



This is a repository copy of *IWAVE—an adaptive filter approach to phase lock and the dynamic characterization of pseudo-harmonic waves*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/230678/>

Version: Published Version

Article:

Daw, E.J. orcid.org/0000-0002-3780-5430, Hollows, I.J. orcid.org/0000-0002-3404-6459, Jones, E.L. et al. (5 more authors) (2022) IWAVE—an adaptive filter approach to phase lock and the dynamic characterization of pseudo-harmonic waves. *Review of Scientific Instruments*, 93 (4). 044502. ISSN: 0034-6748

<https://doi.org/10.1063/5.0070394>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

RESEARCH ARTICLE | APRIL 19 2022

IWAVE—An adaptive filter approach to phase lock and the dynamic characterization of pseudo-harmonic waves

E. J. Daw ; I. J. Hollows ; E. L. Jones; R. Kennedy; T. Mistry; T. B. Edo; M. Fays; L. Sun 



Rev. Sci. Instrum. 93, 044502 (2022)

<https://doi.org/10.1063/5.0070394>



Articles You May Be Interested In

Four-quadrant bidirectional AC-DC converter for plug-in electric vehicle charger with grid reactive power support

AIP Conf. Proc. (February 2020)

Hybrid quantum singular spectrum decomposition for time series analysis

AVS Quantum Sci. (May 2023)

Frequency estimation techniques in capacitance-to-frequency conversion measurement

Rev. Sci. Instrum. (January 2020)



MCL
MAD CITY LABS INC.

Closed Loop Nanopositioning Systems with Picometer precision, Low noise and High stability

Force Microscopy and Single Molecule Microscopy Instruments for Quantum, Materials, and Bioscience

Custom Design and Innovative Solutions for the Nanoscale World

Think Nano® | Positioning | Microscopy | Solutions



IWAVE—An adaptive filter approach to phase lock and the dynamic characterization of pseudo-harmonic waves

Cite as: Rev. Sci. Instrum. 93, 044502 (2022); doi: 10.1063/5.0070394

Submitted: 6 September 2021 • Accepted: 30 March 2022 •

Published Online: 19 April 2022



E. J. Daw,^{1,a)} I. J. Hollows,¹ E. L. Jones,¹ R. Kennedy,¹ T. Mistry,¹ T. B. Edo,^{1,2} M. Fays,³ and L. Sun^{2,4}

AFFILIATIONS

¹ Department of Physics and Astronomy, The University of Sheffield, Hicks Building, Hounsfield Road, Sheffield S3 7RH, United Kingdom

² LIGO Laboratory, California Institute of Technology, Pasadena, California 91125, USA

³ Department of Astrophysics, Geophysics and Oceanography (GEO), Space sciences, Technologies and Astrophysics Research (STAR), Université de Liège, allée du six Août 19, 4000 Liège, Belgium

⁴ OzGrav-ANU, Centre for Gravitational Astrophysics, College of Science, The Australian National University, ACT 2601, Australia

^{a)} Author to whom correspondence should be addressed: e.daw@sheffield.ac.uk

ABSTRACT

We present a novel adaptive filtering approach to the dynamic characterization of waves of varying frequencies and amplitudes embedded in arbitrary noise backgrounds. This method, known as IWAVE (Iterative Wave Action angle Variable Estimator), possesses critical advantages over conventional techniques, making it a useful new tool in the dynamic characterization of a wide range of data containing embedded oscillating signals. After a review of existing techniques, we present the IWAVE algorithm, derive its key characteristics, and provide tests of its performance using simulated and real world data.

© 2022 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/5.0070394>

I. INTRODUCTION

The co-inventor of the maser,¹ Arthur Schawlow, is said to have advised his students “Never measure anything but frequency!” There exist a great variety of techniques for measuring the frequencies of pseudo-harmonic waves, broadly addressing two different classes of problems.

In the first class, the average wave characteristics are estimated across a measurement interval under the implicit assumption that the wave properties are essentially static or when changes in these properties over the measurement interval are not of interest. This problem is addressed by a wide variety of methods, including Welch’s method using discrete Fourier transforms (DFTs),^{2,3} Pisarenko’s method,⁴ MUSIC,⁵ and ESPRIT.⁶

In the second class, the oscillator is constantly evolving, and we seek a time-evolving best estimate of oscillator parameters. This problem is most often solved using phase locked loops (PLLs)⁷ or their hardware realization, lock-in amplifiers.⁸ The myriad

applications of PLLs in science and engineering include, for example, a recent proposal for digital PLLs as an alternative technology for the readout of the photodiode clusters used to stabilize the alignment of mirrors in the future gravitational wave detectors.⁹

In searches for almost continuous wave (CW) signals in gravitational wave detectors, the evolving oscillation of continuous wave (CW) signals is at a very low signal-to-noise ratio (SNR), and so, the frequency evolution of such signals in ground based gravitational wave interferometers cannot be inferred from the raw data as is necessary for successful tracking with a conventional PLL. Instead, coherent matched filtering techniques, such as the \mathcal{F} -statistic used in CW searches,^{10,11} are able to achieve higher sensitivity to these weak signals at the price of substantially greater computational burden. Stack-slide-based semi-coherent algorithms expedite the computation to some extent at the cost of sensitivity by summing the signal power in multiple coherent segments after sliding the segments in the frequency domain to account for the signal phase evolution.^{12–15} More efficient semi-coherent methods, e.g., signal

tracking algorithms based on hidden Markov models,^{16–19} have been developed to tackle the computational challenge as well as to allow for some uncertainties in the signal evolution model.^{20–22}

Several other techniques beyond the conventional PLL have been developed for a range of applications outside the field of gravitational wave research although these algorithms are not adapted for the detection of the very weak CW signals expected from gravitational wave sources. One example is the second order generalized integrator (SOGI-PLL).²³ Developed for the problem of characterizing the characteristics of AC voltages in power lines, SOGI was one of the first methods to successfully address the problem of efficient generation of a copy of an input sinusoid that is out of phase (the so-called quadrature or Q phase) with the input signal. Other more recent papers, for example,^{24–27} have developed the SOGI algorithm for a range of practical applications. A second example is the enhanced PLL (EPLL),^{28,29} which solves the problem of tracking the amplitude of a harmonic wave in addition to its frequency. Further PLL developments are well summarized in Ref. 30.

The IWAVE (Iterative Wave Action-angle Variable Estimator) technique described in this paper is a new type of PLL addressing the dynamic characterization of evolving pseudo-sinusoidal signals. Unlike a conventional PLL, the adaptive element is a filter rather than an oscillator or counter. IWAVE has certain advantages over existing PLLs. First, IWAVE produces a benign output when the PLL is unlocked. In the case where the output of the PLL is being used to control something, for example, some parameter of a gravitational wave detector in a closed loop feedback system, then, in colloquial terms, IWAVE does no harm when it is not working. IWAVE naturally tracks the amplitude as well as frequency using a single feedback loop, unlike EPLL, which requires two loops internally for the amplitude and the phase. IWAVE is initialized using a small set of free parameters corresponding directly to physical oscillator properties—just an initial frequency and a single time constant; other PLL algorithms typically contain many control parameters that do not have a clear physical meaning. IWAVE also has the ability to characterize, simultaneously, multiple oscillations having almost-degenerate frequencies using the cross-subtraction method described in Sec. III C. This last advantage has led to detailed analysis of almost-frequency-degenerate violin modes of fused silica suspension wires in advanced LIGO.³¹ Comparisons of IWAVE with the SOGI and EPLL algorithms are given in the Appendix.

As we discuss in Sec. III B, IWAVE as currently implemented is not sensitive to signals where the ratio of the amplitude of the target wave to the root mean square (rms) of the noise background is less than about 0.3. When applied to broadband gravitational wave data, therefore, IWAVE cannot be expected to detect gravitational wave CW signals at the strengths anticipated for sources described in the literature. However, preprocessing steps to divide the data into narrower frequency bands, thereby significantly reducing the rms of the noise, may yield promising search methods. Studies of these ideas are under investigation and will form the subject of future papers. The potential advantage is that by using the data itself to track the evolving frequency of the oscillation, IWAVE may be significantly less reliant on banks of templates for wave evolution and hence may require significantly less computational resources than existing CW search methods. For now, however, IWAVE represents a simple, well-characterized, and useful technique for analyzing quite complex spaces of evolving oscillators at relatively low computational burden,

which is already being applied to studies of important background oscillations in gravitational wave data. This, then, is a method paper describing how IWAVE works and how it performs and giving some examples of that performance on gravitational wave data.

The IWAVE algorithm has many potential applications beyond gravitational wave science. In the control of brushless electric motors, IWAVE could out-perform standard vector control methods at low rotation rates where the back-emf in the motor windings is weak.³² In radio communications, IWAVE might be used to separate closely spaced channels with frequency evolution and multipath splitting.³³ In heart magnetometry, IWAVE could be used to estimate Fourier coefficients of the cardio-magnetic signal using an electrocardiograph signal to modify the phase evolution as the heart-beat rate evolves.³⁴ In physics, IWAVE could be applied to atomic force microscopy³⁵ or as part of an optical squeezing scheme for the readout of interferometers.³⁶ The authors look forward with anticipation to seeing what other applications we have not noticed.

This paper consists of a description of the IWAVE method in Sec. II, a discussion of the limits of applicability of IWAVE in Sec. III, and an overview of certain other PLL methods in the Appendix. Space constraints have led us to leave out many mathematical steps; a full treatment of the mathematics can be found at Ref. 37. A software library implementing IWAVE in C with wrappers into MATLAB and Python/NumPY is available on a public git repository here.³⁸

II. THE IWAVE METHOD

A. The core algorithm

Before writing down the IWAVE core algorithm, we consider as a starting point the \mathcal{Z} -transform³⁹ of a regularly sampled time series, x_p ,

$$\mathcal{Z}_n(\Omega) = \sum_{p=-\infty}^n x_p e^{(p-n)\Omega}, \quad (1)$$

where p increases with time and $\Omega = w - i\Delta$ has the real part w the reciprocal of one e -folding for the weighting of previous samples and the imaginary part Δ equal to the frequency of the \mathcal{Z} -transform component in radians per sample.

$\mathcal{Z}_n(\Omega)$ obeys an iteration equation,

$$\mathcal{Z}_n(\Omega) = e^{-\Omega} \mathcal{Z}_{n-1}(\Omega) + x_n. \quad (2)$$

An unit amplitude phasor input, $x_n = e^{in\Delta}$, to Eq. (2) results in the output $\mathcal{Z}_n = e^{in\Delta} / (1 - e^{-w})$, which has zero phase shift with respect to the input and a larger amplitude dependent on w . Scaling the input by a factor of $1 - e^{-w}$ leads to an iteration algorithm, which passes phasors at frequency Δ with unit gain and zero phase shift,

$$y_n = e^{-w} e^{i\Delta} y_{n-1} + (1 - e^{-w}) x_n. \quad (3)$$

This iteration algorithm is the core of IWAVE. As we shall see, it responds resonantly at frequency Δ . In the language of signal processing, Eq. (3) is an infinite impulse response (IIR) filter because it generates its n th output using the current input x_n , the previous output y_{n-1} , and a two-input, two-output, multi-input, multi-output (MIMO) filter; y_n and x_n are, in general, complex variables. We will also need the real representation of the transfer function for IWAVE derived from Eq. (3) by writing $x_n = x_n^R + ix_n^I$, $y_n = y_n^R + iy_n^I$,

and $y_{n-1} = z^{-1}y_n$, where z^{-1} is the sample delay operator. In these terms, Eq. (3) can be re-written as

$$\begin{pmatrix} y_n^R \\ y_n^I \end{pmatrix} = \begin{bmatrix} H_{11}(z^{-1}) & H_{12}(z^{-1}) \\ H_{21}(z^{-1}) & H_{22}(z^{-1}) \end{bmatrix} \begin{pmatrix} x_n^R \\ x_n^I \end{pmatrix}, \quad (4)$$

where the elements of the transfer function matrix are

$$H(z^{-1}) = \frac{\begin{bmatrix} 1 - e^{-w} z^{-1} \cos \Delta & -e^{-w} z^{-1} \sin \Delta \\ +e^{-w} z^{-1} \sin \Delta & 1 - e^{-w} z^{-1} \cos \Delta \end{bmatrix}}{\left(\frac{1 - 2e^{-w} z^{-1} \cos \Delta + e^{-2w} z^{-2}}{1 - e^{-w}} \right)}. \quad (5)$$

In this form, the transfer function is seen to be second order in sample delay. We determine the response of the core algorithm to an input consisting of a phasor of arbitrary frequency Θ radians per sample by substituting $x_n = e^{in\Theta}$ into Eq. (3). The response is also a phasor, $y_n = Ae^{i(n\Theta + \Phi)}$, where $A(\Theta)$ and $\Phi(\Theta)$ are the frequency dependent magnitude and phase lag of the output phasor with respect to the input given by

$$A(\Theta) = \frac{1 - e^{-w}}{\sqrt{1 - 2e^{-w} \cos(\Delta - \Theta) + e^{-2w}}}, \quad (6)$$

$$\Phi(\Theta) = \arctan\left(\frac{e^{-w} \sin(\Delta - \Theta)}{1 - e^{-w} \cos(\Delta - \Theta)}\right).$$

Thus, phasors of arbitrary frequency are eigenfunctions of the core algorithm with eigenvalues $Ae^{i\Phi}$. The resonant character of the eigenvalues at frequency Δ can be seen in the denominator of Eq. (5), which is identical to the resonant denominator in the SOGI filter.²³ Figure 1 shows A and Φ as a function of Θ for $\Delta = 1.257$ and four values of w : 1, 0.1, 0.01, and 0.001.

Starting from Eq. (6) and writing $\delta = (\Delta - \Theta) \ll 1$ and $w \ll 1$ so that we are in a limit where the frequency is in the vicinity of a narrow resonance, the magnitude of the filter output can be approximated as

$$A(\delta) \approx \frac{1}{\sqrt{1 + \frac{\delta^2}{w^2}}} \quad (7)$$

so that the peak is approximately Lorentzian in shape with full width at half maximum (FWHM) of $2w$ radians per sample. Using the sampling rate, f_s , in Hz, we give other properties of narrow resonances occurring when $w \ll 1$ in Table I, where τ_s is the sampling period in seconds, τ is the response time in seconds, and Δ_0 is the resonant frequency in radians per sample.

B. Application of IWAVE to a real sinusoidal input

We next discuss the usual case where the input data are a real oscillation at the IWAVE resonant frequency, $x_n = \cos(n\Delta)$. We decompose x_n into two phasors, $x_n = x_f + x_b$, where $x_f = e^{+in\Delta}/2$ and $x_b = e^{-in\Delta}/2$, each of which is an eigenfunction of the core algorithm. Substituting the frequencies $\pm\Delta$ into Eq. (6) and rearranging, we obtain the response

$$y_n = e^{i\frac{\Phi}{2}} \left(\frac{(1+a)}{2} \cos\left(n\Delta - \frac{\Phi}{2}\right) + \frac{i(1-a)}{2} \sin\left(n\Delta - \frac{\Phi}{2}\right) \right), \quad (8)$$




FIG. 1. Magnitude (a) and phase (b) of the response of IWAVE to the phasor input as a function of the phasor frequency in radians per sample for different values of w and $\Delta = 1.257$. A smaller w results in a sharper resonant peak.

TABLE I. Properties of narrow resonances of the IWAVE core algorithm in the limit where $w \ll 1$.

Quantity	Symbol	Formula	Units
Full width at half maximum	FWHM or Γ	$\frac{w f_s}{\pi} = \frac{1}{\pi \tau}$	Hz
Quality factor	Q_f	$\frac{\Delta_0}{2w} = \frac{\Delta_0 \tau}{2\tau_s}$...
Resonant frequency	f_0	$\frac{\Delta_0 f_s}{2\pi}$	Hz
Response time	τ	$\frac{w f_s}{\tau_s}$	s

where

$$a = \frac{1 - e^{-w}}{\sqrt{1 - 2e^{-w} \cos(2\Delta) + e^{-2w}}}, \quad (9)$$

$$\phi = \arctan\left(\frac{e^{-w} \sin(2\Delta)}{1 - e^{-w} \cos(2\Delta)}\right).$$

By inspection, the locus of y_n is an ellipse in the complex plane having semi-major and semi-minor axes $1 + a$ and $1 - a$. Figure 2 shows the result of driving the IWAVE core algorithm with an input $x_n = \cos n\Delta$ for $n \geq 0$. The output starts at the origin, spiraling outward toward a limiting ellipse in the steady state. Notice that the argument of the output y_n is always the same as the phase, $n\Delta$, of the input sinusoid.




FIG. 2. Output of the IWAVE core algorithm for an input $x_n = \cos n\Delta$ for $\Delta = 0.3068$, $0 \leq n < 27$, and $w = 0.3$. The points are the individual outputs y_n , and the solid line is the limiting ellipse discussed in Sec. II B.

Input sinusoids at frequencies other than Δ also result in an elliptical steady state output although with different inclination angles and eccentricities and with smaller overall areas due to the falloff in the magnitude of the response for phasors having frequencies far from Δ .

A matrix transformation can be used to transform the elliptical locus into a circular one with the real part of this circular locus being a sinusoid in-phase with the input wave and the imaginary part having the same amplitude but lagging the input wave by 90° . We refer to these in-phase and out-of-phase components as the D and Q phases, respectively. This transformation can be expressed as a sequence of three elementary operations on the vector whose elements are the real and imaginary parts of y_n : a rotation through an angle of $\frac{\phi}{2}$ about the origin; shears parallel to the real and imaginary axes by factors of $\frac{2}{1+a}$ and $\frac{2}{1-a}$, respectively; and finally a rotation through an angle of $-\frac{\phi}{2}$ about the origin. These three matrices can be combined into a single transformation,

$$\begin{pmatrix} D_n \\ Q_n \end{pmatrix} = \begin{pmatrix} 1 + e^{-w} & \frac{e^{-w} - 1}{\tan \Delta} \\ \frac{e^{-w} - 1}{\tan \Delta} & e^{-w} \left(\frac{(e^w - 1)^2}{\sin^2 \Delta} - 1 \right) + 3 \end{pmatrix} \begin{pmatrix} y_n^R \\ y_n^I \end{pmatrix}. \quad (10)$$

The IWAVE core algorithm followed by this matrix transformation results in the output of both D phase and Q phase copies of the input drive. Thus, IWAVE is an example of what is referred to in signal processing parlance as an orthogonal state generator. Our convention follows other papers in the field in that the D (Q) phase output is in (out of) phase with the input at resonance. As we shall see, the Q phase quadrature can be used to generate an error signal to detect changes in frequency, allowing IWAVE to be




FIG. 3. Symbols for IWAVE acting on two-component phasor or real inputs. The parameters w and Δ may be adjusted to reflect changes in the state of the drive.

used in place of a reference oscillator in a PLL. The sum in quadrature of the two phases, $A_n = \sqrt{D_n^2 + Q_n^2}$, is an estimate of the input signal amplitude. We will use the symbols in Fig. 3 to denote the application of IWAVE either to complex phasor or real sinusoidal inputs.

We next consider the transfer functions from a real sinusoidal drive, an example of which is shown in Fig. 4. Note that the dependence on frequency off-resonance is different for the D and Q outputs. The D transfer function rises linearly in frequency below the resonance and falls linearly in frequency above it, having a phase lead of 90° below the resonance and a phase lag of 90° above it. The Q transfer function has a flat frequency response below the resonance but falls as f^{-2} above it and is in phase with the drive below the resonance but 180° out of phase with the drive above it.




FIG. 4. An example of the transfer functions between a real sinusoidal input and the D and Q phase outputs. Here, the sampling rate was 256 Hz, the resonant frequency was 1 Hz, and the filter quality factor Q was set to 12.3. Notice that the low frequency attenuation in the Q phase output is about 0.08, which is $1/Q$. Subfigures (a) and (c) [(b) and (d)] are the magnitude and phase of the D phase (Q phase) output.

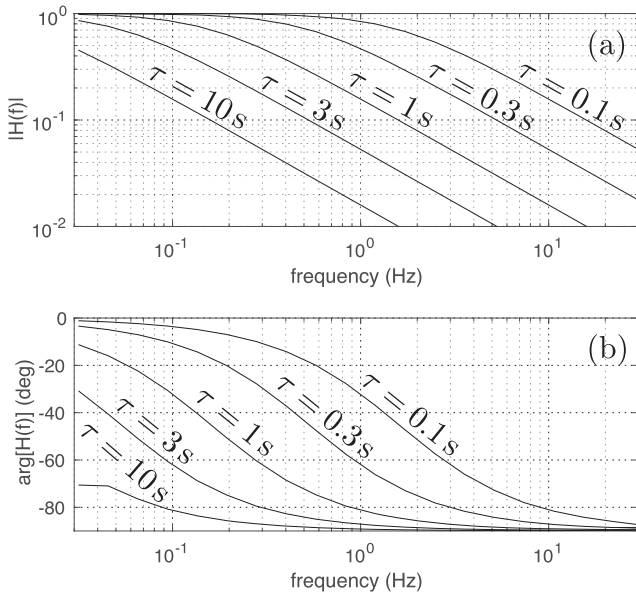


FIG. 5. An example of the response of IWAVE to amplitude modulated signals. The carrier frequency was 60 Hz, and the modulation depth was 10%. Modulation frequencies in the range $33 \text{ mHz} \leq f_{\text{AM}} < 33 \text{ Hz}$ were used, and results are shown for five different values of τ . The 3 dB point for turnover between flat response and proportionality to $1/f_{\text{AM}}$ is $f_{\text{AM}}^{\text{3dB}} = 1/(2\pi\tau)$ in each case. Subfigures (a) and (b) are the magnitudes and phases of the responses, respectively.

This behavior is similar to that of a driven series RLC tank circuit, where the transfer functions from the input voltage to the voltage across the capacitor and resistor are similar to those between the input and the Q and D phase outputs, respectively. The relatively light suppression of low frequency off-resonance signals in the Q phase output can be important, particularly in cases where the quality factor Q_f of the circuit is set low by using a relatively large w coefficient. As in the resonant circuit, the ratio of the resonant to low frequency response is Q_f .

Changes in the amplitude of the incoming wave at the resonance result in a corresponding change in the quadrature sum, $A_n = \sqrt{D_n^2 + Q_n^2}$, but with a response time $\tau = \tau_s/w$ leading to a single pole in the response to amplitude changes at $s = -1/\tau$. Figure 5 shows the response at IWAVE's resonance frequency to an input having a sinusoidally modulated amplitude for a variety of response times.

C. An IWAVE-based phase locked loop

In order to phase lock with IWAVE, we need a measure of departures in the frequency from the frequency Δ of the harmonic wave at the input. We achieve this by exploiting the response time τ of the IWAVE algorithm. Consider a harmonic wave initially at frequency Δ . IWAVE yields both a D_n phase and a Q_n phase copy of this wave at its output. The product of the out-of-phase copy and the input wave, $A^2 \cos(n\Delta) \sin(n\Delta) = (A^2/2) \sin(2n\Delta)$, is a pure harmonic signal at frequency 2Δ . Now, consider an input signal where the wave develops anomalous phase and amplitude disturbances, δ and ϵ , respectively, so that $x_n = A(1 + \epsilon)\cos(n\Delta + \delta)$. For elapsed times significantly less than τ following the onset of these disturbances, the outputs $D_n = A \cos n\Delta$ and $Q_n = A \sin n\Delta$ are unaffected by them. Physically, where the oscillator has an unchanged frequency and amplitude, its complex plane representation precesses in a circle about the origin, and the current IWAVE filter output therefore provides a predictor of the point where such a static oscillator will land the next sample. If the frequency or amplitude evolves, then the combinations $E_n = (x_n - D_n)Q_n$ and $F_n = x_n D_n + Q_n^2 - A_n^2$ are estimates of the departure of the actual evolution of the complex coordinate of the oscillator state from this static model. These combinations can be written in a matrix form as follows:

$$\begin{pmatrix} E_n \\ F_n \end{pmatrix} = \frac{A^2}{2} \begin{bmatrix} -\delta \\ \epsilon \end{bmatrix} + \begin{pmatrix} \cos(2n\Delta) & \sin(2n\Delta) \\ -\sin(2n\Delta) & \cos(2n\Delta) \end{pmatrix} \begin{pmatrix} \delta \\ \epsilon \end{pmatrix}. \quad (11)$$

This equation shows that E_n and F_n each consists of a static offset plus upper sidebands of frequency 2Δ , linear in the phase

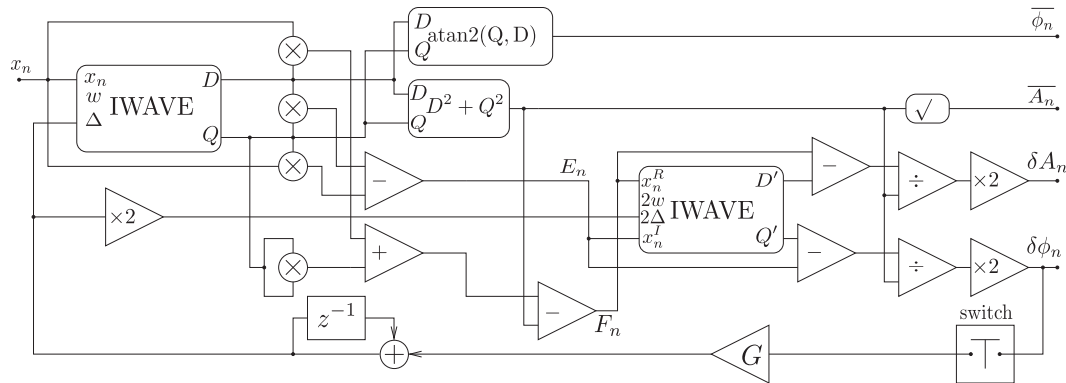


FIG. 6. A schematic of the IWAVE PLL. The error signal after upper sideband filtering is $\delta\phi_n$. The switch closes the feedback loop to adjust Δ and 2Δ for the two IWAVE instances. The use of the complex input IWAVE for attenuation of the 2Δ component in the error signal reduces the computational load compared with the use of a real input IWAVE on the E_n signal alone.

and amplitude offsets, respectively. The upper sidebands, however, appear as components of a rotating phasor in the space spanned by the E_n and F_n signals. We remove them using the complex carrier version of IWAVE, subtracting its outputs from its inputs. We have determined experimentally that using twice the w factor in the 2Δ IWAVE filter yields an acceptable error signal for the detection of phase departures.

Figure 6 is a schematic for the full IWAVE-based phase and amplitude detector. We have only considered here the case where the input data are real and we are tracking harmonic waves. The case where the input data are two-component rotating phasors is simpler as it can be shown that the combination $E'_n = x_n^R Q_n - x_n^I D_n$ contains a pure DC signal in-phase offset with no upper sideband contamination. In addition, shown is the feedback path from the filtered error signal, $\delta\phi_n$, back to Δ through an integrator, discussed in the following.

The response of IWAVE to modulation of the frequency of the carrier is shown in Fig. 7. Knowing the analytic form of the frequency response, we can analyze the closed loop IWAVE PLL to determine an appropriate choice of feedback gain, G . A schematic for the feedback controller is shown in Fig. 8. Any difference between the incoming wave frequency and the IWAVE filter central

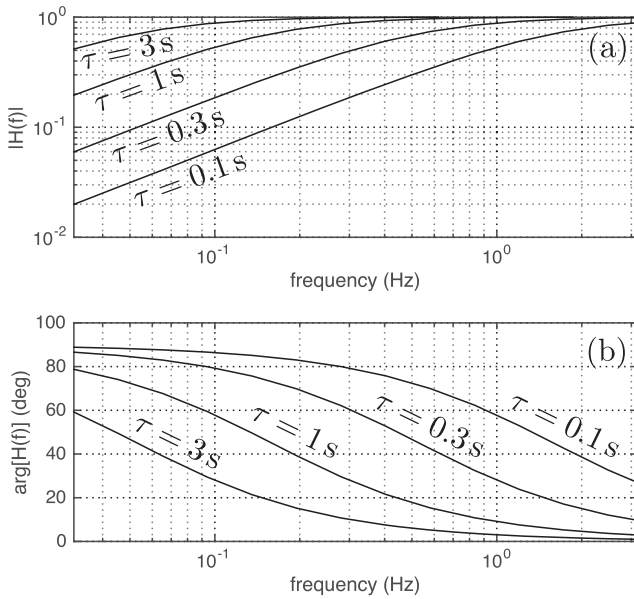


FIG. 7. The transfer function of IWAVE from frequency modulation of the input carrier to the response in $\delta\phi_n$ with the feedback loop open for four different values of τ . The output is insensitive at DC because a step in frequency without feedback moves the input carrier off the IWAVE resonant frequency Δ_0 . Sensitivity increases toward high frequencies because the homodyne detector relies on beats between the Q phase IWAVE output and the carrier, which appear so long as IWAVE has not had the adjustment time, τ , necessary for its outputs to respond to changes in the input signal. Between these two regimes, there is a single pole at frequency $1/(2\pi\tau)$ Hz, which, combined with zero at DC, makes the frequency modulated (FM) response a highpass filter having the same 3 dB point as the lowpass filter associated with the AM response. In each case, the carrier frequency is 60 Hz and the sampling rate is 16 384 Hz. Subfigures (a) and (b) are the magnitude and phase of the transfer function, respectively.

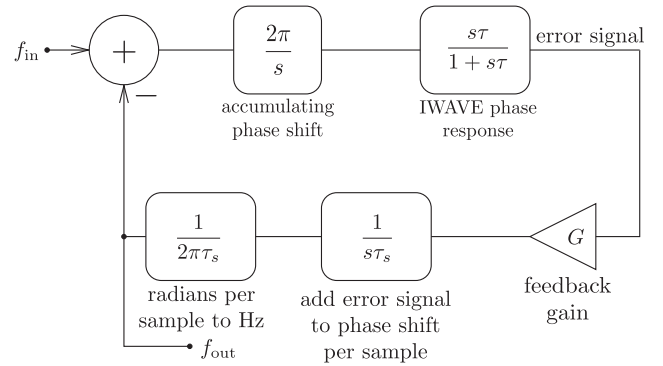


FIG. 8. An s-plane model of the IWAVE PLL.

frequency results in an accumulating phase shift in the homodyne detector. This accumulation of phase is represented by the factor of $2\pi/s$, where $s = 2\pi if$ with f being the signal's frequency. The response of IWAVE to phase, confirmed by the measurements underlying Fig. 7, acts on the accumulated phase to produce the error signal. As shown in Fig. 8, the feedback path from the error signal to a correction in the central frequency of IWAVE consists of an adjustable gain, G , an addition of the gain boosted error signal to the previous value of the phase shift per sample, and a scale factor to convert from radians per sample to frequency in Hz. The closed loop gain is calculated by the usual consistency argument around the loop,

$$f_{\text{out}} = f_{\text{in}} \frac{2\pi}{s} \frac{s\tau}{1+s\tau} G \frac{1}{s\tau_s} \frac{1}{2\pi\tau_s} - f_{\text{out}} \frac{2\pi}{s} \frac{s\tau}{1+s\tau} G \frac{1}{s\tau_s} \frac{1}{2\pi\tau_s} \quad (12)$$

from which we obtain the closed loop transfer function,

$$H(s) = \frac{f_{\text{out}}}{f_{\text{in}}} = \frac{\frac{G}{\tau_s^2}}{s^2 + \frac{s}{\tau} + \frac{G}{\tau_s^2}}. \quad (13)$$

This is the transfer function of a driven damped harmonic oscillator. We want a critically damped response since the control signal will be used to measure the oscillator frequency. For critical damping, we require two coincident real poles, which is achieved if $G = \tau_s^2/(4\tau^2)$. The closed loop transfer function then takes the simpler form

$$H(s) = \left(\frac{\frac{1}{2\tau}}{s + \frac{1}{2\tau}} \right)^2. \quad (14)$$

The response to frequency or phase modulation is therefore flat below the knee frequency of $1/(4\pi\tau)$ Hz, where due to the two poles it has rolled off to -6 dB and drops as $1/f^2$ above that frequency. Larger values of G result in sharper turnover or a resonant peak, corresponding to the underdamped case.

III. LIMITS OF APPLICABILITY OF IWAVE

The time constant, τ , determines the responsiveness of IWAVE to changes in wave frequency and amplitude as well as the noise bandwidth of the filter. Decreasing τ results in a faster response and better ability to stay locked on waves whose frequencies are

changing but also increases the bandwidth of IWAVE to input noise, resulting in a noisier error signal. The optimal τ is small enough so that IWAVE stays locked when the frequency changes but not so small that excessive background noise is admitted by the filter. Section III A is on frequency tracking, and Sec. III B discusses locking in the presence of additive noise and the character of the error signal.

A. Frequency tracking

Consider a wave whose displacement at time t is $h(t) = A \cos(2\pi f(t)t)$ so that the frequency of the wave is changing. The ability of IWAVE to track the evolving wave depends on the value of the response time parameter, τ . Figure 9 shows the results of running IWAVE on a swept sinusoidal wave starting at 20 Hz and increasing linearly in frequency at a variety of rates. At each sweep rate, a variety of values of τ were used. The sum of the squares of the deviation between the actual frequency and the frequency reconstructed by IWAVE over a time interval of 20 s, here referred to as χ^2 , was calculated as a measure of the accuracy of frequency reconstruction.




FIG. 9. Measured χ^2 as defined in the Sec. III A from a simulation where an input swept sinusoid plus additive white Gaussian noise was fed into IWAVE. A range of sweep rates from 0.1 to 2.5 Hz s⁻¹ were used. For each injected wave, IWAVE was run with a range of τ between 0.01 and 5 s. In addition, overlaid are three curves in bold. The curve labeled “measured” is the minimum of $\chi^2(\tau)$ for each df/dt . The curve labeled “theory” is the calculated prediction for the minimum of $\chi^2(\tau)$ from Eq. (15). The curve labeled U.L. is the predicted upper limit on τ above which loss of lock is predicted as discussed in the final paragraph of Sec. III A, given by Eq. (16), vs the measured χ^2 at that τ . The thick black line follows the values of τ where IWAVE loses lock, as demonstrated by the onset of noise in χ^2 , over simulations having a variety of df/dt values.

The value of χ^2 is seen to be a function of τ with a pronounced minimum that is a function of the sweep rate. The χ^2 statistic also becomes noisy at both ends of the range of values of τ . Here, we explain the smooth descent and ascent in χ^2 on either side of the minimum and obtain a formula for the optimum value of τ . We also explain the onset of noise at either end of the range of τ .

Starting at the minimum, χ^2 rises with τ because of the response time of the closed loop transfer function of the servo from Eq. (14). This transfer function is that of two RC lowpass filters in series, each having an exponentially decaying impulse response of timescale 2τ . The response time of the entire transfer function is therefore the time duration of the autocorrelation of this exponentially decaying function. This autocorrelation initially rises linearly in time after the impulse before reaching a maximum at time 2τ and then decaying exponentially. The time between the impulse and the point where the exponential decay reaches $1/e$ of its maximum value is 6τ . This causes the frequency tracking of IWAVE to lag behind that of the input swept sine wave, leading to a frequency discrepancy at any given time of $\Delta f_1 = 6\tau(df/dt)$.

Below the minimum of $\chi^2(\tau)$, the rise in χ^2 with decreasing τ is explained by the frequency response of the core IWAVE algorithm and is best understood by the analogy between the IWAVE algorithm and the characteristics of a damped harmonic oscillator discussed in Sec. II B. At higher values of damping, the frequency response is maximal at a frequency f below the natural frequency of the undamped oscillator, f_0 by an amount given by the relationship $f_0 = \sqrt{f^2 + 1/(2\pi^2\tau^2)}$. This causes a systematic offset between the frequency, f , returned by IWAVE, and the frequency, f_0 , of the input signal. The square of this frequency difference is an additional contribution to the χ^2 statistic.

By squaring and adding the frequency discrepancies arising from these two effects and minimizing with respect to τ , making the assumption that $2\pi^2\tau^2(df/dt)^2 \gg 1$, we arrive at a value of τ that minimizes the χ^2 statistic for frequency tracking,

$$\tau_{\text{opt}} = \frac{1}{\sqrt[6]{288\pi^4 f^2 \left(\frac{df}{dt}\right)^2}}. \quad (15)$$

Figure 9 shows in bold red the values of τ_{opt} corresponding to the measured minimum τ and χ^2 across each of the χ^2 curves vs the measured χ^2 at that minimum along with in bold blue the value of τ_{opt} vs the value of χ^2 at τ_{opt} . There is good agreement between the theoretical τ_{opt} and the value determined from simulations within 29% for the largest value of df/dt studied and within 13% for the smallest one. At larger values of df/dt , τ_{opt} is smaller, so there is more broadband noise in the error signal, and a more sophisticated optimization on τ would lead to a larger optimal value than that predicted by Eq. (15).

We next discuss the breakdown of IWAVE at high values of τ , where the $\chi^2(\tau)$ curves become noisy, indicating loss of lock. The following argument leads to successful prediction of the value of τ where this breakdown occurs. Consider a wave whose displacement at time t is $h(t) = A \cos(2\pi f(t)t)$ so that the frequency of the wave is changing. Assume that IWAVE is locked at time t so that the IWAVE output is $m(t)$ and is equal to $h(t)$. At time $t + \tau$, the wave displacement has evolved to $h(t + \tau) = A \cos(2\pi(f + df)(t + \tau))$. In the time interval $[t, t + \tau]$, the IWAVE output has not had time

to respond to the frequency shift df and therefore takes the form $m(t + \tau) = A \cos(2\pi f(t + \tau))$. The phase shift between the wave and the IWAVE output at time $t + \tau$ is $\Delta\phi \approx 2\pi\tau df$. The error signal for IWAVE is approximated by the integral $\int h(t')m(t')dt'$ over time interval $[t, t + \tau]$. If the phase shift between the incoming wave and the IWAVE output over this time interval is greater than π , the error signal will undergo a sign change causing loss of lock. Therefore, a condition for IWAVE to remain locked is that $\Delta\phi < \pi$ or $\tau\Delta f < 1/2$. Writing $df = \tau \times df(t)/dt$, we arrive at the upper bound that τ should obey

$$\tau < 1/\sqrt{2df/dt} \quad (16)$$

for IWAVE to remain locked. This line is drawn on the χ^2 curves in Fig. 9 in bold black, labeled U.L., vs the measured value of χ^2 at that value of τ in each simulation. The limit tracks the onset of lock loss as demonstrated by large excursions in χ^2 well across different trial values of df/dt .

B. Response to noise

The error signal for the IWAVE PLL is derived from the product of the input data stream and the Q phase output of the IWAVE filter minus the product DQ of the two IWAVE outputs. Where the wave frequency is static, this subtraction removes the upper sideband component at frequency $2f$. When the frequency changes, this causes an additional transient upper sideband component, which is removed using a second IWAVE filter at frequency $2f$ having a response time half that of the primary IWAVE filter. Finally, we divide by the square of the wave amplitude because both the amplitude of the incoming wave and the amplitudes of both IWAVE outputs scale linearly with the wave amplitude. We need to divide this scale out; otherwise, the PLL loop gain will be dependent on the wave amplitude.

This is a form of homodyne detector. Because the error signal incorporates the unfiltered wideband input, the error signal incorporates the broadband noise of the incoming data. The distribution of the error signal therefore reflects the spectral characteristics of the input data over the full Nyquist band. If, for example, the incoming data includes a time domain transient with a broad spectral distribution, this transient will be reflected in the time history of the error signal from IWAVE. If the out of band noise is sufficiently large in amplitude, then IWAVE will lose lock.

We have determined experimentally that at the optimal choice of response time, τ_{opt} given in Eq. (15), then IWAVE will stay locked when the ratio of the wave peak amplitude to the root mean square noise amplitude exceeds 0.3. Future work to improve the performance of IWAVE at lower signal-to-noise ratios could involve, for example, prefiltering the input data to focus on a narrower frequency band about the frequency of interest for the waves under study. The noise content of the error signal, which leads to noise also in the estimate of the IWAVE frequency, is greater at smaller values of τ where the bandwidth of the IWAVE filter resonance is larger. At sufficiently small values of τ , the incursion of noise leads again to loss of lock. This can be seen in Fig. 9 for $\tau < 0.02$. The optimal value of τ is affected by higher noise levels with larger values than that given in Eq. (15) becoming optimal.

The exact value of τ where lock loss occurs and the effects of noise on the optimal value of τ could be determined by a detailed stochastic differential equation analysis and is beyond the scope of this paper. However, a simple scaling argument can be made. Noise in the error signal leads to a random component being added to the phase shift per sample. This means that the reconstructed frequency, which also forms the servo control signal, contains a component that undergoes a random walk and hence grows with the square root of the number of samples. This means that you might expect the onset of loss of lock to occur at τ , which scales as one over the square of the signal-to-noise ratio, so that IWAVE works at τ above a lower limit that goes down by a factor of two when the signal-to-noise ratio is enhanced by a factor of $\sqrt{2}$. This was verified by injecting additive white Gaussian noise on top of the swept sine wave and noting that the threshold for lock loss at low τ reproduced this predicted behavior.

In terms of noise in the error signal, the most significant source is the product of noise in the input data since this noise enters the error signal without bandpassing through the IWAVE filter. However, this broadband noise has an rms amplitude independent of the amplitude of the tracked wave. If the tracked wave drops in amplitude but the rms of the broadband noise at the input does not, then the noise component of the error signal will scale inversely proportional to the amplitude of the line. This is exactly what is seen when IWAVE is run, in practice, on real data. The normalization of the error signal with the amplitude squared is necessary to ensure that the feedback loop gain is independent of the wave amplitude.

The other function of the error signal is to provide an indication of whether or not the PLL is locked. With a non-stationary rms, this is difficult. However, if we scale the error signal by multiplying it by the amplitude of the wave being tracked, this stabilizes the rms of the error signal noise at the level of the line amplitude. If we further divide by the long-term rms of the input data, we then obtain a statistic that has a stable rms of order unity when the servo is locked. Departures from lock manifest themselves as large transient spikes of amplitude greater than ten. This effect will be seen in the discussion of performance on gravitational wave data in Sec. IV. In particular, in Fig. 11, the error signal plotted for each of the four harmonics in the study has been scaled in this way.

C. Use of multiple IWAVE filters in parallel

Multiple IWAVE filters can be applied to multiple harmonic waves in a single data stream. However, when those harmonics are frequencies spaced closely together, this causes crosstalk between the different filters. All harmonics present in the data will enter every IWAVE instance through the error signal. To mitigate this effect, we employ a cross-subtraction scheme, as illustrated in Fig. 10.

The schematic shows only two IWAVE instances, but the technique generalizes to any number of filters. Assume that the two IWAVE filters are locked on line frequencies f_1 and f_2 . The D and Q phase outputs from each IWAVE instance are fed into a phase shifter, which estimates the wave signal one sample in the future at the frequency of the wave tracked by the filter. This wave sample is stored until the next input data sample when it is subtracted from the input data to the other filter. In this way, the input data to each IWAVE filter is purged of the harmonic being followed by the other filter. This technique has been determined




FIG. 10. A schematic showing the cross-subtraction technique for two parallel IWAVE filters. The feedback loop components are omitted for clarity; each IWAVE filter is separately instrumented as shown in Fig. 6.

to successfully lock multiplets of up to 20 filters, leaving the frequency estimates from each filter free of oscillations at the difference between frequencies in the multiplet, the effect seen in the absence of the cross-subtraction technique. The technique works with IWAVE because the two quadrature outputs can be used to generate an output shifted through an arbitrary phase shift and because the outputs are at the same amplitude as the input wave onto which IWAVE is locked.

IV. PERFORMANCE OF IWAVE ON REAL WORLD DATA

We present an example of the application of IWAVE to a set of harmonic noise components of gravitational wave data taken from the LIGO open data center web site.⁴⁰ The data were acquired by the LIGO Hanford interferometer on November 30th, 2016 and consists of 800 s of calibrated strain data from the Hanford interferometer during an 800 s lock stretch. The data were preprocessed with a fourth order Butterworth highpass filter at 30 Hz, followed by four third order Chebyshev type 2 bandpass filters between 5 and 300 Hz, applied in series. Finally, a fine adjustment to the spectrum was made using a single real pole at 10 Hz. The resulting data are dominated by the 20–80 Hz band and are approximately white in that range. It is not a requirement that the input data to IWAVE be whitened, but if a spectral feature is to be successfully tracked by IWAVE, its peak should rise above the noise floor in the surrounding background; whitening ensures that this is the case. Eight harmonic features were identified from a broad power spectrum and were tracked using eight parallel IWAVE instances. These originate from various instrumental sources present in the Hanford instrument at the time when the data were acquired. Violin mode harmonics have been studied in detail by Cumming *et al.*³¹

The results at four of the identified frequencies are shown in Fig. 11. For each harmonic, the frequency, amplitude, and scaled error signal (as described in Sec. III B) are displayed. Amplitudes are in dimensionless strain units, so one represents the 4 km length of the LIGO detector arms. Although the lines are in some cases almost degenerate in frequency, there is no evidence of beats between the reconstructed frequencies. The error signal

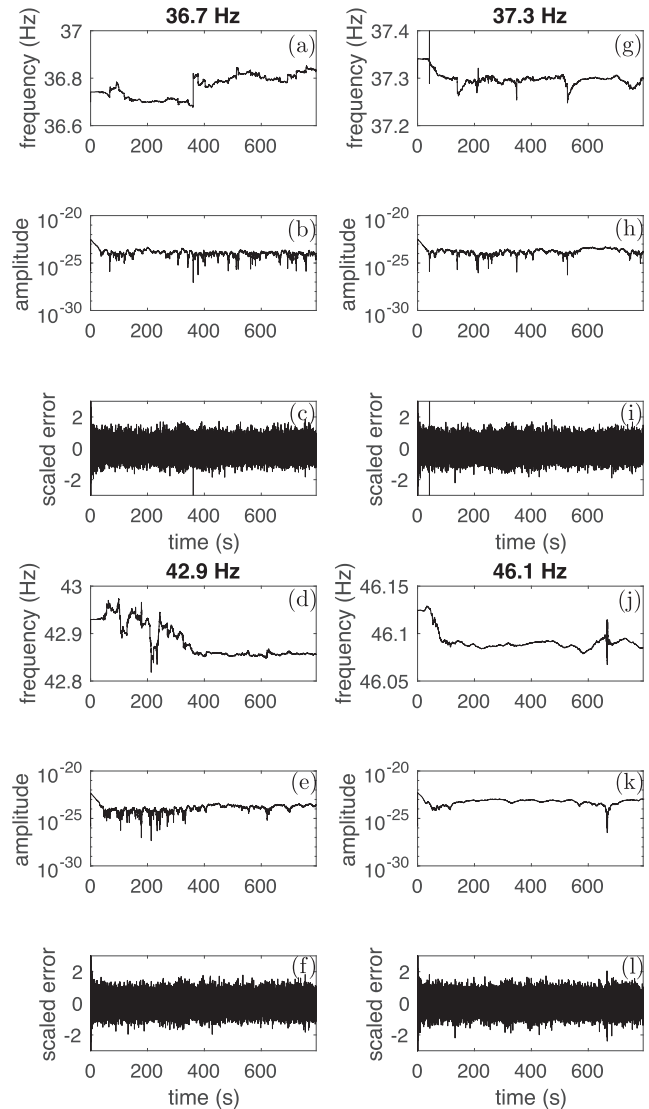


FIG. 11. The results of IWAVE frequency tracking on four pseudo-harmonics present in data from the LIGO Hanford detector. Subfigures (a)–(c), (d)–(f), (g)–(i), and (h)–(l) pertain to harmonics at the nominal frequencies displayed at the top of each group of three sub-plots. At each frequency, the reconstructed frequency, amplitude, and scaled error are plotted. Refer to the discussion in Sec. III B for a description of error signal scaling. Some of the harmonics have frequencies that exhibit significant time evolution; others are more static. In a few cases, IWAVE can be seen to have lost and re-acquired lock. In the case of the 36.7 Hz wave, at about 350 s, a lock loss can be seen in frequency and also via a spike in the scaled error; similarly at about 50 s in the 37.3 Hz wave and at about 680 s in the 46.1 Hz wave.

is roughly static with approximately unit rms although the distribution is non-Gaussian because of the modulation of the input noise by the sinusoidal IWAVE output. A non-statistical transient fluctuation in the modified error signal in the 36.7 Hz line at around 360 s corresponds to a jump in the IWAVE reconstructed frequency, indicating that IWAVE momentarily lost lock

when either the frequency of this sinusoid shifted rapidly or IWAVE jumped between two almost-degenerate harmonics. A second loss of lock can be seen in the 37.3 Hz data at around 30 s accompanying a sudden drop in the amplitude of the harmonic. The eight IWAVE instances all successfully tracked their target harmonics with the reconstructed error signal providing a useful performance indicator.

V. SUMMARY AND FUTURE WORK

We have described IWAVE, a novel orthogonal state generator, resonant filter, and phase locked loop for the dynamic tracking of harmonic waves. The method has a single input parameter, its response time. The algorithm has a low computational load so that many harmonics can be tracked in real time using a single central processing unit (CPU) core. The ability to track multiple closely spaced harmonics means that the method lends itself well to applications where there are dense “forests” of harmonics, such as in LIGO violin mode clusters and communication applications. IWAVE has been applied to LIGO strain data and used to study the character of violin modes in ultralow loss fused silica suspensions. There are many possible applications of the IWAVE method. It is complementary to existing PLL algorithms that we have described in the [Appendix](#). We have supplied software implementations of IWAVE in C with MATLAB and Python wrappers to encourage the community to find other applications and uses.³⁸

The authors can see several directions in which IWAVE could be improved. The error signal is susceptible to broadband noise contamination, and a narrower band alternative would be of benefit in applications with very weak signals although narrowbanding will reduce the responsiveness of the method to frequency changes. There are also applications where feedback is not important, for example, the use of IWAVE for novel resonators that is promising for resonant detectors of weak signals in physics, such as those of dark matter axions.⁴¹ We look forward to seeing what the community finds to do with our harmonic tracking algorithm.

ACKNOWLEDGMENTS

The members of the Sheffield gravitational wave research group acknowledge the support of the Science and Technology Facilities Council under Grant Nos. ST/V005693/1, ST/V001752/1, ST/V001744/1, ST/V001019/1, and ST/R000336/1. I.J.H. was supported by the Hollows Scientific Foundation. L.S. acknowledges the support of the Australian Research Council Centre of Excellence for Gravitational Wave Discovery (OzGrav), Project No. CE170100004, the United States National Science Foundation, and the LIGO Laboratory. M.F. acknowledges the support of the Fonds de la Recherche Scientifique-FNRS, Belgium, under Grant No. 4.4501.19.

This research has made use of data, software, and/or web tools obtained from the Gravitational Wave Open Science Center (<https://www.gw-openscience.org/>), a service of the LIGO Laboratory, the LIGO Scientific Collaboration, and the Virgo Collaboration. The LIGO Laboratory and Advanced LIGO are funded by the United States National Science Foundation (NSF) as well as the Science and Technology Facilities Council (STFC) of the United

Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO600 detector. Additional support for Advanced LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the United States National Science Foundation and operates under cooperative Agreement No. PHY-1764464. Advanced LIGO was built under Grant No. PHY-0823459. Additional support for Advanced LIGO was provided by the Australian Research Council. Virgo is funded through the European Gravitational Observatory (EGO) by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale di Fisica Nucleare (INFN), and the Dutch Nikhef with contributions by institutions from Belgium, Germany, Greece, Hungary, Ireland, Japan, Monaco, Poland, Portugal, and Spain.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

DATA AVAILABILITY

The data that support the findings of this study are available at Ref. 38. The data for the performance tests discussed in Sec. IV can be obtained from the LIGO open data center web site.⁴⁰

APPENDIX: OVERVIEW OF OTHER PLL METHODS

1. SOGI-PLL

The now often-used Generalized Integrator-Based PLL, SOGI-PLL, was introduced in 2006.²³ SOGI is an orthogonal state generator, producing in-phase and quadrature-phase copies of the input wave analogous to the IWAVE *D* and *Q* outputs. These outputs are mixed with two quadratures from a reference oscillator, leading to an error signal in the quadrature-phase output that indicates frequency differences between the SOGI output and the reference. The error signal is fed to a proportional/integral filter. The filter output is added to a frequency offset input ω_n , and the result is used to adjust the coefficients of the SOGI filter to reflect the frequency change as well as to control the reference oscillator. Essentially, the SOGI algorithm is coupled to a conventional phase locked loop with a reference oscillator.

The SOGI algorithm was developed for the field of power grid monitoring, where frequency changes are a small fraction of a nominal constant value, commonly 50, 60 Hz, or harmonics of these. It is not designed for large departures from the frequency set at the ω_n input. The SOGI orthogonal state generator is designed using two s-plane integration stages, shown in the tan SOGI block. These s-plane filters are transformed to the digital domain by, for example, using a Tustin algorithm. The SOGI method does not track the wave amplitude although this can be done in a separate circuit, assuming the frequency of the wave is known, using homodyne detection, for example. The SOGI orthogonal state generator does not have the same transfer functions as the IWAVE one, and this is not surprising given the different methods by which they are obtained although the denominators of the two transfer functions are identical as they both represent a resonant response to a harmonic drive.

2. EPLL

The EPLL although apparently seeming to differ from quadrature signal generation-based PLLs, such as SOGI-PLL, is closely related thereto.²⁸ A conventional PLL is embedded within a second feedback servo that uses the integral of the quadrature phase of the PLL output, multiplied by the unintegrated quadrature phase to synthesize a 2ω signal, which can then be subtracted from the input data, compensating for the 2ω component of the phase detector output inside the PLL. Furthermore, the common amplitudes are also equal to half the amplitude of the sinusoid in the input data. In the EPLL structure, the input data are normalized to the PLL by this amplitude, thereby removing the amplitude dependence of the phase detector.⁴²

Unlike IWAVE, there is a necessity for two servos to be locked at once and several numerical parameters that must be adjusted, including parameters necessary to implement the s-plane design digitally. EPLL is susceptible to interference from other waves present in the input and to DC offsets. The latter two issues are addressed by prefiltering the input to the EPLL. EPLL has been widely implemented because of its comparative simplicity, ability to track wave amplitude, and the availability of the code implementing the method.⁴³

REFERENCES

- ¹A. L. Schawlow and C. H. Townes, *Phys. Rev.* **112**, 1940 (1958).
- ²R. B. Blackman and J. W. Tukey, *The Measurement of Power Spectra* (Dover, 1958).
- ³P. Welch, *IEEE Trans. Audio Electroacoust.* **15**, 70 (1967).
- ⁴V. F. Pisarenko, *Geophys. J. Int.* **33**, 347 (1973).
- ⁵R. Schmidt, *IEEE Trans. Antennas Propag.* **34**, 276 (1986).
- ⁶R. Roy, A. Paulraj, and T. Kailath, *IEEE Trans. Acoust., Speech, Signal Process.* **34**, 1340 (1986).
- ⁷F. M. Gardner, *Phaselock Techniques*, 3rd ed. (Wiley-Interscience, 2005).
- ⁸P. K. Dixon and L. Wu, *Rev. Sci. Instrum.* **60**, 3329 (1989).
- ⁹G. Heinzel, M. D. Álvarez, A. Pizzella, N. Brause, and J. J. E. Delgado, *Phys. Rev. Appl.* **14**, 054013 (2020).
- ¹⁰P. Jaranowski, A. Krolak, and B. F. Schutz, *Phys. Rev. D* **58**, 063001 (1998).
- ¹¹C. Cutler and B. F. Schutz, *Phys. Rev. D* **72**, 063006 (2005).
- ¹²P. R. Brady and T. Creighton, *Phys. Rev. D* **61**, 082001 (2000).
- ¹³V. Dergachev, LIGO Technical Document No. LIGO-T050186, 2005.
- ¹⁴G. Mendell and M. Landry, LIGO Technical Document No. LIGO-T050003, 2005.
- ¹⁵S. Dhurandhar, B. Krishnan, H. Mukhopadhyay, and J. T. Whelan, *Phys. Rev. D* **77**, 082001 (2008).
- ¹⁶A. Viterbi, *IEEE Trans. Inf. Theory* **13**, 260 (1967).
- ¹⁷A. Viterbi, *IEEE Trans. Commun. Technol.* **19**, 751 (1971).
- ¹⁸G. D. Forney, *Proc. IEEE* **61**, 268 (1973).
- ¹⁹B. G. Quinn and E. J. Hannan, *The Estimation and Tracking of Frequency* (Cambridge University Press, 2001), Vol. 9, p. 219.
- ²⁰S. Suvorova, L. Sun, A. Melatos, W. Moran, and R. Evans, *Phys. Rev. D* **93**, 123009 (2016).
- ²¹L. Sun, A. Melatos, S. Suvorova, W. Moran, and R. Evans, *Phys. Rev. D* **97**, 043013 (2018).
- ²²J. Bayley, C. Messenger, and G. Woan, *Phys. Rev. D* **100**, 023006 (2019).
- ²³M. Ciobotaru, R. Teodorescu, and F. Blaabjerg, "A new single-phase PLL structure based on second order generalized integrator," in *2006 37th IEEE Power Electronics Specialists Conference, Jeju, South Korea, 18-22 June 2006* (IEEE, 2006), pp. 1–6.
- ²⁴P. Rodríguez, A. Luna, I. Candela, R. Muijal, R. Teodorescu, and F. Blaabjerg, *IEEE Trans. Ind. Electron.* **58**, 127 (2010).
- ²⁵J. Matas, M. Castilla, J. Miret, L. G. de Vicuña, and R. Guzman, *IEEE Trans. Ind. Electron.* **61**, 2139 (2013).
- ²⁶F. Xiao, L. Dong, L. Li, and X. Liao, *IEEE Trans. Power Electron.* **32**, 1713 (2016).
- ²⁷Z. Xin, X. Wang, Z. Qin, M. Lu, P. C. Loh, and F. Blaabjerg, *IEEE Trans. Power Electron.* **31**, 8068 (2016).
- ²⁸M. Karimi-Ghartemani, *IEEE Trans. Power Electron.* **28**, 4550 (2012).
- ²⁹M. Karimi-Ghartemani, *IEEE Trans. Ind. Electron.* **61**, 1464 (2014).
- ³⁰S. Golestan, J. M. Guerrero, and J. C. Vasquez, *IEEE Trans. Power Electron.* **32**, 9013 (2017).
- ³¹A. V. Cumming, B. Sorazu, E. Daw, G. D. Hammond, J. Hough, R. Jones, I. W. Martin, S. Rowan, K. A. Strain, and D. Williams, *Classical Quantum Gravity* **37**, 195019 (2020).
- ³²J. C. Gamazo-Real, E. Vázquez-Sánchez, and J. Gómez-Gil, *Sensors* **10**, 6901 (2010).
- ³³J. G. Proakis and M. Salehi, in *Digital Communications* (McGraw-Hill, 2008), Chap. 5.
- ³⁴Y. Nakaya and H. Mori, *Clin. Phys. Physiol. Meas.* **13**, 191 (1992).
- ³⁵F. J. Giessibl, *Rev. Mod. Phys.* **75**, 949 (2003).
- ³⁶M. Tse *et al.*, *Phys. Rev. Lett.* **123**, 231107 (2019).
- ³⁷I. J. Hollows, The mathematics of the IWAVE algorithm, https://git.ligo.org/edward.daw/iwave/-/blob/master/documents/paper_rsi/iwavepaper_mathematics.pdf; accessed 4 June 2021.
- ³⁸E. J. Daw, IWAVE git repository, <https://git.ligo.org/edward.daw/iwave>, 2021.
- ³⁹A. V. Oppenheim and R. W. Schaffer, *Digital Signal Processing*, 2nd ed. (Prentice-Hall, 1975).
- ⁴⁰R. Abbott *et al.*, *SoftwareX* **13**, 100658 (2021).
- ⁴¹E. J. Daw, *Nucl. Instrum. Methods Phys. Res., Sect. A* **921**, 50 (2019); [arXiv:1805.11523](https://arxiv.org/abs/1805.11523) [physics.ins-det].
- ⁴²M. Karimi-Ghartemani, *IEEE Trans. Ind. Electron.* **61**, 1464 (2013).
- ⁴³S. Gude and C.-C. Chu, *IEEE Trans. Ind. Appl.* **56**(1), 740–751 (2020).