Exploiting global correlations to detect continuous gravitational waves

Holger J. Pletsch*

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

Bruce Allen†

Max-Planck-Institut für Gravitationsphysik (Albert-Einstein-Institut), Hannover, Germany and
Department of Physics, University of Wisconsin – Milwaukee, USA

Fully coherent searches (over realistic ranges of parameter space and year-long observation times) for unknown sources of continuous gravitational waves are computationally prohibitive. Less expensive hierarchical searches divide the data into shorter segments which are analyzed coherently, then detection statistics from different segments are combined incoherently. Here, we present an improved method for the incoherent step, the Global Correlation Transform (GCT), which exploits global parameter-space correlations in the coherent detection statistic. Application to simulated data shows significant sensitivity improvements compared with previously available methods, increasing the spatial volume probed by more than two orders of magnitude at lower computational cost.

Searching for CW Sources — Direct detection of gravitational waves is the most significant remaining test of Einstein's General Theory of Relativity, and will become an important new astronomical tool.

Rapidly rotating neutron stars are expected to generate continuous gravitational-wave (CW) signals via various mechanisms [1, 2, 3, 4, 5]. Most such stars are electromagnetically invisible, but might be detected and studied via gravitational waves. Recent simulations of neutron star populations [6, 7] suggest that CW sources might eventually be detected with new instruments such as LIGO [8, 9]. World-wide efforts are underway to search for CW signals [10, 11, 12] and observational upper limits already place some constraints on neutron star physics [13, 14].

Because the expected CW signals are weak, sensitive data analysis methods are needed to extract these signals from detector noise. A powerful method is derived in Ref. [15]. This scheme is based on the principle of maximum likelihood detection, which leads to coherent matched filtering. Rotating neutron stars emit monochromatic CW signals, apart from a slowly changing intrinsic frequency. But the terrestrial detector location Doppler-modulates the amplitude and phase of the waveform, as the Earth moves relative to the solar system barycenter (SSB). The parameters describing the signal's amplitude variation may be analytically eliminated by maximizing the coherent matched-filtering statistic [15]. The remaining search parameters describing the signal's phase are the source's sky location, frequency and frequency derivatives, and the resulting coherent detection statistic is called the $\mathcal{F}$-statistic.

This work considers isolated CW sources whose frequency varies linearly with time in the SSB frame. The corresponding phase parameter-space $\mathcal{P}$ is four-dimensional. Standard "physical" coordinates on $\mathcal{P}$ are the frequency $f(t_0)$ at some fiducial time t_0, the frequency's first time derivative $\dot{f}$, and a unit vector $\mathbf{n} =$

$(\cos\delta \cos\alpha, \cos\delta \sin\alpha, \sin\delta)$ on the two-sphere S^2, pointing from the SSB to the source. Here α and δ are right ascension and declination. Thus, a point in parameter space $\mathrm{p} \in \mathcal{P}$ may be labeled by $\mathrm{p} = \{f(t_0), \dot{f}, \mathbf{n}\}$. The $\mathcal{F}$-statistic $\mathcal{F}_{\mathrm{p}}[h]$ is a functional of the detector data set h, and is a function of the point in parameter space $\mathrm{p} \in \mathcal{P}$.

All-sky searches for unknown CW sources using the $\mathcal{F}$-statistic are computationally expensive. For maximum sensitivity, one must convolve the full data set with signal waveforms (templates) corresponding to all possible sources. But the number of templates required for a fully coherent search increases as a high power of the observation time. For one year of data, the computational cost to search a realistic range of parameter space exceeds the total computing power on Earth [15, 16]. Thus a fully-coherent search is limited to much shorter observation times.

Hierarchical semi-coherent search methods address this problem [17, 18, 19]. The data is broken into segments of duration T, where T is much smaller than one year. Each segment is analyzed coherently, computing the $\mathcal{F}$-statistic on a *coarse* grid of templates. Then the $\mathcal{F}$ values from all segments (or statistics derived from $\mathcal{F}$) are incoherently combined using a common *fine* grid of templates, discarding phase information between segments.

Different semi-coherent strategies are currently in use. The Stack-Slide method [17, 18] sums $\mathcal{F}$ values along putative signal tracks in the time-frequency plane. The Hough transform method [19] sums $H(\mathcal{F} - \mathcal{F}_{\mathrm{th}})$ where $\mathcal{F}_{\mathrm{th}}$ is a constant predefined threshold. The Heavyside function $H(\mathrm{x})$ is unity for positive x and vanishes elsewhere. This latter technique is currently used by Einstein@Home [12], a public distributed computing project carrying out the most sensitive blind CW searches.

A central element in these semi-coherent methods is the design of, and link between, the coarse and fine grids. The key quantity is the fractional loss (called *mismatch* $\mathcal{M}$) in

expected $\mathcal{F}$-statistic for a given signal p at a nearby grid point p′. Locally Taylor-expanding $\mathcal{M}$ (to quadratic order) in the differences of the coordinates $\{f(t_0), \dot{f}, \mathbf{n}\}$ of p and p′ defines a signature ++++ metric $\mathrm{d}s^2$ [16, 20, 21, 22]. Current methods consider parameter correlations in $\mathcal{F}$ to *linear* order in T and discard higher orders in T from the metric.

The $\mathcal{F}$-statistic has strong *global* correlations [23, 24] in the physical coordinates $\{f(t_0), \dot{f}, \mathbf{n}\}$, that extend outside the region in which the mismatch is well-approximated by the local metric given above [25]. Recent work [24] has shown that (for a given signal) the region where the expected $\mathcal{F}$-statistic has maximal value may be described by a separate equation for each order of T, when T is small compared to one year. The solutions to each equations is a hypersurface, whose intersections describe the global correlations in $\mathcal{F}$.

For currently used values of T (a day or longer) it is also crucial to consider the fractional loss and global structure of $\mathcal{F}$ to *second-order* in T [24]. For source frequencies above $1\,\mathrm{kHz}$ and for values of T longer than about $60\,\mathrm{h}$, additional orders in T would be needed.

The GCT method — This work exploits the global correlations in the coherent detection statistic $\mathcal{F}$ to construct a significantly improved semi-coherent search technique for continuous gravitational waves. We call the new method the "Global Correlation Transform" (GCT). The GCT technique uses the improved understanding of the global parameter-space correlations to make three important improvements to the incoherent step.

First, in previous approaches the fine grid is obtained by refining the coarse grid in three dimensions, $\dot{f}$ and $\mathbf{n}$. With the GCT, the fine grid is obtained by refining the coarse grid in only one dimension, $\dot{f}$. This greatly reduces the computational cost at *equal detection sensitivity*, although it also reduces the accuracy with which the parameters of a source are estimated. But this is a very profitable trade, because in a hierarchical search the primary goal of the first stages is to discard the uninteresting regions of parameter space. Later follow-up stages use longer coherent integrations to more accurately determine the source parameters.

Second, existing techniques combine the coherent results less effectively than the GCT technique, because they do not use metric information beyond linear order in T. This gives the GCT higher sensitivity at equal computational cost.

Third, the GCT can simultaneously do a Stack-Slide-like summing of $\mathcal{F}$ values and a Hough-like summing of $H(\mathcal{F} - \mathcal{F}_{\mathrm{th}})$, with a lower total computational cost than either one of these methods individually.

For a given CW source with realistic phase parameter values ($f \leq 1\,\mathrm{kHz}$, $|\dot{f}| \leq f/50\,\mathrm{yr}$) and coherent data segment lengths $T \leq 60\,\mathrm{h}$, the global-correlation structure of the $\mathcal{F}$-statistic is well described by the first- and second-

order global-correlation equations [24]:

$$\nu(t) = f(t) + f(t)\dot{\boldsymbol{\xi}}(t) \cdot \mathbf{n} + \dot{f}\,\boldsymbol{\xi}(t) \cdot \mathbf{n}\,,$$
$$\dot{\nu}(t) = \dot{f} + f(t)\ddot{\boldsymbol{\xi}}(t) \cdot \mathbf{n} + 2\dot{f}\,\dot{\boldsymbol{\xi}}(t) \cdot \mathbf{n}\,,$$
$$\text{where } f(t) \equiv f(t_0) + (t - t_0)\dot{f}\,. \tag{1}$$

Here $\boldsymbol{\xi}(t) \equiv \mathbf{r}_{\mathrm{orb}}(t)/c$, with $\mathbf{r}_{\mathrm{orb}}(t)$ denoting the vector from the Earth's barycenter to the SSB, and c the speed of light. Apart from an overall factor, the quantities $\nu(t)$ and $\dot{\nu}(t)$ are called the "global-correlation parameters". They can be interpreted as the source's instantaneous frequency and frequency derivative at the detector, at detector time t.

The global-correlation parameters provide new coordinates (ν and $\dot{\nu}$) on $\mathcal{P}$. It is useful to also introduce new (real-valued) sky coordinates n_{x} and n_{y} (as in [26]):

$$n_{\mathrm{x}}(t) + i\,n_{\mathrm{y}}(t) = f(t)\,\tau_E\,\cos\delta_D\,\cos\delta\,\mathrm{e}^{i[\alpha - \alpha_D(t)]}\,. \tag{2}$$

Here $\tau_E = R_E/c \approx 21\,\mathrm{ms}$ is the light travel time from the Earth center to the detector, and $\alpha_D(t), \delta_D$ are the detector position at time t. The metric separation $\mathrm{d}s^2$ is

$$\mathrm{d}s^2/(2\pi)^2 = \mathrm{d}\nu^2\,T^2/12 + \gamma^2\,\mathrm{d}\dot{\nu}^2\,T^4/720 + \mathrm{d}n_{\mathrm{x}}^2/2$$
$$+ \mathrm{d}n_{\mathrm{y}}^2/2 - \mathrm{d}\nu\,\mathrm{d}n_{\mathrm{y}}\,T/(\pi\ell) + \mathrm{d}\dot{\nu}\,\mathrm{d}n_{\mathrm{x}}\,T^2/(\pi\ell)^2\,. \tag{3}$$

In defining differences in coordinates $\{\nu, \dot{\nu}, n_{\mathrm{x}}, n_{\mathrm{y}}\}$, the time t in Eqs. (1) and (2) is the midpoint of the data segment spanning times $[t - T/2, t + T/2]$, and $\gamma = 1$. To simplify the form of the metric, T is taken to be a positive integer number ℓ of sidereal days.

The new coordinates $\{\nu, \dot{\nu}, n_{\mathrm{x}}, n_{\mathrm{y}}\}$ have important advantages over the original coordinates $\{f, \dot{f}, \mathbf{n}\}$. The metric is explicitly coordinate-independent (showing that $\mathcal{P}$ is flat). In fact, the region around a point p in which the mismatch $\mathcal{M}$ is well-approximated by $\mathrm{d}s^2$ is *much* larger [25].

Consider a segment of data h_{p} which contains a strong CW source with phase parameters p. If the sky separation patch is small enough to neglect the $\mathrm{d}n_{\mathrm{x}}$ and $\mathrm{d}n_{\mathrm{y}}$ terms in Eq. (3), then $\mathcal{F}_{\mathrm{p}'}[h_{\mathrm{p}}]$ is extremized for all p′ that have the *same global-correlation parameters* ν and $\dot{\nu}$ as p. This set of points in $\mathcal{P}$ forms a two-dimensional surface $\mathrm{d}\nu = \mathrm{d}\dot{\nu} = 0$. Thus, for all sources within the sky patch, there exists a different $(f, \dot{f})$ pair with those same values of ν and $\dot{\nu}$. This property is exploited by the GCT algorithm.

An implementation of the GCT — To start, the data set is divided into N segments of length T (potentially including short gaps in operation) labeled by the integer $j = 1, ..., N$. The segments span time intervals $[t_j - T/2, t_j + T/2]$. The detector-time midpoint of segment j is t_j and $t_0 = \frac{1}{N}\sum_{j=1}^{N} t_j$ is the fiducial time.

Every segment is analyzed *coherently* on a coarse grid in phase parameter space $\mathcal{P}$. This grid is constructed so that no point in $\mathcal{P}$ is farther than a specified distance from some coarse-grid point, where the distance measure is defined by the metric above. To simplify the grid construction, large frequency bands are analyzed by breaking them into many narrow sub-bands. For each data segment j, and at each

coarse grid point, the $\mathcal{F}$-statistic is evaluated, and "stats" are obtained. Here, the word "stat" denotes the two-tuple $(\mathcal{F}_j, H(\mathcal{F}_j - \mathcal{F}_{\text{th}}))$.

A typical coarse grid (used for all data segments) is the Cartesian product of a rectangular grid in $f, \dot{f}$ and a grid on the sky-sphere $\mathbf{n} \in S^2$. The spacings in f and $\dot{f}$ are $\Delta f = \sqrt{12m}/(\pi T)$ and $\Delta \dot{f} = \sqrt{720m}/(\pi T^2)$, where m is the one-dimensional metric mismatch parameter [11]. The spacing of the coarse sky grid is chosen so that the dn_{x} and dn_{y} terms in Eq. (3) may be neglected. When orthogonally projected onto the equatorial unit disk, the sky grid is uniform, and contains $\approx 2\pi/(\Delta \varphi)^2$ points, with $\Delta \varphi = \sqrt{2m}/(\pi f \tau_E \cos \delta_D)$.

The *incoherent* step combines the "stats" obtained by the coherent analysis, using a fine grid in $\mathcal{P}$. At each point in the fine grid, a "stat" value is obtained by summing one stat value from each of the N coarse grids. The coarse grid point is the one with the same sky position as the fine grid point, which has the smallest separation in the global correlation parameters, calculated using the metric Eq. (3) above. The final result is a "stat" value at each point on the fine grid. The first element of the stat is the sum of the $\mathcal{F}$-statistic values from the coarse grid points. The second element is a number count, reflecting the number of data segments in which $\mathcal{F}_{\text{th}}$ was exceeded. A detectable CW signal leads to a fine-grid point with a high number count and a large sum of $\mathcal{F}$-statistics.

The spacing of the fine grid is determined from the metric for the fractional loss of the expected $\sum_{j=1}^{N} \mathcal{F}_j$ due to parameter offsets between a putative signal location and a fine grid point at the fiducial time t_0. This may be calculated as in [17], by averaging the coarse-grid metric over the N different segments. Since each coarse-grid metric is no longer calculated at the data-segment midpoints (but at t_0), the coefficients change between segments because of the time-dependence of the parameter-space coordinates. For our choice of t_0 and T, the only additional term in the metric Eq. (3) that does not average to zero is $(t_j - t_0)^2 T^2 d\dot{\nu}^2/12$, and the averaged metric takes a form identical to Eq. (3) but with

$$\gamma^2 = 1 + \frac{60}{T^2 N} \sum_{j=1}^{N} (t_j - t_0)^2, \qquad (4)$$

where the parameter offsets in Eq. (3) are calculated at the fiducial time t_0. Thus, the fine grid may be identical to the coarse grid except that the spacing $\Delta \dot{f}$ is smaller by a factor γ, which is of order N when the number of data segments is large. No further refinement in frequency or sky position is needed. Coherent integration over the total observation time would require refining both $\Delta \dot{f}$ and Δf (increasing points $\propto N^3$), plus similar sky refinements.

GCT versus Hough performance — Monte Carlo simulations are used to illustrate the improved performance of the GCT compared with the conventional Hough transform
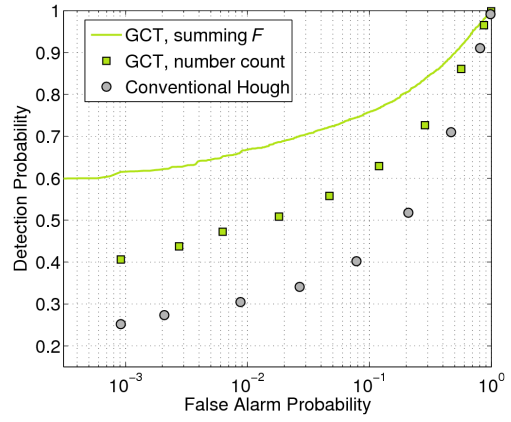

FIG. 1: Receiver operating characteristic curve. The GCT performs better than the conventional Hough transform method.

method. The software tools used are part of LALApps [27] and employ accurate barycentering routines with timing errors below 3μs. To provide a realistic comparison, simulated data sets covered the same time intervals as the input data used for the current (S5R5) Einstein@Home search [12]. Those data, from LIGO Hanford (H1, 4km) and LIGO Livingston (L1, 4km), are not contiguous, but contain gaps when the detectors are not operating. The total time interval spanned is about 264 days, containing 121 data segments of $25\,\text{h}$ duration.

The false alarm probabilities are obtained using $5\,000$ simulated data sets with different realizations of stationary Gaussian white noise, with one-sided strain spectral density $\sqrt{S_h} = 3.25 \times 10^{-22}\,\text{Hz}^{-1/2}$. To find the detection probabilities, different CW signals with fixed strain amplitude h_0 are added. The parameters [15] are randomly drawn from uniform distributions in $\cos(\text{inclination } \iota)$, polarization ψ, initial phase ϕ_0, the entire sky, $f(t_0) \in [100.1, 100.3]\,\text{Hz}$, and $\dot{f} \in [-1.29, -0.711]\,\text{nHz/s}$.

Figure 1 compares the performance of the two methods. The receiver operating characteristic is the detection probability as a function of false alarm probability, at fixed source strain amplitude $h_0 = 6 \times 10^{-24}$. Because the number count (using $\mathcal{F}_{\text{th}} = 2.6$) is discrete, the two "curves" in Fig. 1 consist of discrete points. The GCT (using *either* number counts *or* summed $\mathcal{F}$ as a detection statistic) is superior to the conventional Hough method.

In addition, the GCT is computationally faster. This comparison used *identical* coherent stages ($m = 0.3$, with $2\,981$ coarse-grid points) for the GCT and conventional Hough method. But in the incoherent combination stage, the GCT and the conventional Hough method used *different* fine grids. The GCT fine grid had 506 times as many points as the coarse grid, but the Hough fine grid had $7\,056$ times as many points. In spite of using 14 times fewer fine-grid points, the GCT gave *substantially higher* sensitivity.

Figure 2 shows another comparison of the GCT and Hough method. It compares the detection efficiencies for
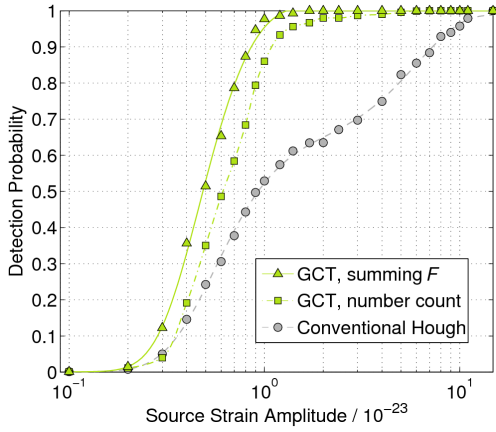
FIG. 2: Probability of detection as a function of source strain amplitude h_0, at a fixed false alarm probability of 1%.

different values of source strain amplitude h_0, at a fixed 1% false alarm probability. As above, each point in Fig. 2 is obtained by analyzing $2\,000$ simulated data sets. Again, the GCT in both modes of operation performs substantially better than the Hough method. For example, compare the source strain amplitude h_0 needed to obtain 90% detection probability. The strain required by the GCT in number-count operation mode is smaller by a factor of about six than the strain needed by the Hough method, making the "distance reach" [15] of the GCT six times larger. This increases the number of potentially detectable sources by more than two orders of magnitude, since the "visible" spatial volume increases as the cube of the distance. In fact the lower computational cost of the GCT would also allow increases in N or T, even further improving the sensitivity.

These results are qualitatively independent of frequency, as confirmed in additional comparisons.

Conclusions — A new semi-coherent technique for detecting continuous gravitational-wave sources has been described. In contrast to previous approaches, the GCT exploits global parameter-space correlations in the coherent detection statistic $\mathcal{F}$, in the subsequent incoherent combination step. For coherent integration times $T \leq 60\,\mathrm{h}$, the global correlations are well-described by the second-order (in T) formulae presented here. The method should also be extendible to longer coherent integration times by including higher orders in T. It could also be extended to search for CW signals from non-isolated sources (i.e. in binary systems) as well as to space-based detectors.

Realistic Monte Carlo simulations show that the GCT is much more sensitive than the Hough transform method (currently the most sensitive CW search technique). The GCT is also computationally simpler, and more efficient.

The LIGO Scientific Collaboration is currently working to deploy the GCT on the Einstein@Home project [12], starting with LIGO S6 data. The combination of new and more sensitive search techniques, and new and more sensitive data, greatly increases the chance of making the first gravitational wave detection of a CW source. The detection of CW signals will provide new means to discover and locate neutron stars, and will eventually provide unique insights into the nature of matter at high densities.

[*] Electronic address: Holger.Pletsch@aei.mpg.de
[†] Electronic address: Bruce.Allen@aei.mpg.de
[1] B. J. Owen, L. Lindblom, C. Cutler, B. F. Schutz, A. Vecchio, and N. Andersson, Phys. Rev. D **58**, 084020 (1998).
[2] G. Ushomirsky, C. Cutler, and L. Bildsten, Mon. Not. Roy. Astron. Soc. **319**, 902 (2000).
[3] C. Cutler, Phys. Rev. D **66**, 084025 (2002).
[4] D. I. Jones and N. Andersson, Mon. Not. Roy. Astron. Soc. **331**, 203 (2002).
[5] B. J. Owen, Phys. Rev. Lett. **95**, 211101 (2005).
[6] B. Knispel and B. Allen, Phys. Rev. D **78**, 044031 (2008).
[7] C. J. Horowitz and K. Kadau (2009), arXiv:0904.1986.
[8] A. Abramovici et al., Science **256**, 325 (1992).
[9] B. C. Barish and R. Weiss, Physics Today **52**, 44 (1999).
[10] B. Abbott et al. (The LIGO Scientific Collaboration), Phys. Rev. D **77**, 022001 (2008).
[11] B. Abbott et al. (The LIGO Scientific Collaboration), Phys. Rev. D **79**, 022001 (2009).
[12] Einstein@Home is a vounteer distributed computing project; see http://einstein.phys.uwm.edu/.
[13] B. Abbott et al. (The LIGO Scientific Collaboration), The Astrophysical Journal Letters **683**, L45 (2008).
[14] B. Abbott et al. (The LIGO Scientific Collaboration), Phys. Rev. Lett. **102**, 111102 (2009).
[15] P. Jaranowski, A. Królak, and B. F. Schutz, Phys. Rev. D **58**, 063001 (1998).
[16] P. R. Brady, T. Creighton, C. Cutler, and B. F. Schutz, Phys. Rev. D **57**, 2101 (1998).
[17] P. R. Brady and T. Creighton, Phys. Rev. D **61**, 082001 (2000).
[18] C. Cutler, I. Gholami, and B. Krishnan, Phys. Rev. D **72**, 042004 (2005).
[19] B. Krishnan, A. M. Sintes, M. A. Papa, B. F. Schutz, S. Frasca and C. Palomba, Phys. Rev. D **70**, 082001 (2004).
[20] R. Balasubramanian, B. S. Sathyaprakash, and S. V. Dhurandhar, Phys. Rev. D **53**, 3033 (1996).
[21] B. J. Owen, Phys. Rev. D **53**, 6749 (1996).
[22] R. Prix, Phys. Rev. D **75**, 023004 (2007).
[23] R. Prix and Y. Itoh, Class. Quant. Grav. **22**, S1003 (2005).
[24] H. J. Pletsch, Phys. Rev. D **78**, 102005 (2008).
[25] Using $\{\nu, \dot{\nu}, n_{\mathrm{x}}, n_{\mathrm{y}}\}$ coordinates the metric $\mathrm{d}s^2$ remains accurate up to $\mathcal{M} = 0.3$. In contrast, in the $\{f, \dot{f}, \mathbf{n}\}$ coordinates $\mathrm{d}s^2$ can yield errors greater than 10% for mismatches as small as 0.001. In part, this is because the metric in $\{f, \dot{f}, \mathbf{n}\}$ varies significantly over the $\mathcal{M} = 0.3$ region.
[26] P. Astone, K. M. Borkowski, P. Jaranowski, and A. Królak, Phys. Rev. D **65**, 042003 (2002).
[27] http://www.lsc-group.phys.uwm.edu/daswg/.