S_1](1 + 2p^2N|\gamma_1|^2R^{-1})$ as

$$P_{\text{DET}}(S_{1,\text{th}}, \theta_{S_1}^2) \approx \int_{S_{1,\text{th}}}^{\infty} \mathcal{N}\{S_1; E_p[S_1], \text{Var}_p[S_1]\} dS_1 \approx \frac{1}{2} \text{erfc}\left\{\left(\frac{S_{1,\text{th}} - E_0[S_1]}{\sqrt{\text{Var}_0[S_1]}} - \theta_{S_1}^2\right) \frac{1}{\sqrt{2 + 4p^2N|\gamma_1|^2R^{-1}}}\right\}. \quad (67)$$

The minimum detectable pulsed fraction is obtained by first inverting Equation (66) to get $S_{1,\text{th}}(P_{\text{FA}}^*)$, which in a second step is substituted in Equation (67) to obtain the threshold S/N $\theta_{S_1}^*$ as

$$\theta_{S_1}^* = \theta_{S_1}(P_{\text{FA}}^*, P_{\text{DET}}^*) \approx [\sqrt{2} \text{erfc}^{-1}(2P_{\text{FA}}^*) - \sqrt{2 + 4p^2N|\gamma_1|^2R^{-1}} \text{erfc}^{-1}(2P_{\text{DET}}^*)]^{1/2}. \quad (68)$$

Finally, using Equation (68), one can convert Equation (65) into the threshold pulsed fraction $p_{\text{scoh},1}^{-1}$, determining the semicoherent sensitivity as

$$p_{\text{scoh},1}^{-1} = \frac{\sqrt{N} R^{-1/4}}{\theta_{S_1}^*} |\gamma_1| = \frac{\sqrt{\mu T} R^{1/4}}{\theta_{S_1}^*} |\gamma_1|, \quad (69)$$

where we used $N = \mu T R$. This reveals the square-root scaling with the coherence window size T and the expected fourth-root scaling with R of the semicoherent sensitivity. Furthermore, using $R = T_{\text{obs}}/T$, we can rewrite the previous equation as

$$p_{\text{scoh},1}^{-1} = \frac{\sqrt{\mu} (T T_{\text{obs}})^{1/4}}{\theta_{S_1}^*} |\gamma_1|. \quad (70)$$

As a comparison, recall that the coherent sensitivity as of Equation (46), $p_{\text{coh},1}^{-1} \propto \sqrt{T_{\text{coh},1}}$ increases with the square root of the coherent integration time $T_{\text{coh},1}$. Here, Equation (70) shows that the semicoherent sensitivity, $p_{\text{scoh},1}^{-1} \propto \sqrt{(T T_{\text{obs}})^{1/2}}$, increases with the square root of the geometric mean of the coherence window size T and the total observation time T_{obs} .

It should be noted that while the semicoherent method allows for the use of short lag windows in order to detect pulsations, there is the additional requirement that there is at least one pair of pulsed photons that arrive within T of each other. This sets a fundamental lower limit on T . However, for typical pulsed fractions and photon arrival rates considered in this work, this lower limit is on the order of only a few hours.

4.2. Grid-point Counting for Semicoherent Search

To optimally construct the search grid for the semicoherent statistic S_1 , it is necessary to re-evaluate the appropriate metric on parameter space. Analog to Equation (32), we define the mismatch for S_1 as the fractional loss in semicoherent S/N squared,

$$\bar{m} = 1 - \frac{\theta_{S_1}^2(\mathbf{u})}{\theta_{S_1}^2(\mathbf{u}_{\text{sig}})} = 1 - \frac{\theta_{S_1}^2(\mathbf{u}_{\text{sig}} + \Delta\mathbf{u})}{\theta_{S_1}^2(\mathbf{u}_{\text{sig}})}. \quad (71)$$

Expanding the mismatch \bar{m} to second order in the offsets $\Delta\mathbf{u}$ as in Equation (33) yields the semicoherent metric tensor \bar{G} ,

$$\bar{m} = \sum_{k,\ell} \bar{G}_{k\ell} \Delta u^k \Delta u^\ell + \mathcal{O}(\Delta u^3). \quad (72)$$

We derive the components $\bar{G}_{k\ell}$ from the phase model in Appendix G analog to the methods described in Pletsch (2010). Following the same steps as in Section 3.3, we find that \bar{G} is also diagonal, and the total number of grid points for a semicoherent step can thus be written as

$$\mathcal{N}_{\text{scoh}} = \frac{1}{16} \mathcal{U} \bar{m}^{-2} \sqrt{\det \bar{G}}, \quad (73)$$

where \bar{m} here represents the maximum mismatch per dimension used for grid construction. As derived in Appendix G, the determinant of the semicoherent metric is

$$\sqrt{\det \bar{G}} = \frac{\pi^4}{12\sqrt{3}} T^3 f^2 r_E^2 R \left[1 - \text{sinc}^2\left(\frac{\Omega_E T}{2\pi}\right)\right]. \quad (74)$$

As in Section 3.3, for practical purposes we construct the grid for the highest frequency searched f_{max} in a given frequency band. Thus, we can rewrite Equation (73) as

$$\mathcal{N}_{\text{scoh}} = \bar{m}^{-2} \frac{\pi^4}{192\sqrt{3}} T^3 f_{\text{max}}^2 r_E^2 R \left[1 - \text{sinc}^2\left(\frac{\Omega_E T}{2\pi}\right)\right] \mathcal{U}, \quad (75)$$

where the proper search volume \mathcal{U} has been defined previously in Equation (35).

To extract the scaling of $\mathcal{N}_{\text{scoh}}$ with T , we use the following approximation:

$$\left[1 - \text{sinc}^2\left(\frac{\Omega_E T}{2\pi}\right)\right] \approx \begin{cases} \frac{\Omega_E^2 T^2}{12}, & T < 0.551\text{yr} \\ 1, & T \geq 0.551\text{yr}, \end{cases} \quad (76)$$

which is illustrated in Figure 5. Hence, using $R = T_{\text{obs}}/T$, one finds that the total number of grid points in the semicoherent search scales as

$$\mathcal{N}_{\text{scoh}} \propto \bar{m}^{-2} \left(\frac{\Omega_E^2}{12}\right)^{\frac{(s-3)}{2}} T^{s-1} T_{\text{obs}} f_{\text{max}}^2, \quad (77)$$

where the exponent s is given by

$$s \approx \begin{cases} 5, & T < 0.551\text{yr} \\ 3, & T \geq 0.551\text{yr} \end{cases}. \quad (78)$$




Figure 5. Scaling of the determinant of the semicoherent metric \bar{G} as function of the coherent window size T (black solid curve). The red dot–dashed curve shows the model for the semicoherent metric determinant from the approximation of Equation (76) used to estimate the computing cost scaling.

(A color version of this figure is available in the online journal.)

4.3. Semicoherent Search Sensitivity at Fixed Computing Cost

In analogy to Section 3.4, here we adopt a similar model for the computational cost of a semicoherent search, which is proportional to the number of search-grid points $\mathcal{N}_{\text{scoh}}$ needed. We again assume that the FFT algorithm is used to compute S_1 over $f_{\text{max}}T$ frequency bins and again split the total mismatch \bar{m}_{tot} into the mismatch due to a frequency offset \bar{m}_f and the mismatch due to offsets in the other parameters \bar{m} . Hence, using Equations (77) and (52), the semicoherent computing cost model $C_{\text{scoh},1}$ is obtained as

$$C_{\text{scoh},1} = K_{\text{scoh},s} \bar{m}^{-3/2} T^{s-1} T_{\text{obs}} \log_2(T f_{\text{max}}), \quad (79)$$

where $K_{\text{scoh},s}$ denotes a constant of proportionality that depends on s ,

$$K_{\text{scoh},s} = K_{\text{FFT}} \frac{\sqrt{2} \pi^3 r_E^2 f_{\text{max}}^2 \mathcal{U}}{96} \left(\frac{\Omega_E^2}{12} \right)^{(s-3)/2}, \quad (80)$$

as well as on the implementation and computing-hardware-dependent constant K_{FFT} as in Equation (45). Analog to Equation (44), here we also assume that the FFT algorithm is used, hence the \log_2 factor in Equation (79). In Section 4.1, we found the sensitivity of the semicoherent search as of Equation (70) can be approximately described by

$$p_{\text{scoh},1}^{-1} = \sqrt{(1 - \langle \bar{m}_{\text{tot}} \rangle)} \mu \frac{|\gamma_1|}{\theta_{S_1}^*} T^{1/4} T_{\text{obs}}^{1/4}, \quad (81)$$

where $\langle \bar{m}_{\text{tot}} \rangle = \langle \bar{m}_f \rangle + 3\xi \bar{m}$ denotes again the total average mismatch of the search grid.

With the sensitivity and computing-cost model at hand, we can now illustrate the increased efficiency that a semicoherent search offers over a fully coherent search. We compare the sensitivity $p_{\text{scoh},1}^{-1}$ of a semicoherent search with coherence window size T over a data set that in total spans the observational time interval T_{obs} to the sensitivity $p_{\text{coh},1}^{-1}$ of a fully coherent search with coherent integration time $T_{\text{coh},1}$, at the same computational cost: $C_{\text{scoh},1} = C_{\text{coh},1}$. For a given computing cost $C_{\text{scoh},1}$ and observational data set spanning T_{obs} , Equation (79) determines T . This value of T can then be used to obtain the sensitivity $p_{\text{scoh},1}^{-1}$ via Equation (81). Similarly, as described in Section 3.5, the given value of $C_{\text{coh},1}$ determines $T_{\text{coh},1}$ and thus provides the corresponding $p_{\text{coh},1}^{-1}$.




Figure 6. Comparison of a semicoherent and fully coherent search sensitivity at equal computing cost and given observational data time spans of $T_{\text{obs}} = 2$ yr (dotted–dashed curves), $T_{\text{obs}} = 5$ yr (solid curves), $T_{\text{obs}} = 10$ yr (dashed curves). The top panel shows the coherence window size T of the corresponding semicoherent search as a function of the sensitivity ratio. The bottom panel shows, for a coherent search, the integration time $T_{\text{coh},1}$, i.e., the subset of T_{obs} that could be fully coherently analyzed with the same computing cost as the semicoherent search with the corresponding T shown in the upper panel. The sensitivity is for $P_{\text{FA}}^* = 10^{-3}$ and $P_{\text{DET}}^* = 0.9$ in each case and a typical pulsed signal with $p = 0.1$ and $|\gamma_1|^2 = 0.35$ (see Figure 2). Since the sensitivity ratio is in all practically relevant cases much greater than unity, the semicoherent search approach more efficient.

(A color version of this figure is available in the online journal.)

The so-obtained ratio of sensitivities $p_{\text{scoh},1}^{-1}/p_{\text{coh},1}^{-1}$ is studied numerically in Figure 6 for realistic computational power available, such as *Einstein@Home*. In both cases, the optimal mismatch parameters are assumed, which are independent of computing cost (see Appendices E and H). As can be seen in the figure, this sensitivity ratio is always greater than unity and increases as T decreases, which is representative of the fact that the sensitivity of a semicoherent search decreases more slowly than that of a coherent search as the available computing power decreases. While this ratio decreases as T (and, therefore, the computing cost) increases, the absolute search sensitivity always increases with T , and so it is still beneficial to use the largest achievable lag-window size T at the available computational power.

Using a simplified approximation for the semicoherent computing cost model of Equation (79) allows us to obtain some analytical insight into the ratio $p_{\text{scoh},1}^{-1}/p_{\text{coh},1}^{-1}$ at fixed computing cost, similar to what has been done in Section 3.5. Ignoring the slowly varying \log_2 term gives the approximate semicoherent computing cost model as

$$C_{\text{scoh},1} \sim K_{\text{scoh},s} \bar{m}^{-3/2} T^{s-1} T_{\text{obs}}. \quad (82)$$

With this simplified model, $C_{\text{scoh},1} = C_{\text{coh},1}$ can be rewritten using the approximation of Equation (49) as

$$\frac{K_{\text{scoh},s} T_{\text{obs}} T^{s-1}}{\bar{m}^{3/2}} = \frac{K_{\text{coh},a} T_{\text{coh},1}^a}{m^{3/2}}. \quad (83)$$

Furthermore, using Equations (45) and (80) to replace $K_{\text{coh},a}$ and $K_{\text{scoh},s}$, we can rewrite Equation (83) as

$$T = \left(\frac{4 \Omega_E \bar{m}^{3/2} T_{\text{coh},1}^6}{5 \sqrt{6} m^{3/2} T_{\text{obs}}} \right)^{1/4}, \quad (84)$$

where we assume $a = 6$ and $s = 5$ since coherent integration times $T_{\text{coh},1}$ less than half a year will be practically feasible in the near future. This relation can then be used to substitute T in the ratio $p_{\text{scoh},1}^{-1}/p_{\text{coh},1}^{-1}$ using Equations (81) and (46), yielding

$$\frac{p_{\text{scoh},1}^{-1}}{p_{\text{coh},1}^{-1}} \approx 2 \frac{\theta_1^*}{\theta_{S_1}^*} \left(\frac{T_{\text{obs}}}{1\text{yr}} \right)^{1/16} \left(\frac{T_{\text{obs}}}{T_{\text{coh},1}} \right)^{1/8}, \quad (85)$$

where we again assumed the optimal mismatch choices for m and \bar{m} (see Appendices E and H) that are independent of computational cost. For $a = 6$ and $s = 5$ these are $m_{\text{opt}} = 0.172$ and $\bar{m}_{\text{opt}} = 0.146$. Hence, as *Fermi*-LAT data spans several years (implying $T_{\text{obs}} \gtrsim 1\text{yr}$) and typically $\theta_1^* \gtrsim \theta_{S_1}^*$, the sensitivity ratio of Equation (85) exceeds unity in all practically relevant cases. This clearly indicates that at fixed computational cost, a semicoherent blind search is always more sensitive than a fully coherent search over the same parameter space.

5. EFFICIENT IMPLEMENTATION OF A MULTISTAGE SEARCH SCHEME

In Section 3, we argued that under computational cost constraints, blind fully coherent searches without harmonic summing are more efficient, i.e., can typically achieve higher search sensitivity. In Section 4, we showed that at fixed computing cost semicoherent searches are more efficient than fully coherent searches to scan wide parameter space.

These considerations motivate a multistage search strategy, in which the first and by far most computationally expensive stage uses the most efficient method (i.e., a semicoherent search) to explore the entire parameter space. In subsequent stages, the most promising candidates are automatically “followed up” in further, more sensitive steps, ultimately using fully coherent search methods. Since the parameter space relevant for these candidates has been previously narrowed down by the first-stage search, the computing cost constraints are relaxed (i.e., the computing cost of the follow-ups is negligible compared to the overall cost of the first stage of the blind search). Hence, then the usage of fully coherent methods offering the highest sensitivity is made possible.

In this multistage search scheme, before statistically significant candidates from the first-stage semicoherent search are followed-up with fully coherent methods, it is advisable to refine the location of each semicoherent candidate by searching, again semicoherently, using a refined grid with a smaller mismatch. We then “zoom in” on each significant candidate by performing a fully coherent search of the local parameter space around the refined location of the semicoherent candidate, using the full observational data time span, T_{obs} . The search-grid construction of each stage is guided by the metric, as described in Appendices B, G, and I.

When searching for weak signals in the presence of noise, this can cause the refined semicoherent candidate to occur at a small but unknown offset from the true signal parameters. This offset depends on the candidate S/N; candidates with higher S/N have a smaller uncertainty region. In order not to miss weak signals, the coherent follow-up has to cover a conservative region in each dimension around the semicoherent candidate location. Since the parameter space that must be searched coherently has been greatly reduced, this step represents a very small fraction of the overall cost of the search. If the ratio of the coherence window size T used in the first stage and T_{obs} is very large, it is more efficient to insert another intermediate zooming stage that

does another semicoherent search with a coherence window size between T and T_{obs} . This would further reduce the parameter space to be searched in the fully coherent step, ensuring that the follow-up remains a negligible fraction compared to overall search. Finally, candidates from this coherent follow-up step are then ranked for further investigation (e.g., by taking into account higher harmonics, or a more complex phase model) according to their false alarm probability.

Since this multistage scheme is designed such that the largest computational burden is associated with the first stage, it is important to optimize this method of calculating the semicoherent test statistic S_1 as much as possible. In the following, we describe various complementary methods which improve the efficiency and sensitivity of a computationally limited semicoherent search.

5.1. Efficient Computation of Semicoherent Test Statistic

For each sky-position grid point of the search region, the barycentric corrections are applied directly to the LAT-registered arrival times t_{LAT} to obtain the corresponding photon arrival times t at the SSB. The semicoherent detection statistic S_1 as of Equation (54) is then computed over the f and \dot{f} ranges. However, directly computing S_1 from Equation (54) is computationally inefficient. Therefore, here we discuss more efficient ways of how to do this.

Making the dependence of S_1 on the search parameters f and \dot{f} explicit for clarity, we rewrite Equation (57) as

$$S_1(f, \dot{f}) = \sum_{j,k=1}^N w_j w_k e^{-i[\phi(t_j; f, \dot{f}) - \phi(t_k; f, \dot{f})]} \hat{W}_T^{\text{rect}}(\tau_{jk}), \quad (86)$$

where the phase differences in terms of f and \dot{f} are given by

$$\begin{aligned} \phi(t_j; f, \dot{f}) - \phi(t_k; f, \dot{f}) &= 2\pi f \tau_{jk} + \pi \dot{f} [(t_j - t_0)^2 - (t_k - t_0)^2] \\ &= 2\pi f \tau_{jk} + \pi \dot{f} [t_j^2 - t_k^2 - 2t_0 \tau_{jk}]. \end{aligned} \quad (87)$$

Thus, S_1 of Equation (86) takes the following form:

$$S_1(f, \dot{f}) = \sum_{j,k=1}^N w_j w_k e^{-\pi i \dot{f} [t_j^2 - t_k^2 - 2t_0 \tau_{jk}]} \hat{W}_T^{\text{rect}}(\tau_{jk}) e^{-2\pi i f \tau_{jk}}, \quad (88)$$

which allows us to utilize the efficiency of the FFT to scan along the f direction. In the following, we describe how to achieve this. First, we construct an equidistant lag series whose separation is the sampling interval $\delta_\tau = 1/(2f_{\text{max}})$, where f_{max} is equal to the Nyquist frequency f_{Ny} . Then, for each pair of times (t_j, t_k) having a lag τ_{jk} smaller than the lag window (i.e., for which $\hat{W}_T^{\text{rect}}(\tau_{jk}) = 1$), we determine the corresponding bin index b of the equidistant lag series via interpolation. While we study the efficiency of different lag-domain-interpolation schemes below, let us assume here nearest-neighbor interpolation for simplicity. Thus, we just round to the nearest lag-bin index b ,

$$b = \text{round} [\tau_{jk}/\delta_\tau]. \quad (89)$$

The FFT performance is generally best for input sizes that are a power of two (radix-2 FFTs). Therefore, we choose T and f_{max} , such that the total number of lag bins $B_T = T/\delta_\tau = 2Tf_{\text{max}}$

is a power of two. We denote the lag-interpolated version of S_1 from Equation (88) by \hat{S} , which can be written using the lag-bin index b as

$$\hat{S}(f, \dot{f}) = \sum_{b=-B_T/2}^{B_T/2} Y_b(\dot{f}) e^{-2\pi i f \delta_\tau b}, \quad (90)$$

where terms depending on \dot{f} and the photon weights have been absorbed into the complex numbers $Y_b(\dot{f})$. More precisely, each $Y_b(\dot{f})$ is the sum of pairwise weight and \dot{f} phase factors, falling into the same lag bin b ,

$$Y_b(\dot{f}) = \sum_{j=1}^N y_j(b; \dot{f}), \quad (91)$$

where

$$y_j(b, \dot{f}) = \begin{cases} w_j w_k e^{-\pi i \dot{f} [t_j^2 - t_k^2 - 2t_0 \tau_{jk}]}, & \text{round}[\tau_{jk}/\delta_\tau] = b, \\ 0, & \text{else.} \end{cases} \quad (92)$$

Note that the so-constructed lag series Y_b has Hermitian symmetry, i.e., $Y_b = Y_{-b}^*$, and therefore \hat{S} remains entirely real valued.

The above expression for \hat{S} in Equation (90) can be seen as a Fourier transform of the complex lag series Y_b , and so \hat{S} can be computed efficiently at many discrete frequencies by exploiting the FFT algorithm, i.e., by calculating

$$\hat{S}_g(\dot{f}) = \sum_{b=-B_T/2}^{B_T/2} Y_b(\dot{f}) e^{-2\pi i g b / B_T}, \quad (93)$$

where the frequency at the g th bin is $f = g/T$. There exist efficient FFT algorithms (Frigo & Johnson 2005) that can be used to evaluate this complex-to-real (c2r) transform of Equation (93) and that only require the positive lag portion of Y_b to be calculated as an input.

The above formulation of the semicoherent detection statistic, \hat{S}_g , is very similar to the D_ℓ statistic described in A06 as the DFT of the discrete autocorrelation function of the (binned) photon arrival times. However, there are some key differences. While further differences are discussed in the following subsections as we encounter them, here we note a first difference between the methods related to the correction of the frequency derivative \dot{f} . When calculating D_ℓ , the frequency derivative is corrected by constructing a new time series in which the photon arrival times are stretched out according to $t_j = \tilde{t}_j + 1/2(\dot{f}/f)\tilde{t}_j^2$. In order to search the $\{f, \dot{f}\}$ parameter space, the ratio \dot{f}/f is increased by small increments. According to this scheme, the search points in the $\{f, \dot{f}\}$ plane lie along straight lines with increasing gradient, intersecting at the origin. As a result, the search-grid-point density is highly non-uniform in the $\{f, \dot{f}\}$ plane, decreasing from low to high search frequencies. The result is that the search parameter space is highly oversampled in the \dot{f} dimension at low frequencies. This sub-optimal grid-point density implies that far more grid points are needed to cover the parameter space. Decreasing the lag-window size to account for this extra computational cost causes a reduction in sensitivity which more than accounts for the decrease in the average mismatch.⁹ Calculating \hat{S}_g in the manner described

⁹ This is because despite the reduced mismatch in the \dot{f} dimension, the contributions of the other three dimensions still remain and dominate the total mismatch that is relevant for the search sensitivity.

above, where the effect of the frequency derivative is accounted for by the complex lag-series, $Y_b(\dot{f})$, allows us to uniformly sample the $\{f, \dot{f}\}$ plane with the optimal average mismatch.

5.2. Frequency Domain Interpolation

When performing a semicoherent search using \hat{S}_g , computed via the FFT as in Equation (93), for a pulsar signal frequency that does not lie exactly at a Fourier frequency (i.e., not at an integer multiple of $1/T$) a loss in signal power (mismatch) will result. To evaluate the response of \hat{S}_g to signals at a non-Fourier frequency, we consider the case when the lag-series contains a pure sinusoid, with amplitude \hat{S}_0 , at a frequency h/T . Including an appropriate normalization factor of $1/B_T$ for the Fourier transform so that

$$Y_b(0) = \frac{\hat{S}_0}{B_T} e^{2\pi i h b / B_T}. \quad (94)$$

This represents the (unlikely) case of a strong signal, in the absence of noise, where the frequency derivative and sky location have been perfectly matched. Using Equation (93) the response at the g th frequency bin is therefore

$$\begin{aligned} \hat{S}_g &= \sum_{b=-B_T/2}^{B_T/2} Y_b(0) e^{-2\pi i g b / B_T} \\ &= \frac{\hat{S}_0}{B_T} \sum_{b=-B_T/2}^{B_T/2} e^{-2\pi i b (g-h) / B_T}. \end{aligned} \quad (95)$$

The above summation over b can be explicitly calculated and is also called the *Dirichlet* kernel, which is given by

$$\mathcal{D}_N(x) = \sum_{b=-N}^N e^{-i b x} = \frac{\sin((N+1/2)x)}{\sin(x/2)}. \quad (96)$$

Using this identity gives rise to rewrite Equation (95),

$$\begin{aligned} \hat{S}_g &= \frac{\hat{S}_0}{B_T} \mathcal{D}_{B_T/2-1}(2\pi(g-h)/B_T) \\ &= \frac{\hat{S}_0}{B_T} \frac{\sin(\pi(g-h)(1-1/B_T))}{\sin(\pi(g-h)/B_T)} \\ &\approx \frac{\hat{S}_0}{B_T} \frac{\sin(\pi(g-h))}{\sin(\pi(g-h)/B_T)} \\ &\approx \hat{S}_0 \text{sinc}(g-h), \end{aligned} \quad (97)$$

where in the approximation made in the third step we assumed that $1/2 \gg 1/B_T$, and in the fourth step we used in addition the following approximation $\sin(\pi(g-h)/B_T) \approx \pi(g-h)/B_T$, since typically for nearby frequency bins $B_T \gg (g-h)$. Therefore, the match is well described by a sinc function for signals at non-Fourier frequencies and is smallest (i.e., greatest mismatch) if the signal lies exactly halfway between two Fourier frequencies. This is shown in Figure 7, which displays the approximated response of Equation (97).

This loss can be reduced by interpolating the Fourier response halfway between two Fourier frequency bins. One method of interpolating the Fourier transform output, known as zero-padding, is to extend the original lag series (or time series) to twice its original length by adding zeros onto the end.




Figure 7. Illustration of frequency domain interpolation. The dashed blue curve shows the relative response (match) of \hat{S}_g at neighboring Fourier bins as a function of the signal frequency offset. The solid black curve represents the overall DFT response. The overlaid dotted–dashed red curve is the overall DFT response match obtained via the frequency domain interpolation as described in the text.

(A color version of this figure is available in the online journal.)

However, this requires calculating a Fourier transform which is twice as long, and therefore more than twice as costly. To avoid increasing the computational cost, we use a more efficient interpolation technique in the frequency domain, also known as “interbinning” (van der Klis 1989; Ransom et al. 2002). Note that Ransom et al. (2002) gives a formulation for calculating interbin amplitudes for real- or complex-to-complex Fourier transforms. However, in our case, where \hat{S}_g is entirely real valued, it is sufficient to calculate interbins by summing the amplitude of neighboring frequency bins,

$$\hat{S}_{g+1/2} = \frac{1}{\sqrt{2}}(\hat{S}_g + \hat{S}_{g+1}). \quad (98)$$

It is also important to emphasize that our chosen normalization differs from that used by van der Klis (1989) and Ransom et al. (2002), where the interbins are normalized to ensure that all of the signal power is recovered in an interbin if the signal lies exactly halfway between two Fourier bins. Instead, here we use a normalization factor of $1/\sqrt{2}$ ensuring that interbins have the same noise variance as the standard Fourier bins (as was first done by Astone et al. 2010). While the method used in Equation (98) results in a mismatch even for signals at the center of an interbin, ensuring that the noise variance is consistent between bins and interbins facilitates semicoherent candidate ranking for follow-up procedures.

The overall response for signals at non-Fourier frequencies before and after interbinning is shown in Figure 7. Using the interbinning method, the average mismatch due to a frequency offset is reduced from ~ 0.13 to ~ 0.075 , while the maximum mismatch is reduced from ~ 0.36 to ~ 0.14 . Thanks to their simplicity, interbins can be calculated very quickly, and so this performance gain comes at negligible extra computing cost (when compared to the dominant FFT computing cost).

5.3. Complex Heterodyning

Searching a wide range of frequencies (i.e., large f_{\max}) using the test statistic \hat{S} would require computing a single FFT of large size, B_T . The length of an FFT that can be computed is limited by the amount of memory accessible. In particular, extending the frequency search band to the millisecond pulsar regime (i.e., near 1 kHz frequencies) would require a large increase in the sampling rate and would potentially require decreasing the lag-window size (and hence the sensitivity of the search) to make the FFT short enough to fit into memory.

To address this problem, we divide the total frequency range into smaller bands of size Δf (that can be efficiently searched in parallel) using complex heterodyning, without sacrificing sensitivity. Using this method, the center frequency, f_H , of a given subband is shifted to DC, which in the lag domain corresponds to multiplying each lag bin by $e^{-2\pi i f_H \delta_\tau b}$. The heterodyned lag series is therefore defined as

$$Y'_b(\dot{f}, f_H) = Y_b(\dot{f}) e^{-2\pi i f_H \delta_\tau b}, \quad (99)$$

and the frequency at the g th bin becomes

$$f = g/T + f_H. \quad (100)$$

One can therefore compute $\hat{S}_g(\dot{f})$ over the subband $[f_H - \Delta f/2; f_H + \Delta f/2]$ via

$$\hat{S}_g(\dot{f}) = \sum_{b=-B_T/2}^{B_T/2} Y'_b(\dot{f}, f_H) e^{-2\pi i g b/B_T}, \quad (101)$$

in the same way as described in Equation (93), but using a sampling interval of only $\delta_\tau = 1/(\Delta f)$. Hence, we can search subbands in the millisecond-pulsar regime, while the FFT size remains at $B_T = T\Delta f$.

5.4. Lag Domain Interpolation

As outlined above, before the FFT can be performed, the lags τ_{jk} have to be binned into an equidistant lag series. Because the lags τ_{jk} will in general not coincide with the lag-bin centers, the nearest-neighbor interpolation of Equation (89) introduces an additional, frequency-dependent loss (mismatch) of signal power across the frequency band analyzed (e.g., van der Klis 1989; Ransom et al. 2002).

The process of binning in lag can be thought of as convolving the lag series with a binning function. By the convolution theorem, the resulting response across the frequency band is the Fourier transform of this convolving function. For \hat{S}_g as derived above, the binning function (for nearest-neighbor interpolation) is a simple rectangular function of width δ_τ , leading to the sinc response in the frequency domain. As a consequence, this results in an average loss (mismatch) in signal power of $\sim 13\%$ across the entire search band, illustrated in Figure 8.

Improved lag domain interpolation can reduce these losses. A given frequency response can be achieved by weighting the lag series bins around each τ_{jk} with an appropriate interpolation function. Ideal (i.e., lossless) interpolation would lead to a frequency response that is a rectangular function: unity within the search band to remove all bias in the spectrum, and zero outside to prevent any noise from being aliased into the band. Therefore, this ideal case of a rectangular frequency response requires a lag interpolation function that is the sinc

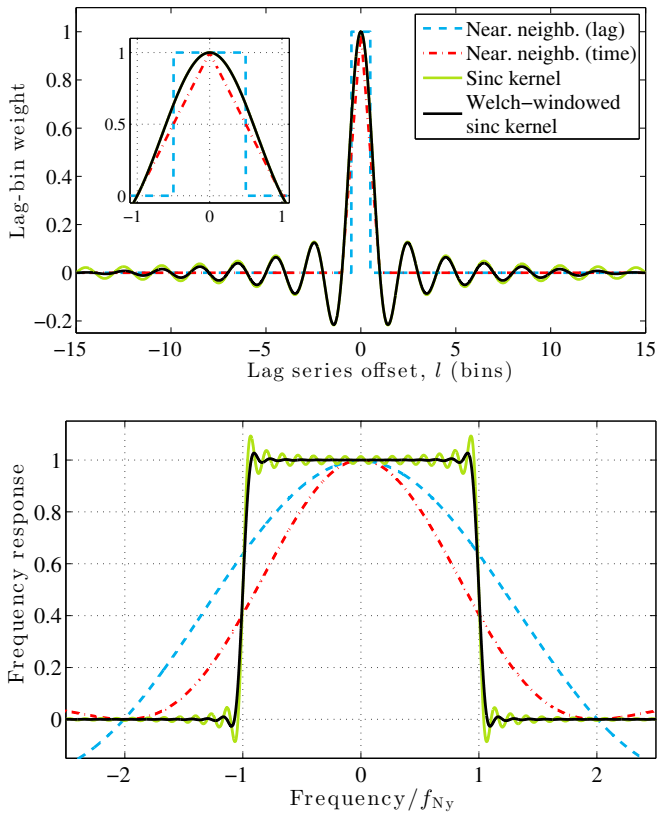


Figure 8. Upper panel: comparison of different lag domain interpolation functions, with the interpolating bin weights calculated over the range of the nearest 15 lag bins on either side of the center. For clarity, the inset shows a zoom of the central region. Lower panel: overall frequency response of each interpolation function from the upper panel, which \hat{S} is multiplied with in the frequency domain. The ideal response would be unity within the search band, i.e., for $-f_{Ny} < (f - f_H) < f_{Ny}$, and zero otherwise. For this specific choice of using the nearest 15 bins for the interpolation, the average mismatch (loss in signal power) across the search band from is $\sim 23\%$ for the rectangular binning function in time, $\sim 13\%$ for the rectangular binning function in lag domain, but only $\sim 1\%$ for the sinc kernel and also for the Welch-windowed sinc kernel that shows reduced Gibbs oscillations.

(A color version of this figure is available in the online journal.)

function. However, this interpolation function has infinite extent in the lag domain and is therefore impossible to realize in practice.

A practical solution is to truncate the sinc function in the lag domain around each τ_{jk} , such that the computational cost of this interpolation remains a negligible fraction of the overall computation time. In fact, one can show that using lag domain interpolation with the sinc function truncated to only the d nearest lag bins for each τ_{jk} is the best d th order approximation in the least squares sense to the ideal (rectangular) response function (e.g., Percival & Walden 1993). As a result, the average loss (mismatch) across the frequency search band is drastically reduced. In the example shown in Figure 8, with a truncated sinc kernel using the $d = 15$ nearest lag bins are on either side reduces this average mismatch to only $\sim 1\%$, as compared to the nearest-neighbor interpolation. Generally, it is often practical to use even more neighboring bins without significantly affecting the computational cost, but reducing the average mismatch even further.

However, as can also be seen in Figure 8, an inconvenient property of the truncated sinc kernel is the Gibbs oscillation throughout the frequency band. These oscillations mean that

the false alarm probabilities of candidates can vary significantly across the frequency band, making it difficult to rank candidate pulsars for follow-up. This problem can be mitigated by multiplying the sinc kernel by another windowing function (Lyons 2004, p. 176). This windowing function is required to be simple (and therefore efficient) to compute, and must still have a reasonably sharp fall-off in frequency near the edges of the bands. We find that the Welch window (an inverted parabola) provides a useful compromise between these requirements. The interpolated lag series, \tilde{Y}_b , is constructed by spreading the original lag-series Y'_b among the first d bins on either side of the nearest bin to a single photon pair with lag τ_{jk} ,

$$\tilde{Y}_{b+l}(\dot{f}, f_H) = Y'_b(\dot{f}, f_H) \text{sinc} \left(b+l - \frac{\tau_{jk}}{\delta_\tau} \right) \times \left[1 - \left(b+l - \frac{\tau_{jk}}{\delta_\tau} \right)^2 \frac{1}{d^2} \right], \quad (102)$$

for $l = 0, \pm 1, \dots, \pm d$. The frequency response of the Welch-windowed sinc kernel is displayed Figure 8. While the average mismatch with the Welch-windowed sinc kernel is comparable to the truncated sinc kernel, the reduced Gibbs oscillation means that the false alarm probabilities of candidates are much more consistent across the frequency band, allowing candidate pulsars to be more easily ranked, albeit with almost no increase in the cost of interpolating the lag-series. Fortunately, the interpolation functions can be efficiently computed using trigonometric look-up tables and recurrence relations. When this efficiency is combined with the typical sparseness of the lag series, the interpolation step remains a negligible fraction of the overall computation time.

Within this framework of lag domain interpolation, another key difference to the A06 method is worth pointing out. In A06, the SSB photon arrival times t_j are binned directly prior to calculating the lags τ_{jk} and the DFT (the D_ℓ in their notation). This implies a rectangular window function in time, which then is convolved with itself leading to a triangular window shape in the lag domain. Hence, the resulting frequency response is effectively the sinc function squared (also shown in Figure 8). This causes significant loss in signal power, especially at the edges of the frequency band, and amounts to a loss of $\sim 23\%$ averaged across the entire frequency band. For comparison, by using the lag domain interpolation technique with the Welch-windowed sinc kernel as presented above, this average loss can be reduced by more than an order of magnitude, from $\sim 23\%$ to $\sim 1\%$, at about the same computational expense.

6. PERFORMANCE DEMONSTRATION

In order to validate the expected sensitivity gain from the improved methods presented in this paper, we perform extensive Monte Carlo simulations. The false alarm probabilities are obtained using simulated data sets with different realizations of 8000 photon arrival times (with unit weights), spanning a realistic observation time of $T_{\text{obs}} = 5$ yr. To find the detection probabilities (for a given false alarm probability), simulated pulsar signals are added, which have the same pulse profile of Gaussian shape whose Fourier coefficient at the fundamental frequency is $|\gamma_1| = 0.82$, and varying pulsed fractions p .

While for computational reasons, the actual parameter space searched in each simulation was chosen smaller than in a real




Figure 9. Comparison of search efficiency of different search methods at fixed computational cost. Shown is the detection probability P_{DET} at increasing pulsed fractions p for the simulated gamma-ray pulsar signals at $P_{\text{FA}} = 10^{-3}$. The solid curves represent fits to each set of data points. Green stars: estimated sensitivity using the A06 method for the same computing cost. Black circles: semicoherent search method, using only nearest-neighbor lag-domain interpolation. Blue crosses: semicoherent search method, using interbinning frequency-domain interpolation, and lag-domain interpolation with a Welch-windowed sinc kernel. Red squares: multistage search method (semicoherent search method using lag-domain and interbinning frequency-domain interpolation, plus fully coherent follow-up). In all cases, the number of simulations was chosen large enough so that the uncertainties of the data points become smaller than the size of the data markers.

(A color version of this figure is available in the online journal.)

search, the main conclusions from these results are unaffected by this. In each simulation, the search covered a frequency bandwidth of 1 Hz and a frequency derivative range of 10^{-13} Hz s $^{-1}$. Each simulation searched the nearest nine sky positions around the signal location, at a uniformly random location on the sky. In the semicoherent search stage, we used a coherence window size of $T = 2^{20}$ s ≈ 12 d.

For further comparison, we also apply the A06 method to the simulated data sets. However, here we obtain a generous sensitivity estimation. This is because the non-uniform sampling of the $\{f, \dot{f}\}$ parameter space (discussed in more detail in Section 5.1) was not accounted for. While this is justifiable for a search for isolated millisecond pulsars, at lower frequencies and larger frequency derivatives (i.e., where most young pulsars are found) this non-optimal sampling requires reducing the lag-window size (and therefore reducing the sensitivity) to achieve the same computational cost.

The results from all simulations are summarized in Figure 9, which shows the detection probability as a function of pulsed fraction for each of the search methods discussed in this paper. From best-fit curves (of typical sigmoid shape) shown in Figure 9, we compare the pulsed fraction required to give a detection probability of 95% at a false alarm probability of 0.1%. We find that this pulsed fraction is around 48% lower for the full multistage method presented here than for the A06 method with approximately the same computational cost. This sensitivity increase is due to several improvements described in previous sections, in particular, use of the parameter space metric to allow optimally spaced grid points; lag- and frequency-domain interpolation to reduce mismatch; and an automated coherent follow-up step to increase sensitivity to weak gamma-ray pulsar signals.

7. CONCLUSIONS

We have presented optimized strategies to improve the efficiency of blind searches for isolated gamma-ray pulsars, whose search sensitivity is computationally limited. Under these conditions, our results confirm that fully coherent searches are generally less efficient than semicoherent searches, as well as that harmonic summing is typically less efficient than searching only for the strongest individual harmonic. We also derived the parameters for most efficient search grids. As motivated by these results, we presented and studied the implementation of a multistage search strategy. We have also presented efficient computation and interpolation techniques for the semicoherent test statistic, offering further important sensitivity gains. Finally, we have conducted realistic simulations which demonstrate the improved performance from our combined advances, providing in a substantial increase in sensitivity (i.e., lowering the minimum detectable pulsed fraction by almost 50%) over previous methods at the same computational cost.

The methods presented here are being implemented with the *Einstein@Home* volunteer computing project to increase the chances of detecting new gamma-ray pulsars among the unidentified LAT sources. While here we have focused on searches for isolated pulsars, the methods also apply to searches for pulsars in binaries, where partial knowledge of the orbit is available from observations at other wavelengths (Pletsch et al. 2012b).

Furthermore, the framework derived in this work in order to obtain an improved understanding of the pulsation search sensitivities underlying the different methods should also be useful for population studies. Specifically, these estimates can facilitate identifying the selection biases in the known gamma-ray pulsar sample, for example, due to the difference in pulse profile shape. In future work, we shall also explore using this framework to improve the efficiency of harmonic summing employing one or more realistic pulse profile templates built from the existing population of known gamma-ray pulsars.

This work was supported by the Max-Planck-Gesellschaft (MPG), as well as by the Deutsche Forschungsgemeinschaft (DFG) through an Emmy Noether research grant PL 710/1-1 (PI: Holger J. Pletsch). We also thank the anonymous referee for suggestions that helped improve the manuscript.

APPENDIX A

DERIVATION OF STATISTICAL PROPERTIES OF COHERENT TEST STATISTIC

From Equation (12) in Section 3.1, the coherent power \mathcal{P}_n can be rewritten as $\mathcal{P}_n = c_n^2 + s_n^2$, where

$$c_n = \sqrt{\frac{2}{N}} \sum_{j=1}^N \cos[n\phi(t_j)], \quad (\text{A1})$$

$$s_n = \sqrt{\frac{2}{N}} \sum_{j=1}^N \sin[n\phi(t_j)]. \quad (\text{A2})$$

Under the null hypothesis $p = 0$, the phases $\phi(t_j)$ are uniformly distributed on $[0, 2\pi]$, and it is straightforward to show that

$$E_0[\cos(n\phi(t_j))] = E_0[\sin(n\phi(t_j))] = 0, \quad (\text{A3a})$$

$$\text{Var}_0[\cos(n\phi(t_j))] = \text{Var}_0[\sin(n\phi(t_j))] = 1/2. \quad (\text{A3b})$$

Since we typically have $N \gg 1$, by appealing to the Central Limit Theorem, the random variables c_n and s_n are normally distributed with zero mean and unit variance,

$$E_0[c_n] = E_0[s_n] = 0, \quad (\text{A4a})$$

$$\text{Var}_0[c_n] = \text{Var}_0[s_n] = 1. \quad (\text{A4b})$$

Hence, \mathcal{P}_n follows a central χ^2 distribution with two degrees of freedom (e.g., Blackman & Tukey 1958). Therefore, the first two moments are $E_0[\mathcal{P}_n] = 2$ and $\text{Var}_0[\mathcal{P}_n] = 4$, as given in Equation (13).

Suppose a pulsed signal is present, $p > 0$, with a pulse profile having the complex Fourier coefficients γ_n as defined by Equation (4). While in this case for the $(1-p)N$ “non-pulsed” photons (i.e., background), Equations (A3) still hold; however, for the pN “pulsed” photons (i.e., not background), one obtains

$$E_p[\cos(n\phi(t_j))] = \Re(\gamma_n), \quad (\text{A5a})$$

$$E_p[\sin(n\phi(t_j))] = -\Im(\gamma_n), \quad (\text{A5b})$$

$$\text{Var}_p[\cos(n\phi(t_j))] = \frac{1}{2} + \frac{\Re(\gamma_{2n})}{2} - \Re(\gamma_n)^2, \quad (\text{A5c})$$

$$\text{Var}_p[\sin(n\phi(t_j))] = \frac{1}{2} - \frac{\Re(\gamma_{2n})}{2} - \Im(\gamma_n)^2. \quad (\text{A5d})$$

Therefore, the random variables c_n and s_n are normally distributed (since $N \gg 1$) with the following mean values and variances:

$$E_p[c_n] = p\sqrt{2N}\Re(\gamma_n), \quad (\text{A6a})$$

$$E_p[s_n] = -p\sqrt{2N}\Im(\gamma_n), \quad (\text{A6b})$$

$$\text{Var}_p[c_n] = 1 + p\Re(\gamma_{2n}) + 2p\Re(\gamma_n)^2, \quad (\text{A6c})$$

$$\text{Var}_p[s_n] = 1 - p\Re(\gamma_{2n}) - 2p\Im(\gamma_n)^2. \quad (\text{A6d})$$

For weak signals (i.e., small pulsed fractions) and typical gamma-ray pulse profiles (see Figure 2), we can approximate these variances as

$$\text{Var}_p[c_n] \approx \text{Var}_p[s_n] \approx 1. \quad (\text{A7})$$

With this approximation, the distribution of \mathcal{P}_n follows a noncentral χ^2 distribution (Groth 1975; Guidorzi 2011) with two degrees of freedom, whose first two moments are

$$E_p[\mathcal{P}_n] \approx 2 + 2p^2N|\gamma_n|^2, \quad (\text{A8a})$$

$$\text{Var}_p[\mathcal{P}_n] \approx 4 + 8p^2N|\gamma_n|^2, \quad (\text{A8b})$$

recovering Equations (14a) and (14b). The noncentrality parameter of that distribution is the second summand in Equation (A8a), $2p^2N|\gamma_n|^2$.

APPENDIX B

COHERENT METRIC

For the purpose of efficient search-grid construction, we exploit a simplified phase model that captures the most dominant effects. It is to be emphasized that we do *not* use this phase model in the actual search when computing the phases at the photon arrival times. Thus, here we assume that the LAT data set spans at least one year, such that the Doppler modulation is dominated by the Earth motion around the SSB.

For very short coherent integration times, the orbital motion of the *Fermi* satellite around the Earth could also introduce further Doppler effects. Comparing this effect to the much larger effect of the Earth’s orbital motion around the sun, which is responsible for the behavior of the metric visible in, e.g., Figure 3, it is clear that this effect would saturate after a small number of orbits. Hence, for coherent integration times of more than a few hours, here it is safe to neglect the rapidly oscillating components of the motion of the *Fermi* satellite around the Earth. Doing so yields the following phase model:

$$\begin{aligned} \phi(t, \mathbf{u}) &= 2\pi f(t - t_0) + \pi \dot{f}(t - t_0)^2 + 2\pi f \frac{\mathbf{n} \cdot \mathbf{r}_E(t)}{c} \\ &= 2\pi f(t - t_0) + \pi \dot{f}(t - t_0)^2 \\ &\quad + 2\pi f r_E [n_x \cos(\Omega_E t) + n_y \sin(\Omega_E t)], \end{aligned} \quad (\text{B1})$$

where n_x and n_y are the components of \mathbf{n} , the unit vector pointing from the SSB to the sky location (α, δ) , projected into the ecliptic plane (using the obliquity of the ecliptic, ϵ),

$$n_x = \cos(\alpha) \cos(\delta), \quad (\text{B2})$$

$$n_y = \cos(\epsilon) \sin(\alpha) \cos(\delta) + \sin(\epsilon) \sin(\delta), \quad (\text{B3})$$

and $\Omega_E = 2\pi/1\text{yr}$, and $r_E = 1\text{AU}/c \sim 500\text{ s}$.

In the presence of a small offset $\Delta\mathbf{u}$ from a signal’s location in parameter space \mathbf{u}_{sig} , we can write the mismatch, $m^{[t_j]}$, in the coherent power in a window of length T , centered on the j th photon as

$$m^{[t_j]} = 1 - \frac{(\theta_{p_1}^2(\mathbf{u}_{\text{sig}} + \Delta\mathbf{u}))^{[t_j]}}{\theta_{p_1}^2(\mathbf{u}_{\text{sig}})} \quad (\text{B4})$$

$$= 1 - |\langle e^{-i\phi(t, \Delta\mathbf{u})} \rangle^{[t_j]}|^2, \quad (\text{B5})$$

where we have replaced the discrete sum of Equation (10) for simplicity by a continuous integral over the coherent integration time T , i.e.,

$$\langle x \rangle^{[t_j]} \equiv \frac{1}{T} \int_{t_j - T/2}^{t_j + T/2} x(t) dt. \quad (\text{B6})$$

Following the derivation in Pletsch (2010), the mismatch can be Taylor expanded up to second order in terms of the parameter offsets, $\Delta\mathbf{u}^k$ to give

$$m^{[t_j]} = \sum_{k, \ell} G_{k\ell}^{[t_j]} \Delta\mathbf{u}^k \Delta\mathbf{u}^\ell + \mathcal{O}(\Delta\mathbf{u}^3). \quad (\text{B7})$$

The coherent metric components are defined as

$$G_{k\ell}^{[t_j]} = \langle \partial_k \phi \partial_\ell \phi \rangle^{[t_j]} - \langle \partial_k \phi \rangle^{[t_j]} \langle \partial_\ell \phi \rangle^{[t_j]}, \quad (\text{B8})$$




Figure 10. Comparison of mismatch in \mathcal{P}_1 (dashed curves) with coherent metric prediction (solid curves). In each panel, the horizontal axis shows the offset from the signal parameters in f (left), \dot{f} (middle), and sky position (right). The sky-location offset is $\sqrt{\Delta n_x^2 + \Delta n_y^2}$, which measures the offset in coordinates (n_x, n_y) in the ecliptic plane. The underlying pulsar signal has been simulated with spin parameters $f = 32$ Hz, $\dot{f} = -10^{-12}$ Hz s $^{-1}$ for a total coherent observation time of $T_{\text{coh}} = 3.4$ yr.

(A color version of this figure is available in the online journal.)

where $\partial_k \phi$ is the partial derivative of the phase at the signal location with respect to the k th component of the parameter offset:

$$\partial_k \phi \equiv \left. \frac{\partial \phi(t; \mathbf{u}_{\text{sig}} + \Delta \mathbf{u})}{\partial (\Delta u^k)} \right|_{\Delta \mathbf{u} = \mathbf{0}}. \quad (\text{B9})$$

Using the simplified phase model of Equation (B1), the metric components for a coherent window, centered on t_j are given by

$$G_{ff}^{[t_j]} = \frac{\pi^2 T^2}{3}, \quad (\text{B10a})$$

$$G_{\dot{f}\dot{f}}^{[t_j]} = \frac{\pi^2 T^4}{180} + \frac{\pi^2 (t_j - t_0)^2 T^2}{3}, \quad (\text{B10b})$$

$$G_{n_x n_x}^{[t_j]} = 2\pi^2 f^2 r_E^2 [1 + \text{sinc}(\Omega_E T / \pi) \cos(2\Omega_E t_j) - 2 \text{sinc}^2(\Omega_E T / 2\pi) \cos^2(\Omega_E t_j)], \quad (\text{B10c})$$

$$G_{n_y n_y}^{[t_j]} = 2\pi^2 f^2 r_E^2 [1 - \text{sinc}(\Omega_E T / \pi) \cos(2\Omega_E t_j) - 2 \text{sinc}^2(\Omega_E T / 2\pi) \sin^2(\Omega_E t_j)]. \quad (\text{B10d})$$

For the specific case of the general expressions above, where $t_j = t_0 = 0$, the metric components for the coherent detection statistic simplify to the following form:

$$G_{ff} = \frac{\pi^2 T^2}{3}, \quad (\text{B11a})$$

$$G_{\dot{f}\dot{f}} = \frac{\pi^2 T^4}{180}, \quad (\text{B11b})$$

$$G_{n_x n_x} = 2\pi^2 f^2 r_E^2 [1 + \text{sinc}(\Omega_E T / \pi) - 2 \text{sinc}^2(\Omega_E T / 2\pi)], \quad (\text{B11c})$$

$$G_{n_y n_y} = 2\pi^2 f^2 r_E^2 [1 - \text{sinc}(\Omega_E T / \pi)]. \quad (\text{B11d})$$

The mismatches predicted by these derived metric components are compared to the measured mismatches in \mathcal{P}_1 for a simulated pulsar signal in Figure 10.

Therefore, the determinant of the coherent metric is found as

$$\begin{aligned} \sqrt{\det G} &= \frac{\pi^4}{\sqrt{135}} T^3 f^2 r_E^2 \\ &\times [1 + \text{sinc}(\Omega_E T_{\text{coh},1} / \pi) - 2 \text{sinc}^2(\Omega_E T_{\text{coh},1} / 2\pi)] \\ &\times [1 - \text{sinc}(\Omega_E T_{\text{coh},1} / \pi)]. \end{aligned} \quad (\text{B12})$$

APPENDIX C

COHERENT METRIC WITH INCOHERENT HARMONIC SUMMING

If a search is performed using the Z_M^2 statistic, i.e., incoherently summing the coherent power \mathcal{P}_n in the first M harmonics, the mismatch, \tilde{m} , becomes

$$\begin{aligned} \tilde{m} &= 1 - \frac{\sum_{n=1}^M \theta_{\mathcal{P}_n}^2(\mathbf{u}_{\text{sig}} + \Delta \mathbf{u})}{\sum_{n=1}^M \theta_{\mathcal{P}_n}^2(\mathbf{u}_{\text{sig}})} \\ &= 1 - \frac{\sum_{n=1}^M |\gamma_n|^2 |e^{-in\phi(t, \Delta \mathbf{u})}|^{[t_j]}|^2}{\sum_{n=1}^M |\gamma_n|^2}. \end{aligned} \quad (\text{C1})$$

Taylor expanding this mismatch to the second order gives the metric components

$$\tilde{m} = \sum_{k, \ell} \tilde{G}_{k\ell}^{[t_j]} \Delta u^k \Delta u^\ell + \mathcal{O}(\Delta \mathbf{u}^3), \quad (\text{C2})$$

which can be expressed using Equation (B8) as

$$\tilde{G}_{k\ell}^{[t_j]} = r^2 G_{k\ell}^{[t_j]}, \quad (\text{C3})$$

where we defined the harmonic refinement factor r from

$$r^2 = \frac{\sum_{n=1}^M |\gamma_n|^2 n^2}{\sum_{n=1}^M |\gamma_n|^2}. \quad (\text{C4})$$

Thus, Equation (C3) indicates that the parameter space must be sampled r times more finely in each dimension when summing the power from M harmonics:

$$\sqrt{\det \tilde{G}} = r^4 \sqrt{\det G}. \quad (\text{C5})$$

The value of this refinement factor r also depends on the signal pulse profile γ_n , which of course is unknown in advance. However, we can consider the two limiting cases. First, for the narrowest possible pulse profile, a Delta function, all coefficients are equal, $|\gamma_n| = 1$, such that

$$r^2 = \frac{1}{M} \sum_{n=1}^M n^2 = \frac{M^2}{3} + \frac{M}{2} + \frac{1}{6}. \quad (\text{C6})$$

Therefore, for $M > 1$, the parameter space must be sampled more finely in each dimension by a factor of approximately $M^2/3$ (to leading order). On the other limiting case, for a sinusoidal pulse profile, where $|\gamma_{n>1}| = 0$, $r = 1$, and thus $\hat{G}_{k\ell}^{[t_j]} = G_{k\ell}^{[t_j]}$, requiring no refinement. Therefore, the range of the harmonic-summing refinement factor is approximately limited to $r \in [1, M]$.

Finally, we would like to point out a further generalization. Suppose a search is performed using the Q_M statistic and a template pulse profile α_n , which is not equal to the Dirac delta function (in this case, Q_M would reduce again to Z_M^2). Then by a straightforward repetition of arguments from the beginning of this section, one obtains the resulting metric tensor $\hat{G}_{k\ell}^{[t_j]}$ for the Q_M test statistic as

$$\hat{G}_{k\ell}^{[t_j]} = \hat{r}^2 G_{k\ell}^{[t_j]}, \quad (\text{C7})$$

where the harmonic refinement factor \hat{r} in this case would be different from Equation (C4), namely,

$$\hat{r}^2 = \frac{\sum_{n=1}^M |\gamma_n|^4 n^2}{\sum_{n=1}^M |\gamma_n|^4}. \quad (\text{C8})$$

APPENDIX D

APPROXIMATE HARMONIC-SUMMING COMPUTING COST

In Section 3.5, we describe an analytical approximation for the computing cost model of incoherent harmonic summing. This approximation is based on ignoring the slowly varying \log_2 factors in Equations (44) and (47). If one then equates $C_{\text{coh},1} = C_{\text{coh},M}$, it follows that $T_{\text{coh},M}$ must be shorter by the factor $(M^2 r^2)^{1/a}$, as given in Equation (50). Here, we study the accuracy of the analytical approximation in terms of the search sensitivity $p_{\text{coh},M}^{-1} \propto \sqrt{T_{\text{coh},M}}$, in comparison to the exact value for $T_{\text{coh},M}$ obtained from numerical evaluation. For a given value of $T_{\text{coh},1}$, we find numerically the exact value of $T_{\text{coh},M}$ such that $C_{\text{coh},1} = C_{\text{coh},M}$. Here, we assume a wide search frequency range, $f_{\text{max}} = 1000$ Hz. The results are displayed in Figure 11, showing that the approximation is accurate to within less than 1% for typical search setups. As can also be seen for the realistic case of $a = 6$, the approximation is generous in favor of the harmonic summing approach; because $T_{\text{coh},M}^{\text{approx}} \gtrsim T_{\text{coh},M}^{\text{exact}}$, the approximation overestimates the true search sensitivity.

APPENDIX E

OPTIMAL MISMATCH IN COHERENT SEARCH

In this section, we use the method of Lagrange multipliers as in Prix & Shaltev (2012) to obtain the optimal average mismatch for a fully coherent search. We use the scalings of the sensitivity $p_{\text{coh},M}^{-1}$ and computing cost $C_{\text{coh},M}$, ignoring the \log_2 FFT scaling factor, from Equations (28) and (49), respectively. In order to find the optimal mismatch at a fixed computing cost C_0 , we search for stationary points of the Lagrange function:

$$\begin{aligned} L(T_{\text{coh},M}, m, M, \lambda) &= p_{\text{coh},M}^{-1} - \lambda(C_{\text{coh},M} - C_0) \\ &= (1 - \langle m_{\text{tot}} \rangle)^{1/2} T_{\text{coh},M}^{1/2} h^*(M) \\ &\quad + \lambda(K'_{\text{coh},a} m^{-3/2} T_{\text{coh},M}^a M^2 r^2(M) - C_0), \end{aligned} \quad (\text{E1})$$




Figure 11. Comparison of the analytical approximation for the harmonic-summing computing cost model (leading to $T_{\text{coh},M}^{\text{approx}}$) to the results obtained from fully numerical evaluation (leading to $T_{\text{coh},M}^{\text{exact}}$), as a function of $T_{\text{coh},1}$ corresponding to the same computing cost $C_{\text{coh},1} = C_{\text{coh},M}$. Since we are interested in the impact on search sensitivity $p_{\text{coh},M}^{-1} \propto \sqrt{T_{\text{coh},M}}$, the vertical axis shows the square root of the ratio. As indicated by the legend the different curves are for different values of scaling exponent a of Equation (41) and number of harmonics summed M .

(A color version of this figure is available in the online journal.)

where λ is a Lagrange multiplier, and we defined $K'_{\text{coh},a} = K_{\text{coh},a} f_{\text{max}}^2$, as well as the function $h^*(M)$ as

$$h^*(M) = \frac{1}{M^{1/4} \theta_M^*} \left[\sum_{n=1}^M |\gamma_n|^2 \right]^{1/2}, \quad (\text{E2})$$

using $*$ to indicate the implicit dependence on P_{FA}^* and P_{DET}^* through θ_M^* . Taking partial derivatives with respect to $T_{\text{coh},M}$, m and M , respectively, yields

$$\frac{\partial L}{\partial T_{\text{coh},M}} = \frac{1}{2} (1 - \langle m_{\text{tot}} \rangle)^{1/2} T_{\text{coh}}^{-1/2} h^*(M) + \frac{a\lambda C_{\text{coh},M}}{T_{\text{coh},M}} = 0, \quad (\text{E3})$$

$$\frac{\partial L}{\partial m} = \frac{1}{2} (1 - \langle m_{\text{tot}} \rangle)^{-1/2} 3\xi T_{\text{coh}}^{1/2} h^*(M) + \frac{3\lambda C_{\text{coh},M}}{2m} = 0, \quad (\text{E4})$$

$$\begin{aligned} \frac{\partial L}{\partial M} &= (1 - \langle m_{\text{tot}} \rangle)^{1/2} T_{\text{coh},M}^{1/2} \frac{\partial h^*(M)}{\partial M} \\ &\quad + \lambda C_{\text{coh},M} \left(\frac{2}{M} + \frac{2}{r(M)} \frac{\partial r}{\partial M} \right) = 0. \end{aligned} \quad (\text{E5})$$

Equating these and rearranging for ξm , we find that the optimal average mismatch for a fully coherent search is

$$3\xi m_{\text{opt}} = \frac{1 - \langle m_f \rangle}{\frac{2a}{3} + 1}. \quad (\text{E6})$$

As we argue in Section 3.4, practical fully coherent searches are computationally limited to integration times $T_{\text{coh},M}$ less than half

a year, implying $a = 6$. If the frequency dimension is interpolated using interbinning, $(m_f) \approx 0.14$, giving $m_{\text{opt}} = 0.172$ for a total average mismatch of $\langle m_{\text{tot}} \rangle = 0.312$. It is noteworthy that this result is independent of the computational cost, the coherent integration time, and the number of harmonics summed.

In principle, one can also rearrange for M to find the optimal number of harmonics, which then requires solving a complicated differential equation. However, the derivatives of the functions $h^*(M)$ defined in Equation (E2) and $r(M)$ defined in Equation (C4) are difficult to obtain for most pulse profiles. Therefore, we followed the approach presented in Section 3.5 to find the optimal M at fixed computing cost, which does not require calculating these derivatives.

APPENDIX F

DERIVATION OF STATISTICAL PROPERTIES OF SEMICOHERENT TEST STATISTIC

From Equation (54), the expectation value of S_1 can be written as

$$E_0[S_1] = E_0 \left[\sum_{j,k} w_j w_k e^{-i(\phi(t_j) - \phi(t_k))} \hat{W}_T^{\text{rect}}(\tau_{jk}) \right]. \quad (\text{F1})$$

In order to evaluate this expectation value, we must take into account terms in the double sum where the photon indexes (j, k) are equal, giving

$$E_0[S_1] = \sum_{j=1}^N w_j^2 \hat{W}_T(0) + \sum_{j \neq k} w_j w_k E_0 \left[e^{-i(\phi(t_j) - \phi(t_k))} \right] \hat{W}_T(\tau_{jk}), \quad (\text{F2})$$

where $\sum_{j \neq k}^N$ denotes a double sum over all photons, excluding terms where $j = k$. Under the null hypothesis, $p = 0$, it holds that

$$E_0[e^{-i\phi(t_j)}] = E_0[e^{i\phi(t_k)}] = 0, \quad (\text{F3})$$

and hence we find that the expectation value of S_1 is simply

$$E_0[S_1] = \sum_{j=1}^N w_j^2 \hat{W}_T(0). \quad (\text{F4})$$

To find the variance of S_1 , we must evaluate

$$E_0[S_1^2] = E_0 \left[\sum_{j,k,l,m} e^{-i(\phi(t_j) - \phi(t_k) + \phi(t_l) - \phi(t_m))} \hat{W}_T(\tau_{jk}) \hat{W}_T(\tau_{lm}) \right]. \quad (\text{F5})$$

Again, taking into account terms where photon indexes are equal, and using Equation (F3), we find that

$$E_0[S_1^2] = \sum_{j=1}^N w_j^4 \hat{W}_T(0)^2 + \sum_{j \neq k} w_j^2 w_k^2 \hat{W}_T(0)^2 + \sum_{j \neq k} w_j^2 w_k^2 \hat{W}_T(\tau_{jk})^2, \quad (\text{F6})$$

and hence the variance of S_1 under the null hypothesis is

$$\begin{aligned} \text{Var}_0[S_1] &= E_0[S_1^2] - E_0[S_1]^2 \\ &= \sum_{j \neq k} w_j^2 w_k^2 \hat{W}_T(\tau_{jk})^2. \end{aligned} \quad (\text{F7})$$

From now on in this section, we will use the rectangular lag-window $\hat{W}_T^{\text{rect}}(\tau_{jk})$ of Equation (56). In addition, we assume binary photon weights for simplicity. In this case, one obtains

$$E_0[S_1] = N, \quad \text{Var}_0[S_1] \approx N^2 R^{-1}. \quad (\text{F8})$$

To derive the moments of the distribution of S_1 in the presence of a perfectly matched signal, we need to distinguish times t_j of non-pulsed photons (i.e., background) from pulsed photons by denoting the latter times as t'_j . We then use the definitions of the Fourier coefficients of the pulse profile to evaluate the expectation values

$$E_p[e^{-in\phi(t'_j)}] = \gamma_n, \quad E_p[e^{in\phi(t'_j)}] = \gamma_n^*. \quad (\text{F9})$$

Evaluating Equations (F1) and (F5), using the expectation values from Equations (F3) and (F9), with a pulsed fraction $p \sim \mathcal{O}(10^{-1})$ and a typical pulse profile γ_n (see Figure 2), gives the first two moments of the distribution of S_1 in the presence of a weak signal as

$$E_p[S_1] \approx N + p^2 N^2 |\gamma_1|^2 R^{-1}, \quad (\text{F10})$$

$$\text{Var}_p[S_1] \approx \frac{N^2}{R} (1 + 2p^2 N |\gamma_1|^2 R^{-1}), \quad (\text{F11})$$

where we have assumed a large number of photons $N \gg 1$, and that R is large enough such that edge effects (e.g., effectively shorter windows near the end of the observational data time span) become negligible.

Again, appealing to the central limit theorem (i.e., assuming that there are many photon pairs within the double sums of Equation (F1)), we can approximate the distribution of S_1 by a normal distribution with the same mean and variance. By comparison with numerical simulations, Figure 12 validates this approximation for the purpose of the sensitivity estimation as presented in Section 4.1.

APPENDIX G

SEMICOHERENT METRIC

To derive the semicoherent metric, we investigate the mismatch in the semicoherent detection statistic in the presence of a strong signal. Starting from Equation (54), using binary photon weights and the rectangular lag window,

$$\begin{aligned} S_1 &= \sum_{j=1}^N \sum_{k=1}^N e^{-i[\phi(t_j) - \phi(t_k)]} \hat{W}_T^{\text{rect}}(\tau_{jk}) \\ &= \sum_{j=1}^N e^{-i\phi(t_j)} \sum_{k=1}^N e^{i\phi(t_k)} \hat{W}_T^{\text{rect}}(\tau_{jk}). \end{aligned} \quad (\text{G1})$$

Again, replacing the sum over k with a continuous integral allows us to write the mismatch as

$$\bar{m} = 1 - \frac{\sum_{j=1}^N e^{-i\phi(t_j, \mathbf{u}_{\text{sig}} + \Delta \mathbf{u})} \langle e^{i\phi(t, \mathbf{u}_{\text{sig}} + \Delta \mathbf{u})} \rangle^{[t_j]}}{\sum_{j=1}^N e^{-i\phi(t_j, \mathbf{u}_{\text{sig}})} \langle e^{i\phi(t, \mathbf{u}_{\text{sig}})} \rangle^{[t_j]}}. \quad (\text{G2})$$




Figure 12. Comparison of empirical and analytically predicted probability density function (PDF) of the semicoherent test statistic S_1 . The blue curves (left) refer to the noise-only case, where S_1 has been calculated from many simulated data sets of $N = 10^4$ unit-weight photons with $R = 60$ to obtain the empirical PDF (solid curve), which is compared to the analytical PDF (dashed). The green curves represent the PDF of S_1 for simulated data sets containing signals with a pulsed fraction of $p = 0.1$ and a profile with $|\gamma_1|^2 = 0.668$, where again the empirical PDF (solid curve) is compared to the analytical PDF (dashed).

(A color version of this figure is available in the online journal.)

Assuming that each coherent window contains the same power (and hence has the same S/N at \mathbf{u}_{sig}), this can be simplified to

$$\bar{m} = 1 - \frac{1}{N} \sum_{j=1}^N e^{-i\phi(t_j, \Delta \mathbf{u})} \langle e^{i\phi(t, \Delta \mathbf{u})} \rangle^{[t_j]}. \quad (\text{G3})$$

Taylor expanding this mismatch around $\Delta \mathbf{u} = 0$ to second order in $\Delta \mathbf{u}$ gives

$$\begin{aligned} \bar{m} &= \frac{i}{N} \sum_{j=1}^N (\partial_k \phi|_{t=t_j} - \langle \partial_k \phi \rangle^{[t_j]}) \Delta u^k \\ &+ \frac{1}{2N} \sum_{j=1}^N (\partial_k \phi|_{t=t_j} \partial_\ell \phi|_{t=t_j} + \langle \partial_k \phi \partial_\ell \phi \rangle^{[t_j]}) \Delta u^k \Delta u^\ell \end{aligned}$$

$$\begin{aligned} &- \frac{1}{N} \sum_{j=1}^N (\partial_k \phi|_{t=t_j} \langle \partial_\ell \phi \rangle^{[t_j]}) \Delta u^k \Delta u^\ell \\ &+ \frac{i}{2N} \sum_{j=1}^N (\partial_k \partial_\ell \phi|_{t=t_j} - \langle \partial_k \partial_\ell \phi \rangle^{[t_j]}) \Delta u^k \Delta u^\ell \\ &+ \mathcal{O}(\Delta u^3), \end{aligned} \quad (\text{G4})$$

where there are implicit sums over repeated indices.

Under the assumption $T \ll T_{\text{obs}}$, the mismatch of Equation (G4) becomes

$$\begin{aligned} \bar{m} &\approx \frac{1}{2N} \sum_{j=1}^N (\langle \partial_k \phi \partial_\ell \phi \rangle^{[t_j]} - \langle \partial_k \phi \rangle^{[t_j]} \langle \partial_\ell \phi \rangle^{[t_j]}) \Delta u^k \Delta u^\ell \\ &= \frac{1}{2N} \sum_{j=1}^N G_{k\ell}^{[t_j]} \Delta u^k \Delta u^\ell. \end{aligned} \quad (\text{G5})$$

Hence, the semicoherent metric components can be found by taking half the average of the coherent metric components of Equations (B10) over all photons in the observation time. Using the approximations given in Pletsch (2010), which are valid under the assumption that the data set spans many years, we find

$$\bar{G}_{ff} = \frac{\pi^2 T^2}{6}, \quad (\text{G6a})$$

$$\bar{G}_{f\dot{f}} = \frac{\pi^2 T^4}{360} \gamma^2, \quad (\text{G6b})$$

$$\bar{G}_{n_x n_x} = \bar{G}_{n_y n_y} = \pi^2 f^2 r_E^2 [1 - \text{sinc}^2(\Omega_E T / 2\pi)], \quad (\text{G6c})$$

where γ is the semicoherent refinement factor (Pletsch & Allen 2009; Pletsch 2010) defined as

$$\gamma^2 = 1 + \frac{60}{N} \sum_{j=1}^N \frac{(t_j - t_0)^2}{T^2}. \quad (\text{G7})$$

The mismatches predicted by these derived metric components are compared to the measured mismatches in S_1 for a simulated pulsar signal in Figure 13. For the purpose of the analytic




Figure 13. Comparison of mismatch in S_1 (dashed curves) with semicoherent metric prediction (solid curves). In each panel, the horizontal axis shows the offset from the signal parameters in f (left), \dot{f} (middle), and sky position (right). The sky-location offset is $\sqrt{\Delta n_x^2 + \Delta n_y^2}$, which measures the offset in coordinates (n_x, n_y) in the ecliptic plane. The underlying pulsar signal has been simulated with parameters $f = 32$ Hz, $\dot{f} = -10^{-12}$ Hz s $^{-1}$ for a total observational data time span of $T_{\text{obs}} = 3.4$ yr and a coherent window size of $T = 524288$ s.

(A color version of this figure is available in the online journal.)

study of the computing cost scaling in this paper, we employ the approximation $\gamma \approx \sqrt{5} T_{\text{obs}}/T = \sqrt{5} R$. Hence, the determinant of the semicoherent metric is obtained as

$$\sqrt{\det \bar{G}} \approx \frac{\pi^4}{4\sqrt{27}} T^3 f^2 r_E^2 R \left[1 - \text{sinc}^2 \left(\frac{\Omega_E T}{2\pi} \right) \right]. \quad (\text{G8})$$

APPENDIX H

OPTIMAL MISMATCH IN SEMICOHERENT SEARCH

Following the same steps as in Appendix E, we can find the optimal average mismatch for a semicoherent search with sensitivity $p_{\text{scoh},1}^{-1}$ at a fixed computing cost C_0 by consideration of the following Lagrange function:

$$\begin{aligned} L(T, \bar{m}, \lambda) &= p_{\text{scoh},1}^{-1} + \lambda(C_{\text{scoh}} - C_0) \\ &= (1 - \langle \bar{m}_{\text{tot}} \rangle)^{1/2} T^{1/4} + \lambda(K'_{\text{scoh}} \bar{m}^{-3/2} T^{(s-1)} - C_0). \end{aligned} \quad (\text{H1})$$

Applying the method of Lagrange multipliers as above, we find that

$$3\xi \bar{m}_{\text{opt}} = \frac{1 - \xi \bar{m}_f}{\frac{4(s-1)}{3} + 1}. \quad (\text{H2})$$

As argued in Section 4.3, an efficient strategy uses coherence window sizes T much less than half a year. In this regime of interest, $s = 5$. Using interbinning to interpolate the frequency spectrum gives $\langle m_f \rangle \approx 0.075$, giving the optimal maximum mismatch in the remaining three parameters as $\bar{m}_{\text{opt}} = 0.146$.

APPENDIX I

SKY-GRID CONSTRUCTION

From the metrics derived above, in Appendices B and G, we know when searching over a grid of sky locations that these grid points should be defined by a uniform grid in the ecliptic plane.

To construct the sky search grid for a source within an angular radius of θ from (α_0, δ_0) , this central point is rotated from equatorial to ecliptic coordinates according to the Earth's axial tilt (using the obliquity of the ecliptic, ϵ) and projected into the ecliptic plane, with Cartesian coordinates (x_0, y_0) ,

$$x_0 = \cos(\alpha_0) \cos(\delta_0), \quad (\text{I1})$$

$$y_0 = \cos(\epsilon) \sin(\alpha_0) \cos(\delta_0) + \sin(\epsilon) \sin(\delta_0). \quad (\text{I2})$$

A square of side length θ on the unit circle is calculated around this point and sampled (using the semicoherent or coherent metric components as appropriate) with spacings

$$\Delta n_x = \Delta n_y = 2\sqrt{m/G_{n_x n_x}}. \quad (\text{I3})$$

These grid points are then projected back onto the unit sphere and rotated into equatorial coordinates for barycentering.

Since a square region is sampled in the ecliptic plane, many of the resulting sky-points lie outside the radius defining the search region on the sky. These points are simply discarded, resulting in the original circular search region on the sky in equatorial coordinates, sampled by a uniform grid defined in the ecliptic plane.

A possible problem arises when the search region crosses the ecliptic equator since when the square is constructed in the ecliptic plane, some points lie outside the unit circle, and therefore cannot be projected onto a unit sphere. This can be overcome by reflecting points, (x, y) , which lie outside the unit circle back into the sphere around the ecliptic longitude, l , of the center of the search region:

$$l = \tan^{-1}(y_0/x_0), \quad (\text{I4a})$$

$$x' = \cos(l) - [x - \cos(l)], \quad (\text{I4b})$$

$$y' = \sin(l) - [y - \sin(l)]. \quad (\text{I4c})$$

The new points (x', y') are then projected into the opposite hemisphere from the central point of the search region, resulting in a grid that covers an area of the sky that wraps around the ecliptic equator.

REFERENCES

- Abdo, A. A., Ackermann, M., Ajello, M., et al. 2009, *Sci*, **325**, 840
 Abdo, A. A., Ajello, M., Allafort, A., et al. 2013, *ApJS*, **208**, 17
 Ackermann, M., Ajello, M., Allafort, A., et al. 2012, *ApJ*, **753**, 83
 Astone, P., Borkowski, K. M., Jaranowski, P., et al. 2010, *PhRvD*, **82**, 022005
 Atwood, W. B., Abdo, A. A., Ackermann, M., et al. 2009, *ApJ*, **697**, 1071
 Atwood, W. B., Ziegler, M., Johnson, R. P., & Baughman, B. M. 2006, *ApJL*, **652**, L49
 Balasubramanian, R., Sathyaprakash, B. S., & Dhurandhar, S. V. 1996, *PhRvD*, **53**, 3033
 Beran, R. J. 1969, *Ann. Math. Stat.*, **40**, 1196
 Bickel, P., Kleijn, B., & Rice, J. 2008, *ApJ*, **685**, 384
 Blackman, R. B., & Tukey, J. W. 1958, *Bell Syst. Tech. J.*, **37**, 485
 Brady, P. R., & Creighton, T. 2000, *PhRvD*, **61**, 082001
 Brady, P. R., Creighton, T., Cutler, C., & Schutz, B. F. 1998, *PhRvD*, **57**, 2101
 Bucccheri, R., Bennett, K., Bignami, G. F., et al. 1983, *A&A*, **128**, 245
 Caraveo, P. A. 2014, *ARA&A*, **52**, 211
 Chandler, A. M., Koh, D. T., Lamb, R. C., et al. 2001, *ApJ*, **556**, 59
 Cutler, C., Gholami, I., & Krishnan, B. 2005, *PhRvD*, **72**, 042004
 de Jager, O. C., Raubenheimer, B. C., & Swanepoel, J. W. H. 1989, *A&A*, **221**, 180
 Frigo, M., & Johnson, S. G. 2005, *Proc. IEEE*, **93**, 216
 Groth, E. J. 1975, *ApJS*, **29**, 285
 Guidorzi, C. 2011, *MNRAS*, **415**, 3561
 Guillemot, L., Johnson, T. J., Venter, C., et al. 2012, *ApJ*, **744**, 33
 Kerr, M. 2011, *ApJ*, **732**, 38
 Lee, K. J., Guillemot, L., Yue, Y. L., Kramer, M., & Champion, D. J. 2012, *MNRAS*, **424**, 2832
 Lyons, R. G. 2004, *Understanding Digital Signal Processing* (2nd ed.; Upper Saddle River, NJ: Prentice Hall Professional Technical Reference)
 Meinshausen, N., Bickel, P., & Rice, J. 2009, *Ann. Appl. Stat.*, **3**, 38
 Nolan, P. L., Abdo, A. A., Ackermann, M., et al. 2012, *ApJS*, **199**, 31
 Owen, B. J. 1996, *PhRvD*, **53**, 6749
 Percival, D. B., & Walden, A. T. 1993, *Spectral Analysis for Physical Applications* (Cambridge: Cambridge Univ. Press)
 Pletsch, H. J. 2010, *PhRvD*, **82**, 042002
 Pletsch, H. J., & Allen, B. 2009, *PhRvL*, **103**, 181102
 Pletsch, H. J., Guillemot, L., Allen, B., et al. 2012a, *ApJ*, **744**, 105
 Pletsch, H. J., Guillemot, L., Allen, B., et al. 2012c, *ApJL*, **755**, L12
 Pletsch, H. J., Guillemot, L., Allen, B., et al. 2013, *ApJL*, **779**, L11
 Pletsch, H. J., Guillemot, L., Fehrmann, H., et al. 2012b, *Sci*, **338**, 1314
 Prix, R. 2007, *CQGra*, **24**, S481
 Prix, R., & Shaltev, M. 2012, *PhRvD*, **85**, 084010
 Ransom, S. M., Eikenberry, S. S., & Middleditch, J. 2002, *AJ*, **124**, 1788
 Ransom, S. M., Ray, P. S., Camilo, F., et al. 2011, *ApJL*, **727**, L16
 Saz Parkinson, P. M., Dormody, M., Ziegler, M., et al. 2010, *ApJ*, **725**, 571
 Stoica, P., & Moses, R. 2005, *Spectral Analysis of Signals* (Upper Saddle River, NJ: Pearson Prentice Hall)
 van der Klis, M. 1989, in *NATO ASI Series, Vol. 262, Timing Neutron Stars*, ed. H. gelman & E. Heuvel (Netherlands: Springer), **27**