An Efficient Algorithm for Computing the Time Resolved Full Sky Cross Power in an Interferometer with Omni-directional Elements

Kipp C. Cannon
Department of Physics, University of Wisconsin Milwaukee,
*PO Box 413, Milwaukee, Wisconsin, 53201-0413**

We derive both time- and frequency-domain algorithms for the fast, approximate, numeric evaluation of the cross-correlation integral that appears in a variety of imaging applications. These algorithms possess a number of useful features. Their operation counts are linear in the amount (duration) of data to be processed, which allows them to be run quickly on small amounts of data, thereby allowing one to resolve time-dependence in the sources being imaged. The algorithms also provide means by which to systematically increase or decrease their accuracy for increases and decreases in speed respectively, that is to say one can choose to obtain the result to lower accuracy and thereby obtain the result more quickly or with more modest computer requirements. The frequency-domain implementation is particularly fast, and on a 32-bit 1.8 GHz Pentium M computer has demonstrated the ability to synthesize full-sky cross-power images for a single baseline up to harmonic order 19, with an SNR of 26 dB, at speeds of over 800 ksample/s.

PACS numbers: 02.70.-c 07.05.Kf 07.60.Ly 42.30.Va 95.75.-z 95.75.Kk 95.75.Mn 95.75.Pq

I. INTRODUCTION

In problems of interferometric imaging, one finds an integral of the form

$$\xi_{1,2}(\hat{s}) = \int_{t-\frac{T}{2}}^{t+\frac{T}{2}} g_1(t' - \vec{r}_1 \cdot \hat{s}) g_2(t' - \vec{r}_2 \cdot \hat{s})\, \mathrm{d}t'. \tag{1}$$

This integral computes the correlation of two functions of time, as a function of a relative delay applied to the two of them. In particular, this form arises when one has omni-directional instruments located at $\vec{r}_1$ and $\vec{r}_2$, the instruments produce time-dependent outputs g_1 and g_2 respectively, and one is interested in computing the "cross power" (correlation) associated with the direction $\hat{s}$ integrated over a time T centred on t. Here, and in what follows, we set the wave speed (for example, c) to 1 — distances are measured in units of time.

For example, in radio interferometry, if g_1 and g_2 are the time-dependent voltages seen in the outputs of two antennas, then $\xi_{1,2}(\hat{s})$ is equivalent to the output of a multiplying correlator [1, Chapter 2]. From this one obtains the complex visibility, from which is synthesized a brightness map of the radio source at which the antennas have been pointed. The integral also appears in synthetic aperture radar and sonar signal processing applications [2].

In gravitational wave astronomy, when the g_i are derived from the strain time series outputs of two gravitational wave antennas, (1) or forms similar to it can be found in the coherent "burst" search algorithms described in [3], [4], and [5], and in the "stochastic" search algorithms described in [6], [7], and [8].

In most applications, for example radio astronomy, sonar, etc., one or more of the following simplifying assumptions can be made:

- The antennas/transducers whose outputs provide the time series are directional, so that it can be assumed that the $g_i(t)$'s contain contributions only from sources in directions $\hat{s}$ within a small range of angles around the nominal direction $\hat{s}_0$ in which the antennas are pointed, $|\hat{s} - \hat{s}_0| \ll 1$.
- The receiving apparatus (or equivalently the received signal) has a relatively narrow bandwidth, $\Delta f / f_0 \ll 1$, so that the autocorrelation of $g_i(t)$ is approximately sinusoidal over the range of time delays of interest.
- Alternatively, the signal may be broad-band but artificial, being the echo of a transmitted signal, and so in that way the autocorrelation is also known in advance. Knowing the autocorrelation of the received signal in advance allows one to avoid explicitly computing (1) for all $\hat{s}$.

*Electronic address: kipp@gravity.phys.uwm.edu

- The angular distribution of power incident from the source field is independent of time, in other words the sources being imaged are not changing.

Gravitational wave astronomy presents a, perhaps unique, challenge in that in the search for gravitational waves using laser interferometer antennas, none of these simplifying assumptions apply. Laser interferometer antennas are relatively broad-band transducers whose bandwidths are comparable to their centre frequencies, they are nearly omni-directional and the Earth is mostly transparent to gravitational waves so Earth-based instruments do not even have the horizon to narrow their field of view, and many of the sources of interest such as supernovae, gamma-ray bursts, and the coalescences of pairs of compact objects, are transient phenomena. On the other hand, Earth-based gravitational wave antennas are immobile and for the most part unsteerable. They also, typically, observe signals with frequencies below ~ 10 kHz, which is much lower than in most other imaging applications, and means that interferometric networks of gravitational wave antennas have comparatively poor directional sensitivity.[1]

The remainder of the paper is the development of a fast numerical implementation of (1) that allows us to obtain an image of the entire sphere, over a broad frequency band, in a time-resolved manner so that transient signals can be detected and localized. We also present tests of the algorithm's performance and convergence for a variety of simulated signals in the presence of Gaussian noise.

II. OVERVIEW

We assume that the input data, $g(t)$, has been discretely sampled at a regular interval Δt, so that we can form a vector whose elements are

$$\vec{g} = \begin{bmatrix} \vdots \\ g(t_j) \\ \vdots \end{bmatrix}, \tag{2}$$

where $t_j = j\Delta t$. We can write the linear transformation

$$g(t) \to g(t - \vec{r}\cdot\hat{s}) \tag{3}$$

as

$$\vec{g}(\hat{s}) = \mathbb{T}(\vec{r}\cdot\hat{s})\cdot\vec{g}, \tag{4}$$

where $\mathbb{T}(\tau)$ is the matrix that retards the time series by the geometric delay τ. The integral (1) can then be written as the inner product

$$\begin{aligned} \xi_{1,2}(\hat{s}) &= \frac{1}{N}\vec{g}_1^{\mathrm{T}}(\hat{s})\cdot\vec{g}_2(\hat{s}) \\ &= \frac{1}{N}\vec{g}_1^{\mathrm{T}}\cdot\mathbb{T}_1^{\mathrm{T}}(\vec{r}_1\cdot\hat{s})\cdot\mathbb{T}_2(\vec{r}_2\cdot\hat{s})\cdot\vec{g}_2, \end{aligned} \tag{5}$$

where $()^{\mathrm{T}}$ denotes the transpose. Each of $\xi_{1,2}$, $\mathbb{T}_1$, and $\mathbb{T}_2$ depends on $\hat{s}$. This dependence can be encoded by projecting them onto some choice of basis, and computing (5) for each of the terms in that expansion separately.

Here, we can see a clear statement of the general problem of the construction of an efficient correlator for computing (1): (i) what choice of basis for the vectors $\vec{g}_i$ makes the product $\mathbb{T}_1^{\mathrm{T}}\cdot\mathbb{T}_2$ as close to diagonal and time independent as possible, and (ii) what choice of basis for the expansion of the angular dependence of $\xi_{ij}(\hat{s})$ minimizes the number of coefficients that need to be computed?

For the case of a network of Earth-based gravitational wave laser interferometer antennas, using an Earth-fixed co-ordinate system for $\hat{s}$ makes $\mathbb{T}_1^{\mathrm{T}}\cdot\mathbb{T}_2$ independent of time, and Fourier transforming $\vec{g}_i$ to the frequency domain makes $\mathbb{T}_1^{\mathrm{T}}\cdot\mathbb{T}_2$ exactly diagonal. Since the discretely-sampled time series $\vec{g}_i$ are necessarily band-limited, $\xi_{ij}(\hat{s})$ is a band-limited function on the sphere and so can be accurately represented by expansion to finite order in spherical harmonics. We shall use a spherical harmonic expansion for $\xi_{ij}(\hat{s})$, although no attempt is made to show that this

[1] Bar antennas are small enough that their orientations can be changed, but because of any gravitational wave antenna's sensitivity to mechanical disturbances this is not something that can be done frequently; perhaps one change of orientation every few months.

represents a general $\xi_{ij}(\hat{s})$ with the smallest number of co-efficients, and infact it will be seen that it certainly does not. However, it is convenient to use spherical harmonics due to their orthogonality, the availability of efficient numerical algorithms for computing them, and the straight-forward framework this basis provides for the systematic lower-order approximation of, and iterative refinement of, $\xi_{ij}(\hat{s})$.

Of course, in an Earth-fixed co-ordinate system, due to the rotation of the Earth a source at a fixed location on the celestial sphere will be seen in different $\hat{s}$ directions at different times. As time passes, point sources will drift through the direction $\hat{s}$, and the integrated cross-power, $\xi_{1,2}(\hat{s})$, becomes the cross power averaged over the sources that passed through $\hat{s}$ during the integration. The effect is similar to a photograph of the night sky taken with a camera on a fixed mount whose shutter has been held open, allowing the stars to leave tracks on the film. If the extended tracks are undesirable, and one wishes to resolve point sources as points, then since the imaging system has a finite angular resolution one can simply choose to leave the shutter open for a length of time short enough that during the "exposure" a point source moves through an angle small compared to the imaging system's angular resolution. Longer integrations are achieved by acquiring a sequence of short integrations, transforming each to a common sky-fixed co-ordinate system, and summing.

III. CONVENTIONS

A. Fourier Transforms

The Fourier transform of the function $g(t)$, and the inverse transform, are

$$g(t) = \int_{-\infty}^{+\infty} \tilde{g}(f) \mathrm{e}^{2\pi \mathrm{i} f t} \, \mathrm{d}f, \tag{6}$$

$$\tilde{g}(f) = \int_{-\infty}^{+\infty} g(t) \mathrm{e}^{-2\pi \mathrm{i} f t} \, \mathrm{d}t. \tag{7}$$

If $g(t)$ is real, then $\tilde{g}(f) = \tilde{g}^*(-f)$, where $()^*$ is the complex conjugate. The Wiener-Khinchin theorem (that the autocorrelation of a function is the Fourier transform of its spectral density) matching these definitions is

$$\int_{-\infty}^{+\infty} g(t'-t) g(t') \, \mathrm{d}t' = \int_{-\infty}^{+\infty} S_g(f) \mathrm{e}^{2\pi \mathrm{i} f t} \, \mathrm{d}f, \tag{8}$$

where the spectral density function is

$$S_g(f) = \left| \tilde{g}(f) \right|^2. \tag{9}$$

The discrete Fourier transform of the N samples $g(t_j) = g(j\Delta t)$, where $j \in \{0, \ldots, N-1\}$, and the inverse transform are

$$g_j = \frac{1}{N} \sum_{k=0}^{N-1} \tilde{g}_k \mathrm{e}^{2\pi \mathrm{i} j k / N}, \tag{10}$$

$$\tilde{g}_k = \sum_{j=0}^{N-1} g_j \mathrm{e}^{-2\pi \mathrm{i} j k / N}. \tag{11}$$

The Kronecker δ is [9, Section 14.6]

$$\frac{1}{N} \sum_{k=0}^{N-1} \mathrm{e}^{-2\pi \mathrm{i}(j_1 - j_2)k/N} = \delta_{j_1 j_2} = \begin{cases} 1, & \text{if } j_1 = j_2, \\ 0, & \text{if } j_1 \neq j_2. \end{cases} \tag{12}$$

B. Spherical Harmonics

The spherical harmonics are

$$\mathrm{Y}_{lm}(\theta, \phi) = \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} \, \mathrm{P}_l^m(\cos\theta) \mathrm{e}^{\mathrm{i} m \phi}, \; l \geq 0, \; |m| \leq l, \tag{13}$$

where $\mathrm{P}_l^m(x)$ are the associated Legendre polynomials,

$$\mathrm{P}_l^m(x) = (-1)^m (1-x^2)^{m/2} \frac{\mathrm{d}^m}{\mathrm{d}x^m} \mathrm{P}_l(x),\ m \geq 0, \tag{14}$$

and where the Legendre polynomial $\mathrm{P}_l(x) = \frac{1}{2^l l!} \frac{\mathrm{d}^l}{\mathrm{d}x^l}(x^2-1)^l$. Associated Legendre polynomials of negative m are related to those of positive m by[2]

$$\mathrm{P}_l^{-m}(x) = (-1)^m \frac{(l-m)!}{(l+m)!} \mathrm{P}_l^m(x),\ m \geq 0. \tag{15}$$

It follows that

$$\mathrm{Y}_{l,-m}(\theta,\phi) = (-1)^m \mathrm{Y}_{lm}^*(\theta,\phi),\ m \geq 0. \tag{16}$$

Spherical harmonics defined in this manner obey the orthogonality relation

$$\int_{\phi=0}^{2\pi} \int_{\theta=0}^{\pi} \mathrm{Y}_{lm}(\theta,\phi)\, \mathrm{Y}_{pq}^*(\theta,\phi) \sin\theta \,\mathrm{d}\theta \,\mathrm{d}\phi = \delta_{lp}\delta_{mq}. \tag{17}$$

The spherical harmonic triple product is

$$\int_{\hat{s}} \mathrm{Y}_{l_1 m_1}(\hat{s})\, \mathrm{Y}_{l_2 m_2}(\hat{s})\, \mathrm{Y}_{lm}^*(\hat{s})\, \mathrm{d}\Omega = (-1)^m \sqrt{\frac{(2l_1+1)(2l_2+1)(2l+1)}{4\pi}} \begin{pmatrix} l_1 & l_2 & l \\ m_1 & m_2 & -m \end{pmatrix} \begin{pmatrix} l_1 & l_2 & l \\ 0 & 0 & 0 \end{pmatrix}, \tag{18}$$

where $\begin{pmatrix} l_1 & l_2 & l \\ m_1 & m_2 & m \end{pmatrix}$ are the Wigner 3-j symbols [13]. Many numerical analysis libraries, such as the GNU Scientific Library [14], contain routines for computing the Wigner 3-j symbols. The 3-j symbol $\begin{pmatrix} l_1 & l_2 & l \\ m_1 & m_2 & m \end{pmatrix} = 0$ if any of the following conditions is *not* met [13]:

$$m_1 + m_2 + m = 0, \tag{19}$$
$$l_1 + l_2 - l \geq 0, \tag{20}$$
$$l_1 - l_2 + l \geq 0, \tag{21}$$
$$-l_1 + l_2 + l \geq 0. \tag{22}$$

We also have the special values [15, Section 27.9]

$$\begin{pmatrix} l_1 & l_2 & l \\ 0 & 0 & 0 \end{pmatrix} = 0, \text{ if } l_1 + l_2 + l \text{ is an odd integer}, \tag{23}$$
$$\begin{pmatrix} l_1 & l_1 & l \\ m_1 & m_2 & m \end{pmatrix} = 0, \text{ if } 2l_1 + l \text{ is an odd integer}. \tag{24}$$

IV. TIME-DOMAIN CORRELATOR

A. Geometric Delay

We begin by constructing a purely time-domain correlator. The first assumption will be that we can choose a co-ordinate system in which $\vec{r}_i$, the position of the i^{th} instrument, is independent of time, and we do all work in that co-ordinate system. Let us denote the geometric delay transformation as the functional $\mathcal{T}$,

$$\mathcal{T}[g_i(t); \vec{r}_i \cdot \hat{s}] = g_i(t - \vec{r}_i \cdot \hat{s}). \tag{25}$$

---

[2] Overall, this definition of $\mathrm{Y}_{lm}(\theta,\phi)$ is equal to those of Arfken [9], Choi et al. [10], Jackson [11], and MathWorld [12] , however in [9] the factor of $(-1)^m$ appearing here in the definition of $\mathrm{P}_l^m(\cos\theta)$ is removed and added to the coefficients multiplying $\mathrm{P}_l^m(\cos\theta)$ in the definition of $\mathrm{Y}_{lm}(\theta,\phi)$, and in [10] the same movement of the $(-1)^m$ occurs but also a factor of $\sqrt{\frac{(2l+1)(l-m)!}{2(l+m)!}}$ is moved from $\mathrm{Y}_{lm}(\theta,\phi)$ into $\mathrm{P}_l^m(\cos\theta)$.

Given a discretely sampled input time series, $g_i(t_k) = g_i(k\Delta t)$, $k \in \{-\infty, \ldots, 0, \ldots, +\infty\}$, if it was obtained from a band-limited function then by the sampling theorem [16] the original continuous-time function is exactly reconstructed by convolving the samples with the sinc function, $\mathrm{sinc}(x) = \sin(\pi x)/(\pi x)$,

$$g_i(t) = \sum_{k=-\infty}^{\infty} g_i(k\Delta t)\,\mathrm{sinc}[(t - k\Delta t)/\Delta t]. \tag{26}$$

Application of the geometric delay functional transforms this to

$$\mathcal{T}\left[g_i(t); \vec{r}_i \cdot \hat{s}\right] = \sum_{k=-\infty}^{\infty} g_i(k\Delta t)\,\mathrm{sinc}[(t - \vec{r}_i \cdot \hat{s} - k\Delta t)/\Delta t]. \tag{27}$$

Constructing a new discretely-sampled time series by evaluating this at $t_j = j\Delta t$, we obtain the components of the matrix form of the geometric delay operator,

$$\mathcal{T}\left[g_i(j\Delta t); \vec{r}_i \cdot \hat{s}\right] = \sum_{k=-\infty}^{\infty} g_i(k\Delta t)\,\mathrm{sinc}[(j\Delta t - \vec{r}_i \cdot \hat{s} - k\Delta t)/\Delta t] \tag{28}$$

$$= \sum_{k=-\infty}^{\infty} \mathbb{T}_{jk}(\vec{r}_i \cdot \hat{s}) g_i(k\Delta t). \tag{29}$$

Therefore

$$\mathbb{T}_{jk}(\vec{r}_i \cdot \hat{s}) = \mathrm{sinc}(j - k - \vec{r}_i \cdot \hat{s}/\Delta t), \tag{30}$$

where $j, k \in \{-\infty, \ldots, 0, \ldots, +\infty\}$. Although in general this matrix is non-zero everywhere, within each row the magnitude of the elements decays as the inverse of their distance from the diagonal. We will approximate $\mathbb{T}$ by choosing some number of elements, $N_{\mathcal{T}}$, to retain within each row, centred on the diagonal, and set the rest to zero. Therefore, the j^{th} sample of the delayed time series for the direction $\hat{s}$ is

$$g_i(j\Delta t; \hat{s}) = \mathcal{T}\left[g_i(j\Delta t); \vec{r}_i \cdot \hat{s}\right] \approx \sum_{k=j-(N_{\mathcal{T}}-1)/2}^{j+(N_{\mathcal{T}}-1)/2} \mathbb{T}_{jk}(\vec{r}_i \cdot \hat{s}) g_i(k\Delta t), \tag{31}$$

where $N_{\mathcal{T}}$ is a positive odd integer.

From (30), we see that the elements of the geometric delay matrix are oscillatory functions on the sphere with some finite frequency, and so we choose to project them onto the basis of spherical harmonics with the expectation that we will be able to achieve a good approximation to the elements by retaining some finite (and small) number of coefficients. Let us denote the l-m^{th} coefficient of the expansion into spherical harmonics of the dependence of $\mathbb{T}_{jk}(\vec{r}_i \cdot \hat{s})$ on $\hat{s}$ as $\mathbb{T}_{jk}^{(lm)}(\vec{r}_i)$, so

$$\mathbb{T}_{jk}(\vec{r}_i \cdot \hat{s}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \mathbb{T}_{jk}^{(lm)}(\vec{r}_i)\,\mathrm{Y}_{lm}(\hat{s}), \tag{32}$$

where

$$\mathbb{T}_{jk}^{(lm)}(\vec{r}_i) = \int_{\hat{s}} \mathbb{T}_{jk}(\vec{r}_i \cdot \hat{s})\,\mathrm{Y}_{lm}^{*}(\hat{s})\,\mathrm{d}\Omega. \tag{33}$$

Since the instrument location $\vec{r}_i$ is independent of time, $\mathbb{T}_{jk}^{(lm)}(\vec{r}_i)$ is a constant, and since $\mathbb{T}_{jk}(\vec{r}_i \cdot \hat{s})$ is real-valued,

$$\mathbb{T}_{jk}^{(l,-m)}(\vec{r}_i) = (-1)^m \left[\mathbb{T}_{jk}^{(lm)}(\vec{r}_i)\right]^* \tag{34}$$

so it is only necessary to record half of them. The coefficients in the harmonic expansion of $\mathbb{T}_{jk}(\vec{r}_i \cdot \hat{s})$ can be obtained to satisfactory accuracy through numerical integration of (33).

Substituting (32) into (31), we can introduce

$$g_i^{(lm)}(j\Delta t) = \sum_{k=j-(N_{\mathcal{T}}-1)/2}^{j+(N_{\mathcal{T}}-1)/2} \mathbb{T}_{jk}^{(lm)}(\vec{r}_i) g_i(k\Delta t) \tag{35}$$

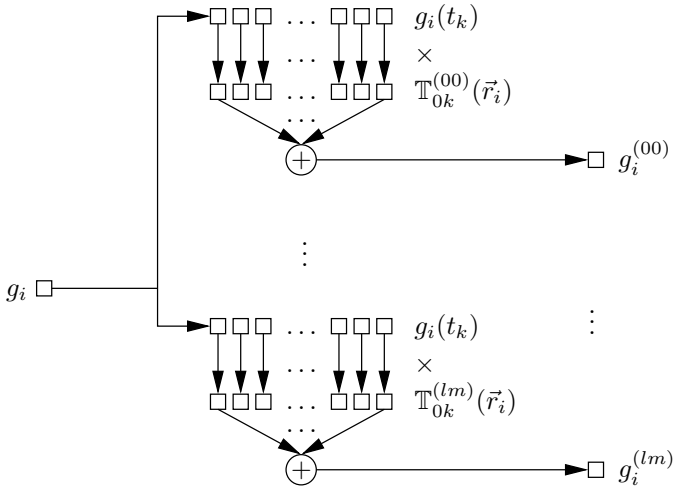

FIG. 1: The time-domain geometric delay algorithm. One-by-one, the input samples of the g_i time series are loaded into shift registers, one register for each spherical harmonic coefficient being computed. After each sample is shifted into a register, the inner product is computed of the shift register's contents with a second register holding the l-m^{th} coefficients of the $j = 0$ row of $\mathbb{T}_{jk}(\vec{r}_i)$. This yields the l-m^{th} coefficient of the (direction-dependent) delayed time series. Each input sample yields exactly one output sample for each coefficient in the spherical harmonic decomposition of the delayed time series, and in this way we construct a description of the delayed time series valid for all directions on the sky. There is a start-up transient that lasts while the shift registers are being populated, after which time the first valid output samples correspond to the time of the central sample in the shift registers.

as the l-m^{th} coefficient of the expansion in spherical harmonics of the $\hat{s}$-dependence of $g_i(j\Delta t; \hat{s})$. This can be thought of as a time series of vectors, in which each vector-valued sample provides the coefficients describing the instantaneous value of a function on the sphere: if one chooses a sample j, and a direction θ, ϕ, and sums the series one obtains the instantaneous value of the time series $g_i(t)$ time-delayed for the chosen direction. This time series of vectors, one time series for each instrument, constitutes the output of the geometric delay section of the time-domain correlator. A graphical depiction of the geometric delay algorithm is shown in Figure 1.

B. Network Power

Having obtained the direction-dependent, geometric-delayed, time series from each of the input time series, we now take their product sample-by-sample to obtain the instantaneous (direction-dependent) cross power in the baseline. The samples whose products we are computing are provided to us as harmonic expansions of functions on the sphere, so to compute the product of two of them we can either

1. transform to the spatial domain, evaluate the product explicitly, and transform back to the harmonic domain, or

2. compute the product in the harmonic domain by convolving the coefficients.

The first approach requires access to a fast spherical harmonic transform (FSHT) algorithm, for example the FSHT of Healy et al. [17]. For functions expanded to order l, they describe algorithms for transforming to and from the spatial domain whose asymptotic complexity scales as $\mathrm{O}(l^2 \log^2 l)$.

We choose the second approach, to compute the product in the harmonic domain, but in the following way. If the co-ordinate origin is placed at the midpoint of the baseline separating the pair of instruments, and oriented so that the baseline lies along the $\hat{z}$ axis, so that $\vec{r}_i = \pm\frac{d}{2}\hat{z}$ for $d = |\vec{r}_1 - \vec{r}_2|$, then $\vec{r}_i \cdot \hat{s} = \pm\frac{d}{2}\cos\theta$ is azimuthally symmetric. Therefore, $\mathbb{T}^{(lm)}_{jk}(\vec{r}_i) = 0$, for $m \neq 0$. This reduces the number of coefficients in an expansion of order l from $(l+1)^2$ to $l+1$, and allows the product to be computed in $\mathrm{O}(l^3)$ operations with a small leading coefficient. If the network of instruments contains only a single baseline, then it is always possible to do this. If the network consists of multiple baselines, then this procedure can be performed for each baseline separately, but there will be a timing error between the baselines on account of their mid-points being in different locations. For example, an impulsive event will be seen

as an increase in power in the baseline farthest from the source slightly in advance of the other baseline.[3] The timing error is no greater than the time required for a wave to propagate from one baseline's mid-point to the next, and if this timing error is small compared to the length of time for which the power will be integrated, then it can be ignored. For Earth-bound gravitational-wave antennas, the power should be integrated over time intervals significantly longer than about 40 ms, otherwise the azimuthally-symmetric approximation should not be used.[4]

Denoting the product of $g_1(t_j;\hat{s})$ and $g_2(t_j;\hat{s})$ on the sphere as

$$G(t_j;\hat{s}) = g_1(t_j;\hat{s})g_2(t_j;\hat{s}), \tag{36}$$

and replacing all three quantities with harmonic expansions, we have

$$\sum_{p=0}^{\infty}\sum_{q=-p}^{p} G^{(pq)}(t_j)\,\mathrm{Y}_{pq}(\hat{s}) = \left[\sum_{l_1=0}^{\infty}\sum_{m_1=-l_1}^{l_1} g_1^{(l_1m_1)}(t_j)\,\mathrm{Y}_{l_1m_1}(\hat{s})\right]\left[\sum_{l_2=0}^{\infty}\sum_{m_2=-l_2}^{l_2} g_2^{(l_2m_2)}(t_j)\,\mathrm{Y}_{l_2m_2}(\hat{s})\right]. \tag{37}$$

If both sides are multiplied by $\mathrm{Y}^*_{lm}(\hat{s})$ and integrated over the sphere, then the orthogonality of the spherical harmonics in (17) used on the left-hand side, and the triple product in (18) on the right, one obtains

$$G^{(lm)}(t_j) = \sum_{l_1=0}^{\infty}\sum_{m_1=-l_1}^{l_1}\sum_{l_2=0}^{\infty}\sum_{m_2=-l_2}^{l_2} g_1^{(l_1m_1)}(t_j)g_2^{(l_2m_2)}(t_j)(-1)^m\sqrt{\frac{(2l_1+1)(2l_2+1)(2l+1)}{4\pi}}\begin{pmatrix} l_1 & l_2 & l \\ m_1 & m_2 & -m\end{pmatrix}\begin{pmatrix} l_1 & l_2 & l \\ 0 & 0 & 0\end{pmatrix}. \tag{38}$$

This is called the solution to the "forward problem" (in the "inverse problem", one knows the coefficients of G and g_1, and obtains those for g_2). If $g_1(t_j;\hat{s})$ and $g_2(t_j;\hat{s})$ are band-limited functions on the sphere, so that $g_1^{(lm)}(t_j) = 0$ for $l > l_{g_1}$, and $g_2^{(lm)}(t_j) = 0$ for $l > l_{g_2}$, then by factoring the sums in (38) and imposing the conditions in (19)–(24) it can be shown that

$$G^{(lm)}(t_j) = (-1)^m\sqrt{\frac{2l+1}{4\pi}}\sum_{l_1=\max\{0,l-l_{g_2}\}}^{\min\{l_{g_1},l+l_{g_2}\}}\sqrt{2l_1+1}\sum_{\substack{l_2=|l-l_1|,\\ l_1+l_2+l\ \text{even}}}^{\min\{l_{g_2},l+l_1\}}\sqrt{2l_2+1}\begin{pmatrix} l_1 & l_2 & l \\ 0 & 0 & 0\end{pmatrix}\times \sum_{m_1=\max\{-l_1,m-l_2\}}^{\min\{+l_1,m+l_2\}} g_1^{(l_1m_1)}(t_j)g_2^{(l_2,(m-m_1))}(t_j)\begin{pmatrix} l_1 & l_2 & l \\ m_1 & m-m_1 & -m\end{pmatrix}. \tag{39}$$

From the limits on the sum over l_1, we see that $G^{(lm)}(t_j) = 0$ for $l > l_G = l_{g_1} + l_{g_2}$. Evaluating (39) is a little better than $\mathrm{O}(l_G^5)$, but if $g_1^{(lm)}(t_j) = 0$ for $m \neq 0$, and $g_2^{(lm)}(t_j) = 0$ for $m \neq 0$, then it can be shown that (39) reduces to

$$G^{(l0)}(t_j) = \sqrt{\frac{2l+1}{4\pi}}\sum_{l_1=\max\{0,l-l_{g_2}\}}^{\min\{l_{g_1},l+l_{g_2}\}}\sqrt{2l_1+1}\sum_{\substack{l_2=|l-l_1|,\\ l_1+l_2+l\ \text{even}}}^{\min\{l_{g_2},l+l_1\}}\sqrt{2l_2+1}\begin{pmatrix} l_1 & l_2 & l \\ 0 & 0 & 0\end{pmatrix}^2 g_1^{(l_10)}(t_j)g_2^{(l_20)}(t_j), \tag{40}$$

which is $\mathrm{O}(l_G^3)$. This is the algorithm we have tested. The evaluation strategy for both (39) and (40) is to fix choices of l_{g_1} and l_{g_2} and then pre-compute a "multiplication plan" as follows. Each coefficient in the result, G^{l0}, is the sum of terms of the form $X_{l_1l_2}g_1^{(l_10)}g_2^{(l_20)}$. The coefficients $g_1^{(l_10)}$ and $g_2^{(l_20)}$ are stored in arrays, so for each output coefficient G^{l0} to be computed, we construct an array of triples of the form $\{\text{offset } 1, \text{offset } 2, X\}$ describing the terms in the sum for that coefficient. The sum is evaluated by iterating over the array of triples, and from each element in the array using the two offsets to retrieve the correct $g_1^{(l_10)}$ and $g_2^{(l_20)}$ from their respective arrays and then adding the product of them and the X value from the triple to an accumulating result for G^{l0}. Each array of triples can be sorted so that the coefficients of the g_i arrays are accessed in a cache-friendly manner. The algorithm is depicted graphically in Figure 2.

---

[3] If done correctly the power increase would be seen simultaneously, but the farther baseline has been moved closer to the source by treating it and the other baseline's midpoints as co-located, so the farther baseline will register a power increase earlier than it should.

[4] For some networks the number may be smaller. For example, for the network consisting of the three LIGO instruments the integration need only be longer than 5 ms before the azimuthally-symmetric approximation starts becoming valid.

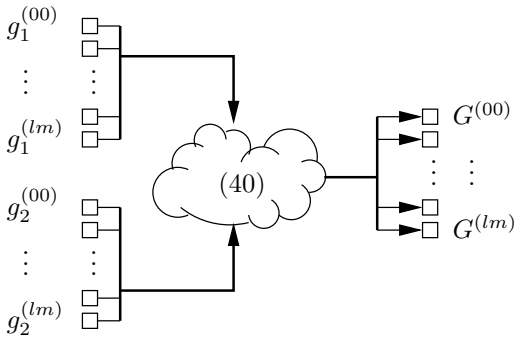

FIG. 2: The time-domain correlator's algorithm for computing the harmonic expansion of the instantaneous baseline "power". At each time step, two complete sets of harmonic coefficients, one for each input time series, are combined via (40), and turned into a set of harmonic coefficients for $G(t_j; \hat{s})$.

C. Summary

The time-domain correlator consists of (35) applied to each input time series to perform the geometric delay, followed by (40) to compute the cross power, followed by a mid-point rule integration to obtain the power over an interval of time.

V. FREQUENCY-DOMAIN CORRELATOR

In the time-domain correlator described above, the geometrically-delayed time series in (27) is obtained by a convolution with a sinc interpolating kernel. This convolution can be implemented using FFT techniques for a significant increase in speed. This modification leads to a frequency-domain correlator, whose input data is provided as frequency series rather than time series.

We begin by considering two discretely-sampled input time series of N samples each, $g_i(k\Delta t)$ for $k \in \{0, \ldots, N-1\}$, and we assume they are periodic. Observing in (30) that the elements $\mathbb{T}_{jk}$ depend only on $j-k$, we can write them as $\mathbb{T}_{jk} = \mathbb{T}_{j-k}$, and replacing time with an integer index, $k\Delta t \to k$, we can write (31) in the form

$$\begin{bmatrix} g_i[0;\hat{s}] \\ \vdots \\ g_i[j;\hat{s}] \\ \vdots \\ g_i[N-1;\hat{s}] \end{bmatrix} = \begin{bmatrix} \mathbb{T}_0(\vec{r}_i\cdot\hat{s}) & \mathbb{T}_{-1}(\vec{r}_i\cdot\hat{s}) & \cdots & \cdots & \mathbb{T}_{+1}(\vec{r}_i\cdot\hat{s}) \\ & & \vdots & & \\ & \cdots & \mathbb{T}_{j-k}(\vec{r}_i\cdot\hat{s}) & \cdots & \\ & & \vdots & & \\ \mathbb{T}_{-1}(\vec{r}_i\cdot\hat{s}) & \cdots & \cdots & \mathbb{T}_{+1}(\vec{r}_i\cdot\hat{s}) & \mathbb{T}_0(\vec{r}_i\cdot\hat{s}) \end{bmatrix} \begin{bmatrix} g_i[0] \\ \vdots \\ g_i[k] \\ \vdots \\ g_i[N-1] \end{bmatrix}. \tag{41}$$

The $\mathbb{T}$ matrix elements' single integer index is $(j-k) \in \{-\frac{1}{2}(N_{\mathcal{T}}-1), \ldots, 0, \ldots, +\frac{1}{2}(N_{\mathcal{T}}-1)\}$, where $N_{\mathcal{T}}$ is an odd integer, and the periodicity of the input data has been used to wrap the rows of the matrix. The elements of the geometric delay matrix are $\mathbb{T}_{j-k} = 0$ for $|j-k| > \frac{1}{2}(N_{\mathcal{T}}-1)$.

In the time-domain case, we clipped the rows of $\mathbb{T}$ to $N_{\mathcal{T}}$ non-zero elements centred on the diagonal simply for the practical need to limit the size of the computation. Here, since the data are assumed periodic, in principal we can compute all the elements of $\mathbb{T}$ and perform the geometric delay exactly. However, we know that in reality we will use data that are not periodic, and have simply been chopped into short segments, transformed to the frequency domain one segment at a time, and fed into the correlator. The periodicity implied by the frequency-domain correlator can be thought of as introducing noise at the boundaries of the input segments. Zeroing the geometric delay matrix's elements for which $|j-k| > \frac{1}{2}(N_{\mathcal{T}}-1)$ introduces numerical errors in the interpolation, but also gives the geometric delay transform a finite impulse response in the time domain. This confines the noise due to the assumption of periodicity to a few samples near the start and end of the output time series. Infact, exactly the first and last $\frac{1}{2}(N_{\mathcal{T}}-1)$ samples are affected by the wrap-around artifacts, and the rest are algebraically identical to the values that are obtained in the purely time-domain correlator. Recall that the time-domain correlator also corrupts the first and last $\frac{1}{2}(N_{\mathcal{T}}-1)$ samples, but in the time-domain case the correlator can process the samples as a continuous stream, and so even in a finite computer the corrupted samples can be made an arbitrarily small fraction of the total number.

Continuing, we write (41) more compactly as

$$g_i[j;\hat{s}] = \sum_{k=0}^{N-1} \mathbb{T}_{j-k}(\vec{r}_i\cdot\hat{s})g_i[k], \tag{42}$$

multiply both sides by $\mathrm{e}^{-2\pi\mathrm{i}jm/N}$, and sum over j to obtain

$$\begin{aligned}\sum_{j=0}^{N-1} g_i[j;\hat{s}]\mathrm{e}^{-2\pi\mathrm{i}jm/N} &= \sum_{j=0}^{N-1}\sum_{k=0}^{N-1} \mathbb{T}_{j-k}(\vec{r}_i\cdot\hat{s})g_i[k]\mathrm{e}^{-2\pi\mathrm{i}jm/N} \\ &= \sum_{k=0}^{N-1} g_i[k]\sum_{j=0}^{N-1} \mathbb{T}_{j-k}(\vec{r}_i\cdot\hat{s})\mathrm{e}^{-2\pi\mathrm{i}jm/N} \\ &= \sum_{k=0}^{N-1} g_i[k]\mathrm{e}^{-2\pi\mathrm{i}km/N}\left(\sum_{j=0}^{N-1} \mathbb{T}_{j-k}(\vec{r}_i\cdot\hat{s})\mathrm{e}^{-2\pi\mathrm{i}(j-k)m/N}\right).\end{aligned} \tag{43}$$

Since $\mathbb{T}_{j-k}$ is periodic, from (11) we recognize the quantity in parentheses on the right-hand-side as $\tilde{\mathbb{T}}_m(\vec{r}_i\cdot\hat{s})$, the inverse DFT of $\mathbb{T}_{j-k}(\vec{r}_i\cdot\hat{s})$,

$$\tilde{\mathbb{T}}_m(\vec{r}_i\cdot\hat{s}) = \sum_{j-k=0}^{N-1} \mathbb{T}_{j-k}(\vec{r}_i\cdot\hat{s})\mathrm{e}^{-2\pi\mathrm{i}(j-k)m/N}. \tag{44}$$

Making that substitution, factoring $\tilde{\mathbb{T}}_m(\vec{r}_i\cdot\hat{s})$ out of the sum, and recognizing the remaining terms in the sum as $\tilde{g}_i[m]$, the inverse DFT of $g_i[k]$, we find that

$$\sum_{j=0}^{N-1} g_i[j;\hat{s}]\mathrm{e}^{-2\pi\mathrm{i}jm/N} = \tilde{\mathbb{T}}_m(\vec{r}_i\cdot\hat{s})\tilde{g}_i[m], \tag{45}$$

which is the convolution theorem. It was important to do this clearly in order to understand that in this result, $\tilde{\mathbb{T}}_m(\vec{r}_i\cdot\hat{s})$ is the inverse DFT of $\mathbb{T}_{j-k}$, but the index $j-k$ increases to the *left* across a row of the matrix. If one computes the values of $\mathbb{T}_{j-k}$ for one row of the matrix, and stores them in an array in order of increasing k (the normal order for components of a matrix), then one will need to frequency invert the inverse DFT of *that* array to obtain the $\tilde{\mathbb{T}}_m(\vec{r}_i\cdot\hat{s})$ that appear here.

Both sides of (45) are a vector whose components are indexed by m. In (45), setting first $i=1$ for instrument 1, then setting $i=2$ for instrument 2 and taking the complex conjugate of this second equation, and then taking the inner product of both sides of the $i=1$ and $i=2$ equations yields

$$\sum_{m=0}^{N-1}\left(\sum_{j_1=0}^{N-1} g_1[j_1;\hat{s}]\mathrm{e}^{-2\pi\mathrm{i}j_1m/N}\sum_{j_2=0}^{N-1} g_2[j_2;\hat{s}]\mathrm{e}^{2\pi\mathrm{i}j_2m/N}\right) = \sum_{m=0}^{N-1}\left(\tilde{\mathbb{T}}_m(\vec{r}_1\cdot\hat{s})\tilde{g}_1[m]\right)\left(\tilde{\mathbb{T}}_m(\vec{r}_2\cdot\hat{s})\tilde{g}_2[m]\right)^*. \tag{46}$$

Dividing by N, and re-organizing the sums, this is equivalent to

$$\sum_{j_1=0}^{N-1}\sum_{j_2=0}^{N-1} g_1[j_1;\hat{s}]g_2[j_2;\hat{s}]\left(\frac{1}{N}\sum_{m=0}^{N-1}\mathrm{e}^{-2\pi\mathrm{i}(j_1-j_2)m/N}\right) = \frac{1}{N}\sum_{m=0}^{N-1}\tilde{g}_1[m]\left(\tilde{\mathbb{T}}_m(\vec{r}_1\cdot\hat{s})\tilde{\mathbb{T}}^*_m(\vec{r}_2\cdot\hat{s})\right)\tilde{g}^*_2[m]. \tag{47}$$

The quantity in parentheses on the left-hand-side is the Kronecker δ in (12), $\delta_{j_1j_2}$, and so

$$\frac{1}{N}\sum_{j=0}^{N-1} g_1[j;\hat{s}]g_2[j;\hat{s}] = \frac{1}{N^2}\sum_{m=0}^{N-1}\tilde{g}_1[m]\left(\tilde{\mathbb{T}}_m(\vec{r}_1\cdot\hat{s})\tilde{\mathbb{T}}^*_m(\vec{r}_2\cdot\hat{s})\right)\tilde{g}^*_2[m], \tag{48}$$

where we have divided by a second factor of N. The left-hand-side of (48) is the discrete form of the cross-correlation integral, (1), and so we have derived the frequency-domain correlator.

Alternatively, one can recognize that in (45), the left-hand side is the inverse DFT, (11), of the geometrically-delayed time series. It has been left as a sum over the time-domain samples so that in (48) we can see clearly how the left-hand

side is the integrated cross power. However, writing the left-hand side of (45) as $\tilde{g}_i[m;\hat{s}]$, the Fourier transform of the direction-dependent geometrically delayed time series, (48) becomes

$$\frac{1}{N}\sum_{m=0}^{N-1}\tilde{g}_1[m;\hat{s}]\tilde{g}_2^*[m;\hat{s}] = \frac{1}{N^2}\sum_{m=0}^{N-1}\tilde{g}_1[m]\left(\tilde{\mathbb{T}}_m(\vec{r}_1\cdot\hat{s})\tilde{\mathbb{T}}_m^*(\vec{r}_2\cdot\hat{s})\right)\tilde{g}_2^*[m]. \tag{49}$$

The left-hand sides of (48) and (49) are related by Parseval's theorem: the square integral can be computed in the time domain, (48), or the frequency domain, (49). The quantity $\tilde{g}_1[m;\hat{s}]\tilde{g}_2^*[m;\hat{s}]$ is then recognized as the cross power in the baseline for the time interval spanned by the original time series, in the frequency bin m. And, of course, it is a vector of harmonic coefficients describing the angular distribution of that quantity on the sky. Therefore, the frequency-domain correlator can, itself, be used to do a direction-resolved time-frequency decomposition of the sky.

The evaluation strategy for (48) is as follows. Again, as in the time-domain case, we write the direction-dependence ($\hat{s}$-dependence) of the delay matrix $\mathbb{T}(\vec{r}_i\cdot\hat{s})$ by expanding it as the harmonic series in (32). We begin by computing the top row of the matrix in (41), the one in which the element whose index is $j-k=0$ is in the first position.[5] Each element in the row is a vector containing the coefficients of the harmonic expansion of that element on the sphere. We inverse DFT the row, where the DFT is of the related coefficients in the vectors: the monopole coefficients for the row are taken as a time series, inverse DFT'ed, and so on. The inverse DFT is then frequency-inverted to give the components of $\tilde{\mathbb{T}}_m(\vec{r}_i\cdot\hat{s})$, $m\in\{0,\ldots,N-1\}$. Each element is a vector containing the coefficients of its harmonic expansion on the sphere. All of this is repeated for both instruments, and then element-by-element the product of the two arrays is computed using either (39) or (40) as appropriate (depending on whether or not azimuthal symmetry is being used). That product is $\tilde{\mathbb{T}}_m(\vec{r}_1\cdot\hat{s})\tilde{\mathbb{T}}_m^*(\vec{r}_2\cdot\hat{s})$, with each element m expanded as a vector of harmonic coefficients, and we store the entire thing for repeated use during correlation.

The correlation algorithm is then nearly trivial. For each segment of N samples, the two input time series are inverse DFT'ed to give $\tilde{g}_1[m]$ and $\tilde{g}_2^*[m]$, and the two of them multiplied together, element-by-element. The inner product of the resulting vector and the vector of pre-computed $\tilde{\mathbb{T}}_m(\vec{r}_1\cdot\hat{s})\tilde{\mathbb{T}}_m^*(\vec{r}_2\cdot\hat{s})$ components yields the baseline's direction-dependent integrated cross power (up to a factor of $1/N^2$), in the form of a single vector of the coefficients of its expansion as a series of spherical harmonics.

A. Summary

The frequency domain correlator is given by (48), with $\tilde{\mathbb{T}}_m(\vec{r}_i\cdot\hat{s})$ given by (44), which is evaluated once for each coefficient in an expansion of $\mathbb{T}_{j-k}(\vec{r}_i\cdot\hat{s})$ in spherical harmonics. The price that must be paid for using the frequency-domain correlator is

- One must choose in advance the length of data to be processed. It cannot be processed in a "streaming mode" as with the purely time-domain correlator, and this can make it difficult to search for transient events of unknown durations.

However, the frequency-domain correlator offers the following benefits over the time-domain correlator:

- Typically the correlator is preceded by some form of data conditioning or filtering that is naturally performed in the frequency domain (for example, normalizing to a noise spectrum), and so the input data is often available as a frequency series for free.
- As stated above, the geometric delay convolution is significantly faster in the frequency domain than in the time-domain, not only because the transformation becomes "diagonal", but also because knowing the length of the transform allows the two geometric delay matrices to be multiplied in advance.
- Since the delay matrices are pre-multiplied, this mostly eliminates the motivation for the "azimuthal symmetry" approximation, allowing arbitrary instrument locations to be used, which in turn eliminates any concern over inter-baseline timing offsets.
- Finally, there is the ability of the frequency-domain correlator to provide a time-frequency decomposition of the cross-correlated power, which may be a more desirable output for the correlator to provide.

---

[5] In practise, it can be easier to work with a different row, for example the row whose left-most non-zero element is in the first position. One can choose any row to transform to the frequency domain, but then rotate the phases so that the result is equal to the inverse DFT of the top row.

VI. CHANGE OF BASIS

The output of either of the correlators described above is expressed in the co-ordinate system of $\hat{s}$, which we've chosen to align with the baseline. It is likely one would prefer to know the result in a different co-ordinate system, one that is obtained by a rotation from the co-ordinate system of the correlator, for example a sky-fixed co-ordinate system. Consider a rotation matrix, $\mathbb{R}$, that rotates one set of co-ordinate basis vectors, $\{\hat{e}_i\}$, to a new set of co-ordinate basis vectors, $\{\hat{e}'_i\}$, by

$$\hat{e}'_i = \mathbb{R}\hat{e}_i. \tag{50}$$

It can be shown that a spherical harmonic rotated on the sphere by $\mathbb{R}$ can be written as a linear combination of the nonrotated spherical harmonics of the same order [18]. That is,[6]

$$\mathrm{Y}'_{lm}(\theta,\phi) = \sum_{m'=-l}^{l} D^l_{mm'}\,\mathrm{Y}_{lm'}(\theta,\phi), \tag{51}$$

where the function Y'_{lm} on the left is the standard l-m^{th} spherical harmonic in the rotated co-ordinate system written as a function of the co-ordinates θ and ϕ in the original co-ordinate system.

Given a function on the sphere expressed as a series of spherical harmonics, $g(\hat{s}) = \sum_{l=0}^{\infty}\sum_{m=-l}^{l} g_{lm}\,\mathrm{Y}_{lm}(\hat{s})$, we wish to know the coefficients, g'_{lm}, for the expansion of the function $g'(\hat{s})$ that is the original function rotated on the sphere by $\mathbb{R}$. This problem is solved by first writing $g'(\hat{s})$ with respect to the transformed harmonics, Y'_{lm}, and then substituting (51) into each term of the sum and solving for g'_{lm}, by which process one finds that

$$g'_{lm'} = \sum_{m=-l}^{l} g_{lm} D^l_{mm'}. \tag{52}$$

Note the difference, relative to (51), in the index over which the sum is performed.

The complex matrix $\mathbb{D}^l$, whose components are $D^l_{mm'}$, is called a "Wigner D matrix". A D matrix is a function of the rotation, $\mathbb{R}$, and there is an extensive body of literature on the problem of the quick and accurate numerical calculation of its components. There is no need to reproduce this information here, only to say that stable and accurate algorithms exist that can compute a set of $\mathbb{D}^l$ matrices upto some order l in linear time in the total number of components, that is each component can be computed in O(few) operations. The total number of components upto order l scales as $\mathrm{O}(l^3)$, which sets the scaling for both the time required to compute the matrices and the time required to apply them. For a specific implementation, we have used the algorithm presented in [10], and find it to be to fast and accurate, and easily implemented from the description in that paper.[7]

VII. LINEARITY OF THE CORRELATOR

Here, we establish the conditions under which the cross-correlation integral (1) is linear in the sources of incident radiation. By "linear in the sources", it is meant that if when the incident radiation is the result of sources A we measure $\xi^A_{ij}(\hat{s})$, and when the incident radiation is the result of sources B we measure $\xi^B_{ij}(\hat{s})$, then when the incident radiation is the result of sources A and B together, we will measure $\xi^{A+B}_{ij}(\hat{s}) = \xi^A_{ij}(\hat{s}) + \xi^B_{ij}(\hat{s})$. When this is true, the correlator's response to a source is not influenced by the presence or otherwise of other sources elsewhere on the sky — the brightness image of the sky is the sum of the images from each source individually.

We assume that both the wave equation for the radiation field and the instrument transducers are linear (including in the "transducers" any data conditioning and/or filtering that is applied to the data before providing it as input to the correlator), and so sources A yield correlator inputs $g^A_i(t)$, sources B yield correlator inputs $g^B_i(t)$, and sources $A+B$ yield correlator inputs $g^{A+B}_i(t) = g^A_i(t) + g^B_i(t)$, for detector i.

The geometric delay functional in (25) is both linear and invertible, and therefore possess a Green's function, $G_{\mathcal{T}}(t,t';\vec{r}_i\cdot\hat{s})$, such that[8]

$$\mathcal{T}[g(t);\vec{r}_i\cdot\hat{s}] = \int_{t'=-\infty}^{+\infty} G_{\mathcal{T}}(t,t';\vec{r}_i\cdot\hat{s})g(t')\,\mathrm{d}t'. \tag{53}$$

---

[6] The components $D^l_{mm'}$ are defined here transposed with respect to the m indexes relative to [10, equation (1.1)].
[7] Note however that we believe that their equation (5.5) is incorrect by an over-all sign on the right-hand side.
[8] Infact, $G_{\mathcal{T}}(t,t';\tau) = \delta(t-\tau-t')$.

From the transform's linearity,

$$\begin{aligned}\xi_{ij}^{A+B}(\hat{s}) &= \int \mathcal{T}\left[g_i^{A+B}(t);\vec{r}_i\cdot\hat{s}\right]\mathcal{T}\left[g_j^{A+B}(t);\vec{r}_j\cdot\hat{s}\right]\,\mathrm{d}t\\ &= \xi_{ij}^{A}(\hat{s}) + \xi_{ij}^{B}(\hat{s}) + \int \mathcal{T}\left[g_i^{A}(t);\vec{r}_i\cdot\hat{s}\right]\mathcal{T}\left[g_j^{B}(t);\vec{r}_j\cdot\hat{s}\right]\,\mathrm{d}t + \int \mathcal{T}\left[g_i^{B}(t);\vec{r}_i\cdot\hat{s}\right]\mathcal{T}\left[g_j^{A}(t);\vec{r}_j\cdot\hat{s}\right]\,\mathrm{d}t. \end{aligned} \tag{54}$$

We now assume

1. that the sources A and the sources B emit radiations that are not correlated with one another, that is that
$$\left\langle g_i^A(t_1) g_j^B(t_2)\right\rangle = 0, \tag{55}$$
2. and that the $g_i(t)$ are stationary and ergodic, that is that measurements of statistical properties, such as $\langle g_i(t)\rangle$, are independent of the choice of t, and that ensemble averages are equivalent to time averages.

The assumption of ergodicity allows the integrals over time in (54) to be replaced with ensemble averages. For example,

$$\int \mathcal{T}\left[g_i^{A}(t);\vec{r}_i\cdot\hat{s}\right]\mathcal{T}\left[g_j^{B}(t);\vec{r}_j\cdot\hat{s}\right]\,\mathrm{d}t = \left\langle \mathcal{T}\left[g_i^{A}(t);\vec{r}_i\cdot\hat{s}\right]\mathcal{T}\left[g_j^{B}(t);\vec{r}_j\cdot\hat{s}\right]\right\rangle. \tag{56}$$

Consider, for a moment, $\left\langle \mathcal{T}\left[g_i^{A}(t_1);\vec{r}_i\cdot\hat{s}\right]\mathcal{T}\left[g_j^{B}(t_2);\vec{r}_j\cdot\hat{s}\right]\right\rangle$. Using (53) and then (55),

$$\begin{aligned}\left\langle \mathcal{T}\left[g_i^{A}(t_1);\vec{r}_i\cdot\hat{s}\right]\mathcal{T}\left[g_j^{B}(t_2);\vec{r}_j\cdot\hat{s}\right]\right\rangle &= \left\langle \int_{t_1'=-\infty}^{\infty} G_{\mathcal{T}}(t_1,t_1';\vec{r}_i\cdot\hat{s}) g_i^A(t_1')\,\mathrm{d}t_1' \cdot \int_{t_2'=-\infty}^{\infty} G_{\mathcal{T}}(t_2,t_2';\vec{r}_j\cdot\hat{s}) g_j^B(t_2')\,\mathrm{d}t_2' \right\rangle\\ &= \int_{t_1'=-\infty}^{\infty}\int_{t_2'=-\infty}^{\infty} G_{\mathcal{T}}(t_1,t_1';\vec{r}_i\cdot\hat{s}) G_{\mathcal{T}}(t_2,t_2';\vec{r}_j\cdot\hat{s}) \left\langle g_i^A(t_1') g_j^B(t_2')\right\rangle\,\mathrm{d}t_1'\,\mathrm{d}t_2'\\ &= 0, \end{aligned} \tag{57}$$

but if this is true for any t_1 and t_2 then it is true for $t_1 = t_2 = t$, and so

$$\left\langle \mathcal{T}\left[g_i^{A}(t);\vec{r}_i\cdot\hat{s}\right]\mathcal{T}\left[g_j^{B}(t);\vec{r}_j\cdot\hat{s}\right]\right\rangle = 0. \tag{58}$$

Similarly for A and B reversed, and so

$$\xi_{ij}^{A+B}(\hat{s}) = \xi_{ij}^{A}(\hat{s}) + \xi_{ij}^{B}(\hat{s}), \tag{59}$$

and the correlator is linear in the sources.

In gravitational wave astronomy, the assumptions 1 and 2 above are likely to not always hold. For example, if two binary compact object coalescences occur in different parts of the sky at nearly the same time, and the component masses in the one are similar to those in the other, then assumption 1 will be violated (the waveforms will be correlated). Of course, on average we don't expect events like these to be synchronized in different parts of the sky, but when they happen to be synchronized the correlator will not respond linearly to them — the correlator will exhibit "source confusion".

VIII. ANGULAR RESOLUTION

The image synthesis algorithm is linear in the sources, so we can answer the question of the angular resolution of $\xi_{ij}(\hat{s})$ by considering the response of a single pair of instruments to a single point source. We'll consider, specifically, sources far from the network of detectors. Without loss of generality, we place instrument 1 at $\vec{r}_1 = +\frac{d}{2}\hat{z}$, and instrument 2 at $\vec{r}_2 = -\frac{d}{2}\hat{z}$. Any other baseline can be transformed to this one by a rotation, and an irrelevant (the sources being at infinity) translation of the mid-point. If a single point source is located at colatitude θ_s, and emits radiation seen as the function $g(t)$ at the co-ordinate origin, then the detector outputs are

$$g_1(t) = g\left(t + \frac{d}{2}\cos\theta_s\right), \tag{60a}$$

$$g_2(t) = g\left(t - \frac{d}{2}\cos\theta_s\right). \tag{60b}$$

Since $\vec{r}_1 \cdot \hat{s} = \frac{d}{2}\cos\theta$, and $\vec{r}_2 \cdot \hat{s} = -\frac{d}{2}\cos\theta$, the cross-correlation integral in (1) is

$$\xi_{1,2}(\theta;\theta_s) = \int_{-\infty}^{+\infty} g\left[t - \frac{d}{2}(\cos\theta - \cos\theta_s)\right] g\left[t + \frac{d}{2}(\cos\theta - \cos\theta_s)\right] \mathrm{d}t. \tag{61}$$

Defining

$$\zeta(\theta;\theta_s) = d(\cos\theta - \cos\theta_s), \tag{62}$$

and making the change of variables $t \to t - \zeta/2$,

$$\xi_{1,2}(\zeta) = \int_{-\infty}^{+\infty} g(t-\zeta)g(t)\,\mathrm{d}t. \tag{63}$$

The quantity ζ is an angular co-ordinate for the sky that measures direction in units of the relative signal delay in the two instruments; from (62), $\zeta = 0$ corresponds to the celestial colatitude of the source. Application of the Wiener-Khinchin theorem in (8) to (63), and assuming $g(t)$ to be real, so that $S_g(f) = S_g(-f)$,

$$\xi_{1,2}(\zeta) = 2\int_0^{\infty} S_g(f)\cos(2\pi f\zeta)\,\mathrm{d}f. \tag{64}$$

This is an expression of the well-known relationship between the angular resolution of an imaging system, the diameter of its aperture, and the spectrum of the radiation being imaged.

Since $\xi_{1,2}(\zeta)$ and $S_g(f)$ are Fourier conjugates of one-another, we observe that in general the higher the frequency content of the the spectral density function, $S_g(f)$, the more narrowly peaked is $\xi_{1,2}(\zeta)$ about $\zeta = 0$ — that is, the better the detector network is able to identify the direction to the source of the incident radiation. For a flat spectral density function with a centre frequency of f_0 and a bandwidth of Δf,

$$S_g(f) = \begin{cases} S_g & \text{if } f_0 - \frac{1}{2}\Delta f \le |f| \le f_0 + \frac{1}{2}\Delta f, \\ 0 & \text{otherwise,} \end{cases} \tag{65}$$

the correlation function is

$$\xi_{1,2}(\zeta) = S_g \Delta f\,\mathrm{sinc}\,(\zeta\Delta f)\cos(2\pi\zeta f_0)\,. \tag{66}$$

For a monochromatic spectrum with frequency f_0,

$$S_g(f) = S_g\delta(|f| - f_0), \tag{67}$$

the correlation function is

$$\xi_{1,2}(\zeta) = S_g\cos(2\pi\zeta f_0)\,. \tag{68}$$

These results can be used to test the time- and frequency-domain correlators by comparing their numerical outputs to these analytic results for injections with these spectra.

The correlation function for a monochromatic source in (68) has peaks occurring at a frequency of f_0 in ζ. From (62), $\zeta(\theta)$ varies most rapidly as a function of celestial colatitude in the vicinity of $\theta = \frac{\pi}{2}$, so placing the source at $\theta_s = \frac{\pi}{2}$ in order for there to be a peak in $\xi_{1,2}(\zeta)$ there, we find that $\xi_{1,2}(\zeta)$'s zero-crossings at $\zeta = \pm\frac{1}{4f_0}$ correspond to celestial colatitudes of

$$\theta = \frac{\pi}{2} \pm \sin^{-1}\frac{1}{4f_0 d}. \tag{69}$$

Since we are considering digital correlators, the data have been discretely sampled at an interval Δt, so the highest possible frequency component in the correlator's input spectrum is the Nyquist frequency at $1/(2\Delta t)$. Inserting this frequency into (69), we find that the correlator's input spectrum can result in the cross-correlated power exhibiting zero-crossings on the sphere no closer to one-another than

$$\Delta\theta = 2\sin^{-1}\frac{\Delta t}{2d}. \tag{70}$$

IX. PARAMETER SELECTION

A. Size of $\mathbb{T}$

In approximating $\mathbb{T}$ by setting to zero elements far from the diagonal, we wish to know how many elements, $N_{\mathcal{T}}$, should be retained. Before doing so, we should answer the question of what those elements should be, for although the "sinc weights" are the correct choice in the all-time case, they may not be the best choice in the case of an $N_{\mathcal{T}}$-point interpolation. For functions analytic in a strip of the complex plane containing the real line, and decaying to 0 no slower than exponentially in the vicinity of positive and negative infinity, retaining the original "sinc weights" as the coefficients of the $N_{\mathcal{T}}$-point interpolation is nearly optimal, that is it can be shown that no choice of coefficients can be made that will reduce the maximum error of the interpolation below that of the sinc interpolation's by more than an amount that decays exponentially in $N_{\mathcal{T}}$ [19], [20]. Stationary Gaussian noise does not satisfy the optimality conditions since it does not decay, but if we consider only the signal part of the correlator's input, and assume that the signal obeys the interpolation's optimality criteria (which for a physical signal is almost certainly the case), we can argue in favour of the continued use of sinc interpolation based on its ability to near-optimally reconstruct the non-noise component of the correlator's input.

The error bound for sinc interpolation is that there exists a constant, C'_p, such that for an $N_{\mathcal{T}}$-point interpolation the error is

$$E \leq C'_p \sqrt{N_{\mathcal{T}}} \exp\left(-\frac{\pi}{2}\sqrt{\frac{N_{\mathcal{T}}}{p/(p-1)}}\right), \tag{71}$$

where $1 < p \leq \infty$ is a measure of the decay rate of the interpolated function, and is not important here [20]. The constant C'_p is not stated a priori, so it is not possible to choose $N_{\mathcal{T}}$ to obtain a desired error bound E, but (71) can tell us by how much $N_{\mathcal{T}}$ should be increased in order to reduce an observed error by a desired amount.

Nevertheless, an instrument's distance from the origin (network's phase centre) places a lower bound on $N_{\mathcal{T}}$. For an instrument at $\vec{r}$, to delay-correct its output to the network's phase centre can require a delay with a magnitude as large as $|\vec{r}|$. The part of the sky closest to the instrument corresponds to a delay correction of $+|\vec{r}|$, and the part of the sky farthest from the instrument a delay correction of $-|\vec{r}|$. If the sample period is Δt, then the range of delays implies a minimum of

$$N_{\mathcal{T}} \geq 2\left\lceil \frac{|\vec{r}|}{\Delta t} \right\rceil + 1 \tag{72}$$

elements in the geometric delay matrix, with (an even number of) additional elements added as needed to improve accuracy. In (72), the ceiling of the ratio of the instrument's distance from the network's phase centre to the sample period is taken in order to obtain an integer, and 1 is added to make the total odd.

B. Expansion of $\mathbb{T}_{jk}(\vec{r}_i \cdot \hat{s})$

In representing the direction dependence of $\mathbb{T}_{jk}(\vec{r}_i \cdot \hat{s})$ by expanding it as a series in $\mathrm{Y}_{lm}(\hat{s})$, we wish to know the order at which the expansion can be truncated. Rather than carefully deriving error bounds, we'll solve this by simply finding the harmonic order having the same density of zero-crossings on the sphere as the target function, with the knowledge that having not done this carefully we will likely need to adjust the result up or down a little. From (30), $\mathbb{T}_{jk}(\vec{r}_i \cdot \hat{s})$ goes through zero whenever $\vec{r}_i \cdot \hat{s}/\Delta t$ is an integer. If γ is the angle between $\vec{r}_i$ and $\hat{s}$, then $\vec{r}_i \cdot \hat{s}/\Delta t = |\vec{r}_i|\,\Delta t^{-1} \cos\gamma$ is changing most rapidly near $\gamma = \pm\pi/2$ where $|\vec{r}_i|\,\Delta t^{-1}\cos\gamma \approx \mp\,|\vec{r}_i|\,\Delta t^{-1}\gamma$, respectively. Therefore, the smallest change in direction, $\Delta\gamma$, that sees $\vec{r}_i \cdot \hat{s}/\Delta t$ change by 1, in other words the angle between the most closely-spaced zero-crossings of the matrix element $\mathbb{T}_{jk}(\vec{r}_i \cdot \hat{s})$, is $\Delta\gamma = \Delta t/\,|\vec{r}_i|$. Therefore, we should expect a decomposition of $\mathbb{T}_{jk}(\vec{r}_i \cdot \hat{s})$ into spherical harmonics to require expansion upto order $l_{\mathcal{T}} \sim \lceil \pi/\Delta\gamma \rceil$, or

$$l_{\mathcal{T}} \sim \left\lceil \frac{\pi\,|\vec{r}_i|}{\Delta t} \right\rceil. \tag{73}$$

C. Expansion of $\xi_{ij}(\hat{s})$

We wish to know the order at which a spherical harmonic decomposition of the sky can be truncated, and again we'll do this by finding the harmonic order whose zero-crossings occur at the same spatial frequency as the target

function. From (70), we see that a pair instruments producing data sampled at a rate Δt cannot yield an observed angular distribution of cross power having zero-crossings more closely spaced than $\Delta\theta = 2\sin^{-1}\Delta t/(2d)$, where d is the distance between the two instruments. Of all pairs of instruments in the network, the pair with the greatest distance between them, $d = d_{\max}$, yields the most closely-spaced zero-crossings, $\Delta\theta_{\min}$. For the spherical harmonic $\mathrm{Y}_{lm}(\hat{s})$, zero crossings are separated by an angle of about π/l. Therefore, we expect a decomposition of the correlation function $\xi_{1,2}(\hat{s})$ into spherical harmonics to require expansion upto order $l_\xi \sim \lceil \frac{\pi}{\Delta\theta_{\min}} \rceil$, or

$$l_\xi \sim \left\lceil \frac{\pi}{2\sin^{-1}\Delta t/(2d_{\max})} \right\rceil . \tag{74}$$

X. TESTS

Both the time-domain and frequency-domain correlators described above have been implemented in C using double-precision arithmetic and tested for speed and accuracy on a 32-bit Pentium-M based system. The tests have been designed to be relevant to the application of the algorithm to gravitational wave interferometry, and so the choices of baseline length, sample rate, etc., are made in that context.

A. Accuracy

The numerical accuracy of the implementation is tested by injecting Gaussian white noise and comparing the correlator's output to the analytic result in (66). The comparison is performed by numerically expanding the analytic result in spherical harmonics, then subtracting from that the vector of harmonic coefficients produced by the correlator. This gives the vector of harmonic coefficients for the error in the correlator output,

$$\Delta\xi^{(lm)} = \xi^{(lm)}_{\text{exact}} - \xi^{(lm)}. \tag{75}$$

Taking the inner product of that vector with its own complex conjugate gives the square integral over the sphere of the error in the cross power,

$$\int_{\hat{s}} \left|\xi_{\text{exact}}(\hat{s}) - \xi(\hat{s})\right|^2 \,\mathrm{d}^2\Omega = \sum_{l,m} \Delta\xi^{(lm)} \left(\Delta\xi^{(lm)}\right)^* , \tag{76}$$

so dividing by 4π rad^2 and taking the square root gives the RMS error in the cross power integrated over the sphere,

$$\Delta\xi_{\text{RMS}} = \sqrt{\frac{1}{4\pi}\sum_{l,m} \Delta\xi^{(lm)} \left(\Delta\xi^{(lm)}\right)^*} . \tag{77}$$

We can divide this by the maximum value of ξ_{exact} on the sphere (which equals the power incident upon the network from the source) to measure the RMS error in the correlator's cross power as a fraction of the injected power,

$$\Delta\bar{\xi}_{\text{RMS}} = \frac{1}{\max\xi_{\text{exact}}} \sqrt{\frac{1}{4\pi}\sum_{l,m} \Delta\xi^{(lm)} \left(\Delta\xi^{(lm)}\right)^*} . \tag{78}$$

The sum over l in (78) can be split into two parts to make a more precise statement about the origin of the error in the correlator's output. If the correlator produces harmonic coefficients for orders $l \le l_\xi$, then we can write

$$\Delta\bar{\xi}_{\text{RMS}} = \sqrt{\Delta\bar{\xi}^2_{\mathcal{T}\text{RMS}} + \Delta\bar{\xi}^2_{l_\xi\text{RMS}}} , \tag{79}$$

where

$$\Delta\bar{\xi}_{l_\xi\text{RMS}} = \frac{1}{\max\xi_{\text{exact}}} \sqrt{\frac{1}{4\pi}\sum_{\substack{m \\ l>l_\xi}} \Delta\xi^{(lm)} \left(\Delta\xi^{(lm)}\right)^*} , \tag{80}$$

$$\Delta\bar{\xi}_{\mathcal{T}\text{RMS}} = \frac{1}{\max\xi_{\text{exact}}} \sqrt{\frac{1}{4\pi}\sum_{\substack{m \\ l\le l_\xi}} \Delta\xi^{(lm)} \left(\Delta\xi^{(lm)}\right)^*} . \tag{81}$$

						Time Domain		Frequency Domain		
Δt	$N_{\mathcal{T}}$	$l_{\mathcal{T}}$	l_ξ	θ_s	$\Delta\bar{\xi}_{l_\xi\text{RMS}}$	$\Delta\bar{\xi}_{\mathcal{T}\text{RMS}}$	$\Delta\bar{\xi}_{\text{RMS}}$	$\Delta\bar{\xi}_{\mathcal{T}\text{RMS}}$	$\Delta\bar{\xi}_{\text{RMS}}$	Figure
512^{-1} Hz^{-1}	7	9	17	0 rad	−23 dB	−10 dB	−10 dB	−10 dB	−10 dB	3(a)
				$\pi/4$ rad	−24 dB	−11 dB	−11 dB	−11 dB	−11 dB	3(b)
				$\pi/2$ rad	−21 dB	−14 dB	−14 dB	−14 dB	−14 dB	3(c)
	127	13	19	0 rad	−29 dB	−27 dB	−26 dB	−27 dB	−26 dB	4(a)
				$\pi/4$ rad	−32 dB	−27 dB	−26 dB	−27 dB	−26 dB	4(b)
				$\pi/2$ rad	−28 dB	−27 dB	−26 dB	−27 dB	−26 dB	4(c)
4096^{-1} Hz^{-1}	43	65	129	0 rad	−31 dB	−11 dB	−11 dB	−11 dB	−11 dB	
				$\pi/4$ rad	−30 dB	−10 dB	−10 dB	−10 dB	−10 dB	
				$\pi/2$ rad	−25 dB	−10 dB	−10 dB	−10 dB	−10 dB	
	163	69	131	0 rad	−34 dB	−26 dB	−26 dB	−26 dB	−26 dB	5(a)
				$\pi/4$ rad	−32 dB	−27 dB	−27 dB	−27 dB	−27 dB	5(b)
				$\pi/2$ rad	−28 dB	−28 dB	−26 dB	−28 dB	−26 dB	5(c)

TABLE I: Correlator fractional errors (RMS over the sphere) for a Gaussian white noise injection sampled at a variety of frequencies, for a variety of correlator parameter choices, and for a variety of baseline orientations. In all cases, the baseline is 10 ms long and centred on the origin. For each sample rate, the first entries in the table correspond to the correlator parameter values in (72), (73), and (74). The column labeled θ_s is the angle between the baseline's axis and the point source of injected noise.

These two errors are added in quadrature to obtain the total fractional RMS error. The first, $\Delta\bar{\xi}_{l_\xi\text{RMS}}$, is the sum of contributions to the error from coefficients the correlator does not compute, and so is the contribution to the fractional RMS error due to l_ξ being too small. The second, $\Delta\bar{\xi}_{\mathcal{T}\text{RMS}}$, is the error due to the correlator computing incorrectly those harmonic coefficients it does produce — the error due to the delay transformation $\mathcal{T}$ being inaccurate. Round-off errors are negligible, so two things contribute to inaccuracies in the delay transformation: the interpolation matrix having too few elements, $N_{\mathcal{T}}$, and the matrix elements being expanded to too low an harmonic order, $l_{\mathcal{T}}$. In the numerical tests done here, no attempt has been made to isolate these two effects although in principle it is possible to do so since the error in each delay matrix element due to its truncation at finite harmonic order can be measured directly.

We now consider a single baseline in the x-y plane, 10 ms in length, centred on the origin. We place a source of stationary Gaussian white noise with a variance of 1 (arbitrary units) on the x axis. The fractional RMS errors for a variety of baseline orientations and sampling rates are summarized in Table I, and plots of the correlator's output can be seen in Figures 3, 4, and 5. For each sample rate listed in Table I, the first values listed for the three correlator parameters $N_{\mathcal{T}}$, $l_{\mathcal{T}}$, and l_ξ, are the values obtained from (72), (73), and (74) respectively. Recalling that in all three cases those formulae were our estimate of the minimum required value, it is not surprising to see somewhat poor accuracy in the correlator's output — errors at about the 10% level. Nevertheless, Figure 3 shows that as we had hoped, these values for the parameters *do* succeed in capturing the character of the solution: maxima and minima are not far from the correct locations and have not unreasonable amplitudes, and in particular the maximum associated with the injection is recovered quite well. It is interesting to observe in Table I that for the "minimum" parameter values the error is dominated by the poor quality of the geometric delay transformation. This is perhaps expected since stationary Gaussian noise fails the convergence criteria for finite sinc interpolation.

For the 512 Hz Gaussian noise, increasing the harmonic order of the cross power, l_ξ, by 2 gets the phases of the maxima and minima correct, and experimentation shows that increasing the harmonic order of the geometric delay matrix elements by 4 and adding 120 elements to each row of the matrix balances the errors so that the error associated with no one of the parameters dominates. These choices might be thought of as the minimum "reasonable" values, in that they are the minimum required to obtain what is essentially the correct solution: the correct number of maxima and minima, and errors balanced across the various sources. With these parameters, the numerical errors in the correlator hover around −26 dB relative to the peak cross power in the baseline. The parameter values can be increased for greater accuracy, but speed suffers, and 26 dB of signal-to-noise is probably considered adequate in most applications.

A second set of tests was performed with the sample rate increased to 4096 Hz. The parameter values specified in (72)–(74) again result in numerical errors at about the 10% level. This time, however, rather than experimenting with other parameter values, the adjustments made in the 512 Hz case were simply applied here again to test the appropriateness of these adjustments as a general "rule of thumb": l_ξ was increased by 2, $l_{\mathcal{T}}$ was increased by 4, and $N_{\mathcal{T}}$ was increased by 120. As can be seen in Table I, this modification again reduces the errors to the −26 dB level

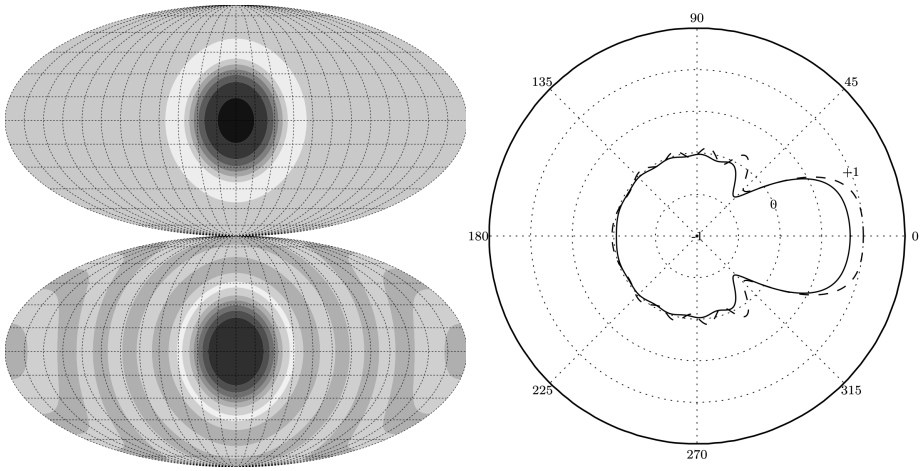

(a)

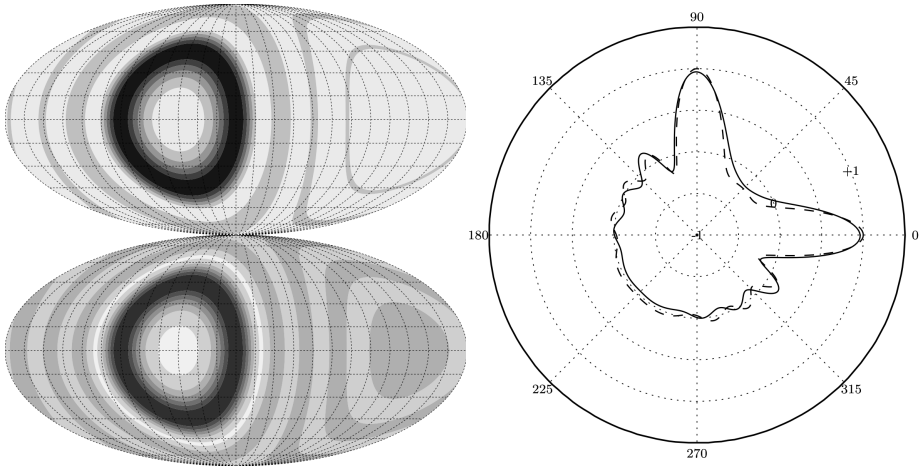

(b)

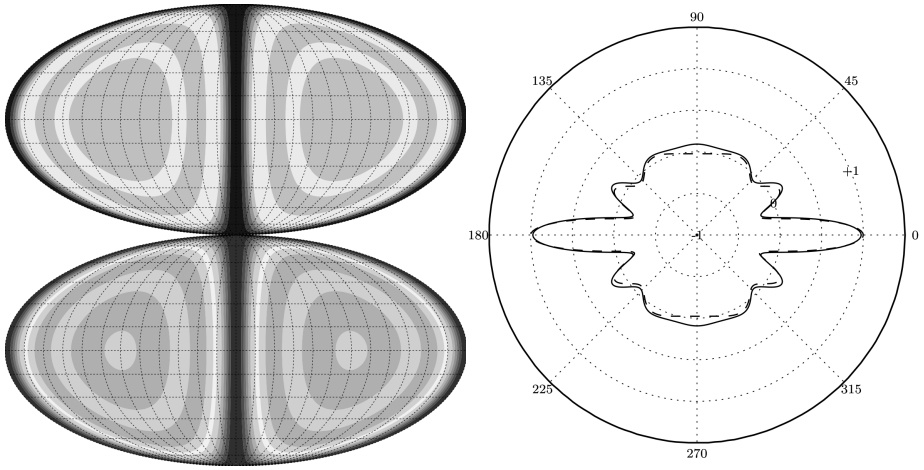

(c)

FIG. 3: Plots for the tests in Table I of the "minimal" parameter values. Each of the elliptical plots is a Mollweide projection of the sky as seen by an observer at the co-ordinate origin (baseline phase centre). The equator in these plots is the x-y plane, the z axis points towards the top of each image and the x axis to the centre. The upper projection in each pair is a contour plot of the correlator's output, and the lower projection is a contour plot of the exact result; the same contour levels are used for both plots, with darker shades indicating greater cross-power. The circular images show detailed comparisons of the correlator's output to the exact solution in the x-y plane. The correlator's output is shown as the solid curve, and the exact solution in dashes, with the angle in degrees from the x axis marked around the outside. These parameter choices aparently capture the character of the solution, but not the details.

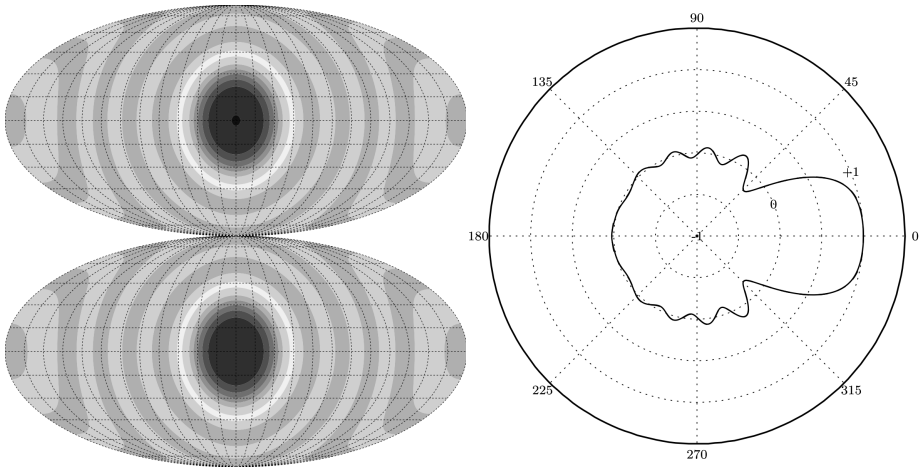

(a)

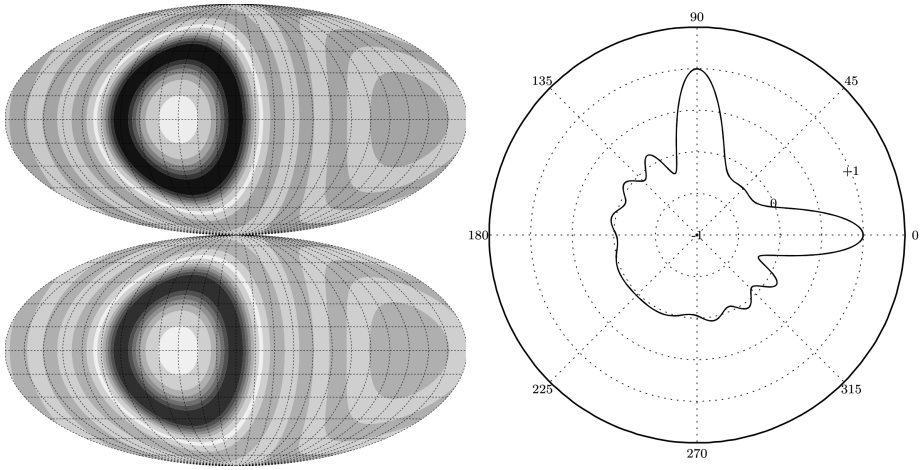

(b)

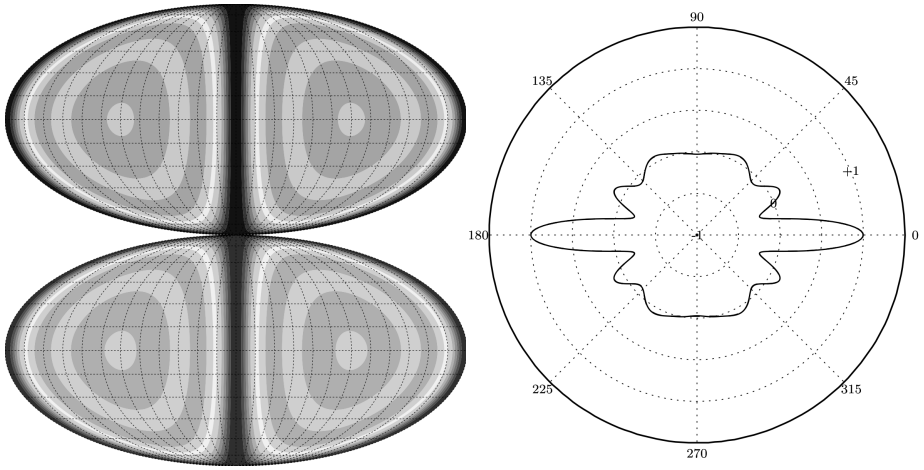

(c)

FIG. 4: Plots for the tests in Table I, continued. For these choices of correlator parameters, the numerical and exact solutions are nearly indistinguishable in these plots.

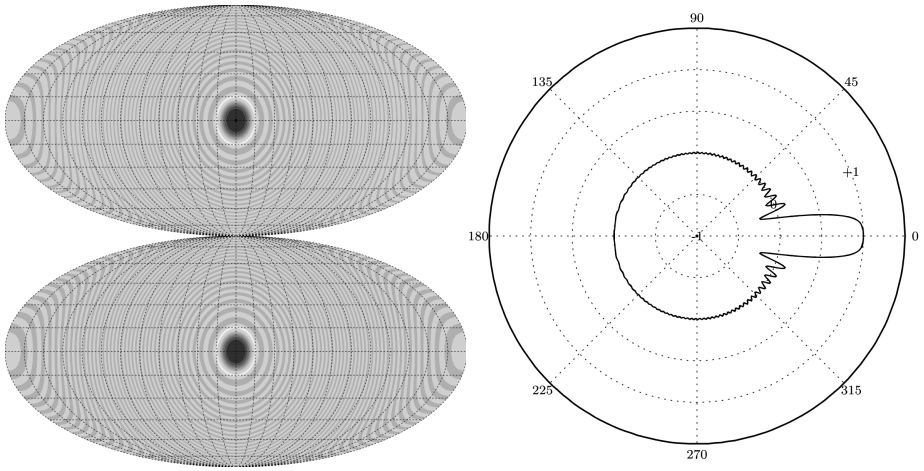

(a)

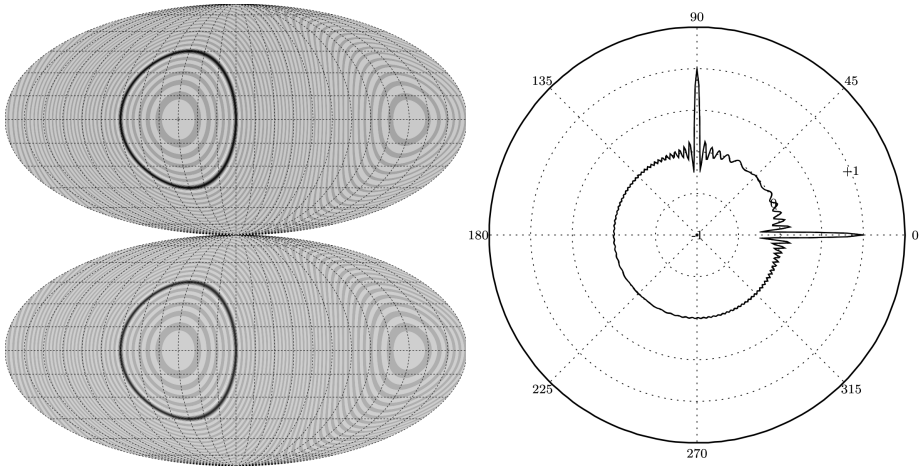

(b)

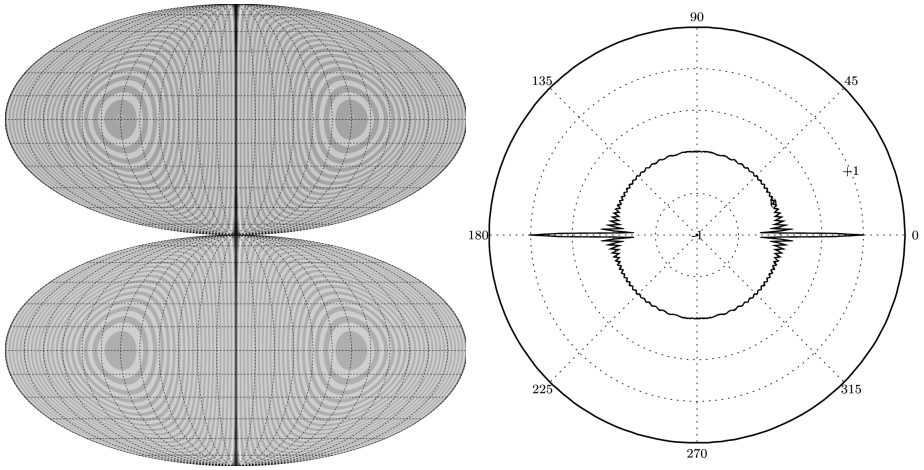

(c)

FIG. 5: Plots for the tests in Table I, continued. The sample rate is increased to 4096 Hz. The correlator's output is nearly indistinguishable from the exact solution.

Processor	Word Size	Frequency	l_ξ	Correlator Type	Correlator Speed
Pentium M	32 bits	1.8 GHz	19	F.D.	822251 samples/s
				T.D.	18094.5 samples/s
			131	F.D.	128038 samples/s
				T.D.	149.5 samples/s

TABLE II: Correlator speed. All tests were performed on GNU/Linux systems using the GCC C compiler and compiler flags appropriate for the hardware. The correlator's speed is the rate at which the correlator processes input samples for a single baseline (1 sample for a baseline = 1 sample each from two instruments).

and balances them across the various contributions. Figure 5 confirms that the correlator recovers the correct solution well. Given that using the parameter values from (72)–(74) with the modifications

$$N_\mathcal{T} \to N_\mathcal{T} + 120 \tag{82}$$

$$l_\mathcal{T} \to l_\mathcal{T} + 4 \tag{83}$$

$$l_\xi \to l_\xi + 2 \tag{84}$$

yield equivalent accuracy over sample rates spanning nearly an order of magnitude, and also balance the error contributions well, it seems it is reasonable to use this as general rule of thumb in tuning the correlator.

B. Speed

The speed of the algorithms is measured by measuring the wall-clock time required for the correlators to process some number of samples of stationary, unit-variance, Gaussian white noise. In both the time- and frequency-domain correlators, there is a large start-up transient associated with the harmonic analysis of the transformation matrix coefficients. In an Earth-fixed co-ordinate system these matrices need never be recomputed, so this time is subtracted from the run time, but the time required to divide the instrument time series into segments, correlate each segment individually, rotate the results to a sky-fixed co-ordinate system using D matrices, and then sum the snap-shots is included as the price to be paid for time-independant delay transformations. Not included is the time required to generate the injection data since a real system does not need to do this, but also not included is any time that might be required for I/O. In the case of the frequency domain correlator, the time required to convert the input time series to a frequency series *is* included in the measured execution time, although in practice this often comes for free for example when the correlator is preceded by frequency-domain data conditioning.

The results of the speed measurements are summarized in Table II. In all cases, the modified correlator parameters in (82)–(84) were used. The frequency domain correlator's operation count appears to be approximately linear in the harmonic order of the output, which is as expected due to the ability to precompute almost everything in that case (resulting in a fixed operation count per output coefficient). The time-domain correlator appears to be a little worse than quadratic in the harmonic order. This is better than expected since the time-domain correlator is dominated by the harmonic product in (40), which we had estimated to be cubic. Clearly the time-domain correlator becomes impractical for larger harmonic orders, but its scaling at large l could be improved by replacing the harmonic product in (40) with a fast harmonic transform algorithm. This has not been tested.

It is important to understand that in both the time- and frequency-domain cases, the number of samples the correlator can process per second is determined entirely by the harmonic order of the cross-power that is desired. The processing speed depends on the bandwidth of the input data and other parameters of the interferometer network only in so far as the harmonic order of the cross-power depends on those parameters. For example, a low-frequency radio interferometer consisting of a single baseline sampled at 5×10^5 samples/s has a 1000-fold wider bandwidth than the 512 Hz example above, but if the interferometer elements are moved 1000 times closer, to a separation of a few kilometres, then the harmonic order of the cross-power is the same, and the frequency-domain correlator tested above could produce a full-sky image in better than real-time on the test system.

XI. CONCLUSIONS

We have derived both time- and frequency-domain algorithms for the fast, approximate, numeric evaluation of the cross-correlation integral that appears in a variety of imaging applications. These algorithms possess a number of

useful features. Their operation counts are linear in the amount (duration) of data to be processed, which allows them to be run quickly on small amounts of data, thereby allowing one to resolve time-dependence in the sources being imaged. The algorithms also provide means by which to systematically increase or decrease their accuracy for increases and decreases in speed respectively, that is to say one can choose to obtain the result to lower accuracy and thereby obtain the result more quickly or with more modest computer requirements. The frequency-domain implementation is particularly fast, and on a 32-bit 1.8 GHz Pentium M computer has demonstrated the ability to synthesize full-sky cross-power images for a single baseline upto harmonic order 19, with an SNR of 26 dB, at speeds of over 800 ksample/s.

XII. ACKNOWLEDGEMENTS

I would like to thank Prof. Wayne Cannon and Prof. Jolien Creighton for their interest, discussions, and input on this work. I would also like to thank Prof. Cheol Choi, Dr. Joseph Ivanic, Prof. Mark Gordon, and Prof. Klaus Ruedenberg for graciously providing the source code for the D matrix generator described in [10].

---

[1] A. R. Thompson, J. M. Moran, and J. George W. Swenson, *Interferometry and Synthesis in Radio Astronomy.* John Wiley & Sons, Inc., second ed., 2001.
[2] S. Haykin, J. H. Justice, N. L. Owsley, J. L. Yen, and A. C. Kak, *Array Signal Processing.* Prentice-Hall, Inc., 1985.
[3] W. G. Anderson, P. R. Brady, J. D. E. Creighton, and Éanna É. Flanagan, "Excess power statistic for detection of burst sources of gravitational radiation," *Physical Review* **D63** (February, 2001) 042003, `arXiv:gr-qc/0008066`.
[4] Y. Gürsel and M. Tinto, "Near optimal solution to the inverse problem for gravitational-wave bursts," *Physical Review* **D40** (December, 1989) 3884–3938.
[5] S. Klimenko, S. Mohanty, M. Rakhmanov, and G. Mitselmakher, "Constraint likelihood analysis for a network of gravitational wave detectors," *Physical Review* **D72** (December, 2005).
[6] B. Allen and A. C. Ottewill, "Detection of anisotropies in the gravitational-wave stochastic background," *Physical Review* **D56** (July, 1997) 545–563, `arXiv:gr-qc/9607068`.
[7] S. W. Ballmer, "A radiometer for stochastic gravitational waves," *Classical and Quantum Gravity* **23** (April, 2006) S179–S185.
[8] N. J. Cornish, "Mapping the gravitational-wave background," *Classical and Quantum Gravity* **18** (October, 2001) 4277–4291, `astro-ph/0105374`.
[9] G. Arken, *Mathematical Methods for Physicists.* Academic Press, Inc., 3rd ed., 1985.
[10] C. H. Choi, J. Ivanic, M. S. Gordon, and K. Ruedenberg, "Rapid and stable determination of rotation matrices between spherical harmonics by direct recursion," *Journal of Chemical Physics* **111** (November, 1999) 8825–8831.
[11] J. D. Jackson, *Classical Electrodynamics.* John Wiley & Sons, Inc., third ed., 1998.
[12] E. W. Weisstein, "MathWorld — a Wolfram web resource." `http://mathworld.wolfram.com`.
[13] M. Rotenberg, R. Bivins, N. Metropolis, and J. K. W. Jr., *The 3-J and 6-J Symbols.* MIT Press, 1959.
[14] M. Galassi, J. Davies, J. Theiler, B. Gough, G. Jungman, M. Booth, and F. Rossi, *GNU Scientific Library Reference Manual.* Network Theory, Ltd., 2nd ed., 2006.
[15] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions With Formulas, Graphs, and Mathematical Tables.* National Bureau of Standards, 9th ed., 1970.
[16] M. Unser, "Sampling — 50 years after Shannon," *Proceedings of the IEEE* **88** (April, 2000) 569–587.
[17] D. M. H. Jr., D. Rockmore, P. J. Kostelec, and S. S. B. Moore, "FFTs for the 2-sphere," *The Journal of Fourier Analysis and Applications* **9** (2003), no. 4, 341–385.
[18] A. Messiah, *Quantum Mechanics: Volume II.* North Holland Pub. Co., 12th ed., 1981.
[19] F. Stenger, "Summary of sinc numerical methods," *Journal of Computational and Applied Mathematics* **121** (September, 2000) 379–420.
[20] M. Sugihara, "Near optimality of the sinc approximation," *Mathematics of Computation* **72** (June, 2002) 767–786.