



# Deconvolving Pulsar Signals with Cyclic Spectroscopy: A Systematic Evaluation

Timothy Dolch<sup>1,2</sup> , Dan R. Stinebring<sup>3</sup> , Glenn Jones<sup>4,5</sup>, Hengrui Zhu<sup>3</sup> , Ryan S. Lynch<sup>6</sup> , Tyler Cohen<sup>7</sup> , Paul B. Demorest<sup>8</sup> , Michael T. Lam<sup>9,10</sup> , Lina Levin<sup>11</sup> , Maura A. McLaughlin<sup>12,13</sup> , and Nipuni T. Palliyaguru<sup>14</sup>

<sup>1</sup> Department of Physics, Hillsdale College, 33 E. College Street, Hillsdale, MI 49242, USA

<sup>2</sup> Eureka Scientific, Inc., 2452 Delmer Street, Suite 100, Oakland, CA 94602-3017, USA

<sup>3</sup> Department of Physics and Astronomy, Oberlin College, Oberlin, OH 44074, USA

<sup>4</sup> Rigetti Computing, Inc., 775 Heinz Avenue, Berkeley, CA 94710, USA

<sup>5</sup> Columbia Astrophysics Laboratory, Columbia University, NY 10027, USA

<sup>6</sup> Green Bank Observatory, P.O. Box 2, Green Bank, WV 24944, USA

<sup>7</sup> Department of Physics, New Mexico Institute of Mining and Technology, 801 Leroy Place, Socorro, NM 87801, USA

<sup>8</sup> National Radio Astronomy Observatory, 1003 Lopezville Road, Socorro, NM 87801, USA

<sup>9</sup> School of Physics and Astronomy, Rochester Institute of Technology, Rochester, NY 14623, USA

<sup>10</sup> Laboratory for Multiwavelength Astronomy, Rochester Institute of Technology, Rochester, NY 14623, USA

<sup>11</sup> Jodrell Bank Centre for Astrophysics, University of Manchester, Manchester, M13 9PL, UK

<sup>12</sup> Department of Physics and Astronomy, West Virginia University, P.O. Box 6315, Morgantown, WV 26506, USA

<sup>13</sup> Center for Gravitational Waves and Cosmology, West Virginia University, Chestnut Ridge Research Building, Morgantown, WV 26505, USA

<sup>14</sup> Department of Physics & Astronomy Texas Tech University, Lubbock, TX 79409, USA

Received 2020 August 23; revised 2021 March 28; accepted 2021 March 31; published 2021 May 28

## Abstract

Radio pulsar signals are significantly perturbed by their propagation through the ionized interstellar medium. In addition to the frequency-dependent pulse times of arrival due to dispersion, pulse shapes are also distorted and shifted, having been scattered by the inhomogeneous interstellar plasma, affecting pulse arrival times. Understanding the degree to which scattering affects pulsar timing is important for gravitational-wave detection with pulsar timing arrays (PTAs), which depend on the reliability of pulsars as stable clocks with an uncertainty of  $\sim 100$  ns or less over  $\sim 10$  yr or more. Scattering can be described as a convolution of the intrinsic pulse shape with an impulse response function representing the effects of multipath propagation. In previous studies, the technique of cyclic spectroscopy has been applied to pulsar signals to deconvolve the effects of scattering from the original emitted signals, increasing the overall timing precision. We present an analysis of simulated data to test the quality of deconvolution using cyclic spectroscopy over a range of parameters characterizing interstellar scattering and pulsar signal-to-noise ratio (S/N). We show that cyclic spectroscopy is most effective for high S/N and/or highly scattered pulsars. We conclude that cyclic spectroscopy could play an important role in scattering correction to distant populations of highly scattered pulsars not currently included in PTAs. For future telescopes and for current instruments such as the Green Bank Telescope upgraded with the ultrawide bandwidth receiver, cyclic spectroscopy could potentially double the number of PTA-quality pulsars.

*Unified Astronomy Thesaurus concepts:* Millisecond pulsars (1062); Pulsars (1306); Radio pulsars (1353); Interstellar medium (847); Interstellar plasma (851); Radio astronomy (1338); Gravitational waves (678)

## 1. Introduction

Millisecond pulsars (MSPs) are excellent tools for long-period gravitational-wave (GW) detection. GWs are perturbations in the spacetime metric predicted by general relativity, first indirectly detected in a binary neutron star (Hulse & Taylor 1975) and then directly by the Laser Interferometer Gravitational-Wave Observatory (Abbott et al. 2016). Using pulsar timing for GW detection requires measurements of time-of-arrival (TOA) perturbations from many MSPs. Long-term monitoring of MSPs distributed widely across the sky (pulsar timing arrays; PTAs) is a robust method for nanohertz (i.e., light-year wavelength) GW detection (Sazhkin 1978; Foster & Backer 1990). The PTA method of detecting GWs with frequencies in the range  $10^{-9}$ – $10^{-7}$  Hz is complementary to laser interferometers, which probe GW frequencies of  $>10^{-6}$  Hz. PTAs are sensitive to GWs because metric disturbances induce a variation of the observed pulse arrival time; hence, the TOA will be delayed or advanced relative to a predicted arrival time. PTA collaborations include the North American Nanohertz Observatory for Gravitational Waves (NANOGrav; McLaughlin 2013), the European Pulsar Timing Array; Kramer & Champion 2013), and the Parkes Pulsar

Timing Array; Hobbs 2013; Manchester et al. 2013), working collectively as the IPTA (International Pulsar Timing Array; Hobbs et al. 2010; Manchester & IPTA 2013). The GW sources to which PTAs are likely the most sensitive are merging supermassive black hole binaries (Detweiler 1979; Hellings & Downs 1983, Arzoumanian et al. 2018). The NANOGrav Collaboration uses the Green Bank Telescope (GBT) at the Green Bank Observatory and formerly used the 300 m William E. Gordon telescope at the Arecibo Observatory (AO).

Rapidly rotating neutron stars, observed as pulsars, are remarkably stable astrophysical clocks. Acquiring a useful TOA data set for the purpose of GW detection involves collecting pulse arrival times every 1–4 weeks from each pulsar in the PTA for a decade or more. In the case of NANOGrav (Alam et al. 2021a, 2021b),  $\sim 20$  minute observations are taken every 1–3 weeks, depending on the pulsar. Pulses are averaged or “folded” at the appropriate pulse period, from which we obtain an average pulse profile shape over the observation’s duration. Each folded set of pulses is assigned an arrival time in frequency-dependent sub-bands. The folded pulse profiles are then compared to a timing model that accounts for the many phenomena that influence arrival times. The list includes many

timing perturbations from the ionized interstellar medium, or IISM, between Earth and the pulsar (Stinebring 2013; Dolch 2018), such as variations in dispersion measure (DM; a line-of-sight integral of the electron density in the IISM), scattering, and interstellar scintillation (ISS). These effects are all chromatic, unlike frequency-independent effects such as spin noise intrinsic to the pulsar and any GW signals present during observations.

Significant physical processes not explicitly included in a timing model will lead to significantly nonzero timing residuals, the differences between the measured arrival times and those predicted by the timing model. After accounting for these astrophysical and systematic effects, the final residuals may reveal evidence for GWs, or provide upper limits on their magnitude (Siemens et al. 2013). Modeling chromatic timing effects also provides invaluable ancillary data for studying turbulence, lensing, and other plasma structures in the IISM (Lam et al. 2018, Stinebring et al. 2019b).

Building on studies of PSR B1937+21 (Demorest 2011; Walker et al. 2013), and of other simulations (Palliyaguru et al. 2015), this paper explores the conditions under which scattering can be deconvolved using a signal processing technique known as cyclic spectroscopy (CS). Section 2 describes some specific aspects of the IISM to which deconvolution can be applied. Section 3 derives a figure of merit that can be used to predict how well coherent deconvolution will work for a given pulsar. In Section 4 we describe simulated data sets that represent a variety of pulsar observations. Section 5 shows the results of the simulations, with comparisons to observations. Section 6 makes predictions about the number of pulsars that could become PTA quality after applying CS. We discuss implications for GW detection and for radio astronomy in general in Section 7.

## 2. The Transfer Function of the IISM

To a high degree of accuracy, the ionized ISM acts as a linear filter on the radio waves passing through it (Hankins 1971). Using standard terminology from signal processing, we can characterize this either by the complex-valued voltage impulse response function (IRF), denoted  $h(t)$ , or equivalently, through the Fourier transform (FT) of  $h(t)$ , the transfer function  $H(\nu)$  (Hankins & Rickett 1975).

The dominant source of frequency-dependent delay due to the IISM is from plasma dispersion. This is time-dependent for a particular line of sight because of the combined motion of the pulsar, the observer, and the intervening medium. The observed DM at any given moment introduces a TOA delay  $\propto \text{DM}/\nu^2$ , where  $\nu$  is the observed radio frequency. The effects of DM variations on timing are typically represented by a single time-varying number, the broadband DM, although a more complete model would incorporate a frequency-dependent DM (Cordes et al. 2016). The transfer function of the IISM is dominated by cold plasma dispersion, which can be modeled as an all-pass quadratic chirp filter. In modern pulsar timing observations, the majority of this dispersive effect is removed by applying the inverse filter function to the received voltage data in real time, using a fixed value of the DM. This process is referred to as coherent dedispersion (Hankins & Rickett 1975). Stochastic and systematic variations about the fixed DM are then incorporated into the timing model.

### 2.1. Interstellar Scattering of Pulsar Signals

In addition to dispersion delay, the radio signal is scattered by inhomogeneities in the IISM. Since the unscattered angular size of the pulsar is exceedingly small, multipath scattering results in a variety of interference effects that can be used to assess and potentially mitigate the effects of scattering delay. The scattering is traditionally described by two regimes: diffractive and refractive (Rickett 1990) although the demarcation is not always sharp. Diffractive interstellar scintillation, DISS, refers to the short-timescale (minutes to hours) variation as the observer moves through the random diffraction pattern caused by the scattered ray bundle. This stochastic pattern can be easily visualized by plotting the pulsar spectrum as a function of time, resulting in a *dynamic spectrum*. Although random in nature, this pattern has a characteristic width in frequency and time resulting in bright islands of power known as *scintles* surrounded by regions of relatively weaker signal.

Larger scale inhomogeneities in the IISM can cause the ray bundle to expand or contract or, in extreme cases, break into multiple components that would represent multi-imaging events if the necessary angular resolution was available. These refractive interstellar scintillation effects have a timescale of days to weeks, governed by the length of time it takes for the multipath ray bundle to move transversely to the line of sight by its width, thereby guaranteeing a fresh (uncorrelated) signal path. It has been conventional to model the IISM as a single phase-changing screen with a Kolmogorov density structure (Cordes et al. 2006), and we take that approach here. Increasingly, however, pulsar scintillation studies are yielding a different picture of the IISM: one in which any sight line crosses multiple relatively thin scattering screens with highly nonstationary statistics (Stinebring et al. 2019a).

The IRF of the IISM is changing with time because of motion of the line of sight through the medium. It is useful, therefore, to define a long-time average of the intensity IRF. We will refer to this as the pulse broadening function (PBF), denoted as  $\bar{h}_I(t) = \langle |h(t)|^2 \rangle$ , with the angle brackets denoting a time average. A Kolmogorov spectrum of electron density fluctuations in a thin screen results in a PBF that is close in functional form to a one-sided exponential, although in practice the kernel has a broader wing (Coles et al. 2010). The amplitude of the IRF fluctuates about the square root of the mean PBF. The instantaneous IRF depends on the configuration of the IISM at a given time during an observation. The centroid of a PBF corresponds to a pulsar's characteristic scattering timescale, or  $\tau_s$ , a quantity often cited in the literature (Levin et al. 2016). The corresponding scintle width, or diffractive bandwidth  $\Delta\nu_d$  (Rickett 1990), is

$$\Delta\nu_d = \frac{C_1}{2\pi\tau_s}, \quad (1)$$

where  $C_1$  is a dimensionless constant of order unity. The effect of scattering is strongly frequency dependent:  $\tau_s \propto \nu^{-4.4}$ . The diffractive timescale, the timescale on which the scintillation pattern changes (also known as the scintillation timescale), scales as  $\propto \nu^{-2.2}$  (Cordes & Shannon 2010) and is typically minutes to hours for PTA observations. (The actual frequency dependencies for any given line of sight may differ significantly; our method in this paper does not rely on any particular exponent.)

A first-order correction of scattering delays as a function of time can be accomplished by simply subtracting  $\tau_s$  from TOAs. This reduces a scattering delay to a single number instead of an entire function. The  $\tau_s$  value is typically obtained through an autocorrelation function analysis of the scintillation structure, given that  $\tau_s$  is inversely proportional to the scintillation bandwidth. Levin et al. (2016) measures the long-term  $\tau_s$  for NANOGrav pulsars, while Turner et al. (2020) provide an updated analysis on the recent NANOGrav 12.5 yr data set. However, the correction by this single parameter  $\tau_s$  does not account for the distortion of the IRF shape away from the one-sided exponential ensemble average PBF, which may lead to higher order timing errors. On the other hand, by determining the IRF as a function of time and frequency through the cyclic spectrum, we can remove both the delay and the distortion of the pulse profile by a deconvolution.

Coherently deconvolving the IRF from the measured voltage waveform intrinsic to the pulsar provides a more accurate correction by separating the intrinsic pulsar signal from the scattering transfer function of the IISM (Jones et al. 2013). Several papers demonstrate a method for doing this using CS (Demorest 2011; Walker et al. 2013). This technique relies on computing the cyclic spectrum of the voltage waveform, which is defined for cyclostationary signals—e.g. any noise modulated by a periodic envelope (Gardner 1991; Antoni 2007). Pulsar emission is a particularly good example of a cyclostationary signal.

The deconvolution algorithm discussed in the following sections will be referred to as WDS, in reference to Walker et al. (2013). We note that the approach explored here and pioneered in early CS papers relies on a minimum of assumptions about the problem: the pulse profile and IRF do not change over the time span analyzed, and each pulse consists of uncorrelated noise (Hankins & Rickett 1975). In principle, no assumptions are needed about the distribution of scattering material along the line of sight or the inhomogeneity spectrum of that material. This is in contrast to intensity-only techniques that have been employed in the past (e.g., Bhat et al. 2004; Löhmer et al. 2004; Geyer & Karastergiou 2016; Geyer et al. 2017) that require specific assumptions about the functional form of the scattering kernel.

## 2.2. The Signal Model

When a pulsed signal arrives at a telescope, the standard practice is to detect pulses by dedispersing and folding data in real time. The full electric field (E-field) information is typically discarded to keep the recorded data volume manageable. Pulsar data therefore tends to be an average of pulse intensity profiles (for a particular polarization),  $I(\nu, t) = |E(\nu, t)|^2$ . For coherent WDS deconvolution, the phase information must be preserved in a cyclic spectrum. In practice, the receiver band is divided into frequency channels with a digital filterbank and the complex voltage time series corresponding to  $E(\nu, t)$ , for each frequency channel  $\nu$ , is recorded to disk. Through the methods outlined here, the ultimate aim of deconvolution will be to understand how the E-field of the original pulse train is changed by its interaction with the IISM (Walker & Stinebring 2005; Walker et al. 2008).

To describe the process of CS deconvolution, we take a simple model for the E-field of a pulsar signal passing through

the IISM and telescope:

$$E(t) = [p(t)N(t)] * h(t) + n_{\text{sys}}(t), \quad (2)$$

where  $p(t)$  is the original (unconvolved) pulse profile at time  $t$  mod  $P$ ,  $N(t)$  is the intrinsic modulated pulsar noise,  $h(t)$  is the IRF,  $P$  is the pulse period, and  $n_{\text{sys}}(t)$  is the sky and receiver noise present in the system, uncorrelated across pulse periods. This can also be written as

$$E(t) = X(t) * h(t) + n_{\text{sys}}(t), \quad (3)$$

in which  $X(t) = p(t)N(t)$ . In the frequency domain the signal model becomes

$$E(\nu) = [p(\nu) * N(\nu)]H(\nu) + n_{\text{sys}}(\nu) \quad (4)$$

$$= X(\nu)H(\nu) + n_{\text{sys}}(\nu), \quad (5)$$

using  $X$  instead of  $p * N$  because the convolution occurs upon emission at the pulsar.

The *cyclic spectrum* of  $E(t)$  is

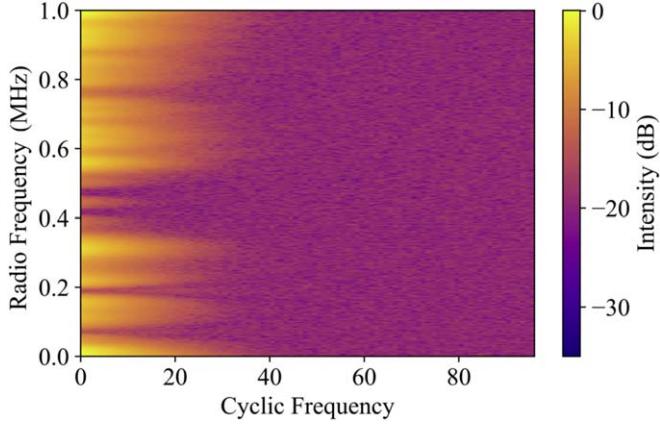
$$S_E(\nu, \alpha_k) = \langle E(\nu + \alpha_k/2)E^*(\nu - \alpha_k/2) \rangle, \quad (6)$$

where  $\nu$  is the radio frequency at when the signal is measured and  $\alpha_k = k/P$  is the cyclic frequency, also known as the *modulation frequency*. The integer  $k$  also refers to the number of pulse periods in the *time domain* that we would shift the signal and its conjugate before taking their product. The mean here is taken over an integer number of pulses. The resulting folded cyclic spectrum is complex valued with amplitude and phase for each  $(\nu, \alpha_k)$  pair, and is not defined for nonperiodic signals.

Encapsulating the complete effect of the IISM on the intrinsic (pre-ISM) pulsed signal requires preserving the E-field phase information. If the IISM were not present, we would have  $h(t) = 1$ . For every short integration time on the order of the scintillation timescale we can construct a cyclic spectrum as a function of  $\nu$  and  $\alpha_k$  instead of an intensity spectrum as a function of  $\nu$  only. The cyclic spectrum can be rewritten in the following form:

$$\begin{aligned} S_E(\nu, \alpha_k) &= \langle H(\nu + \alpha_k/2)H^*(\nu - \alpha_k/2) \\ &\quad \times X(\nu + \alpha_k/2)X^*(\nu - \alpha_k/2) \rangle \\ &= \langle H(\nu + \alpha_k/2)H^*(\nu - \alpha_k/2) \rangle S_x(\alpha_k), \end{aligned} \quad (7)$$

where  $S_x(\nu, \alpha_k)$  is the FT of the folded intrinsic pulse profile, which we will write as  $S_x(\alpha_k)$ , assuming that there is negligible pulse profile evolution with radio frequency across the relevant observing band. Because of this assumption, we calculate cyclic spectra over fairly narrow bandwidths, meaning that variations in radio frequency in that spectra will be due to the IISM alone. The shifts by  $\pm\alpha_k/2$  represent the range of possible phase differences induced in the E-field from the IISM. Due to the periodic modulation (i.e., the intrinsically cyclostationary nature) of the intrinsic pulse, the modulus of  $S_x(\alpha_k)$  is also called the harmonic profile. The cyclic spectrum,  $S_E(\nu, \alpha_k)$ , is then a 2D array that contains both folded E-field amplitude and E-field phase information. *Phase* in this context specifically refers to the electromagnetic phase induced only by the propagation of the intrinsic pulses through the IISM, not the pulse phase. The FT of the cyclic spectrum along the  $\alpha_k$ -axis yields the cyclic periodogram, which shows the pulsar signal



**Figure 1.** Idealized, simulated example: magnitude of the cyclic spectrum for a scattered pulsar signal. The color scaling is logarithmic in power with an arbitrary reference level. Power is seen as a function of both cyclic frequency  $\alpha_k$  (hertz) and radio frequency  $\nu$  (megahertz), chosen in this simulated example to be a 1 MHz band near 435 MHz.

averaged over pulse phase and radio frequency. The cyclic periodogram can produce finer frequency resolution than a standard filterbank for periodic signals.

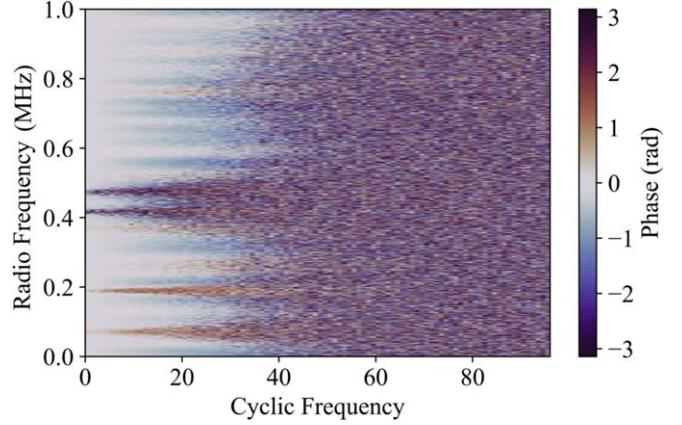
### 3. The Cyclic Merit, a Metric of Deconvolution Quality

Regardless of any later analysis one wishes to perform, the cyclic spectrum is itself an efficient way of collecting and storing the phase information from a given pulsar observation without resorting to large baseband data sets. The cyclic spectrum is also a tool for IISM deconvolution: the form of the spectrum expressed in Equation (7) suggests that  $H$  can be separated from  $S_x(\alpha_k)$  in principle. In practice this separation is nontrivial.

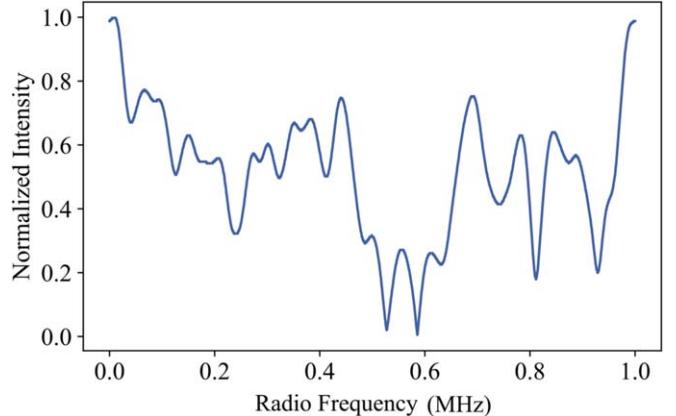
Detailed derivations of relevant CS quantities and their expected noise properties are provided in Demorest (2011) and Walker et al. (2013). In order to predict the fidelity of the IRF determined by the WDS algorithm, we derive a figure of merit related to the signal-to-noise ratio (S/N) of the cyclic spectrum.

#### 3.1. A Description of the Cyclic Spectrum by Example

We first consider some properties of the cyclic spectrum in the presence of interstellar scattering. For pulse-modulated noise, we consider the cyclic spectrum of a scattered pulse. In this example, the cyclic spectrum corresponds to a 1.6 ms pulse period (the same as PSR B1937+21) with a  $\sim 5\%$  duty cycle scattered with a  $4 \mu\text{s}$  scattering tail. These values represent a typical MSP that NANOGrav would routinely observe, with the period chosen to match that of PSR B1937+21. The amplitude of the cyclic spectrum is simply  $|H(\nu)| |S_x(\alpha_k)|$  (see Figure 1). The examples shown in Figures 1 and 2 have a pulse profile S/N of 70. In the blue regions of Figure 2, the phase of the cyclic spectrum  $\Phi_{S_E}$  remains coherent with  $\alpha_k$ . Deconvolving IISM effects using CS hinges on this phase structure, which represents the degeneracy between  $H(\nu + \alpha_k/2)$  and  $H(\nu - \alpha_k/2)$  being broken. In contrast, a nonperiodic source such as a quasar (in this noiseless case) would have a zero-valued cyclic spectrum in amplitude everywhere except for the  $\alpha_k = 0$  column, which is the standard spectrum of the source (Figure 3), and would have an entirely incoherent complex phase across radio frequency.



**Figure 2.** Idealized, simulated example: phase of the cyclic spectrum for a scattered pulsar signal. The color scaling is  $\pm\pi$  in phase. Phase is seen as a function of both cyclic frequency  $\alpha_k$  (hertz) and radio frequency  $\nu$  (megahertz), chosen in this simulated example to be over 1 MHz near 435 MHz. The zero-to-negative phase transition within the scintles is the phase slope resulting from the IISM. The cyclic spectrum makes this phase slope apparent. Figure 1 is scaled logarithmically; by contrast, the phase structure in this figure, scaled linearly, goes out to much higher harmonics than does the amplitude.



**Figure 3.** The radio frequency spectrum for the same simulated cyclic spectrum shown in Figures 1 and 2. This is the zero-valued column of Figure 1 along the radio frequency axis. Each peak is a *scintle* due to ISS.

#### 3.2. A Figure of Merit for Deconvolution by CS

The WDS algorithm searches for both the true (unscattered) pulse profile and for the IRF that best explains the measured CS. The algorithm performs this search by taking an initial guess at the IRF, usually a delta function, which corresponds to the specification  $\partial H(\nu)/\partial\nu = 0$  (no scintles) and a constant phase. The  $M$  complex samples that make up the IRF are then varied as a  $2M$ -dimensional parameter space, holding the assumed intrinsic profile constant. While this methodology inherently introduces many possible degeneracies, Walker et al. (2013) show that the cyclic spectrum structure due to scattering will not be significantly covariant with the intrinsic profile. The intrinsic profile can then be determined by iteratively solving several consecutive cyclic spectra, while refining the estimate of the intrinsic profile with each subsequent iteration (the *outer loop*). The assumed intrinsic profile (at any step of the outer loop that iterates through successive cyclic spectra) and the IRF that yielded the best-fit cyclic spectrum represent estimates of the intrinsic pulse profile and the IRF of the IISM. The WDS procedure in Walker et al. (2013) yields an IRF for the original MSP B1937+21 at 430 MHz, in which the best-fit IRF clearly

has an even more finely resolved scattering structure than that shown by earlier methods (Stinebring et al. 2001) in which astronomical unit-sized scattering structures in the IISM were observed. The intrinsic profile is consistent between two nearby frequencies treated independently, showing a high degree of convergence.

In our implementation of the WDS algorithm, we do not use the outer loop. This is because each simulation is limited to a single cyclic spectrum made from a single realization of an IRF. One could simulate a realistic sequence of IRFs that evolve with time to more fully characterize the WDS algorithm, but doing so is beyond the scope of this paper. In real observations, once a best-fit intrinsic profile is found, it can be used in repeated epochs, as it will generally vary much less significantly than the IRF on a diffractive timescale, likely making the inner loop more critical.

We find that WDS fits for the cyclic spectrum amplitude  $|S_E(\nu, \alpha_k)|$  in a relatively small number of iterations. This is not surprising because  $|S_E(\nu, 0)|$  is proportional to  $|H(\nu)|$ , providing a good initial estimate of the magnitude. As we move out to higher  $\alpha_k$  in  $|S_E(\nu, \alpha_k)|$ , the functional form across radio frequency is narrowed and overtaken by noise. Fitting the phase of  $S_E(\nu, \alpha_k)$  requires considerably more effort by the nonlinear optimization at the heart of WDS, for the following important reason. Radio frequency channels between scintles have little to no signal at low values of  $\alpha_k$ , which results in regions of phase ambiguity. Qualitatively observing the evolution of the fitting process in our runs of the code, the complex phase of  $H(\nu)$  often appears to converge long after its complex amplitude. See Walker et al. (2013) for a more detailed explanation.

The transfer function phase due to interfering E-field phases from the cyclic spectrum can be obtained as follows. In order to deconvolve the IISM along every ray path, we need E-field phase information contained in the phase of the transfer function  $H$  of the data we are fitting. Taking the complex phase of Equation (5), we have

$$\Phi_{S_E}(\nu, \alpha_k) = \Phi_H(\nu + \alpha_k/2) - \Phi_H(\nu - \alpha_k/2) + \phi_{S_x}(\alpha_k) \quad (8)$$

$$= \alpha_k \frac{\Phi_H(\nu + \alpha_k/2) - \Phi_H(\nu - \alpha_k/2)}{\alpha_k} + \phi_{S_x}(\alpha_k), \quad (9)$$

where  $\Phi_E$  is the phase of the cyclic spectrum at  $\nu$  and  $\alpha_k$ ,  $\Phi_H$  is the phase of the transfer function  $H(\nu, \alpha_k)$ , and  $\phi_{S_x}$  is the phase of the intrinsic profile harmonics. In the limit  $\alpha_k \ll \Delta\nu$ , where  $\Delta\nu$  is the radio frequency channelization, we have

$$\phi_{S_E}(\nu, \alpha_k) \approx \alpha_k \frac{d\Phi_H}{d\nu} + \phi_{S_x}(\alpha_k) \quad (10)$$

$$= \frac{k}{P} \frac{d\Phi_H}{d\nu} + \phi_{S_x}(\alpha_k). \quad (11)$$

The problem then becomes solving for  $\Phi_H$  and  $\phi_{S_x}$  in Equation (11). The WDS algorithm arrives at a solution through optimization techniques. The likelihood of extracting useful phase information from the cyclic spectrum can be found by calculating a figure of merit, which is the expected average value of  $\Phi$  divided by the uncertainty in that estimate (see the

Appendix for derivation and details):

$$m_{\text{cyc}} = \frac{\Phi}{\delta\Phi} = 2\pi \frac{\tau_s W_e}{P^2} (\text{S/N}) \sqrt{\sum_k k^2 a_k}, \quad (12)$$

where  $\tau_s$  is the scattering time,  $P$  is the pulsar period,  $W_e$  is the equivalent width<sup>15</sup> of the average pulse profile, S/N is the averaged S/N of the maximum of the pulse profile divided by the off-pulse,<sup>16</sup> and  $a_k \equiv A_k/A_0$ , where the  $A_k$  values are the amplitudes of the FT of the intensity pulse profile (note the importance of CS for MSPs, with  $m_{\text{cyc}} \propto P^{-1}$ , assuming  $W_e \propto P$ ).

For  $m_{\text{cyc}} \gg 1$ , the cyclic spectrum should provide enough information for WDS to successfully deconvolve the IRF to a reasonably high accuracy. Testing this assumption with simulations is critical because it is not obvious how specific features of the algorithm might limit recovery, such as phase wrapping, in which the phase of  $H(\nu)$  might be lost between scintles. Many other possible subtleties of WDS described in Walker et al. (2013) could also reduce the effectiveness of the algorithm.

A high S/N is clearly a criterion for successful deconvolution. In one sense, a long scattering tail provides more phase slope (see Figure 2) that can be leveraged by the deconvolution algorithm if it rises above the phase noise as we move toward higher harmonics or  $\alpha_k$  values. However, a highly scattered pulsar may still be more difficult for the WDS algorithm to deconvolve, given the phase-wrapping considerations just mentioned. We now proceed to describe tests of the effectiveness of the WDS algorithm on simulated data sets and evaluate the results in relation to the cyclic figure of merit.

#### 4. Deconvolving the Simulated Data Sets

We generated a suite of artificial IRFs and recovered them using WDS, looking at these results by varying the S/N of the de-scattered pulse profile and the simulated  $\tau_s$  in order to determine the region in the  $\tau_s$ -S/N parameter space in which CS is effective. WDS was re-implemented in Python in the publicly available PYCYC code.<sup>17</sup>

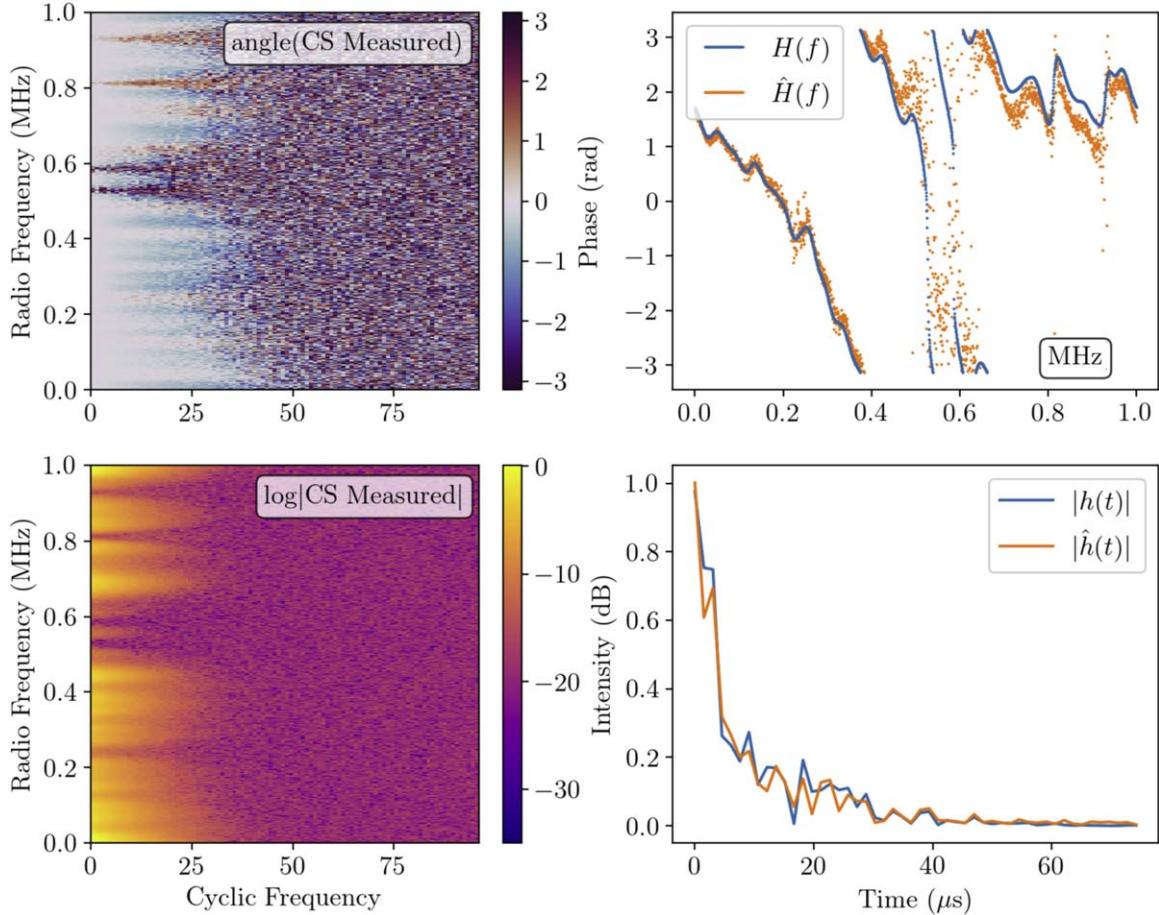
##### 4.1. Setup and Properties of Simulation Suite

We assumed 2048 frequency bins over a bandwidth of 1 MHz, in order to ensure resolved scintles; thousands of channels of frequency resolution across several scintles is routinely possible after employing CS, as with the periodic spectrum of PSR B1937+21 in Demorest (2011). The very narrow radio frequency channelization assumed here is not strictly necessary for the success of CS deconvolution; what is necessary is a large enough  $d\Phi_H/d\nu$  to accumulate CS phase while avoiding obscuration by random variations in  $\phi_{S_x}(\alpha_k)$  as in Section 3.2. The simulated pulsar has a duty cycle for its intrinsic pulse shape of about 5%, similar to PSR B1937+21 (Kramer et al. 1998).

<sup>15</sup>  $W_e$  is the width of a top-hat pulse with the same maximum value as a folded profile (Lorimer & Kramer 2004).

<sup>16</sup> CS does not require specifying the arbitrary *on* and *off* pulse regions used in standard radio pulsar analysis. Instead, the  $\alpha_k$  structure contains generalized information about the strength of the repeated signal across the profile. We assume, however, that the pulse profile has been smoothed to the optimal sharpness width,  $W_s$ , defined in the Appendix.

<sup>17</sup> <https://github.com/gitj/pycyc>



**Figure 4.** Simulated recovery of an IRF for an artificial pulsar signal with period 1.6 ms and scattering time constant  $\tau_s = 4 \mu\text{s}$ ,  $\text{S/N} = 70$ . The top left plot shows the phase of the cyclic spectrum as a function of cyclic frequency in hertz, including a stable phase due to scattering, similar to the toy model in Figure 2, with the color scaling representing complex phase. The bottom left plot shows the amplitude of the cyclic spectrum as a function of cyclic frequency in hertz, similar to the toy model in Figure 1, with a color scaling logarithmic in power. The upper right figure shows the complex phases of the simulated (blue) and recovered (orange) transfer functions, or FTs of the IRFs. The x-axis shows frequency (in megahertz). The bottom right plot shows the simulated (blue) and recovered (orange) impulse responses in the time domain, for only the first 75 of 2048 samples, scaled so that the peak value is 1 dB. See Section 3 for further details. The plots are generated from the PYCYC code.

We generated transfer functions each with characteristic width  $\tau_s$  by using the relevant routines in PYCYC to multiply the exponential IRF envelope by complex Gaussian noise. We then added noise directly to the simulated cyclic spectrum to simulate the presence of radiometer noise. In both stages the noise was white and Gaussian in each of the real and imaginary components. The one-sided exponential with multiplicative noise becomes the IRF, which the WDS algorithm will recover. An entirely realistic pulse profile would include two components of white noise in the time domain, one additive, corresponding to radiometer noise, and one component of amplitude-modulated noise (AMN), emanating from the emission mechanism from the pulsar itself. Doing so is beyond the scope of this paper, as such an arrangement would be too expensive computationally. Adding artificial noise directly to the cyclic spectrum is slightly unrealistic because each element of the 2D cyclic spectrum array receives an independently generated noise value. In reality, the noise and signal values in a cyclic spectrum are not independent from bin to bin because the values result from the correlation products of the measured voltage data. The purpose of the simulation, however, is to explore the conditions under which a phase slope can rise above the noise present in a cyclic spectrum, regardless of the detailed characteristics of that noise.

Figure 4 shows an example simulation of a realistic IRF and its resulting scintillation structure after a converged WDS fitting process. The bottom right panel shows the time-domain magnitude of the simulated IRF (blue) and the recovered IRF (orange). The signal bandwidth is 1 MHz and the pulse period is 1.6 ms in the artificial pulsar signal. The number of frequency channels is chosen to reduce computation time. It is not necessary to simulate a range of bandwidths because the important quantity to vary is the number of scintles across the band. In the example shown in the figure,  $\tau_s$  is the  $1/e$  timescale in the blue curve in the bottom right. For a 1.6 ms pulsar,  $5 \mu\text{s}$  of scattering is a small fraction of the pulsar’s period, which is why the noisy one-sided exponential is only visible in the first  $\sim 20$  of the 2048 time bins of the IRF. The upper right panel shows the complex phase of the transfer function, or FT of the IRF. As demonstrated in Section 2, the slope of the phase of the transfer function is proportional to the scattering timescale, or the centroid of the IRF in the bottom right panel. The bottom left and upper left panels show, respectively, the amplitude and phase of the complex cyclic spectrum array, to which WDS was applied and from which the results in the upper right and lower right panels were extracted. The color scale of the amplitude plot is in logarithmic units, and the scaling of the phase plot ranges from  $-\pi$  to  $\pi$ . In the

amplitude of the cyclic spectrum, the magnitude of  $H(\nu)$  is discernible along the vertical direction. Each scintle decays with increasing harmonic in the horizontal direction. In the phase of the cyclic spectrum a slope is apparent in the horizontal direction. The degree to which this phase stability is measurable across harmonics becomes a diagnostic for how well IISM deconvolution is possible.

The cyclic merit  $m_{\text{cyc}}$  of the example in Figure 4 is  $\sim 4$ . While the quality of recovery in this case is reasonably good by eye, the low cyclic merit value implies that this example is an outlier among iterations that would generally not allow for reliable CS deconvolution.

The upper right panel of Figure 4 also demonstrates the inherent difficulty in the WDS fitting process caused by phase ambiguities. In the example shown, the lowest quality fit around 0.6 MHz is caused by a low S/N in that region of the cyclic spectrum, which is visible in the upper left panel. For a pulsar signal scattered significantly more than this example, the number of phase wraps may become so large that the phase of  $H(\nu)$  may become indistinguishable from noise in the low-amplitude regions of the bandwidth.

Our goal is to find the quality of IRF reconstruction in both complex amplitude and phase. We define 36 *cells* in a two-parameter grid of simulated pulsar signals, varying pulse profile S/N and scattering timescale parameters, to see how the WDS algorithm behaves in practice for a variety of possible astrophysical signals.

We chose scattering times of 1, 2, 4, 8, 16, 32, 64, 128, and 256  $\mu\text{s}$  for the  $\tau_s$  of a B1937-like pulsar (period 1.6 ms). We used full-bandwidth intrinsic pulse profile S/N (peak to off-pulse rms) values of 20, 70, 650, and 2600. For each combination of  $\tau_s$  and S/N, we ran at least 60 simulations with different random number seeds determining both the radiometer noise and the realization of the scattering variations modulating the one-sided exponential. Each simulation had a limit of 3000 objective function evaluations by the nonlinear optimizer. Convergence occurs when successive function values differ by a factor of approximately the machine  $\epsilon$ , or the smallest numerical step possible on a given processor, which is an extremely high standard of convergence. Most runs converged long before 3000 objective function evaluations.

#### 4.2. Measuring the Quality of IRF Recovery

Once WDS determined a best-fit IRF for a particular simulation, we calculated the quality of recovery by comparing the input and output IRFs.<sup>18</sup> For the  $>60$  simulations in a cell, we define a goodness-of-fit metric  $\Delta$  as

$$\Delta \equiv \frac{1}{3\tau_s} \sum_{i=0}^J \left( \frac{h_{\text{input}}(t_i)}{\sum_{j=0}^J h_{\text{input}}(t_j)} - \frac{h_{\text{output}}(t_i)}{\sum_{j=0}^J h_{\text{output}}(t_j)} \right)^2, \quad (13)$$

where  $J$  is determined such that  $J\delta t = 3\tau_s$ , where  $\delta t$  is the sampling interval ( $t_J = 3\tau_s$ ). By construction, the simulated

<sup>18</sup> In reality, the true complex IRF will have an amplitude of  $\sqrt{2/\tau_s}$  guaranteeing its normalization to 1, but here we are only interested in recovering the functional form of the complex IRF. Once a best-fit IRF and (via an outer loop) intrinsic profile are obtained, their convolution can always be rescaled to match the flux of the measured profile.

IRFs (each  $h_{\text{input}}$  above), are one-sided exponential functions with  $1/e$  widths of  $\tau_s$  and multiplicative random noise but otherwise noiseless. The multiplicative noise is a simulated astrophysical signal with each IRF representing a different realization of the IISM. The radiometer noise, the observational noise added to the simulated cyclic spectra, is present in both the recovered intrinsic pulse profile and the recovered IRF. Normalizing IRFs to a value of 1 is difficult in the presence of the additive noise acquired in the WDS fitting process, in which each iteration of  $h$  depends on most recent best-fit cyclic spectrum because the normalization will be influenced by the noise present in  $h(t)$  when  $t \gg \tau_s$ . We instead perform our normalization only out to  $3\tau_s$ , where the IRF signal is present, and then take the sum-of-squares value out to  $3\tau_s$  as well. We finally divide by  $3\tau_s$  to compare the goodness-of-fit per time bin, so as to be able to compare simulations with different  $\tau_s$  values. The parameter  $\Delta$  is taken to be the typical quality of reconstruction for a given  $\tau_s$  and S/N ratio. In the deconvolution example of Figure 4,  $\log_{10}\Delta = -3.1$ , which is  $\sim 1.3\sigma$  above the mean of our simulated  $\log_{10}\Delta$  values (lower  $\log_{10}\Delta$  representing better fit quality), consistent with the  $m_{\text{cyc}}$  range that represents varying recovery quality as stated in 4.1.

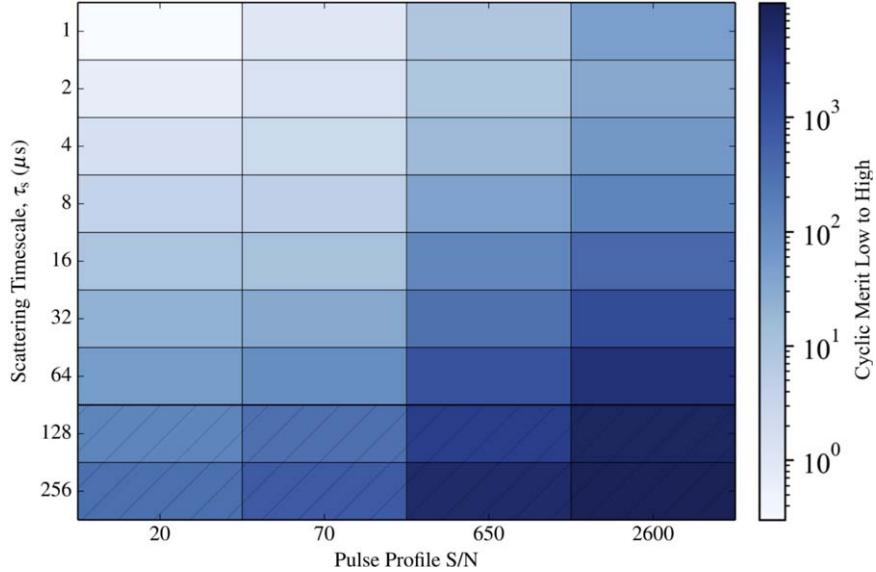
The quantity  $\Delta$  used in this work is different than the *demerit* used in Walker et al. (2013), a quantity that compares the data to a model cyclic spectrum at each step of the fitting process. The WDS algorithm minimizes the demerit. Here, we are instead comparing a simulated IRF to the recovered best-fit IRF. In Appendix B of Walker et al. (2013), the authors show that quality of fit (the positive curvatures or second derivatives of demerit) is proportional to  $F^2$ , where  $F$  is the total pulsed flux in a pulsar signal. For a fixed S/N in our simulations, where S/N is the peak to off-pulse rms, increasing  $\tau_s$  is equivalent to increasing  $F$  in the suite of simulated pulsars. For our high  $\tau_s$ -valued and/or high S/N-valued individual simulations, then, the metric  $\Delta$  should be closely related to the best-fit demerits from the WDS algorithm. However, the  $\tau_s$  and S/N values at which the recovered IRFs become unreliable is not straightforward, which is why we propose  $m_{\text{cyc}}$  as a new metric, to be verified in the following section with simulations.

## 5. Results and Comparisons with Observations

Using the simulations, quality measurements, and theoretical cyclic merit from the previous two sections, we compare a suite of simulations to test the dependence of  $m_{\text{cyc}}$  on  $\tau_s$  and S/N. We also compare several observations to the theoretical cyclic merit predictions, establishing a fiducial metric for the reliability of routine deconvolution of the IISM.

### 5.1. Measuring the Simulated Cyclic Merit Relationship

Results from the simulations are summarized in Figure 5. Each cell in the grid corresponds to a combination of  $\tau_s$  and S/N ratio. The color scale corresponds to the logarithm of the  $m_{\text{cyc}}$  values described in Section 3. The  $\Delta$  values from each corresponding simulated and recovered IRF pair were calibrated to the theoretical  $m_{\text{cyc}}$  values by the procedure to be described in this section, and then rescaled to arrive at the merit-based color scaling in Figure 5. Note that while the vertical axis shows  $\tau_s$  values that are powers of two, the S/N values in the horizontal axis increase by a factor of  $\sim 3.5$  except from 70–2600, which increases by almost two of the



**Figure 5.** Quality of IRF recovery for simulated pulsars using the WDS algorithm (Section 5.1). Higher values in colorbar mean better quality of recovery. The values in each cell were originally goodness-of-fit metrics  $\Delta$ , then reassigned to cyclic merit values  $m_{\text{cyc}}$  (Equation (12)), as described in Section 5.2. Values of  $\tau_s$  are in microseconds. Each cell corresponds to  $>60$  simulations, with the color corresponding to the median  $m_{\text{cyc}}$  value. Simulations in the crosshatched region are those in which scintles are not resolved by the fixed number of radio frequency channels across the simulations performed here, in which case the recovered IRFs are broad one-sided exponentials. The simulations verify that the effectiveness of CS deconvolution is proportional to both pulse profile S/N and scattering timescale  $\tau_s$ .

logarithmic intervals. We did not simulate the missing middle column to save computing time; the gap is taken into account in our linear regression procedure.

For the B1937+21-like artificial pulsar used in the simulations (similar in period and in main pulse width, with no interpulse in simulations), we relate our goodness-of-fit metric  $\Delta$  to  $m_{\text{cyc}}$  by

$$\Delta = A m_{\text{cyc}}^{-\mu} \quad (14)$$

because we expect an inversely monotonic relationship between  $\Delta$  and  $m_{\text{cyc}}$  across many orders of magnitude in  $\tau_s$  and S/N, but we do not know the details of how the two will correspond at low  $m_{\text{cyc}}$ , as discussed in Section 4.2. We assume  $\mu$  is a constant with respect to  $\tau_s$  and S/N. Although we already have a theoretical expectation for how  $m_{\text{cyc}}$  will depend on  $\tau_s$  and S/N from Equation (A11), we parameterize possible power-law dependencies as

$$m_{\text{cyc}} \propto \tau_s^\beta (\text{S/N})^\gamma. \quad (15)$$

We are interested in verifying that the simulation results demonstrate that  $\beta = 1$  and  $\gamma = 1$ , which would be consistent with the derivation in the Appendix. Rewriting this in logarithmic form, Equation (15) becomes

$$\log_{10} m_{\text{cyc}} = \beta \log_{10} \tau_s + \gamma \log_{10} (\text{S/N}) + C, \quad (16)$$

where  $C$  is a constant derived from pulsar-specific parameters in Equation (A11). We can then substitute Equation (14) into Equation (16) to obtain

$$-\frac{1}{\mu} \log_{10} \Delta = \beta \log_{10} \tau_s + \gamma \log_{10} (\text{S/N}) + D, \quad (17)$$

where  $D$  depends on both  $A$  and  $C$ . Proceeding, we note that because

$$\log_{10} \Delta = -\mu \beta \log_{10} \tau_s - \mu \gamma \log_{10} (\text{S/N}) - \mu D, \quad (18)$$

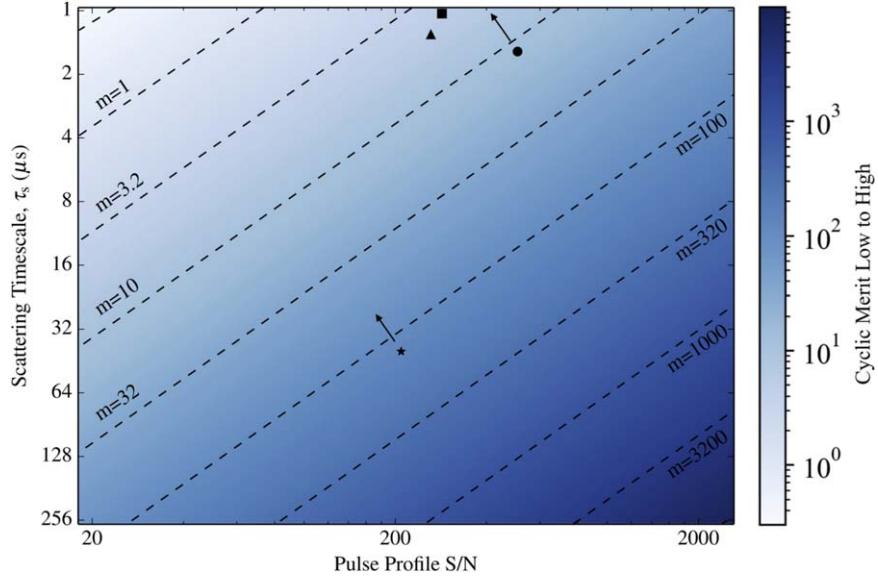
we cannot measure  $\beta$  and  $\gamma$  from the results of our simulations because each will be covariant with  $\mu$ . However, we can measure the ratio  $\gamma/\beta$ , which should be 1, by least-squares linear regression. This is equivalent to measuring the direction of the gradient of  $\log_{10} \Delta$  but not its magnitude. Because we have derived  $\beta = 1$  and  $\gamma = 1$  in the Appendix, the gradient of  $\Delta$  should point in the (1, 1) direction in  $\log_{10} \tau_s$ - $\log_{10} (\text{S/N})$  space. We ignore the constant  $\mu D$  in our fitting.

The distribution of the  $\Delta$  values within each cell was approximately Gaussian. Across cells, the median  $\Delta$  values appear approximately coplanar in Figure 5 by eye. An exception is the upper right cell, which is a local minimum compared to the surrounding cell. We examined the individual input and output IRFs and found that because the S/N was very high in this region, most  $\Delta$  values were very near zero. The statistical distribution of the  $\Delta$  values was different from the apparently Gaussian distribution in other cells, most likely because the statistics are determined more by the numerical properties of the computational process than by the simulated IISM properties. In the following fit, we neglect this outlier upper right cell.

A two-parameter unweighted linear regression of  $\log_{10} \Delta$  with  $\log_{10} \tau_s$  and  $\log_{10} (\text{S/N})$  yields slopes of  $-\mu \beta = -1.86 \pm 0.07$  and  $-\mu \gamma = -1.92 \pm 0.07$ , respectively, with an  $r^2$  regression statistic of 0.98. The slope ratio  $\gamma/\beta$  is therefore  $0.98 \pm 0.05$ . With these nearly coplanar values and the expected gradient direction, we compute the  $m_{\text{cyc}}$  values from Equation (12) for a B1937+21-like pulsar and assign those values to the cells, shown in the color scale of Figure 5.

The bottom crosshatched two rows of Figure 5 are those in which the scintles are not resolved, and in which, as a result, the success of the WDS algorithm in those cells is not representative of the full capability of the algorithm.

The radio frequency channel size here is 1 MHz/2048 = 0.488 kHz, which is held constant over all the simulations so that  $\tau_s$  and S/N are the only variables changed. By



**Figure 6.** Theoretical quality of impulse response function recovery (Section 5.2) from Equation (12). Four examples of CS deconvolution on real pulsars are also shown: PSR J2317+1439 at 327 MHz (triangle), PSR B1937+21 at 430 MHz (star; from Walker et al. 2013), PSR J1713+0747 at 327 MHz (square), and PSR B1937+21 at 1410 MHz (circle). Of the four real pulsar examples, the upper three real pulsars did not successfully recover an IRF, which is consistent with their low cyclic merit ( $m_{\text{cyc}}$ ) values. The predicted cyclic merit values shown in the color scale are for a generic pulsar with the period and equivalent width of PSR B1937+21, and the pulse profile shape of the PSR B1937+21 main pulse only. Pulsars J1713+0747 and J2317+1439 therefore have  $m_{\text{cyc}}$  values that differ slightly from those shown. The arrows next to both observations of PSR B1937+21 indicate that in reality, the interpulse reduces the cyclic merit by a factor of  $\sim 2$ , resulting in three of the four observations with  $m_{\text{cyc}} < 10$ , those three that we also found to not regularly converge on a best-fit IRF. Note that a pulsar's S/N is entirely dependent on the telescope sensitivity and on the folding time. Upcoming telescopes will move more pulsars to the right. Inclusion of significantly scattered pulsars not currently used in timing programs will further populate the bottom of the plot with high  $m_{\text{cyc}}$  pulsars.

Equation (1), for  $\tau_s = 64\mu\text{s}$ ,  $\Delta\nu_d = 2.5\text{ kHz}$ , and we have  $\sim 5$  channels per scintle. At  $\tau = 128\mu\text{s}$  we have one or two channels per scintle, making resolving the scintles difficult. Because the scintillation structure across radio frequency corresponds to the multiplicative noise structure in an IRF, unresolved scintles are equivalent to a one-sided exponential IRF with no additional structure. The situation of unresolved scintles often occurs when observing pulsars at radio frequencies  $< 100\text{ MHz}$ ; scattering timescales are measured by taking an intrinsic pulse shape derived from higher frequencies and convolving with one-sided exponentials until the model profiles match the data (Bansal et al. 2019). The crosshatched rows represent simulations with no IISM information gained over the method of successive convolutions with smooth, one-sided exponential IRFs. However, we are not aware of any reasons why WDS would perform poorly for the highly scattered pulsars represented in the bottom two rows of Figure 5 if we had used a greater number of channels. The slope ratio  $\gamma/\beta$  for Figure 5, removing both the upper right outlier and the crosshatched cells, is  $1.02 \pm 0.08$  with an  $r^2$  regression statistic of 0.97, consistent with the prediction from Equation (12). The S/N in Figure 5 is integrated over the band, not the single-channel S/N.

### 5.2. Comparisons of Figure of Merit with Observations

In Figure 6, we show  $m_{\text{cyc}}$  predictions using, as in Figure 5, the period and main pulse  $W_e$  for PSR B1937+21 as representative quantities for a generic MSP. We then compare to data.

The observations used here, in addition to the 430 MHz data on PSR B1937+21 used in Demorest (2011) and Walker et al. (2013), were taken from three separate campaigns at Arecibo

on pulsars PSR J1713+0747 (4.6 ms period) at 327 MHz, PSR B1937+21 (1.6 ms period) at 1.4 GHz, and PSR J2317+1439 (3.4 ms period) at 327 MHz. All three data sets were baseband. The PSR J1713+0747 data (AO P2627, PI: Palliyaguru) were taken at 430 MHz on 2011 September 19 with Mock spectrometers as backend receivers and were 40 MHz wide in bandwidth, with a 10 MHz slice used for the deconvolution attempt. The PSR B1937+21 data (AO P2676, PI: Dolch) were taken on 2012 September 19 and the PSR J2317+1439 data (AO P2824, PI: Stinebring) were taken on 2013 August 2, both with the PUPPI (Puerto Rican Ultimate Pulsar Processing Instrument) backend (DuPlain et al. 2008) with a 200 MHz and a 50 MHz bandwidth, respectively. Data were folded into cyclic spectra using the `-cyclic` option in the DSSPSR package (Demorest 2011; van Straten & Bailes 2011). The deconvolutions used 25 MHz of the available bandwidth unless stated otherwise, which simplified the data set because of band gaps between channels. The cyclic spectrum was accumulated in integrations such that the duration of each integration was less than the diffractive timescale, so that the IRF would be stable throughout each integration. The S/N of each data set was high enough that scintles were visually apparent in dynamic spectra formed from the data. For most subintegrations of data, the WDS procedure was applied and no converging IRF was found. We attempted a number of different starting conditions, varying the shift in phase between input profile and the measured profile, but in none of these cases did the algorithm yield a coherent solution. The same stopping and convergence criteria were used on both the data and the simulations. We found one well-fitting IRF in a subintegration of the PSR B1937+21 data at 1.5 GHz, but based on the cyclic merit values in Figure 5, we expect an occasional convergence

by chance that may nor may not represent the true underlying state of the IISM during that subintegration.

The examples of deconvolution attempts on real data are plotted in Figure 6 along with the converging example from (Demorest 2011) and Walker et al. (2013). The S/N values for the data new to this paper were computed using a 25 MHz bandwidth. The three non-converging examples are in a region of  $\tau_s$ -S/N space predicted to have poorer deconvolution quality than the converging example by an order of magnitude. The  $\tau_s$  values and each pulsar’s measured frequency power law used here are from the NANOGrav 12.5 yr data set (Turner et al. 2020), unless otherwise stated. We recognize these values are representative of highly variable IISM structures along a pulsar’s line of sight (the variability motivating CS deconvolution) but many reported values in the literature differ. The  $a_k$  values (Equation (12)) used were also from the NANOGrav 12.5 yr data set (Alam et al. 2021a). The  $m_{\text{cyc}}$  value for PSR J1713+0747 at 430 MHz ( $\tau_s = 1.0 \mu\text{s}$ ; S/N = 285) was 4.3. For PSR B1937+21 at 1.4 GHz<sup>19</sup> ( $\tau_s = 1.6 \mu\text{s}$ ; S/N = 506),  $m_{\text{cyc}} = 12.2$ . PSR B1937+21 at 430 MHz had  $m_{\text{cyc}} = 125$  ( $\tau_s = 40 \mu\text{s}$ ; S/N = 209, Walker et al. 2013; Demorest 2011, using the  $-3.66$  frequency power-law dependence of  $\tau_s$  in Ramachandran et al. 2006 at a frequency of 1273 MHz where our baseband observations were located). Finally, PSR J2317+1439 at 327 MHz ( $\tau_s = 1.3 \mu\text{s}$ ; S/N = 262) had  $m_{\text{cyc}} = 5.1$ .

In reality, these four pulsar observations will correspond to different color scales in Figure 6 because they have different values of  $P$ ,  $W$ , and  $\sum_k k^2 a_k$  in Equation (12), whereas the color scale of Figure 5 is based on PSR B1937+21. The  $m_{\text{cyc}}$  values for PSR J1713+0747 and PSR J2317+1439 are slightly higher and lower, respectively, due to the sharper pulse of the second producing more  $a_k$  terms in Equation (12). Our stated  $m_{\text{cyc}}$  values for PSR B1937+21 are greater than the true  $m_{\text{cyc}}$  values by a factor of  $\sim 2$  because the presence of the interpulse reduces  $\sum_k k^2 a_k$  in Equation (12). The arrows in Figure 6 point to the true  $m_{\text{cyc}}$  contour line where the two PSR B1937+21 observations should reside. For PSR B1937+21 at 430 MHz (Walker et al. 2013), the actual  $m_{\text{cyc}}$  value is 66. Because our simulations were done on a single-component pulse profile representing a generic MSP, in Figure 6 we still plot the  $m_{\text{cyc}}$  value from PSR B1937+21’s main component  $m_{\text{cyc}}$  only. The results remain consistent with the concept of  $m_{\text{cyc}}$  as a unitless detection *signal* of a scattering tail. Taking the interpulse into account, the three unsuccessful CS deconvolution attempts have  $m_{\text{cyc}} < 10$  and the successful attempt has  $m_{\text{cyc}} = 66$ , which is why we use  $m_{\text{cyc}} = 10$  as an absolute minimum requirement for CS-enhanced pulsar timing in Section 6.

The predictions in Figure 6, verified by the simulations in Figure 5, suggest that despite the large amount of phase wrapping, a pulsar in the bottom left of Figure 6—highly scattered but low S/N—could be successfully IISM deconvolved. The large amount of scattering would mean a high duty cycle in the scattered pulse profile. The de-scattered profile in such a case would both be scattering corrected in terms of its TOAs (by eliminating variations in the IRF), as well as turned into a better precision-timed pulsar with its low-duty-cycle unscattered pulse profile. However, this possibility would need case-by-case evaluation: as we move toward lower frequencies,  $\tau_s$  increases, but the diffractive timescale decreases (Cordes & Shannon 2010), limiting the

maximum integration time. A smaller  $\Delta\nu_d$  also lowers the S/N per scintle. A correct application of the figure of merit (Figure 6) to pulsar observations should take these considerations into account.

A more detailed simulation will be necessary to evaluate the following cases, which depend on the pulsar and the observation in question:

1. If the pulsar flux becomes comparable to the telescope’s system equivalent flux-density, then the on-pulse noise (due to the pulsar’s own AMN) will significantly exceed the off-pulse noise, unlike our simulations, which assume the same noise properties across all phases of the profile.
2. If the number of pulses averaged to obtain a cyclic spectrum is too low, then the effects of pulse phase jitter will be important (Lam et al. 2019), and the correlations between  $S_E(\nu, \alpha_k)$  bins in the cyclic spectrum array will become significant. Pulse phase jitter—the intrinsic white noise distributed TOAs of single pulses—cannot be reduced by telescope sensitivity, as the TOA uncertainties are intrinsic to the pulsar. The TOA uncertainties can be reduced by longer integration times, but this is not helpful for mitigating scattering in real time, with the longest integration time on the order of a diffractive timescale. For highly scattered pulsars, the diffractive timescale may in some cases become so low that the necessary S/N for CS deconvolution is not possible.<sup>20</sup>

As long as a particular pulsar observation is free of the above two properties, Figure 6 can be expected to reasonably predict whether CS deconvolution is feasible.

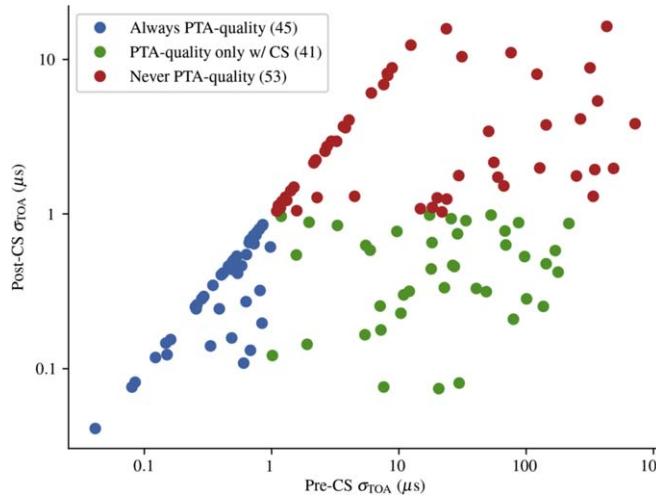
Figure 6 does not imply that CS is totally inapplicable to pulsar observations with  $m_{\text{cyc}} < 10$ . Even if deconvolution is infeasible, a cyclic spectrum still contains a detailed measurement of the presence of scattering, in addition to providing an extremely fine-frequency resolution. In some cases, CS deconvolution may be possible even when  $m_{\text{cyc}} < 10$ . What we have tested here is a straightforward, untailored application of the WDS algorithm. In practice, well-measured long-term PBFs for particular pulsars could serve as ideal initial conditions, instead of the delta functions used here. An ideal frequency binning could also change the behavior of the algorithm. Such approaches would model a line of sight and repeatedly update that model, much in the same way that pulsar timing models are repeatedly updated. The cyclic merit concept, then, must be understood as a measure of how reliably the WDS algorithm can quickly and routinely deconvolve, as in a real-time data processing pipeline.

## 6. Predictions for Pulsar Populations

Figure 6 suggests that many pulsars in addition to PSR B1937+21 could have high  $m_{\text{cyc}}$  values, most of which are not currently included in PTAs. We estimate the number of pulsars that could be successfully deconvolved in a future version of Figure 6 using the example of the GBT ultrawide bandwidth (UWB) receiver currently under construction. The GBT’s UWB receiver, similar to the UWB receiver developed by CSIRO for the Parkes Radio Telescope, spans the 0.7–4 GHz range.

<sup>19</sup> Here we use the  $\tau_s$  from the one well-fitting IRF.

<sup>20</sup> NANOGrav Memo #4: <http://nanograv.org/assets/files/memos/NANOGrav-Memo-004.pdf>



**Figure 7.** The improvement in timing precision for a simulated Galactic MSP population from the `PsRPopPy` code observed with the GBT’s UWB receiver, currently under construction. CS deconvolution has the potential to double the number of PTA-quality MSPs with  $\sigma_{\text{TOA}} < 1 \mu\text{s}$  at the GBT using the UWB receiver. The (blue) *Always PTA-quality MSPs* have similar properties to current NANOGrav pulsars. The (green) *PTA-quality only with CS* are those MSPs that could be included in the NANOGrav PTA because of the improved UWB sensitivity.

We simulated<sup>21</sup> the full Galactic MSP population using the `PsRPopPy` library (Bates et al. 2014) with the expected UWB system temperatures and gains across the bandwidth,  $T_{\text{obs}} = 15$  minutes (corresponding to typical DISS timescales), and assuming a Gaussian pulse profile shape. The pulsars had the requirement that  $\tau_s < P/2$ , because more highly scattered pulsars are unlikely to be detected in untargeted surveys. We also assumed that any pulsar with a large value of  $m_{\text{cyc}}$  would have sufficiently high S/N to be detected in future surveys. Finally, we calculated the rms timing residual,  $\sigma_{\text{TOA}}$ , expected (Lam et al. 2018) for two cases: (1) traditional processing, in which scattering is not mitigated for any MSP, and (2) CS processing, in which we set  $\tau_s = 0$  for any MSP with  $m_{\text{cyc}} \geq 10$  at 0.8 GHz (equivalent to completely mitigating all IISM noise at a typical NANOGrav timing frequency). We used  $m_{\text{cyc,full}}$  from Equation (A12), because we assume that the high-resolution dynamic spectra from CS will resolve scintles, enabling the usage of a bandwidth over which we would measure  $m_{\text{cyc}}$  that varies with  $\tau_s$ . The timing rms due to variations in the IRF is assumed to be 50% of  $\tau_s$  (Lam et al. 2018). The quantity  $\sigma_{\text{TOA}}$  also includes the rms from timing residuals from phenomena such as long-term DM variations.

The results are shown in Figure 7. We find that  $\sim 40$  MSPs that would otherwise not be of PTA quality (i.e.,  $\sigma_{\text{TOA}} > 1 \mu\text{s}$ ; the pulsars shown in green) become PTA quality with the use of CS and the GBT UWB receiver. This is compared to 45 MSPs that are always of PTA quality using the UWB receiver, even without the use of CS (the pulsars shown in blue). These population simulations suggest that CS deconvolution could double the number of PTA-quality MSPs, using the example of the GBT as one of many telescopes undergoing a receiver upgrade. We also find that the mean  $\sigma_{\text{TOA}}$  improves from  $\sim 45 \mu\text{s}$  without CS to  $\sim 2 \mu\text{s}$  with CS. We note that this simulated population includes pulsars similar to current

NANOGrav pulsars as well as pulsars either not yet discovered or not yet included as PTA pulsars due to high scattering and/or low S/N. The off-diagonal blue points in Figure 7 likely represent pulsars not yet discovered (see Section 7). The blue points along the diagonal represent a population similar to NANOGrav pulsars, both discovered and yet to be discovered, for which CS deconvolution does not improve timing.

Several caveats about the cyclic merit values used to create Figure 7 bear mentioning. Because detailed profile information is difficult to extract from the population simulations, we only use the  $a_1$  term in Equation (A12), which may *underestimate* the cyclic merit of each pulsar in the simulated population, and thus the number that would improve to PTA quality. On the other hand, we also assumed that  $m_{\text{cyc}} > 10$  implies complete scattering removal; the true degree of removal will depend on the details of the timing pipeline used (see Section 7). The threshold  $m_{\text{cyc}}$  might in practice also be larger than 10; all we can say definitively is that  $m_{\text{cyc}} \gg 1$  for scattering removal (with motivation for a threshold of 10 from our data in Section 5.2). Although we calculated  $\sigma_{\text{TOA}}$  at 0.8 GHz,  $m_{\text{cyc}}$  may also need to be  $> 10$  at more than one frequency across the UWB receiver’s bandwidth in order to improve timing precision, possibly reducing the number of pulsars correctable to PTA quality. As these caveats are difficult to quantify, we do not attempt to refine them further in this work. At the very least, the PTA-quality only with CS designation in Figure 7 suggests a *possible* population of pulsars with good timing quality that has yet to be exploited. The artificial pulsars with a factor of  $\sim 100$  improvement are not known PTA pulsars that would improve by such a significant factor (note the smaller improvement factors for the Always PTA-quality pulsars). The pulsars with a factor of  $\sim 100$  improvement in  $\sigma_{\text{TOA}}$  are rather those that in the absence of CS deconvolution would likely have never had a full timing solution prior to CS pulse sharpening. The improvements in Figure 7 must be understood with this last consideration especially in mind.

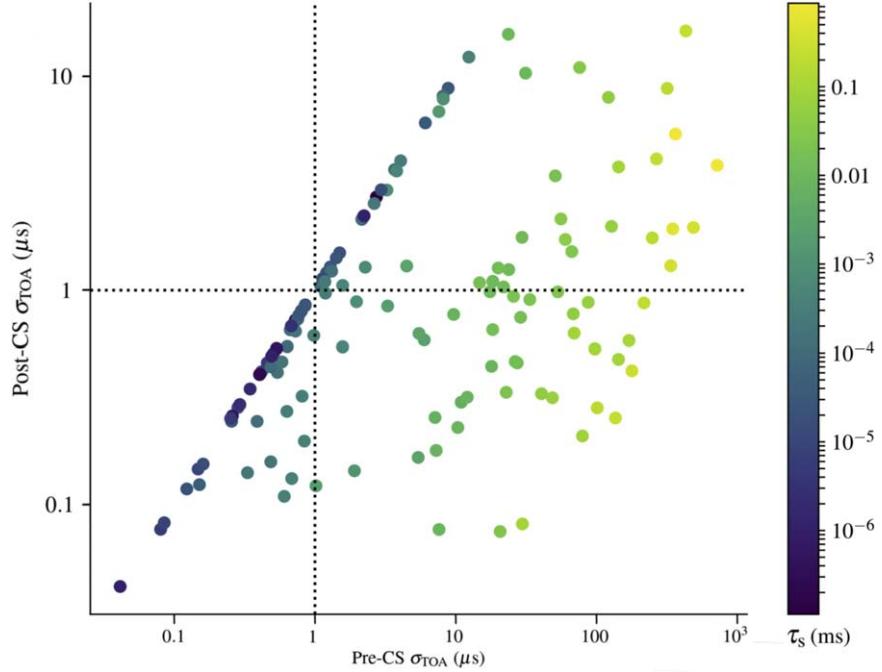
Figures 8, 9, and 10 show some interesting properties of the predicted populations. Pulsars with the largest improvement (Figure 8) tend to be those with large  $\tau_s$  values but also narrow intrinsic pulse widths. The probability density function of pulse widths is shown in Figure 9. We measured width by  $W_{\text{eff}}/P$ , where  $1/W_{\text{eff}}$  is the mean-squared derivative of the pulse profile (Downs & Reichley 1983; Lam et al. 2016). The distribution in Figure 9 tends toward narrower pulse widths for the PTA-quality only with CS subpopulation, suggesting that their narrow intrinsic widths in particular will enable better timing. While these have notably narrow pulse widths, the lognormal distribution from which widths are drawn matches the width distribution from the NANOGrav 12.5 yr data set. As Figure 10 shows, the distribution of the PTA-quality only with CS pulsars in Galactic coordinates tends toward the Galactic center, with the highest timing residual rms improvements corresponding to the most distant pulsars. Bright, intrinsically narrow, and highly scattered pulsars may very well acquire low timing rms values by means of CS.

## 7. Conclusions and Next Steps

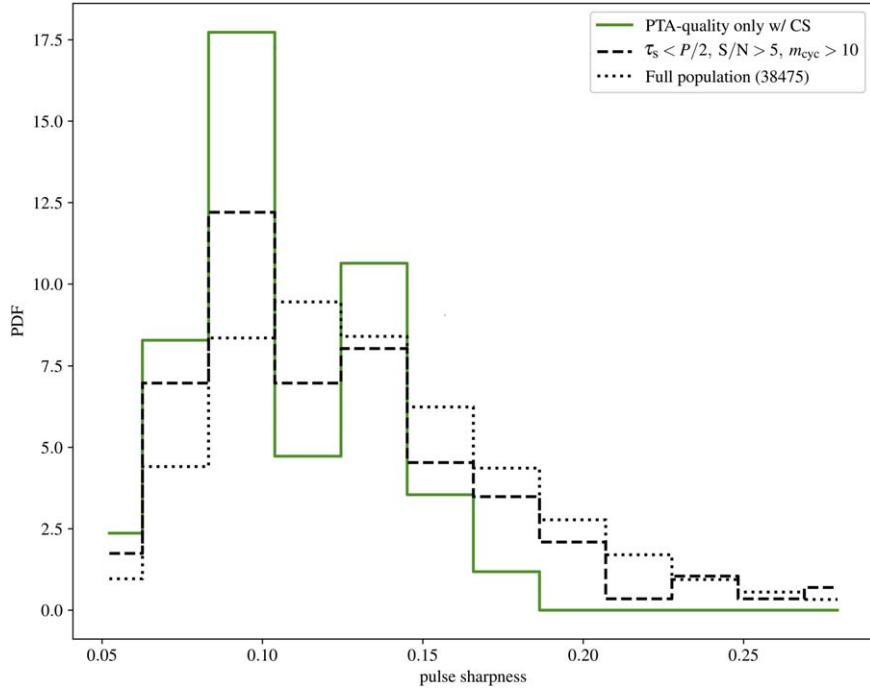
To reiterate, we expect several possible CS deconvolution benefits for PTA pulsars:

1. *Improved timing precision of current PTA pulsars* through scattering correction at some or all radio

<sup>21</sup> Simulation details and results are available at <http://www.aoc.nrao.edu/~tcohen/research/popsynth.shtml>.



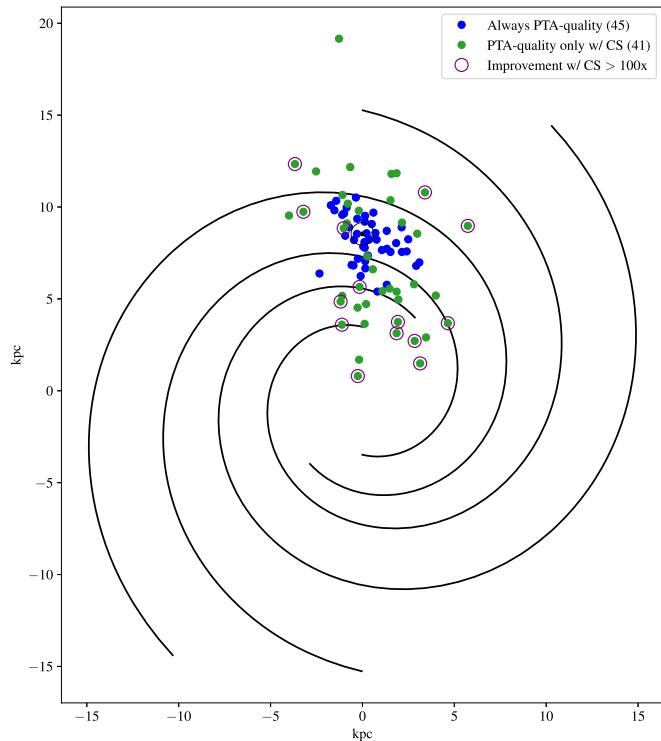
**Figure 8.** The same simulated pulsars are plotted as in Figure 7, but with the color scale representing the  $\tau_s$  value of each simulated pulsar. The pulsars with the greatest timing improvements tend to be those with the longest scattering tails, but the majority of the PTA-quality only with CS pulsars have scattering tails with  $\tau_s$  values of  $\sim 10\text{--}100\ \mu\text{s}$ , in accordance with Figure 5. The dotted lines correspond to the classification regions from Figure 7.



**Figure 9.** Probability density functions of the pulse sharpness measurements,  $W_{\text{eff}}/P$ , for the intrinsic (unscattered) pulse shapes from all the simulated pulsars in Figure 7. Both the PTA-quality only with CS and the entire population of simulated pulsars have sharper intrinsic pulses, enabling good timing quality if CS deconvolution is applied.

frequencies across a wide bandwidth. For any particular pulsar currently with PTA-quality timing precision, the timing correction may be one of many small second-order corrections, but at least for some current PTA pulsars, an S/N improvement with a new receiver and/or telescope may enable a more significant correction (blue points in Figure 7, allowing a movement to the right in Figure 6).

2. *Future pulsar discoveries*, likely to be more distant and highly scattered, that can *become* PTA quality through removal of scattering variations. CS deconvolution also applies to known pulsars currently with low timing precision, IISM-deconvolved to have PTA-quality timing precision.
3. *Enabling observations at lower frequencies* not presently used for pulsar timing where pulsars are intrinsically



**Figure 10.** Galactic spatial distribution of simulated pulsar set from Figure 7. The solar system is the black circled dot. Those pulsars acquiring PTA quality after CS deconvolution tend to be more distant than those PTA pulsars with no or small timing improvements after CS deconvolution. The small subset with extremely large timing improvements are the most distant of the PTA-quality only with CS pulsars, consistent with their significant scattering from Figure 9.

brighter, given the approximate relationship  $\tau_s \propto \nu^{-4.4}$ . Low-frequency telescopes and instruments could become more important for pulsar timing.

A procedure for using CS to improve pulsar timing precision is described in Palliyaguru et al. (2015). The approach uses a timing model, fitting the WDS-recovered IRF with a one-sided exponential. The results are encouraging, although as the paper states, work still remains to implement a full pulse-sharpening procedure. Our present work demonstrates that the pulse sharpening itself will likely be effective if  $m_{\text{cyc}} > 10$  for a particular pulsar; demonstration of a full timing pipeline involving CS deconvolution is a work in progress.

Several additional procedures will need to be incorporated into a full CS timing pipeline. Scattering correction methods are most relevant to low radio frequencies. Because high timing precision requires a DM measurement, which in turn requires high precision across a wide range of frequencies, scattering variations (and thus the changing IRF) would be highly covariant with any DM variations. If scattering variations are significant enough, then scattering delays can be partially absorbed into the DM measurement. Such an effect would, of course, represent a real source of error not otherwise taken into account, and CS would be useful for characterizing that error, but improved timing precision is not guaranteed. If scattering is high enough, then DM measurements could become contaminated (Levin et al. 2016). For most pulsars, the scintillation timescale, the timescale on which IRF realizations change, is significantly less than the typical timescale for DM variations, reducing the possibility of covariance.

For PSR B1937+21, CS deconvolution is unsuccessful only for the 1.4 GHz data presented here; at 430 MHz, Demorest (2011) and Walker et al. (2013) successfully obtained an IRF. Being heavily subject to red noise (Kaspi et al. 1994; Lam et al. 2017, 2016), PSR B1937+21 is not a useful contributor for most expected GW signals, given the necessity for stability on decades' timescales. It remains included in the list of routinely timed NANOGrav pulsars in order to compare arrival times between telescopes, and because of its very high short-term timing precision, making it sensitive to burst sources on timescales of approximately months (Finn & Lommen 2010). Such bursts could arise from nearby individual supermassive black hole binaries in highly elliptical orbits. PSR B1937+21's short-term GW sensitivity could be improved by routine CS deconvolution.

In the case of NANOGrav, the PTA's pulsars typically have  $\tau_s$  values  $< 10 \mu\text{s}$  (Levin et al. 2016; Turner et al. 2020), corresponding to low cyclic merit values. This is not surprising because the pulsars were selected for high timing precision. CS, however, remains an effective technique for obtaining extremely fine-frequency resolution, and should still be performed as a routine data product. Finer frequency resolution helps to better excise radio frequency interference (RFI) using the RFI itself as the cyclostationary *signal*, and can also provide higher frequency-resolution dynamic spectra, either for IISM studies or as a diagnostic tool for timing. For example, highly scattered epochs could be measured with fine-resolution cyclic periodograms (as demonstrated by Archibald et al. 2014) and then appropriately down weighted in a GW detection pipeline. Preliminary high-resolution dynamic spectra using CS were also created as part of the IPTA's 24 hr global campaign on PSR J1713+0747 (Dolch et al. 2014).

The practical difficulty of CS as a routine data product is the large volume of baseband data. Real-time CS backend receivers are currently being pursued in order to create cyclic spectra and/or IISM-deconvolved timing data as a direct data product, so that baseband data will not need to be saved for post-processing.

Within the current 12.5 yr NANOGrav data set (Alam et al. 2021a, 2021b), PSR J1903+0327 is an example of a high-DM NANOGrav pulsar that would likely benefit from such a routine procedure. Using NANOGrav 12.5 yr data set profiles, values from the ATNF pulsar catalog,<sup>22</sup> and the assumption of increasing flux at low frequencies with a  $-1.4$  power-law dependence, we estimate  $m_{\text{cyc}} \sim 4$  at 820 MHz for PSR J1903+0327. The true scattering timescale has not yet been estimated from resolved scintles. Improved frequency resolution with CS and/or improved telescope sensitivity could show a higher merit value. The calculation assumes an the AO sensitivity, but given the improved system temperature of the UWB receiver, this is still a pulsar worth exploring with the upgraded GBT. Other distant, high-DM pulsars for which good non-NANOGrav timing solutions already exist could be added to the array. However, these added pulsars must not be scattered so highly that low-frequency scintles are unresolvable (Cordes & Shannon 2010); in the NANOGrav 12.5 yr data set scattering measurements (Turner et al. 2020), all pulsars either have  $\tau_s < 1 \mu\text{s}$  at 820 MHz or have unresolvable scintles at the observed frequencies. Routine, fine-frequency dynamic spectra from CS may resolve currently unresolved scintles for PTA

<sup>22</sup> <https://www.atnf.csiro.au/people/pulsar/psrcat/>

pulsars such as PSR J1903+0327, after which, CS pulse sharpening may be applicable.

Other possible benefits of CS deconvolution are likely to occur for other telescopes not yet mentioned (Dolch 2019). The Canadian HI Mapping Experiment (CHIME; Bandura et al. 2014; Ng et al. 2020), although primarily searching for high-redshift baryon acoustic oscillations, observes transiting pulsars daily from 400–800 MHz, typical frequencies in which  $\tau_s$  is large enough that CS deconvolution could be successfully applied on many pulsars. Combined with the high cadence, CHIME’s sensitivity to short-timescale GW bursts would improve with routine CS deconvolution. The telescope is not steerable, but a pulsar-specialized backend receiver extracts pulsar timing data from the large volume of information constantly obtained at the instrument’s many receivers (CHIME/Pulsar Collaboration et al. 2020). CHIME can obtain high-time-sampled, low-frequency data as well as make discoveries of pulsars in its decl. range. The future Square Kilometer Array (SKA), and the current pathfinders—Australian Square Kilometre Array Pathfinder (Johnston et al. 2009) and MeerKAT (Gibbon et al. 2013)—are expected to discover a significant fraction of all the visible pulsars in the Galaxy (Cordes et al. 2004). A Next-Generation Very-Large Array (ngVLA; Chatterjee 2018) has also been proposed for the northern hemisphere. Both telescopes are expected to have a factor of  $\sim 10$  improvement in sensitivity over AO, improving the  $m_{\text{cyc}}$  for many currently fainter pulsars. As already discussed in Section 6, the GBT UWB receiver will enable simultaneous, highly sensitive pulsar observations across 3.3 GHz. In Figure 6, these future telescopes or upgrades would move the example pulsars at least an order of magnitude to the right on the diagram. As a low- $\sigma_{\text{TOA}}$  example, PSR J1713+0747 would be located in the upper right region of the plot instead of the upper left. While CS deconvolution would require many computing nodes across a large bandwidth, such computing resources will be a part of the full SKA or ngVLA. With these telescopes, many pulsars will have S/N similar to that of PSR B1937+21 now. These changes could heavily populate the diagnostic diagram in this paper with pulsars that are high S/N enough and/or highly scattered enough for feasible routine CS deconvolution.

In addition to the successful application of CS to Low Frequency Array data (Archibald et al. 2014), other low-frequency radio telescopes (e.g., 200 MHz and below) could benefit from routine CS data products, for the sake of both fine channelization data products and for deconvolution. While these frequencies do not usually allow for precision timing, the significant scattering tails present could be utilized vis-à-vis CS deconvolution to significantly improve timing precision, or to arrive at stable timing solutions otherwise unavailable before CS deconvolution. The Long-Wavelength Array (LWA), recently expanded to two stations with a significant sensitivity increase, is a relevant example. The LWA is already monitoring many pulsars (Stovall et al. 2015), several of which are currently NANOGrav MSPs.<sup>23</sup> Future low-frequency telescopes such as SKA-low in the southern hemisphere and the proposed LWA-Swarm telescope (Dowell & Taylor 2018; Taylor et al. 2019) may have sensitivities for a significant S/N boost in Figure 6. The Expanded LWA (ELWA; Taylor et al. 2017; Ruan et al. 2020) already uses recently installed 50–80 MHz dipoles on the

VLA together with the two LWA stations. Even if low-frequency timing does not become feasible, these low-frequency telescopes aided with CS can complement timing observations with DM measurements and with monitoring of CS-resolved secondary spectra (the parabolic 2D FTs of dynamic spectra containing scattering structure information) as done with PSR B1937+21 in Walker et al. (2013). Finally, CS has been successfully applied to data from the VLA Low Band Ionospheric and Transient Experiment (VLITE; Polisensky et al. 2016) project in order to provide very high-frequency resolution for RFI excision (M. Kerr, private communication). For the transient search aspect of VLITE, the improved frequency resolution of periodic RFI (Demorest 2011), and therefore the improved quality of RFI excision, could aid the search for fast radio bursts or other radio transients. The Breakthrough Listen project, along these lines, has searched for technosignature periodicities in association with fast radio bursts (Price et al. 2019), by creating cyclic spectra, shifting at multiples of small time intervals and seeing if power appears in harmonics of a particular period. Beyond transient radio astronomy, as in radio spectral line observations, routine CS would also help excise RFI.

CS may also serve as an important pulsar searching tool. The presence of scattering in a CS phase slope could serve as an additional pulsar searching metric, along with trial DMs and periods, for an astrophysically pulsating signal. Searches near the highly scattered Galactic center might particularly benefit from CS as a routine and/or real-time data product.

We thank the referee for the many helpful suggestions that have been incorporated into this paper. We thank J. M. Cordes, T. T. Pennucci, M. Kerr, and the NANOGrav Collaboration for helpful discussions and suggestions, the Noise Budget Working Group in particular. We thank the AO staff for their assistance during the observations reported in this work, in particular P. Perillat, J. S. Deneva, H. Hernandez, and the telescope operators. We are grateful for the use of the Bowser computing cluster at WVU with the assistance of N. E. Garver-Daniels. From Hillsdale College, T.D. gratefully acknowledges computing resources from the Department of Physics, summer leave funds, and professional development funds for conducting this research. The authors were partially supported through the National Science Foundation (NSF) PIRE program award number 0968296. Members of the NANOGrav Collaboration (T.D., D.S., R.L., P.D., M.L., and M.M.) acknowledge support from the NSF Physics Frontiers Center (PFC) award number 1430284. T.D. acknowledges NANOGrav seed-funding for meetings where this research was conducted. AO is a facility of the NSF operated under cooperative agreement by UCF in alliance with Yang Enterprises, Inc. and Universidad Metropolitana. The Green Bank Observatory (GBO) is a facility of the NSF operated under cooperative agreement by Associated Universities, Inc. The GBT UWB receiver project is funded in part by the Gordon and Betty Moore Foundation through Grant GBMF7576 to Associated Universities Inc. to support the work of the GBO and the NANOGrav PFC. T.D., and M.L. acknowledge NSF AAG award number 2009468.

## Appendix

### Derivation of the Cyclic Merit

According to Equation (10), the phase of the cyclic spectrum has two contributions: one term induced by the derivative of

<sup>23</sup> <https://lwa10g.alliance.unm.edu/>

the phase of the IISM transfer function and the other intrinsic to the pulse profile. By the shift theorem in Fourier analysis, a translation in the time domain induces a phase slope in the frequency domain. Since the time-domain impulse response function is characterized by a delay  $\tau_s$ , we can estimate the phase derivative of the transfer function by  $-2\pi\tau_s$ ,

$$\Phi_{S_E}(\nu, \alpha_k) = \frac{k}{P} \frac{d\Phi_H}{d\nu} + \phi_{S_x}(\alpha_k) \quad (A1)$$

$$\approx -2\pi \frac{k\tau_s}{P} + \phi_{S_x}(\alpha_k) \quad (A2)$$

$$\approx k b + \phi_{S_x}(\alpha_k), \quad (A3)$$

where  $b = 2\pi\tau_s/P$  is the average cyclic phase slope, also equal to the expected phase at  $k = 1$ .

The spectrum of the intrinsic pulse profile can be determined accurately by the WDS algorithm, since the Hessian is diagonal with respect to the set of parameters describing the profile (Walker et al. 2013). Hence, in general, the second term in Equation (A3) does not complicate the phase retrieval process. At larger modulation frequencies, the CS amplitude is lost to the noise as in Figure 1 near  $\alpha_k = 40$ , which causes a transition to incoherence in phase as in Figure 2. We can estimate the phase noise in a cyclic spectrum at the  $k$ th harmonic to be

$$\Delta\Phi_k \equiv \sigma_{\text{noise}}/A_k, \quad (A4)$$

where  $A_k$  is the frequency-averaged amplitude of the cyclic spectrum at harmonic  $\alpha_k$ . By a weighted linear fit, the phase slope  $b$  can be estimated by

$$b = \frac{1}{\Delta_w} \sum_k w_k k \Phi_E(k), \quad (A5)$$

where the least-squares weights are  $w_k = 1/[\delta\Phi_E(k)]^2$ , and  $\Delta_w \equiv \sum_k w_k k^2$  does not depend on the measured  $\Phi_E(k)$  values. Then, the uncertainty in the phase slope  $b$  is

$$\delta b = \frac{1}{\Delta_w} \sqrt{\sum_k w_k^2 k^2 [\delta\Phi_E(k)]^2} \quad (A6)$$

$$= \frac{1}{\Delta_w} \sqrt{\sum_k \frac{k^2}{[\delta\Phi_E(k)]^2}}, \quad (A7)$$

where all the sums are from 1 to  $N_{\text{max}}/2 + 1$  and where  $N_{\text{max}}$  is the number of time samples in the folded profile. This can be further simplified by expressing the uncertainty of the cyclic spectrum phase in terms of observational parameters from the radiometer equation:

$$\delta\Phi_E(k) \approx \Delta\Phi_k = \frac{\Delta\Phi_0}{a_k} = \frac{1}{a_k} \frac{T_{\text{sys}}}{gS_0 \sqrt{B\mathcal{T}_{\text{obs}}}}. \quad (A8)$$

Here,  $T_{\text{sys}}$  is the system temperature,  $g$  is the telescope gain (units: Kelvin per Jansky),  $S_0$  is the pulsar continuum flux (single polarization),  $B$  is the bandwidth over which the phase information is integrated (ideally channelized to the diffractive bandwidth,  $\Delta\nu_d$ ), and  $\mathcal{T}_{\text{obs}}$  is the observing time in the integration. Taken together, this allows us to define a quality metric, the *cyclic merit*, for CS phase retrieval, which we express as the S/N of the cyclic spectrum phase

$$m_{\text{cyc}} = \frac{b}{\delta b} = \frac{2\pi\tau_s}{P} \frac{gS_0}{T_{\text{sys}}} \sqrt{B\mathcal{T}_{\text{obs}} \sum_k k^2 a_k}. \quad (A9)$$

The above metric can then be written in terms of the folded pulse profile S/N

$$m_{\text{cyc}} = \frac{2\pi\tau_s W_e}{P^2} (\text{S/N}) \sqrt{\sum_k k^2 a_k}, \quad (A10)$$

where  $W_e$  is the equivalent width of the pulse profile, and

$$(\text{S/N}) \equiv \frac{gS_0}{T_{\text{sys}}} \frac{P}{W_e} \sqrt{B\mathcal{T}_{\text{obs}} \left( \frac{W_e}{P} \right)} = \frac{gS_0}{T_{\text{sys}}} \sqrt{B\mathcal{T}_{\text{obs}} \left( \frac{P}{W_e} \right)}. \quad (A11)$$

Here, we have assumed that the profile has been smoothed to a time resolution of  $W_e$  for an optimal trade-off between pulse sharpness and S/N. In principle, the smoothing should be done to a time resolution of the sharpness width,  $W_s \equiv \int dt |U'(t)|^2$ , where  $U(t)$  is a unit amplitude average pulse profile, and the prime indicates a time derivative (Cordes & Shannon 2010). In practice, the quantitative difference between  $W_e$  and  $W_s$  is expected to be small for realistic pulse profiles.

Equation (A9) is the cyclic merit for determining the *average* phase slope, which is a useful estimator of  $\tau_s$ . If a full reconstruction of the transfer function is required, we must estimate the cyclic spectrum phase over a diffractive bandwidth. This requires setting  $B = \Delta\nu_d \approx 1/(2\pi\tau_s)$  in Equation (A9) from which we obtain a cyclic merit for a full reconstruction of the transfer function

$$m_{\text{cyc,full}} = \frac{gS_0}{P\mathcal{T}_{\text{sys}}} \sqrt{2\pi\tau_s \mathcal{T}_{\text{obs}} \sum_k k^2 a_k}, \quad (A12)$$

where we see a square-root dependence on the scattering delay.

## ORCID iDs

Timothy Dolch <https://orcid.org/0000-0001-8885-6388>  
 Dan R. Stinebring <https://orcid.org/0000-0002-1797-3277>  
 Hengrui Zhu <https://orcid.org/0000-0001-9027-4184>  
 Ryan S. Lynch <https://orcid.org/0000-0001-5229-7430>  
 Tyler Cohen <https://orcid.org/0000-0001-7587-5483>  
 Paul B. Demorest <https://orcid.org/0000-0002-6664-965X>  
 Michael T. Lam <https://orcid.org/0000-0003-0721-651X>  
 Lina Levin <https://orcid.org/0000-0002-2034-2986>  
 Maura A. McLaughlin <https://orcid.org/0000-0001-7697-7422>  
 Nipuni T. Palliyaguru <https://orcid.org/0000-0002-4828-0262>

## References

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016, *PhRvL*, **116**, 061102  
 Alam, M. F., Arzoumanian, Z., Baker, P. T., et al. 2021a, *ApJS*, **252**, 4  
 Alam, M. F., Arzoumanian, Z., Baker, P. T., et al. 2021b, *ApJS*, **252**, 5  
 Antoni, J. 2007, *MSSP*, **21**, 597  
 Archibald, A. M., Kondratiev, V. I., Hessels, J. W. T., & Stinebring, D. R. 2014, *ApJL*, **790**, L22  
 Arzoumanian, Z., Baker, P. T., Brazier, A., et al. 2018, *ApJ*, **859**, 47  
 Bandura, K., Addison, G. E., Amiri, M., et al. 2014, *Proc. SPIE*, **9145**, 22  
 Bansal, K., Taylor, G. B., Stovall, K., & Dowell, J. 2019, *ApJ*, **875**, 146  
 Bates, S. D., Lorimer, D. R., Rane, A., & Swiggum, J. 2014, *MNRAS*, **439**, 2893  
 Bhat, N. D. R., Cordes, J. M., Camilo, F., Nice, D. J., & Lorimer, D. R. 2004, *ApJ*, **605**, 759  
 Chatterjee, S. 2018, in ASP Conf. Ser. 517, Science with a Next Generation Very Large Array, ed. E. Murphy (San Francisco, CA: ASP), **751**  
 CHIME/Pulsar Collaboration, Amiri, M., Bandura, K. M., et al. 2020, arXiv:2008.05681

Coles, W. A., Rickett, B. J., Gao, J. J., Hobbs, G., & Verbiest, J. P. W. 2010, *ApJ*, **717**, 1206

Cordes, J. M., Kramer, M., Lazio, T. J. W., et al. 2004, *NewAR*, **48**, 1413

Cordes, J. M., Rickett, B. J., Stinebring, D. R., & Coles, W. A. 2006, *ApJ*, **637**, 346

Cordes, J. M., & Shannon, R. M. 2010, arXiv:1010.3785

Cordes, J. M., Shannon, R. M., & Stinebring, D. R. 2016, *ApJ*, **817**, 16

Demorest, P. B. 2011, *MNRAS*, **416**, 2821

Detweiler, S. 1979, *ApJ*, **234**, 1100

Dolch, T. 2019, in 2019 United States National Committee of URSI National Radio Science Meeting (USNC-URSI NRSRM) (New York: IEEE), 1

Dolch, T., Lam, M. T., Cordes, J., et al. 2014, *ApJ*, **794**, 21

Dolch, T., NANOGrav Collaboration, Chatterjee, S., et al. 2018, *JPhCS*, **957**, 012007

Dowell, J., & Taylor, G. B. 2018, *JAI*, **7**, 1850006

Downs, G. S., & Reichley, P. E. 1983, *ApJS*, **53**, 169

DuPlain, R., Benson, J., & Sessoms, E. 2008, *Proc. SPIE*, **7019**, 70191A

Finn, L. S., & Lommen, A. N. 2010, *ApJ*, **718**, 1400

Foster, R. S., & Backer, D. C. 1990, *ApJ*, **361**, 300

Gardner, W. A. 1991, *ISPM*, **8**, 14

Geyer, M., & Karastergiou, A. 2016, *MNRAS*, **462**, 2587

Geyer, M., Karastergiou, A., Kondratiev, V. I., et al. 2017, *MNRAS*, **470**, 2659

Gibson, T. B., Rotich, E. K., Kourouma, H. Y. S., et al. 2013, *Proc. SPIE*, **9008**, 90080P

Hankins, T. H. 1971, *ApJ*, **169**, 487

Hankins, T. H., & Rickett, B. J. 1975, in Methods in Computational Physics. Vol. 14—Radio Astronomy, ed. B. Alder, S. Fernbach, & M. Rotenberg (New York: Academic), 55

Hellings, R. W., & Downs, G. S. 1983, *ApJL*, **265**, L39

Hobbs, G. 2013, *CQGra*, **30**, 224007

Hobbs, G., Archibald, A., Arzoumanian, Z., et al. 2010, *CQGra*, **27**, 084013

Hulse, R. A., & Taylor, J. H. 1975, *ApJL*, **195**, L51

Johnston, S., Feain, I. J., & Gupta, N. 2009, in ASP Conf. Ser., **407**, The Low-Frequency Radio Universe, ed. D. J. Saikia, D. A. Green, Y. Gupta, & T. Venturi (San Francisco, CA: ASP)

Jones, G., Cordes, J., Demorest, P. B., et al. 2013, in 2013 US National Committee of URSI National Radio Science Meeting (USNC-URSI NRSRM) (New York: IEEE), 1

Kaspi, V. M., Taylor, J. H., & Ryba, M. F. 1994, *ApJ*, **428**, 713

Kramer, M., & Champion, D. J. 2013, *CQGra*, **30**, 224009

Kramer, M., Xilouris, K. M., Lorimer, D. R., et al. 1998, *ApJ*, **501**, 270

Lam, M. T., Cordes, J. M., Chatterjee, S., et al. 2016, *ApJ*, **819**, 155

Lam, M. T., Cordes, J. M., Chatterjee, S., et al. 2017, *ApJ*, **834**, 35

Lam, M. T., Ellis, J. A., Grillo, G., et al. 2018, *ApJ*, **861**, 132

Lam, M. T., McLaughlin, M. A., Arzoumanian, Z., et al. 2019, *ApJ*, **872**, 193

Lam, M. T., McLaughlin, M. A., Cordes, J. M., Chatterjee, S., & Lazio, T. J. W. 2018, *ApJ*, **861**, 12

Levin, L., McLaughlin, M. A., Jones, G., et al. 2016, *ApJ*, **818**, 166

Löhmer, O., Mitra, D., Gupta, Y., Kramer, M., & Ahuja, A. 2004, in How Does the Galaxy Work?, ed. E. J. Alfaro, E. Pérez, & J. Franco (Dordrecht: Springer), 327

Lorimer, D. R., & Kramer, M. 2004, *Handbook of Pulsar Astronomy* (Cambridge: Cambridge Univ. Press)

Manchester, R. N., Hobbs, G., Bailes, M., et al. 2013, *PASA*, **30**, 17

Manchester, R. N. & IPTA 2013, *CQGra*, **30**, 224010

McLaughlin, M. A. 2013, *CQGra*, **30**, 224008

Ng, C., Pandhi, A., Naidu, A., et al. 2020, *MNRAS*, **496**, 2836

Palliyaguru, N., Stinebring, D., McLaughlin, M., Demorest, P., & Jones, G. 2015, *ApJ*, **815**, 89

Polisensky, E., Lane, W. M., Hyman, S. D., et al. 2016, *ApJ*, **832**, 60

Price, D. C., Foster, G., Geyer, M., et al. 2019, *MNRAS*, **486**, 3636

Ramachandran, R., Demorest, P., Backer, D. C., Cognard, I., & Lommen, A. 2006, *ApJ*, **645**, 303

Rickett, B. J. 1990, *ARA&A*, **28**, 561

Ruan, D., Taylor, G. B., Dowell, J., et al. 2020, *MNRAS*, **495**, 2125

Sazhin, M. V. 1978, *SvA*, **22**, 36

Siemens, X., Ellis, J., Jenet, F., & Romano, J. D. 2013, *CQGra*, **30**, 224015

Stinebring, D. 2013, *CQGra*, **30**, 224006

Stinebring, D. R., Chatterjee, S., Clark, S. E., et al. 2019b, *BAAS*, **51**, 492

Stinebring, D. R., McLaughlin, M. A., Cordes, J. M., et al. 2001, *ApJL*, **549**, L97

Stinebring, D. R., Rickett, B. J., & Koch Ocker, S. 2019a, *ApJ*, **870**, 82

Stovall, K., Ray, P. S., Blythe, J., et al. 2015, *ApJ*, **808**, 156

Taylor, G., Dowell, J., Malins, J., et al. 2017, arXiv:1708.00090

Taylor, G., Dowell, J., Pihlström, Y., et al. 2019, *BAAS*, **51**, 2

Turner, J. E., McLaughlin, M. A., Cordes, J. M., et al. 2020, arXiv:2012.09884

van Straten, W., & Bailes, M. 2011, *PASA*, **28**, 1

Walker, M. A., Demorest, P. B., & van Straten, W. 2013, *ApJ*, **779**, 99

Walker, M. A., Koopmans, L. V. E., Stinebring, D. R., & van Straten, W. 2008, *MNRAS*, **388**, 1214

Walker, M. A., & Stinebring, D. R. 2005, *MNRAS*, **362**, 1279