



Memorandum 009

Correction Methods for Interstellar Pulse Broadening

J. M. Cordes

Authored Apr 23, 2012 - Nov 1, 2012

Typos Corrected Aug 19, 2023

Additional Comments Jan 17, 2025

Correction Methods for Interstellar Pulse Broadening

J. M. Cordes

Astronomy Department, Cornell University

`cordes@astro.cornell.edu`

Apr 23, 2012 - Nov 1, 2012; typos corrected Aug 19, 2023; comments Jan 17, 2025

ABSTRACT

These are notes on pulsar signal models and methods for deconvolving scattering broadening or otherwise mitigating scattering delays for pulsar timing. They were written in 2012 so there are consequently no recent references. The document is in slide format for the most part without a comprehensive narrative between topics. Two retrospective comments: The applicability of the mean shift regime is limited, as discussed in *Fundamental Noise Processes in Pulsar Timing* (JMC et al. 2025 in preparation). However, phase retrieval using the Hilbert transform appears to be a promising avenue of research for scattering corrections.

1. Outline

Outline:

1. Basics: DISS and AMN definitions
2. Frequency domain quantities
3. Intrinsic AMN fluctuations (examples in figures)
4. Statistics of the frequency ACF
5. Signal flow diagrams and setting up inverse problems
6. Timing perturbations from the ISM
7. Correction regimes for scattering
8. Mean shift regime
9. Requirements and methods for incoherent deconvolution
10. Requirements for coherent deconvolution (phase space)
11. Methods for coherent deconvolution (including cyclic spectroscopy, phase retrieval using Hilbert transform)

Notation: As is common, $h(t)$ is used to denote the impulse response of a linear filter. There should be no confusion with the dimensionless strain produced by gravitational waves (GWs) since the context of the discussion here is different and more general than for GW applications.

Scattering and Scintillation in the ISM

Terms:

Diffraction (aka scattering) causes multipath propagation of a pulsar signal that leads to pulse broadening and scintillation. Diffraction is caused by very small scales $\ell \lesssim 10^6$ km.

Refraction from larger scale structures causes TOA variations and modification of the diffraction pattern.

Diffraction angle:

$$\theta_d = \frac{\lambda}{2\pi\ell_d},$$

where ℓ_d is the diffraction scale (the spatial size of an intensity maximum in the diffraction pattern).

Pulse broadening: the field and intensity broadening functions are

$$h(t) \quad \text{and} \quad p(t) = |h(t)|^2.$$

Pulse broadening time: is defined as the 1/e time scale of the intensity PBF $p(t)$, usually obtained from long-term averages

$$\tau_d \approx \frac{D\theta_d^2}{2c}.$$

Scintillation bandwidth: is defined as the HWHM of the autocorrelation function of the intensity spectrum $I(\nu)$

$$\Delta\nu_d = \frac{C_1}{2\pi\tau_d},$$

where $C_1 \approx 1$ is a constant that depends on the distribution and wavenumber spectrum of scattering material.

Scintillation time scale, the characteristic time for intensity variations. Using Equations C4-C6 from CR98 for the case where the effective transverse velocity is dominated by the pulsar velocity, we have

$$\Delta t_d = \frac{\ell_d}{V_{\text{eff}}} \approx \frac{\lambda}{2\pi V_{p\perp}} \left(\frac{D}{2c\tau_d} \right)^{1/2} = \frac{c}{2\pi\nu V_{\text{eff}}} \left(\frac{D}{2c\tau_d} \right)^{1/2} \approx \frac{108 \text{ sec}}{\nu V_{100}} \left(\frac{D}{2c\tau_d} \right)^{1/2},$$

for ν in GHz, V_{eff} in units of 100 km s^{-1} , D in kpc, and τ_d in μs .

Amplitude Modulated Noise (AMN)

For an amplitude $a(t)$ that modulates complex white, Gaussian noise $m(t)$ with statistics

$$\langle m(t) \rangle = 0 \quad \text{and} \quad \langle m(t)m^*(t') \rangle = M_0 \Delta(t - t'),$$

where $\Delta(t)$ is a delta-function like quantity with $\Delta(0) = 1$. The pulsar wavefield is

$$\varepsilon_i(t) = a(t)m(t).$$

The measured wavefield is the sum of the pulsar wavefield modified by the impulse response $h(t)$ for scattering in the ISM and additive, complex noise $n(t)$

$$\varepsilon(t) = h(t) * \varepsilon_i(t) + n(t) = h(t) * [a(t)m(t)] + n(t).$$

The additive noise is radiometer noise and has statistics similar to those for $m(t)$ with M_0 replaced by N_0 , i.e. $\langle n(t)n^*(t') \rangle = N_0 \Delta(t - t')$.

The Fourier transform of ε is

$$\tilde{\varepsilon}(\nu) = \tilde{h}(\nu) [\tilde{a}(\nu) * \tilde{m}(\nu)] + \tilde{n}(\nu).$$

where the forward FT is defined with a kernel $\exp(-2\pi i \nu t)$.

AMPSN: A more realistic model for the pulsar signal replaces $m(t)$ with complex, polarized shot noise. Observations of the Crab pulsar support this model.

SAMPSN: A further extension includes scintillations as done above.

Frequency Domain Quantities

The spectrum of the pulsar signal is

$$S_i(\nu) = |\tilde{\varepsilon}_i(\nu)|^2$$

and the field autocorrelation function (ACF) is

$$\Gamma_{\tilde{\varepsilon}_i}(\nu, \alpha) = \tilde{\varepsilon}_i(\nu + \alpha/2) \tilde{\varepsilon}_i^*(\nu - \alpha/2).$$

For AMN the ensemble average correlation function is

$$\begin{aligned} \langle \Gamma_{\tilde{\varepsilon}_i}(\nu, \alpha) \rangle &= \iint dt_1 dt_2 a(t_1) a(t_2) \langle m(t_1) m^*(t_2) \rangle e^{-2\pi i[(\nu + \alpha/2)t_1 - (\nu - \alpha/2)t_2]} \\ &= M_0 \int dt_1 a^2(t_1) e^{-2\pi i \alpha t_1} \\ &\equiv M_0 \tilde{A}(\alpha), \end{aligned}$$

where $A(t) \equiv a^2(t)$ is the intensity modulation that can be aperiodic or periodic. The ensemble average spectrum is then

$$\langle S_i(\nu) \rangle = \langle \Gamma_{\tilde{\varepsilon}_i}(\nu, 0) \rangle = M_0 \tilde{A}(0). \quad (1)$$

We include the ISM by using the pulse-broadening function $h(t)$ defined below. The PBF $h(t)$ changes on the diffractive scintillation time scale. We can therefore calculate the average ACF over an ensemble of the AMN while considering, in effect, a single realization of $h(t)$. We could just as well consider $h(t)$ to be deterministic. Excluding additive noise for now, we get

$$\langle \Gamma_{\tilde{\varepsilon}}(\nu, \alpha) \rangle = \langle \Gamma_{\tilde{\varepsilon}_i}(\nu, \alpha) \rangle \Gamma_{\tilde{h}}(\nu, \alpha) = M_0 \tilde{A}(\alpha) \Gamma_{\tilde{h}}(\nu, \alpha).$$

For a **periodic amplitude modulation** with period P that extends over a time span $T = N_p P$,

$$a(t) = \sum_{j=0}^{N_p-1} a_1(t - jP),$$

we have

$$\tilde{A}(\nu) = P^{-1} \tilde{A}_1(\nu) \delta(\nu - k/P).$$

The frequency correlation function is discretely sampled at harmonics $\alpha_k = k/P$,

$$\langle \Gamma_{\tilde{\varepsilon}}(\nu, \alpha) \rangle = P^{-1} M_0 \tilde{A}_1(\alpha_k) \Gamma_{\tilde{h}}(\nu, \alpha_k) \delta(\alpha - \alpha_k).$$

Re-do using DFT and definition of discrete sinc function put into appendix

Intrinsic AMN Fluctuations

The AMN signal yields two kinds of fluctuations. First are those arising from the noise process, $m(t)$. The second kind is from stochastic variations of the amplitude modulation, $a(t)$, including the phase jitter described above but also from amplitude variations of individual pulses.

The AMN signal is broadband but has a characteristic frequency scale $\sim W_A^{-1}$ that results from the convolution of $\tilde{a}(\nu)$ with $\tilde{m}(\nu)$, where W_A is a characteristic width of $A(t)$. For a spectrum obtained using an FFT of length $T \gg W_A$, the following inequalities hold:

$$\frac{1}{T} \ll \frac{1}{W_A} \ll B.$$

E.g. for a total bandwidth $B = 100$ MHz and a pulse width $W_A = 100 \mu s$, a spectrum obtained from a $T = 1$ s long FFT we have

$$1 \text{ Hz} \ll 10 \text{ kHz} \ll 100 \text{ MHz}.$$

The spectrum $S_i(\nu)$ of the AMN signal has 100% fluctuations because it is a χ_2^2 random variable. The frequency domain ACF $\Gamma_{\tilde{\varepsilon}_i}(\delta\nu)$ also has 100% fluctuations when calculated from a single Fourier transform.

Both the spectrum and the ACF are correlated over a frequency scale $\sim 1/W_A$. The frequency correlation scale can be seen in plots given below.

When averaged over N_b data blocks, fluctuations in the spectrum and the ACF decrease by $N_b^{-1/2}$.

Properties of $h(t)$

$h(t)$ will typically have the form of an envelope function that multiplies a noise-like process. The noise process is actually a consequence of the particular path lengths that the radiation field has from diffraction and is therefore **persistent** over time scales much less than the scintillation time Δt_{ISS} .

It is therefore useful to think of two times: one describes variations on time scales of order the inverse bandwidth (ns) and the other “epochal” time is on the time scale of Δt_{ISS} . The time “t” in $h(t)$ refers to the fast time scale and the epochal time dependence is implicit.

For other cases it may turn out that there are only a few ray paths on which radiation arrives and the impulse response can be written as a sequence of delta functions:

$$h(t) = \sum_j a_j \delta(t - t_j).$$

Figure 1 shows an example PBF.

Other useful quantities are the Fourier transform of $h(t)$, $\tilde{h}(\nu)$, and the frequency ACF,

$$\Gamma_{\tilde{h}}(\nu, \delta\nu) = \tilde{h}(\nu + \delta\nu/2) \tilde{h}^*(\nu - \delta\nu/2).$$

The spectrum is simply

$$H(\nu) = |\tilde{h}(\nu)|^2 = \Gamma_{\tilde{h}}(\nu, 0). \quad (2)$$

The FT of $H(\nu)$ is, by the Wiener-Khinchine theorem, the temporal ACF of $h(t)$,

$$\Gamma_h(t, \tau) = h(t) h^*(t + \tau).$$

Example PBF




Fig. 1.— Example pulse broadening function consisting of a one-sided exponential function that multiplies complex Gaussian noise, $h(t) = e^{-t/\tau}U(t)n(t)$ where U is the unit step function and $n(t)$ is a noise process that is the same for macroscopic times less than the scintillation time Δt_{ISS} . The top panel shows the real and imaginary parts of $h(t)$ along with the envelope function shown as dashed lines. The middle panel shows the phase, and the bottom panel shows the intensity PBF, $p(t) = |h(t)|^2$, along with the envelope function shown as a dashed line.

AMN Components, Example 1: Single Pulse

Figure 2 shows the components of the scattered AMN signal in both the time and frequency domains. Additional examples are shown at the end of the document.

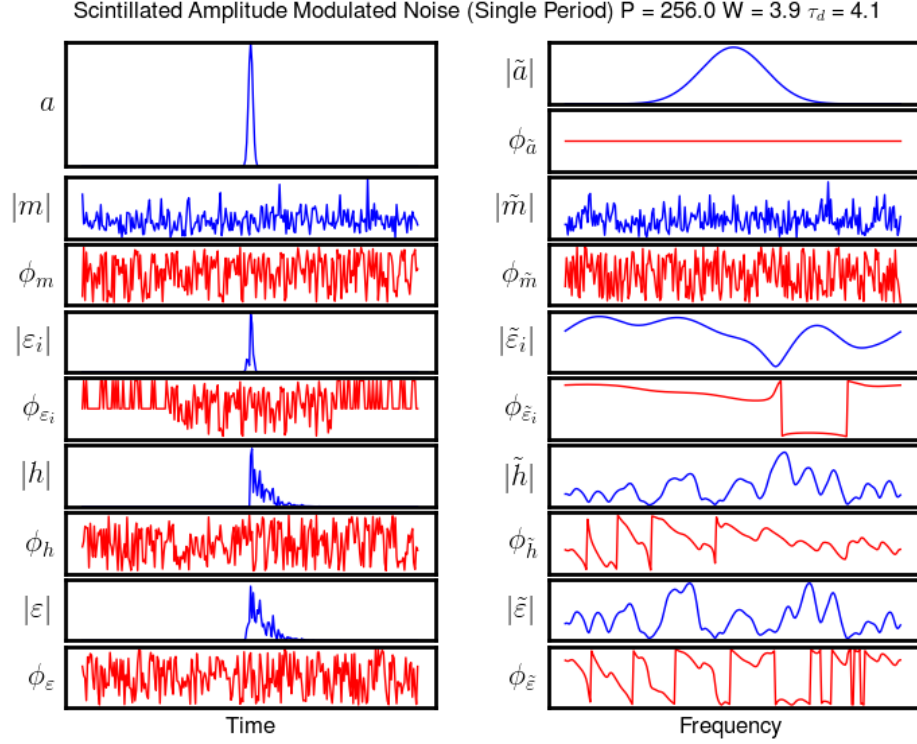



Fig. 2.— Components of scintillated amplitude modulated noise in the time and frequency domain. Amplitudes are plotted in blue, phases in red.

Signal Flow Diagram for Pulse Broadening

The field PBF is related to the ACF of its Fourier transform and to the intensity PBF according to the following diagram. Two-sided arrows denote Fourier transforms while downward arrows show irreversible processes (squared magnitude and correlation) that remove information:



Terminology:

$h(t)$ is the DISS impulse response or field PBF

$p(t) = |h(t)|^2$ is the intensity PBF.

The Fourier transform of a single data block would yield a spectrum $I(\nu) = |\tilde{h}(\nu)|^2$ that has 100% fluctuations because it has χ_2^2 statistics.

It can be shown that the ACF $\Gamma_{\tilde{h}}$ also has 100% variations.

Ensemble averages of p and $\Gamma_{\tilde{h}}$ are estimated by averaging these quantities over multiple data blocks.

Inverse problem: The goal is to determine h when we have, say, the ACF of its FT, $\Gamma_{\tilde{h}}(\delta\nu)$. This is the same as determining \tilde{h} from the magnitude of *its* FT, which is $|h| = \sqrt{p(t)}$.

In practice statistical measures include the pulsar signal along with the PBF. For sake of discussion, if we ignore the pulsar contribution, we could use the following approach: (I think this isn't correct because \tilde{h} is not “causal” (i.e. one-sided) in the frequency domain the way h IS causal in time.)

1. Estimate $\Gamma_{\tilde{h}}(\delta\nu)$ from data.
2. Calculate $p = |h|^2$ from the IFT of $\Gamma_{\tilde{h}}$.
3. Determine the minimum-phase of the FT using the Hilbert transform (HT), $\phi_{\text{HT}} = \text{HT} \{ \log \sqrt{p} \}$
4. Calculate the real and imaginary parts of $h_{\text{HT}}(t)$ as $(h_r, h_i)_{\text{HT}} = \sqrt{p}(\cos \phi_{\text{HT}}, \sin \phi_{\text{HT}})$.
5. Update the HT solution using any additional constraints and by least-squares optimization.

Example Fourier Relationships for a Realization of $h(t)$




Fig. 3.— Fourier relationships for an example PBF consisting of an exponential envelope that multiplies complex Gaussian noise. The time constant for the ensemble average $p(t)$ is $\tau = 7.1$ (arbitrary units). Top left: the real and imaginary parts of $h(t)$. Top right: DFT of $h(t)$, $\tilde{h}(\nu)$. Bottom right: the ACF of $\Gamma_{\tilde{h}}(\delta\nu)$. Bottom left: the directly calculated intensity PBF, $p(t)$, and that calculated from the inverse DFT of the ACF, $\hat{p}(t)$.

Reciprocal Signal Flow Diagram for Pulse Broadening

An alternative diagram describes the reciprocal case where the correlation and squared-magnitude operations are applied in the opposite domains than previously:

$$\begin{array}{ccc}
 \boxed{h(t)} & \longleftrightarrow & \boxed{\tilde{h}(\nu)} \\
 \text{ACF} \quad \downarrow & & \downarrow \quad |\cdot|^2 \\
 \boxed{\Gamma_h(\tau) = h(t)h^*(t + \tau)} & \longleftrightarrow & \boxed{H(\nu) = |\tilde{h}(\nu)|^2}
 \end{array}$$

Inverse problem: The goal is to determine h , as before, but now we have the the magnitude of its FT, \sqrt{H} . The numerical problem is the same, determining a function from its ACF or from the magnitude of its FT.

Relationship to the Secondary Spectrum

The alternative signal-flow diagram can be extended to include the secondary spectrum. In this case we consider a perfectly coherent signal from the pulsar (i.e. a delta function in time) so that the secondary spectrum $H_2(\tau)$ is a function of just the scattering:

$$\begin{array}{ccc}
 \boxed{h(t)} & \longleftrightarrow & \boxed{\tilde{h}(\nu)} \\
 \text{ACF} \quad \downarrow & & \downarrow \quad |\cdot|^2 \\
 \boxed{\Gamma_h(\tau) = h(t)h^*(t+\tau)} & \longleftrightarrow & \boxed{H(\nu) = |\tilde{h}(\nu)|^2} \\
 & & \text{Spectrum of } h(t) \\
 |\cdot|^2 \quad \downarrow & & \downarrow \quad \text{ACF} \\
 \boxed{H_2(\tau) = |\Gamma_h(\tau)|^2} & \longleftrightarrow & \boxed{\Gamma_H(\delta\nu)} \\
 \text{Secondary Spectrum} & &
 \end{array}$$

Inverse problem: A new problem presents itself: from the secondary spectrum we can get the magnitude of the ACF of h , from which we can try to determine the complex ACF, $\Gamma_h(\tau)$ and other quantities, in turn. The advantage is that Γ_h can be averaged over many scintillation time scales.

Signal Flow Diagram for Pulsar Signal + ISM

$$\begin{array}{ccc}
 \boxed{\varepsilon(t) = h(t) * \varepsilon_i(t)} & \longleftrightarrow & \boxed{\tilde{h}(\nu)\tilde{\varepsilon}_i(\nu)} \\
 |\cdot|^2 \downarrow & & \downarrow \text{ACF} \\
 \boxed{\begin{array}{l} I(t) = |\langle \varepsilon(t) |^2 \rangle \\ \langle I(t) \rangle = A(t) * \langle p(t) \rangle \end{array}} & \longleftrightarrow & \boxed{\begin{array}{l} \Gamma_{\tilde{\varepsilon}}(\delta\nu) = \Gamma_{\tilde{h}}(\delta\nu)\Gamma_{\tilde{\varepsilon}_i}(\delta\nu) \\ \langle \Gamma_{\tilde{\varepsilon}}(\delta\nu) \rangle = \tilde{A}(\delta\nu)\tilde{p}(\delta\nu) \end{array}}
 \end{array}$$

SAMN Intensity Statistics Example 1

Figure 4 shows intensity-like quantities and their second moments in the time and frequency domains. The frequency domain quantities have been averaged over 10^3 realizations of the intrinsic pulsar signal while holding fixed the interstellar field PBF. The average spectrum $\langle S_i(\nu) \rangle$ (shown as a dotted line) is converged to within a fractional variation $\sim 10^{-3/2}$ of a constant value. Similarly, the spectrum $\langle S(\nu) \rangle$ of the scattered wavefield has converged to $H(\nu)$ to the same precision.

Other examples can be found in Figures 39 to 41 for other averages of 10, 100 and 10^4 realizations.

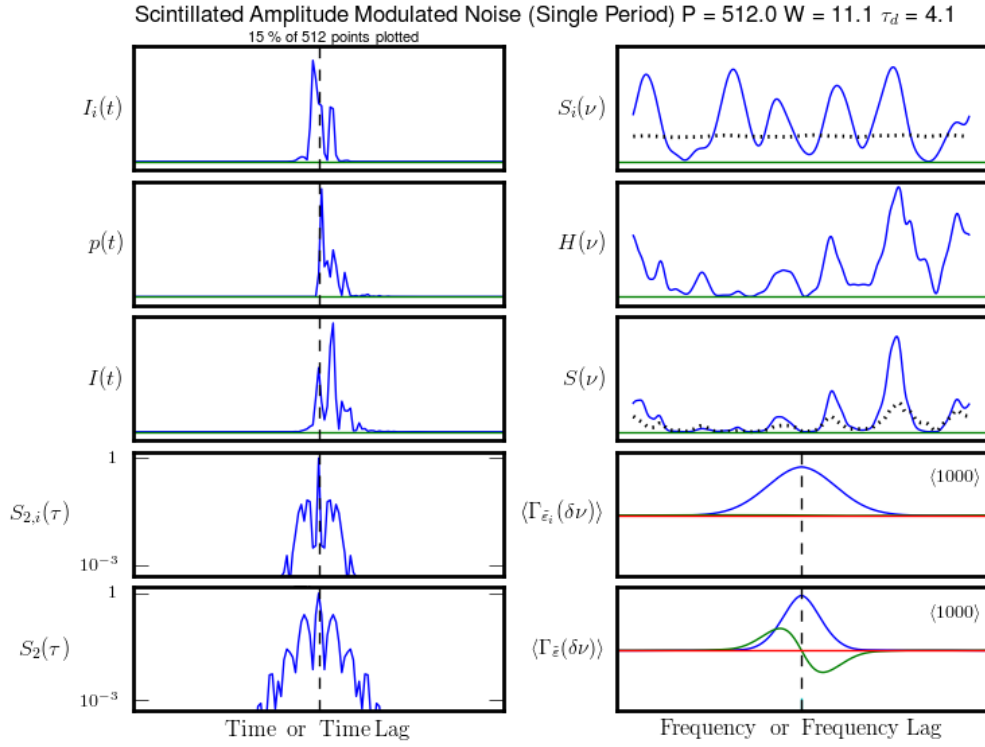


Fig. 4.— Time domain intensity-like quantities are shown on the left. Frequency-domain quantities are on the right. The dotted lines in the plots for $S_i(\nu)$ and $S(\nu)$ are averages over 10^3 realizations. The ISM transfer function $H(\nu)$ is kept fixed for all realizations, so the average of $S(\nu)$ tends toward the shape of $H(\nu)$. The ACFs Γ_{ε_i} and Γ_{ε} are averaged over 10^3 realizations, as indicated.

Fluctuations of Spectra and ACFs

The SAMN signal has a spectrum

$$S(\nu) = |\tilde{\varepsilon}(\nu)|^2$$

that has an ensemble mean

$$\langle S(\nu) \rangle = \langle |\tilde{\varepsilon}(\nu)|^2 \rangle \equiv \langle \Gamma_{\tilde{\varepsilon}}(0) \rangle \equiv M_0 H(\nu) \tilde{A}(0).$$

The variance of $S(\nu)$ follows by considering $\tilde{\varepsilon}$ to be a complex Gaussian process, which it is by virtue of applying the CLT to the FT:

$$\text{Var} \{S(\nu)\} = \langle |\tilde{\varepsilon}(\nu)|^4 \rangle - \langle |\tilde{\varepsilon}(\nu)|^2 \rangle^2 = \langle S(\nu) \rangle^2.$$

$\Gamma_{\tilde{\varepsilon}}(\delta\nu)$ is another second moment whose variance is given by a fourth moment of the field. As before the ensemble mean of the ACF is (excluding the contribution from additive noise)

$$\langle \Gamma_{\tilde{\varepsilon}}(\delta\nu) \rangle = M_0 \Gamma_{\tilde{h}}(\delta\nu) \tilde{A}(\delta\nu)$$

and its second moment is

$$\langle |\Gamma_{\tilde{\varepsilon}}(\delta\nu)|^2 \rangle = |\langle \Gamma_{\tilde{\varepsilon}}(\delta\nu) \rangle|^2 + \langle S(\nu + \delta\nu/2) \rangle \langle S(\nu - \delta\nu/2) \rangle.$$

The variance is therefore

$$\text{Var} \{ \Gamma_{\tilde{\varepsilon}}(\delta\nu) \} = \langle S(\nu + \delta\nu/2) \rangle \langle S(\nu - \delta\nu/2) \rangle.$$

and the fractional error is

$$\frac{[\text{Var} \{ \Gamma_{\tilde{\varepsilon}}(\delta\nu) \}]^{1/2}}{\Gamma_{\tilde{\varepsilon}}(\delta\nu)} = \frac{[\langle S(\nu + \delta\nu/2) \rangle \langle S(\nu - \delta\nu/2) \rangle]^{1/2}}{M_0 H(\nu) \tilde{A}(0)} = \frac{[H(\nu + \delta\nu/2) H(\nu - \delta\nu/2)]^{1/2}}{H(\nu)},$$

which is 100% at the origin.

Variance of Frequency-domain Quantities (no-ISM case)

Consider the frequency-domain ACF of the intrinsic field $\varepsilon(\nu)$ combined with additive noise, $\tilde{n}(\nu)$, excluding for now the effects of the ISM:

$$\Gamma_{\tilde{\varepsilon}}(\nu, \alpha) = \tilde{\varepsilon}(\nu + \alpha/2)\tilde{\varepsilon}^*(\nu - \alpha/2).$$

The ensemble mean of the ACF is

$$\langle \Gamma_{\tilde{\varepsilon}}(\nu, \alpha) \rangle = M_0 \tilde{A}(\alpha) + N_0 N_t \Delta(\alpha),$$

where $M_0 = \langle |m(t)|^2 \rangle$, $N_0 = \langle |n(t)|^2 \rangle$, $\Delta(\alpha)$ is a Kronecker delta like function with $\Delta(0) = 1$. To calculate the ACF's variance we need the second moment of $\Gamma_{\tilde{\varepsilon}}(\nu, \alpha)$,

$$\langle |\Gamma_{\tilde{\varepsilon}}(\nu, \alpha)|^2 \rangle = \langle \varepsilon(\nu + \alpha/2)\varepsilon^*(\nu - \alpha/2)\varepsilon^*(\nu + \alpha/2)\varepsilon(\nu - \alpha/2) \rangle.$$

To simplify, we assume complex Gaussian statistics for both $m(t)$ and $n(t)$ so that the fourth moment of each of these can be expanded into the product of second moments. We consider an FT of N_t time samples defined as

$$\tilde{A}(\nu) = \sum_t A(t) e^{-2\pi i \nu t}.$$

We relate the zero-frequency Fourier amplitude to the mean of A as

$$\tilde{A}(0) = N_t \bar{A}.$$

The variance is then

$$\begin{aligned} \text{Var} [\Gamma_{\tilde{\varepsilon}}(\nu, \alpha)] &\equiv \langle |\Gamma_{\tilde{\varepsilon}}(\nu, \alpha)|^2 \rangle - |\langle \Gamma_{\tilde{\varepsilon}}(\nu, \alpha) \rangle|^2 \\ &= \langle S_i(\nu + \alpha/2) \rangle \langle S_i(\nu - \alpha/2) \rangle + \langle S_n(\nu + \alpha/2) \rangle \langle S_n(\nu - \alpha/2) \rangle \\ &\quad + \langle S_i(\nu + \alpha/2) \rangle \langle S_n(\nu - \alpha/2) \rangle \\ &\quad + \langle S_i(\nu - \alpha/2) \rangle \langle S_n(\nu + \alpha/2) \rangle \\ &= N_t^2 (M_0 \bar{A} + N_0)^2. \end{aligned}$$

We define the signal-to-noise ratio of the correlation function as

$$\left(\frac{S}{N} \right)_{\Gamma_{\tilde{\varepsilon}}} = \frac{\Gamma_{\tilde{\varepsilon}}(\nu, \alpha)}{\text{Var}^{1/2} [\Gamma_{\tilde{\varepsilon}}(\nu, \alpha)]} = \frac{\tilde{A}(\alpha)/\tilde{A}(0) + (N_0/M_0 \bar{A})\Delta(\alpha)}{1 + N_0/M_0 \bar{A}}.$$

The S/N does not depend on frequency in this model (which assumes a flat spectrum for the pulsar signal). At zero lag ($\alpha = 0$) the S/N is unity as expected for analysis of a single data block of a complex Gaussian process (which each of $\tilde{\varepsilon}_i$ and \tilde{n} and their sum are). When summed over N_b data blocks, the S/N is larger by a factor $\sqrt{N_b}$. If the ratio of the mean noise level N_0 to average pulse amplitude \bar{A} is large, the S/N is accordingly smaller. However, the $\sqrt{N_b}$ increase in S/N still applies.

Variance of Frequency-domain Quantities (with ISM)

It is easy to include the ISM into the ACF variance calculation. For deconvolution purposes we need the ACF calculated over a time less than the DISS time scale. By definition the field PBF is (nearly) constant in epochal time for times much less than the DISS time, so in effect $h(t)$ can be considered deterministic. Using the ACF defined previously,

$$\Gamma_{\tilde{h}}(\nu, \alpha) = h(\nu + \alpha/2)h^*(\nu - \alpha/2),$$

the ACF of the scintillated AMN signal has ensemble-average mean

$$\begin{aligned}\langle \Gamma_{\tilde{\varepsilon}}(\nu, \alpha) \rangle &= \Gamma_{\tilde{h}}(\nu, \alpha)\Gamma_{\tilde{\varepsilon}_i}(\nu, \alpha) + \Gamma_{\tilde{n}}(\nu, \alpha) \\ &= \Gamma_{\tilde{h}}(\nu, \alpha)M_0\tilde{A}(\alpha) + N_0N_t\Delta(\alpha).\end{aligned}$$

It is useful to normalize the FT $\tilde{A}(\alpha)$ by $\tilde{A}(0)$ and to use the arithmetic mean $\bar{A} = \int dt A(t) = N_t^{-1}\tilde{A}(0)$ and then write

$$\langle \Gamma_{\tilde{\varepsilon}}(\nu, \alpha) \rangle = N_t \left[M_0\bar{A}\Gamma_{\tilde{h}}(\nu, \alpha)\tilde{A}(\alpha)/\tilde{A}(0) + N_0\Delta(\alpha) \right]$$

The variance is then

$$\begin{aligned}\text{Var} [\Gamma_{\tilde{\varepsilon}}(\nu, \alpha)] &\equiv \langle |\Gamma_{\tilde{\varepsilon}}(\nu, \alpha)|^2 \rangle - |\langle \Gamma_{\tilde{\varepsilon}}(\nu, \alpha) \rangle|^2 \\ &= |\Gamma_{\tilde{h}}(\nu, \alpha)|^2 \langle S_i(\nu + \alpha/2) \rangle \langle S_i(\nu - \alpha/2) \rangle + \langle S_n(\nu + \alpha/2) \rangle \langle S_n(\nu - \alpha/2) \rangle \\ &\quad + H(\nu + \alpha/2) \langle S_i(\nu + \alpha/2) \rangle \langle S_n(\nu - \alpha/2) \rangle \\ &\quad + H(\nu - \alpha/2) \langle S_i(\nu - \alpha/2) \rangle \langle S_n(\nu + \alpha/2) \rangle \\ &= N_t^2 \left\{ |\Gamma_{\tilde{h}}(\nu, \alpha)|^2 (M_0\bar{A})^2 + N_0^2 + [H(\nu + \alpha/2) + H(\nu - \alpha/2)] M_0N_0\bar{A} \right\}\end{aligned}$$

The S/N of the ACF is

$$\begin{aligned}\left(\frac{S}{N} \right)_{\Gamma_{\tilde{\varepsilon}}} &= \frac{\Gamma_{\tilde{\varepsilon}}(\nu, \alpha)}{\text{Var}^{1/2} [\Gamma_{\tilde{\varepsilon}}(\nu, \alpha)]} \\ &= \frac{\Gamma_{\tilde{h}}(\nu, \alpha)\tilde{A}(\alpha)/\tilde{A}(0) + (N_0/M_0\bar{A})\Delta(\alpha)}{\left\{ |\Gamma_{\tilde{h}}(\nu, \alpha)|^2 + (N_0/M_0\bar{A})^2 + [H(\nu + \alpha/2) + H(\nu - \alpha/2)] N_0/M_0\bar{A} \right\}^{1/2}}.\end{aligned}$$

At zero lag, $\Gamma_{\tilde{h}}(\nu, 0) = H(\nu)$ and the S/N is unity, again as expected. At lags near zero ($\alpha = 0^+$), we have

$$\left(\frac{S}{N} \right)_{\Gamma_{\tilde{\varepsilon}}(0^+)} = \frac{1}{1 + N_0/M_0\bar{A}H(\nu)}.$$

The S/N is unity when the signal is much stronger than the additive noise.

Timing Perturbations

TOAs are perturbed in several ways by the ISM, including dispersion, scattering, and refraction. The latter two are closely related to diffractive and refractive interstellar scintillation (DISS, RISS).

Dispersion: Dispersive arrival times vary deterministically with frequency. Even though there are stochastic variations in time, we assume that DM at any epoch is known well enough that there is no perturbation from any error.

Pulse broadening: Since the pulse-broadening function is *causal*, convolution of it with the intrinsic pulsar signal always leads to a delay that should be corrected. In addition there are applications that can benefit from deconvolution of the PBF to obtain the intrinsic pulse shape. We therefore distinguish between TOA correction and deconvolution since the former does not necessarily require the latter.

The TOA perturbation from $h(t)$ is twofold:

Mean delay: Because $h(t)$ is causal, there is always a positive shift of the arrival time. For an interstellar medium with stationary statistics, the ensemble average of $p(t) = |h(t)|^2$ is a stable function with a well defined mean time $\langle t \rangle$.

Finite-scintle variations: The noiselike aspect of $h(t)$ yields variations about the mean delay that are related to the finite number of scintles that are included in any data set used to calculate the TOA.

Refraction perturbations: Refraction causes angle-of-arrival variations on times scales of days to months or longer. These produce additional timing perturbations as discussed in Cordes & Shannon (2010).

Pulse Broadening Function

Pulse broadening from multipath scattering in the ISM is described by a pulse-broadening function (PBF) that is convolved with the emitted pulse shape to produce the measured pulse.

We identify three time-averages for PBFs, based on the amount of time averaging. The “instantaneous” PBF is constant over a time scale less than a scintillation time Δt_d . The PBF applicable to a long time average spanning many scintillation times at a given epoch is a “snapshot” average (Narayan & Goodman 1989) that may be statistically precise but will not have converged to the ensemble average because slow, refractive modulations will alter the PBF on time scales up to years and longer associated with the broad distribution of length scales in electron-density irregularities. We find it convenient to separate the long, refractive modulations from the shorter term variations in the PBF and therefore will refer to ensemble averages that pertain to a wavenumber spectrum that is effectively truncated at wavenumbers less than the refraction or multipath wavenumber (e.g. Cordes et al. 1986) and then acknowledge the existence of time variations in the PBF, especially the mean pulse-broadening time, caused by larger-scale refractive scales.

The characteristic pulse-broadening time is $\tau_d \sim D\theta_d^2/2c$ where θ_d is the diffraction angle and D is an effective distance. The timing perturbation is related to the width of the PBF, which is strongly wavelength dependent ($\propto \lambda^4$). The pulse broadening times in Figure 5 demonstrate that they are a strong function of both DM. The empirical fit from Bhat et al. (2004) is $\log \tau_d = -6.46 + 0.154 \log \text{DM} + 1.07(\log \text{DM})^2 - (3.86 \pm 0.16) \log \nu$, with a scatter ± 0.65 in $\log \tau_d$.

We write the PBF as $p_d(t) = \bar{p}_d(t) + \delta p_d(t)$. where \bar{p}_d is the ensemble-average shape¹ expected from a particular medium and consider it to be a slow function of epoch. By contrast, δp_d encapsulates fast departures from the ensemble average. Defining $p_d(t)$ to have unit area, the characteristic broadening time is

$$\bar{\tau}_d = \int dt t p_d(t) = \tau_d + \delta\tau_d, \quad (3)$$

and its ensemble average as τ_d .

The scintillation bandwidth and the mean pulse broadening time $\bar{\tau}_d$ are related by²

$$2\pi\bar{\tau}_d\Delta\nu_d = C_1, \quad (4)$$

¹This might also be called a “snapshot” average using the language of Narayan & Goodman (1989), which takes into account that a statistically precise measurement at a given epoch is still not a good ensemble average because there are long-term refraction effects. Our approach is to treat refraction effects as a separate modulation imposed on a statistically precise average at a given epoch.

²Our definitions of $\Delta\nu_d$ and $\bar{\tau}_d$ follow Cordes & Rickett (1998), i.e. $\Delta\nu_d$ is the HWHM of the intensity ACF vs. frequency lag and $\bar{\tau}_d$ is the mean delay of the PBF. This differs from Lambert & Rickett (1999), who use alternative definitions of τ_d in a similar expression and thus quote different values of C_1 than we use here.

where C_1 depends on all the properties of the medium (Cordes & Rickett 1998; Lambert & Rickett 2000), such as the wavenumber spectrum, its thickness and location along the (LOS) and on its transverse extent (Cordes & Lazio 2001). For simple media, values of C_1 can be calculated. For a *thin screen* unbounded transverse to the LOS and with a circular Gaussian angular scattering function, the PBF is a one-sided exponential function with $1/e$ time scale τ_d . In this case the mean time delay is also τ_d so that $C_1 = 1$. For media that are plausibly relevant to pulsar scattering, C_1 can vary by a factor of nearly two (Cordes & Rickett 1998; Lambert & Rickett 1999).

Scattering Correction Regimes

Different views on correction: The TOA perturbation from DISS is associated with the *frequency structure* in the scattered wavefield. This structure decorrelates on a scintillation time scale, Δt_{ISS} (by definition). The perturbation amplitude has a mean component and a variation about the mean. The mean perturbation is related to the pulse-broadening time, τ_d . The variations result from the changing number of scintles in the spectrum combined with τ_d . One philosophy is to correct the pulse shape and TOA on time scales smaller than Δt_{ISS} even if the total integration time spans multiples of Δt_{ISS} . Another approach is to average over multiple Δt_{ISS} and deconvolve the intensity profile. This approach may be less exact but likely will have better S/N.

There are **three regimes** for the correctability of the delay:

1. **Mean shift regime:** In this regime the pulse broadening time is much smaller than the pulse width: $\tau_d \ll W_A$. The TOA shift is simply the mean of the pulse broadening function (Cordes et al. 1990; Hemberger & Stinebring 2008; Cordes & Shannon (2010))

$$\bar{t} = \frac{\int dt tp(t)}{\int dt p(t)}. \quad (5)$$

If the shape of p is somehow known, a determination of τ_d at each epoch from a variety of methods is all that is needed. Generally, however, the shape is not known as studies of secondary spectra have shown.

2. **Coherent deconvolution regime:** In this regime, it is possible to determine the field PBF $h(t)$ from the wavefield ACF. This determination is not unique but it may be sufficient to correct TOAs and deconvolve the pulse shape. Inference of $h(t)$ must be done on a time scale smaller than the scintillation time Δt_{ISS} , which limits the signal-to-noise ratio.
3. **Incoherent deconvolution regime:** For pulsar-frequency combinations where the scattering is very large (e.g. $\tau_d \sim \Delta t_{\text{ISS}}$), it is not possible to determine $h(t)$. However, the pulse can be deconvolved using a template bank of possible intensity PBFs to determine (nonuniquely) both $p(t)$ and the intrinsic pulse shape. An advantage of this approach is that the deconvolution can be done on pulses with long integration times $\gg \Delta t_{\text{ISS}}$ that provide high signal-to-noise ratios.

Non-uniqueness issues: As always, additive noise disallows any unique determination of the PBF. More fundamental, however, is that the starting point is often a correlation function of one of the PBFs or its Fourier transform. It is well known that many functions can have the same correlation function, so the inverse problem does not have a unique solution. However, for PBFs the allowed solutions are restricted by **causality** of the PBF and by the **positivity** of the deconvolved pulse shape.

Minimum-delay PBFs: If the PBFs have the minimum delay property, determination of $h(t)$ or $p(t)$ from one of the relevant correlation functions can be done uniquely.

Mean Shift Regime I

Conceptually, the mean shift defined in Eq. 5 is related to the ACF of h as follows:

$$\Gamma_{\tilde{h}}(\delta\nu) \equiv \tilde{p}(\delta\nu) = \int dt e^{-2\pi i \delta\nu t} p(t).$$

Taking the derivative with respect to $\delta\nu$, we have

$$\partial_{\delta\nu} \Gamma_{\tilde{h}}(\delta\nu) = -2\pi i \int dt e^{-2\pi i \delta\nu t} t p(t).$$

Assuming for simplicity that $p(t)$ is normalized to unit area, we then have

$$\bar{t}_p = \frac{i}{2\pi} \partial_{\delta\nu} [\Gamma_{\tilde{h}}(\delta\nu)]_{\delta\nu=0}.$$

Since p is real, $\Gamma_{\tilde{h}}$ is Hermitian and the derivative at the origin is purely imaginary, allowing us to write

$$\bar{t}_p = -\frac{1}{2\pi} \partial_{\delta\nu} \text{Im} [\Gamma_{\tilde{h}}(\delta\nu)]_{\delta\nu=0}.$$

Because p is also causal, the derivative of the imaginary part is negative, yielding $\bar{t}_p \geq 0$.

If $p(t)$ has arbitrary area, then

$$\bar{t}_p = \frac{i}{2\pi} \partial_{\delta\nu} [\ln \Gamma_{\tilde{h}}(\delta\nu)]_{\delta\nu=0}.$$

For the actual pulsar signal, we obtain the frequency ACF

$$\Gamma_{\tilde{\varepsilon}}(\delta\nu) = \Gamma_{\tilde{h}}(\delta\nu) \tilde{A}(\delta\nu) = \tilde{p}(\delta\nu) \tilde{A}(\delta\nu).$$

Using the same approach as before, the TOA shift associated with $\Gamma_{\tilde{\varepsilon}}$ is

$$\bar{t}_{\varepsilon} = -\frac{1}{2\pi} \partial_{\delta\nu} \text{Im} [\Gamma_{\tilde{h}}(\delta\nu)]_{\delta\nu=0} = \bar{t}_p + \bar{t}_A,$$

where the shift associated with the true pulse shape is

$$\bar{t}_A = -\frac{1}{2\pi} \partial_{\delta\nu} \text{Im} [\tilde{A}(\delta\nu)]_{\delta\nu=0}.$$

The pulse shift \bar{t}_A is arbitrary with respect to how the time origin for $A(t)$ is defined. For a symmetric profile delayed to a time t_0 , we would have $\bar{t}_A = t_0$.

For an asymmetric pulse shape and an arbitrary delay t_0 , the total delay obtained from the derivative of $\Gamma_{\tilde{\varepsilon}}(\delta\nu)$ is

$$\bar{t}_{\varepsilon} = \bar{t}_p + \bar{t}_A + t_0.$$

The Mean-Shift Regime II

A case of interest for nearby pulsars or observations at high frequencies is where the pulse broadening time is small compared to the pulse width, $\tau_d \ll W$ but where the strong scintillation regime applies. This requires that the scintillation bandwidth $\Delta\nu_d \ll \nu$ or, equivalently, $\tau_d\nu \gg 1$. We specify the mean-shift regime as

$$\frac{C_1}{2\pi\nu} \ll \tau_d \ll W \quad \text{or} \quad 0.2 \text{ ns } \nu_{\text{GHz}}^{-1} \ll \tau_d \ll 1 \text{ ms } W_{\text{ms}}. \quad (6)$$

Under these conditions the TOA perturbation is simply the pulse broadening time calculated from Eq. 5 using the instantaneous PBF (not the ensemble-average PBF). The TOA perturbation also depends on the number of scintles contributing to an observation. In many cases it is large and the TOA perturbation converges to some mean value. This requires that $\Delta\nu_d \ll B$, where B is the bandwidth, or that $\Delta t_{\text{ISS}} \ll T$, where T is the total time of the observation. Figure 5 yields the range of DM for which the shift approximation applies: $\tau_d < 1 \text{ ms}$ for DM $\lesssim 100, 250$ and 400 pc cm^{-3} at 0.4, 1 and 2 GHz, respectively.

One approach to timing precision is to require that τ_d be smaller than some specified rms timing error, σ_{max} , which might be 100-ns or less for pulsar timing array applications but could be much larger for other timing purposes (such as pulsars in the Galactic center). If $\tau_d < \sigma_{\text{max}}$, pulse broadening might be ignored entirely. Alternatively, it might be corrected to a fractional precision ϵ_τ , yielding a residual error $\epsilon_\tau \tau_d < \sigma_{\text{max}}$. Using the empirical fit for τ_d vs. DM shown in Figure 5, we can then solve for DM(ν) for different σ_{max} . Results are shown in Figure 5. The smallest pulse broadening occurs in the lower-right portion of the diagram whereas objects with large DMs are strongly affected by pulse broadening. Many pulsars have steep spectra, so the choice of frequencies for a given object requires a compromise.

When the shift approximation holds, the TOA can be corrected by subtracting $\bar{\tau}_d$ from nominal values (Hemberger & Stinebring 2008).

However, because by definition $\bar{\tau}_d$ is small, we cannot measure it directly at the frequency of the TOA measurements. There are several ways to estimate it, however.

Correlation Approach

The timing delay τ_d can be estimated by using Eq. 4, where the characteristic scintillation bandwidth $\Delta\nu_d$ is calculated as the half-width at half maximum of the intensity correlation function $\Gamma_{\delta I}(\delta\nu)$ that is calculated from the dynamic spectrum $I(\nu, t)$. This approach requires knowledge of the shape of the PBF so the type of medium must be known to determine the appropriate value for C_1 . Alternatively, a characteristic time τ_{S_2} can be calculated from the secondary spectrum, the squared-magnitude of the Fourier transform of the dynamic spectrum (e.g. Stinebring et al. 2001), which is related to the autocorrelation function of the PBF. The relationship of τ_{S_2} to τ_d is also medium dependent and requires knowledge of their ratio that is analogous to knowing C_1 . The correction of TOAs for the mean PBF delay is therefore problematic unless ancillary constraints

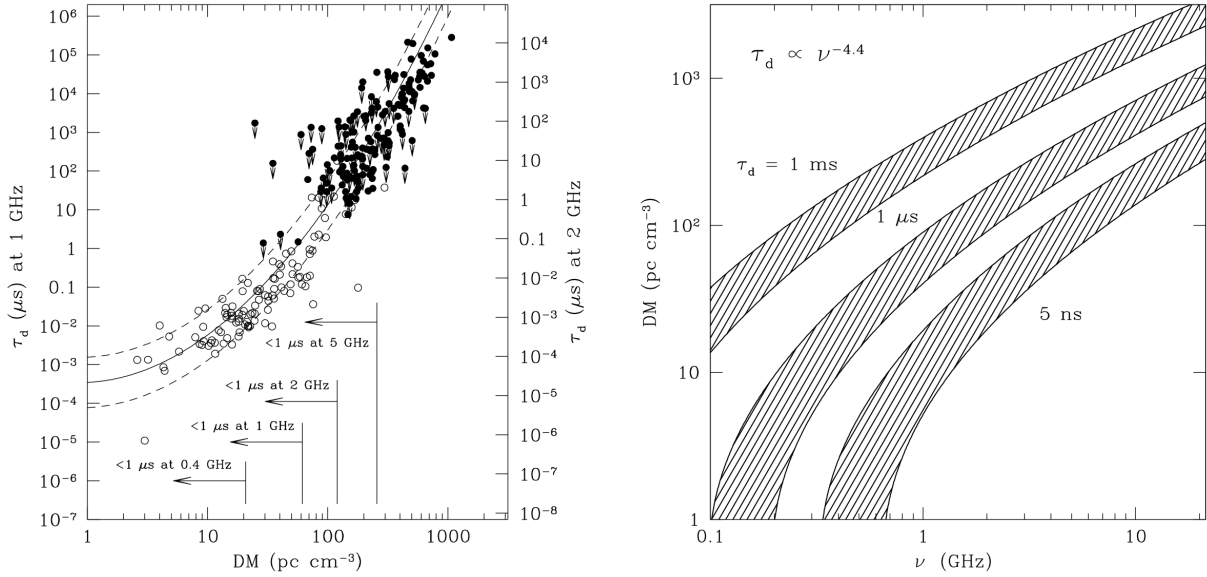


Fig. 5.— (Left) Pulse broadening time plotted against DM for pulsar data used to construct the NE2001 model (Cordes & Lazio 2002). The solid curve shows the parabolic fit given in the text for $\log \tau_d$, and the dashed curves show $\pm 1.0\sigma$ deviations from the fit. Open circles are inferred values using scintillation bandwidth measurements; filled circles are direct measurements of τ_d . Downward arrows denote upper limits. All measurements have been scaled to 1 GHz assuming $\tau_d \propto \nu^{-4.4}$ and the right-hand scale for 2 GHz assumes the same scaling. The strong scattering regime is assumed, which requires $\tau_d \gg C_1/2\pi\nu$ or $\tau_d \gg 2 \times 10^{-4} \mu s$ at 1 GHz. At lower frequencies, the entire range of τ_d in the plot satisfies the requirement. Leftward arrows designate maximum values of DM for which $\tau_d < 1$ ms at 0.4, 1 and 2 GHz using the nominal parabolic fit. (Right) Contours of constant pulse broadening time, τ_d , using the fit shown in Figure 5 including the $\pm 1\sigma$ uncertainty in the fit, which results in the widths of the shaded regions. A $\tau_d \propto \nu^{-4.4}$ scaling with frequency is used. This scaling is steeper than appropriate for some high DM objects but is typical for low-DM pulsars.

on the scattering physics are known, namely whether a thin screen or extended-medium applies, whether the scattering region is bounded transverse to the LOS, what the wavenumber spectrum is, etc. The secondary spectrum provides much of this information (Stinebring et al. 2001; Hemberger & Stinebring 2008).

The broadening time can be estimated from the scintillation bandwidth obtained at a frequency ν' ,

$$\hat{\tau}_d(\nu') = \frac{\hat{C}_1}{2\pi \widehat{\Delta\nu_d}(\nu')}, \quad (7)$$

where quantities requiring an estimate or an assumed value are hatted. The broadening time can

be scaled to other frequencies using

$$\hat{\tau}_d(\nu) = \hat{\tau}_d(\nu') \left(\frac{\nu}{\nu'} \right)^{-\hat{X}_{\text{PBF}}}. \quad (8)$$

The estimation error of the chromatic term and thus on t_∞ expands into terms dependent on individual errors, δC_1 , $\delta \Delta \nu_d$, and δX_{PBF} :

$$\delta \tau_d = \tau_d - \hat{\tau}_d \approx \tau_d \left[\frac{\delta C_1}{C_1} - \frac{\delta \Delta \nu_d}{\Delta \nu_d} - \delta X_{\text{PBF}} \ln \frac{\nu}{\nu'} \right]. \quad (9)$$

Based on estimates of scintillation bandwidths in the literature, the range of possible values of C_1 , and the error in the exponent, δX_{PBF} , the combined error in $\delta \tau_d$ is unlikely to be less than 10% with this approach. However, a 10% error in removing pulse broadening implies that the shaded bands in Figure 5 will shift to the left by about a factor of 1.8 in frequency (0.26 in $\log \nu$).

Secondary Spectrum Approach

Phase Retrieval Methods

A second, empirical approach requires much less *a priori* knowledge about the scattering medium. It calculates the complete PBF using only a measurement of the ACF of the PBF or, equivalently, the magnitude of the PBF's Fourier transform. The ACF of the PBF can be obtained from the secondary spectrum, for example. Walker et al. (2008) apply two-dimensional phase retrieval to the secondary spectrum and demonstrate the general principle. For the one-dimensional ACF with n_τ lags, there are n_τ functions that can produce the measured ACF. Many of these are acausal or include negative values and can therefore be ruled out. Those that remain include the function with minimum delay and its time reverse, which has maximum delay. Methods exist to calculate the *minimum-delay* PBF from the ACF, which is the unique member in the large family of PBFs consistent with the ACF that has its amplitude most concentrated toward the origin in a mean-square sense (Scargle 1981). While physically a minimum-delay solution is not demanded by scattering and refracting geometries, it may be the most probable result given that scattering and refraction angles have distributions peaked at zero. Also it may provide a more accurate correction compared to use of an idealized form for the PBF used to provide a value for C_1 in the correlation method. Application of inversion techniques has been explored in unpublished work (JMC). Another approach (P. Demorest & M. Walker, private communication) investigates the phase of the wavefield directly to determine the PBF.

Incoherent Deconvolution

When the scattering time τ_d is too large for the mean-shift regime to apply and where the signal-to-noise required for coherent deconvolution is not achieved, the PBF can be deconvolved incoherently from the measured pulse shape. Incoherent deconvolution is tricky because the PBF is unknown and so it must be assumed (e.g. Kuz'min & Izvekova 1993) or deduced from the data (Bhat et al. 2004). To be useful, the pulse-broadening time τ_d must be large enough to be “discernible,” so we define this as meaning that it is larger than some fraction ϵ of the pulse width W . Defining also the sample time across the pulse profile (i.e. the sample interval in pulse phase expressed in temporal units), Δt_p we have the inequalities

$$B^{-1} \ll \Delta t_p \lesssim \epsilon W \lesssim \tau_d.$$

Defining the duty cycle $f_{\text{DC}} = W/P$ we define appropriate regions in the frequency-time plane by requiring

$$\tau_d \gtrsim \epsilon W = \epsilon f_{\text{DC}} P.$$

Figure 6 shows demarkation lines using this relation for 0.1, 0.4, 1.4, and 2 GHz assuming a constant duty cycle $f_{\text{DC}} = 0.05$. If a period-dependent duty cycle is used, $f_{\text{DC}} \propto P^{-1/2}$, the curves are not dramatically different except for small periods where the duty cycle is large, requiring a larger broadening time to satisfy the condition. Since the MSPs used in precision timing have pulse widths much narrower than predicted by the scaling that applies to long-period pulsars, the fixed duty-cycle curves appear appropriate.




Fig. 6.— Regions in the DM-Period plane where it is efficacious to incoherently deconvolve the pulse-broadening function. The region above each line indicates the region for the frequency labeling each line.

Requirements for Coherent Deconvolution I

The primary observable is the wavefield ε (e.g. the baseband signal) and the primary statistic for extracting the field PBF $h(t)$ and the unscattered pulse is the ACF Γ_ε .

Determining $h(t)$ or equivalently $\tilde{h}(\nu)$ from the ACF of \tilde{h} or from the intensity PBF $p(t) = |h(t)|^2$ is the classic problem of determining a function from the magnitude of its Fourier transform.

There are two kinds of requirements:

1. Integration times, pulse widths, scattering times and bandwidths must satisfy certain relationships. These in turn can put constraints on the periods P and dispersion measures DM that are amenable to coherent deconvolution.
2. There needs to be adequate signal-to-noise ratio so that the ACF Γ_ε can be well estimated.

Requirements for Coherent Deconvolution II

Using $\Gamma_\varepsilon(\alpha)$ to determine $h(t)$ has the following requirements:

1. By assumption, scattering matters because the strong-scattering regime applies:

$$\Delta\nu_d \ll \nu \quad \text{or} \quad \boxed{\tau_d \gg \frac{1}{2\pi\nu}}.$$

2. A more restrictive constraint arises by requiring that a scintle be no wider than the bandwidth, $2\Delta\nu_d \lesssim B$, or

$$\boxed{\tau_d \gtrsim \frac{1}{\pi B}}$$

3. The data span T used to calculate an FT combined with the number of data blocks N_b needed to improve statistical variations of Γ_ε from pulsar self noise must satisfy

$$\boxed{N_b T \ll \Delta t_d \approx \frac{c}{\nu V_{\text{eff}}} \left(\frac{D}{2c\tau_d} \right)^{1/2}}$$

Keeping statistical variations smaller than a fractional error $\sigma_{\Gamma_{\varepsilon, \text{max}}} / |\Gamma_\varepsilon|$ implies the number of blocks must satisfy

$$\boxed{N_b \geq \left(\frac{|\Gamma_\varepsilon|}{\sigma_{\Gamma_{\varepsilon, \text{max}}}} \right)^2}$$

4. The frequency resolution needs to resolve scintles of width $\sim \Delta\nu_d$:

$$\Delta\nu = \frac{1}{T} \ll \frac{1}{2\pi\tau_d} \quad \text{or} \quad \boxed{T \gg 2\pi\tau_d}$$

5. The ACF needs to be sampled adequately. Since it factors into $\tilde{A}(\alpha_k)$ and $\Gamma_{\tilde{h}}(\alpha_k)$, we require

$$\begin{aligned} \Delta\alpha = \frac{1}{P} \ll W_A^{-1} \quad & \text{or} \quad \boxed{P \gg W_A} \\ & \text{and} \\ \Delta\alpha \ll \frac{1}{2\pi\tau_d} \quad & \text{or} \quad \boxed{P \gg 2\pi\tau_d} \end{aligned}$$

6. To have sufficient frequency resolution, the FT can be calculated over one or more pulse periods

$$\boxed{T \geq P}$$

Requirements for Coherent Deconvolution III

For the purpose of defining regimes in P and DM, we consider the following inequalities:

$$W_A \sim 2\pi\tau_d \ll P \leq T \ll \frac{\Delta t_d}{N_b}$$

or putting in explicit DM and frequency dependences,

$$2\pi\tau_d(\text{DM}, \nu) \ll P \leq T \ll \frac{c}{2\pi\nu V_{\text{eff}} N_b} \left(\frac{D}{2c\tau_d(\text{DM}, \nu)} \right)^{1/2}.$$

We replace the extreme inequalities “ \ll ” with a factor $F = 5$ and we use a nominal number of blocks $N_b = 10^4$ to yield 1% errors in the ACF $\Gamma_{\hat{\varepsilon}}$:

$$10\pi\tau_d \left(\frac{F}{5} \right) \lesssim P \lesssim 10^{-4.7} \left(\frac{5}{F} \right) \left(\frac{10^4}{N_b} \right) \times \frac{c}{2\pi\nu V_{\text{eff}} N_b} \left(\frac{D}{2c\tau_d(\text{DM}, \nu)} \right)^{1/2}.$$

To evaluate the constraints we use the empirical fit for pulse broadening times from Bhat et al. (2004) assuming $\tau_d \propto \nu^{-4}$,

$$\log \tau_d(\text{ms}) = -6.59 + 0.129 \log \text{DM} + 1.02(\log \text{DM})^2 - 4 \log \nu.$$

There is significant scatter about this expression that should be included in the analysis. That will come later.

We therefore have two constraints on period as a function of τ_d (which depends implicitly on DM) that define a region in the P-DM plane where deconvolution is possible.

Equality of the two constraints on P , $2\pi\tau_d = \Delta t_{\text{ISS}}/N_b$ corresponds to equality of the pulse broadening time to the diffractive scintillation time scale divided by the time for one block, as applies to low-frequency observations of high-DM pulsars:

$$\tau_d = \left[\frac{cD}{2(2\pi)^4(\nu V_{\text{eff}} N_b)^2} \right]^{1/3} = \frac{67 \text{ ms}}{(\nu V_{100})^{2/3}} \frac{D^{1/3}}{N_b^{2/3}} = \frac{144 \mu\text{s}}{(\nu V_{100})^{2/3}} \left(\frac{10^4}{N_b} \right)^{2/3}.$$

Somewhat counterintuitively, a larger pulse broadening time can be tolerated for more distant, lower velocity pulsars observed at low frequencies.

Figures 7 to 10 show regimes at four frequencies for $N_b = 10, 100, 10^3$ and 10^4 blocks.

Coherent Deconvolution Regimes for $N_b = 10$

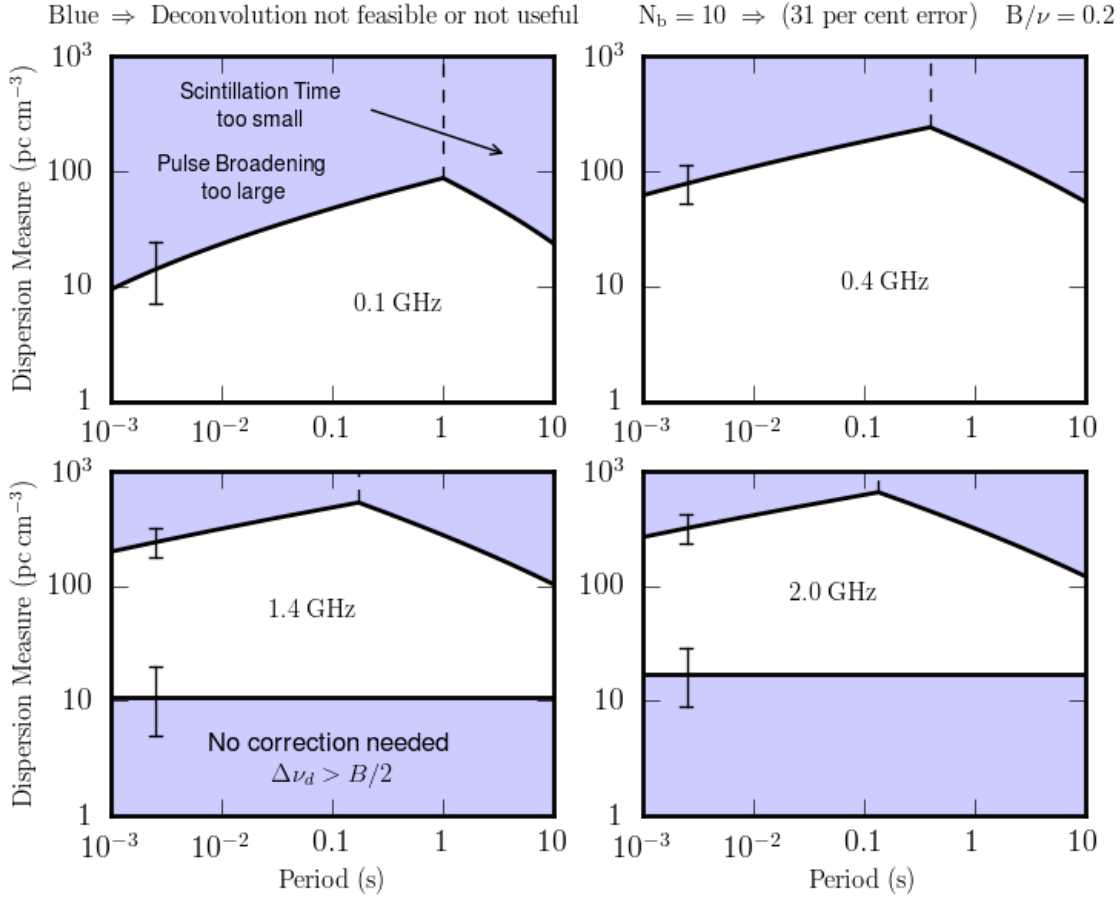


Fig. 7.— Regions in P –DM space where it is advantageous to attempt coherent deconvolution and the conditions for coherent deconvolution are satisfied. This figure is for an analysis of $N_b = 10$ blocks, which corresponds to 31% error in the average correlation function. The blue shaded regions correspond to where either the condition $2\pi\tau_d \ll P$ or $P \ll \Delta t_d/N_b$ is not satisfied. At 1.4 and 2 GHz, the blue region at the bottom of each frame arises because the total scattering delay is small enough, $\tau_d \lesssim 2\nu^{-1}$ ns, so that deconvolution is not necessary given larger contributions to timing errors. In all four frames, the apex of the allowed (clear) region corresponds to $2\pi\tau_d = \Delta t_d/N_b$. The vertical dashed line marks the period of the apex. The vertical error bars indicate the range of the associated curves corresponding to the range of pulse broadening seen at a given DM. A fractional bandwidth $B/\nu = 0.2$ has been assumed.

Coherent Deconvolution Regimes for $N_b = 100$

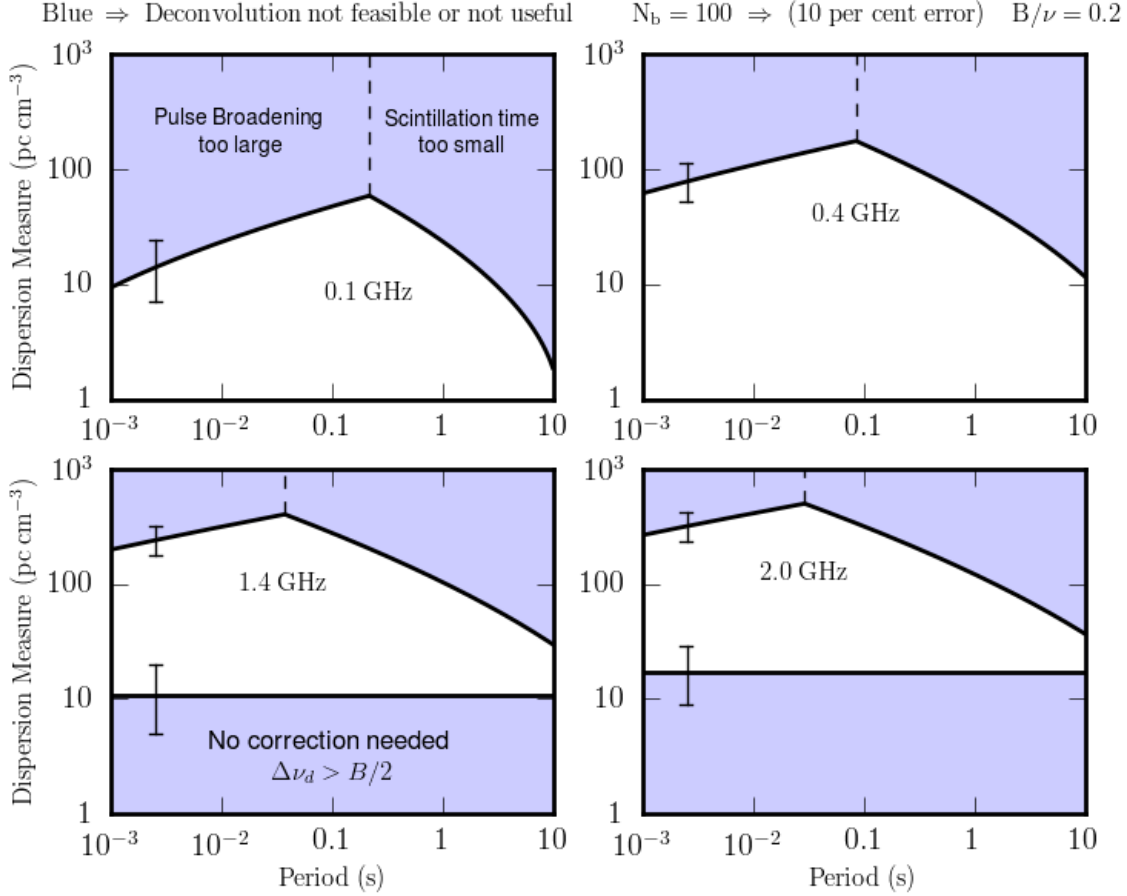


Fig. 8.— Regions in P –DM space where it is advantageous to attempt coherent deconvolution and the conditions for coherent deconvolution are satisfied. This figure is for an analysis of $N_b = 100$ blocks, which corresponds to 10% error in the average correlation function. The blue shaded regions correspond to where either the condition $2\pi\tau_d \ll P$ or $P \ll \Delta t_d/N_b$ is not satisfied. At 1.4 and 2 GHz, the blue region at the bottom of each frame arises because the total scattering delay is small enough, $\tau_d \lesssim 2\nu^{-1}$ ns, so that deconvolution is not necessary given larger contributions to timing errors. In all four frames, the apex of the allowed (clear) region corresponds to $2\pi\tau_d = \Delta t_d/N_b$. The vertical dashed line marks the period of the apex. The vertical error bars indicate the range of the associated curves corresponding to the range of pulse broadening seen at a given DM. A fractional bandwidth $B/\nu = 0.2$ has been assumed.

Coherent Deconvolution Regimes for $N_b = 10^3$

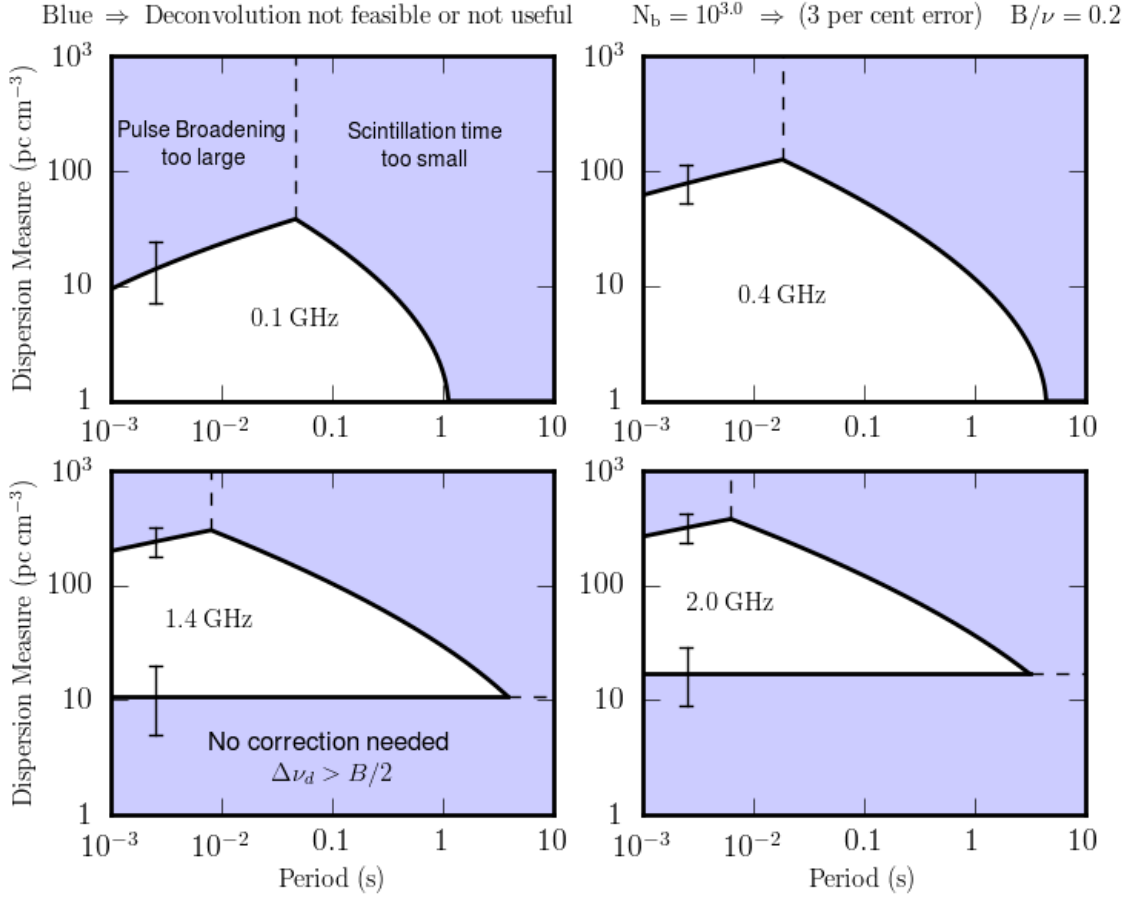


Fig. 9.— Regions in P –DM space where it is advantageous to attempt coherent deconvolution and the conditions for coherent deconvolution are satisfied. This figure is for an analysis of $N_b = 10^3$ blocks, which corresponds to 3% error in the average correlation function. The blue shaded regions correspond to where either the condition $2\pi\tau_d \ll P$ or $P \ll \Delta t_d/N_b$ is not satisfied. At 1.4 and 2 GHz, the blue region at the bottom of each frame arises because the total scattering delay is small enough, $\tau_d \lesssim 2\nu^{-1}$ ns, so that deconvolution is not necessary given larger contributions to timing errors. In all four frames, the apex of the allowed (clear) region corresponds to $2\pi\tau_d = \Delta t_d/N_b$. The vertical dashed line marks the period of the apex. The vertical error bars indicate the range of the associated curves corresponding to the range of pulse broadening seen at a given DM. A fractional bandwidth $B/\nu = 0.2$ has been assumed.

Coherent Deconvolution Regimes for $N_b = 10^4$

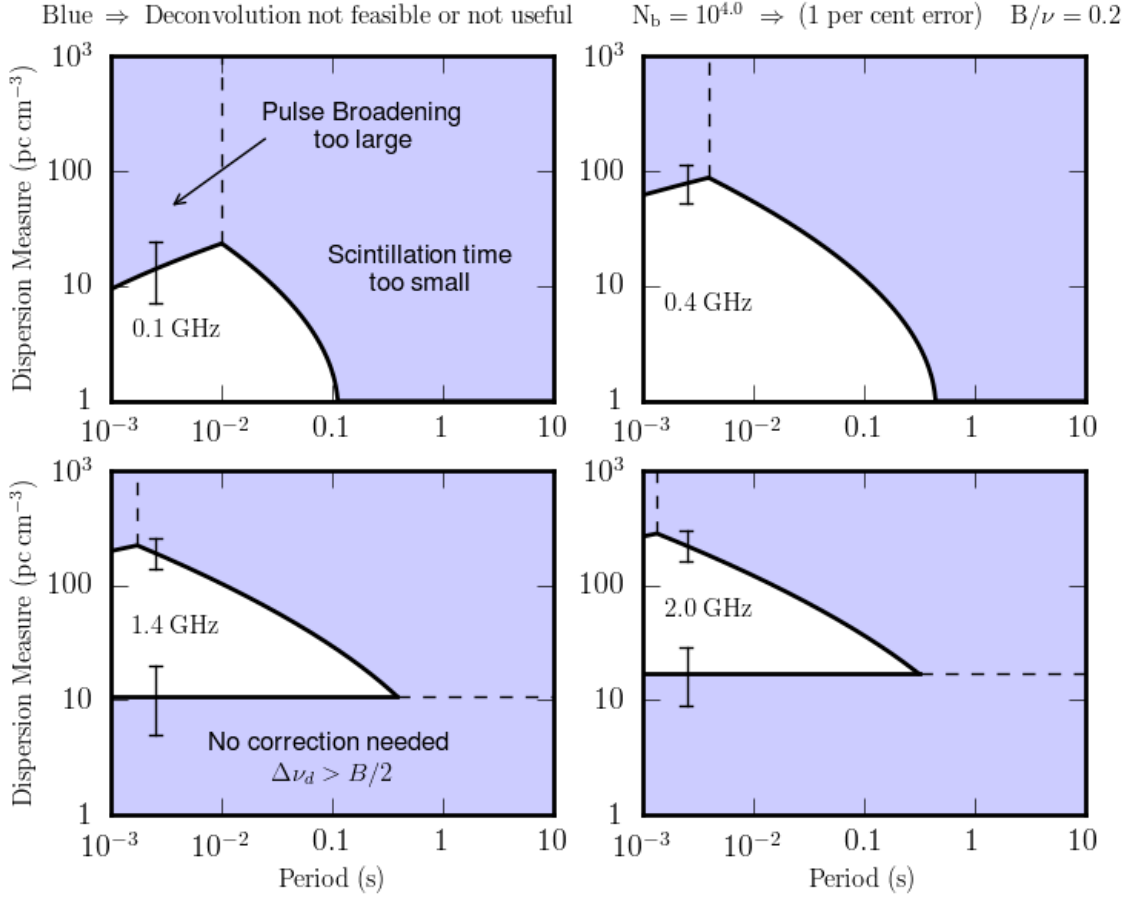


Fig. 10.— Regions in P –DM space where it is advantageous to attempt coherent deconvolution and the conditions for coherent deconvolution are satisfied. This figure is for an analysis of $N_b = 10^4$ blocks, which corresponds to 1% error in the average correlation function. The blue shaded regions correspond to where either the condition $2\pi\tau_d \ll P$ or $P \ll \Delta t_d/N_b$ is not satisfied. At 1.4 and 2 GHz, the blue region at the bottom of each frame arises because the total scattering delay is small enough, $\tau_d \lesssim 2\nu^{-1}$ ns, so that deconvolution is not necessary given larger contributions to timing errors. In all four frames, the apex of the allowed (clear) region corresponds to $2\pi\tau_d = \Delta t_d/N_b$. The vertical dashed line marks the period of the apex. The vertical error bars indicate the range of the associated curves corresponding to the range of pulse broadening seen at a given DM. A fractional bandwidth $B/\nu = 0.2$ has been assumed.

Methods for Coherent Deconvolution

1. Solve for $h(t)$ given the spectrum $S(\nu)$. This could make use of phase retrieval with the HT combined with other constraints such as a maximum entropy term in the cost function for the deconvolved pulse profile to enforce positivity and identifying zeros in the z-transform of $h(t)$.
2. Solve for $h(t)$ using the frequency ACF $\Gamma_\varepsilon(\nu, \alpha)$. The maximum entropy cost can be used here too.

Idealized cases: First we see how well $h(t)$ can be determined from the magnitude of its Fourier transform $H(\nu)$. This corresponds to the case where the pulsar signal is a delta function.

Another idealized case consists of determining $h(t)$ from measurements of the correlation function $\Gamma_{\tilde{h}}(\nu, \alpha)$.

Realistic cases: Actual data involve determination of the PBF and an arbitrary pulse shape $A(t)$ from measurements of the spectrum

$$S(\nu) = H(\nu) |\tilde{\varepsilon}_i(\nu)|^2$$

and the field ACF

$$\Gamma_\varepsilon(\nu, \alpha) = \Gamma_{\tilde{h}}(\nu, \alpha) \Gamma_{\tilde{\varepsilon}_i}(\nu, \alpha).$$

Both of these can be averaged over as many data blocks N_b as are allowed by the DISS time scale. The larger N_b is, the lower the contribution from self noise to $S(\nu)$ and $\Gamma_\varepsilon(\nu, \alpha)$.

Phase Retrieval for Causal Pulse Broadening Functions

By Fermat’s principle, the PBF is causal³. Treating the PBF as a filter that describes scattering in the ISM, we can consider the following:

Theorem: A general causal filter can be written as the product of the *minimum-phase* filter and a set of single-stage *all-pass* filters.

The power of this theorem is that the field PBF may be determinable uniquely from measurements that provide only partial information, such as the magnitude $|\tilde{h}(\nu)|$ and autocorrelation function $\Gamma_{\tilde{h}(\nu)}(\nu, \delta\nu)$ of the Fourier transform of the PBF, $\tilde{h}(\nu)$.

If the true PBF has the minimum-phase property, it can be determined uniquely. An exponential decaying PBF satisfies this criterion as do other monotonically decaying functions. If the PBF does not have this property, it may still be determined uniquely from perfect data (i.e. with no contribution from pulsar self-noise or additive noise) by combining the minimum-phase solution with all-pass filters, the number of which will depend on how much scattering there is.

The approach explored below is as follows:

1. Determine the minimum-phase ϕ_{HT} from the Hilbert transform of the magnitude of the FT, $|\tilde{h}(\nu)|$, and then determine the PBF $h_{h\text{HT}}(t)$, and the implied mean PB time, \bar{t}_{HT} .
2. Characterize how well \bar{t}_{HT} can be used as an estimate by itself as the *true* PB time.
3. Explore calculation of a better restoration by using additional information from the data to define all-pass (AP) filters that augment the HT solution, in accordance with the decomposition theorem for causal filters. For this, we fit for the parameters of the AP filters by requiring a good fit to the ACF, $\Gamma_{\tilde{h}(\nu)}(\nu, \delta\nu)$, by the corresponding ACF calculated from the HT and AP filters.

³While causality of the PBF is clear in principle, there are practical issues in defining the relevant time origin. For a pulse of arbitrary shape and width, the definition of the arrival time depends on a choice for a fiducial time within the pulse. Dispersion and refraction from large scale structure both cause energy to arrive later than in a vacuum. In principle these can both be removed by exploiting their frequency dependences (ν^{-2} and ν^{-4} , respectively). However, any intrinsic frequency dependence of the pulse shape will couple with interstellar effects to introduce uncertainty in the arrival time.

Idealized Cases for Determining $h(t)$ I.

Phase retrieval from the magnitude $|\tilde{h}(\nu)|$ using only the Hilbert transform: we consider cases where the PBF is a one-sided exponential that modulates complex noise and has a nominal (ensemble average) time scale τ_d (in the intensity). The mean time delay is τ_d but realization-to-realization variations are a factor of two smaller or larger. Using only the Hilbert transform to retrieve the phase for scattering times of only a few samples, we find that the true scattering time (which differs statistically from τ_d) is never smaller than the scattering time calculated from the HT solution, as expected. In some cases, the two are equal while on average the ratio is of order 0.5.

Figure 11 shows a case⁴ where the retrieved PBF is completely consistent with the true PBF for a nominal $\tau_d = 3.1$. A case is shown in Figure 12 where the retrieved PBF is concentrated closer to the origin than the true PBF. Scrutiny of the true and retrieved phases in the figure shows that there are similarities at some frequencies but there are a few frequency ranges where the phase rotation is much more rapid.

Interpretation: The HT solution is the “minimum phase” solution which really means “minimum delay.” The HT solution is consistent with the available data (the magnitude of $\tilde{h}(\nu)$) but concentrates power as close as possible to the time origin. One can see in Figure 12 that the HT phase matches some of the variations in the true phase but the true phase shows more total phase wrap across the band. This suggests that if there is a way to identify where these phase wraps occur in frequency and fit for them, that the true phase variation can be retrieved.

⁴The results in Figures 11-15 use $N = 1024$ -point DFTs with no zero-padding. Since the PBFs used have $1/e$ time scales $\tau_d \ll N$, they decay to zero well within the first half of the time-domain vector, even including the long exponential tail. Bias from the inherent periodicity of the DFT is therefore avoided.




Fig. 11.— Example of phase retrieval using the HT for an exponential PBF where the retrieved and true pulse broadening times are equal. The nominal broadening time is $\tau_d = 3.1$ in arbitrary units so the $1/e$ time in $h(t)$ is $2\tau_d$. From top to bottom the frames show the complex PBF, the magnitude of its FT, the phase of the FT, the minimum-phase calculated from the HT of $\log |\tilde{h}(\nu)|$, and the resulting predicted $h(t)$.




Fig. 12.— Example of phase retrieval using the HT for an exponential PBF where the retrieved pulse broadening time is substantially less than the true time. The nominal broadening time is $\tau_d = 3.1$ in arbitrary units so the $1/e$ time in $h(t)$ is $2\tau_d$. From top to bottom the frames show the complex PBF, the magnitude of its FT, the phase of the FT, the minimum-phase calculated from the HT of $\log |\tilde{h}(\nu)|$, and the resulting predicted $h(t)$. The phase curves in the third and fourth panels from the top are nearly identical except for fast phase rotations in the true phase at frequency indices $\sim 430, 620$, and 800 . These can be modeled with all-pass filters that augment the HT solution to bring the net model and actual phase curve into agreement. The nominal broadening time is $\tau_d = 3.1$ in arbitrary units so the $1/e$ time in $h(t)$ is $2\tau_d$. From top to bottom the frames show the complex PBF, the magnitude of its FT, the phase of the FT, the minimum-phase calculated from the HT of $\log |\tilde{h}(\nu)|$, and the resulting predicted $h(t)$.

Idealized Cases for Determining $h(t)$ II.

Properties of the Hilbert-transform solution for the pulse-broadening function: Figures 13-15 show the mean pulse-broadening (PB) time for the true PBF plotted against the mean PB time using the HT solution along with a scatter plot and histogram of their ratio. The three figures cover progressively larger pulse-broadening times. For small PB times less than one sample, the HT solution is highly likely to be equal to the true PBF. For larger PB times from 5 to 50 samples, the PB time calculated from the HT solution (\bar{t}_{HT}) tends toward a constant fraction of the true PB time (\bar{t}). For intermediate cases where the PB time is one to five samples, the HT solution gives the correct PB time in some cases but not in others. Later we show that in this regime the HT solution can be augmented with a series of all-pass filters that can yield the true PBF.

Correction methods in the mean-shift regime: Based on these results we define three methods whose choice depends on the degree of scattering:

1. *Method 1* for large scattering times (many samples): Use the HT solution and scale the PB time \bar{t}_{HT} on the basis of simulations to get the TOA correction.
2. *Method 2* for small scattering times (less than one sample): Use the PB time calculated from the HT solution as the TOA correction.
3. *Method 3*: For intermediate scattering times of one to about 10 samples, start with the HT solution and augment it with cascaded all-pass filters whose parameters are constrained by the data.




Fig. 13.— Top panel: Mean scattering time from the HT solution plotted against the mean time for the actual PBF. The solid line is a least-squares fit to the plotted data, $0.06717 + 0.6727x$. The dashed line shows equality of the estimated and true times. The true time is never smaller than the estimated time, as expected for the minimum-delay solution. Five different ensemble-mean values from 0.1 to 10 time units were used from 0.1 to 1 time samples. Bottom left: scatter plot of the ratio $R = \bar{t}_{\text{HT}}/\bar{t}$. Bottom right: Histogram of R with its mean and rms values indicated. For the small pulse broadening times shown here, the vast majority of the Hilbert-transform solutions for the PBF are identical to the true PBF.




Fig. 14.— Top panel: Mean scattering time from the HT solution plotted against the mean time for the actual PBF. Five different ensemble-mean values from 0.1 to 10 time units were used from 1 to 5 time samples. The solid line is a least-squares fit to the plotted data, $y = 0.3446 + 0.5989x$. The dashed line shows equality of the estimated and true times. The true time is never smaller than the estimated time, as expected for the minimum-delay solution. Bottom left: scatter plot of the ratio $R = \bar{t}_{\text{HT}}/\bar{t}$. Bottom right: Histogram of R with its mean and rms values indicated. For these larger pulse broadening times, there are fewer cases where the Hilbert-transform solutions for the PBF are identical to the true PBF, as reflected in the smaller mean value of R .

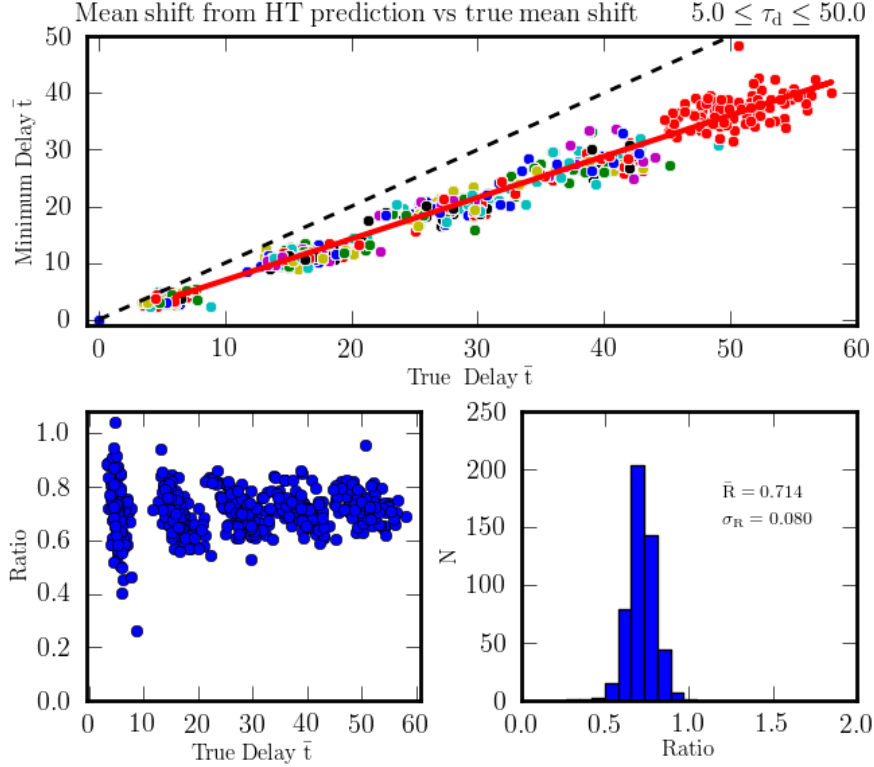


Fig. 15.— Top panel: Mean scattering time from the HT solution plotted against the mean time for the actual PBF. Five different ensemble-mean values from 0.1 to 10 time units were used from 5 to 50 time samples. The solid line is a least-squares fit to the plotted data, $y = -0.2988 + 0.7278x$. The dashed line shows equality of the estimated and true times. The true time is never smaller than the estimated time, as expected for the minimum-delay solution. Bottom left: scatter plot of the ratio $R = \bar{t}_{\text{HT}}/\bar{t}$. Bottom right: Histogram of R with its mean and rms values indicated. For these larger pulse broadening times, none of the the Hilbert-transform solutions for the PBF are identical to the true PBF. However the ratio R clusters around a mean value ~ 0.7 with the degree of scatter decreasing with larger pulse-broadening time, as seen in the bottom left panel. This is associated with the number of scintles contained within the spectrum, which increases with larger pulse-broadening time.

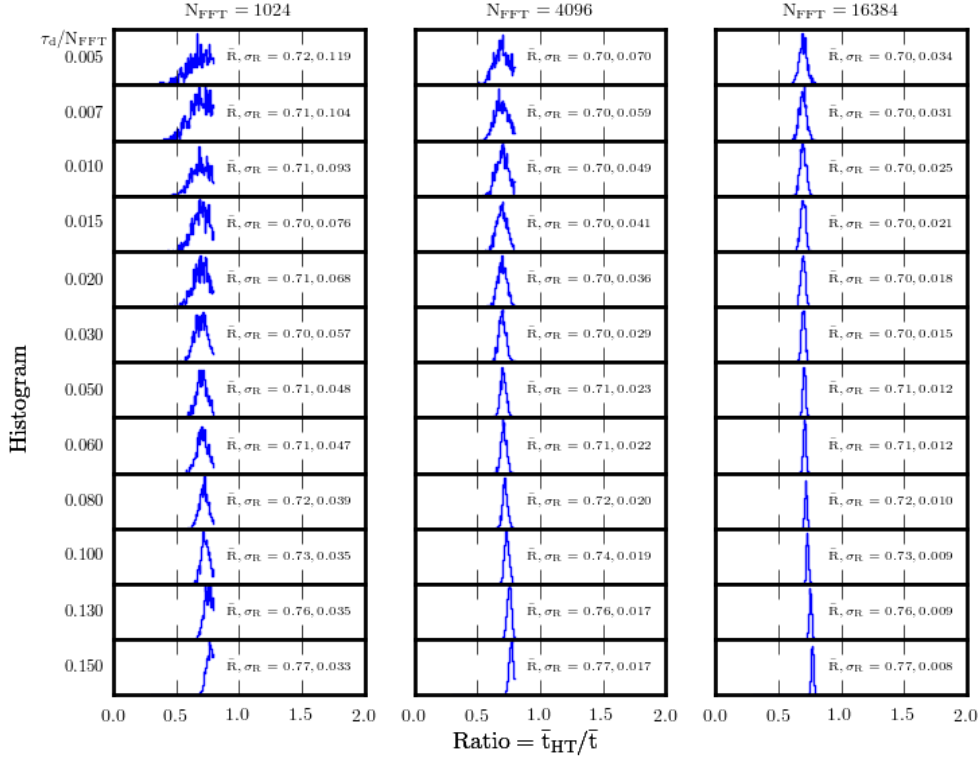


Fig. 16.— Histograms of the ratio $\bar{t}_{\text{HT}}/\bar{t}$ of the mean PB time obtained from the Hilbert-transform solution to the true mean PB time for 1000 realizations in each case. The columns are for different lengths of FFTs (1k, 4k, and 16k) and the rows are for different values of the pulse broadening time, expressed relative to the length of the FFT. Each plot is labelled with the mean and rms ratio.




Fig. 17.— Plot of the ratio $\bar{t}_{\text{HT}}/\bar{t}$ vs. τ_d/N_{FFT} for three FFT lengths (1k, 4k, and 16k) based on averages over 1000 realizations for each plotted point. Error bars on the mean value are visible for some of the points.

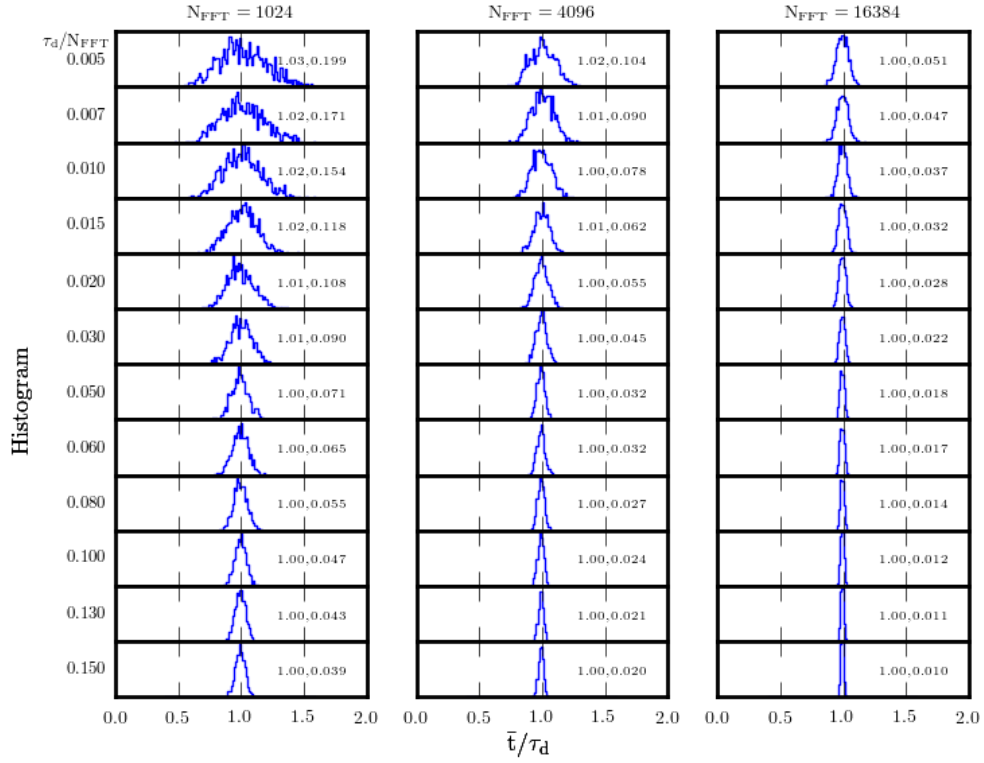


Fig. 18.— Histograms of the ratio \bar{t}/τ_d for different τ_d/N_{FFT} and three FFT lengths (1k, 4k, and 16k). Each frame is labelled with the mean and rms ratio. As expected the mean values are unity or close to it and the rms values decrease in going from the top-left case to the bottom-right case because the number of implied scintles contributing to a given point increases.

Method 1 for Calculating δTOA in the Mean-Shift Regime

A basis for TOA correction: Figure 16 shows histograms of the ratio of the mean PB time calculated using the HT solution to the true PB time on a realization by realization basis⁵. The true PB time for a given realization varies statistically from the ensemble average τ_d because the true PBF is calculated as an exponential envelope that multiplies complex Gaussian noise. The centroids of the histograms are near a ratio $\bar{t}_{\text{HT}}/\bar{t} \approx 0.7$ to 0.8 but there is a small variation vs the ratio τ_d/N_{FFT} . For a fixed value of this latter ratio, the histograms get progressively narrower with increasing N_{FFT} because there are more scintles (correlated intensity maxima) across the spectrum and, evidently, the ratio scales as the inverse square root of this number.

The number of scintles N_{scintles} is proportional to the total bandwidth $B = 1/\Delta t$ divided by the scintillation bandwidth: $N_{\text{scintles}} = \eta B/\Delta\nu_d = (2\pi\eta/C_1)B\tau_d$ for quantities in dimensional units (where the proportionality constant is $\eta \approx 0.2$ and C_1 depends on the shape of the ensemble-average PBF; for an exponential shape, $C_1 = 1$). For τ_d expressed in samples and $C_1 = 1$ this becomes $N_{\text{scintles}} = 2\pi\eta(\tau_d/N_{\text{FFT}})N_{\text{FFT}}$. For fixed τ_d/N_{FFT} , N_{scintles} scales linearly with the FFT length and we expect the histogram width to scale as $\sim N_{\text{scintles}}^{-1/2}$. Similarly for fixed FFT length, N_{scintles} scales linearly with τ_d . Thus the width of the histogram is largest in the top left of Figure 16 and decreases toward the right or downward.

The mean trend of the ratio of pulse-broadening times is plotted against τ_d/N_{FFT} in Figure 17, where each plotted point is based on 1000 realizations of the phase retrieval. For small values of τ_d/N_{FFT} the ratio is larger for smaller FFT lengths but the ratio asymptotes to the same value for $\tau_d/N_{\text{FFT}} \gtrsim 0.1$.

⁵The results shown use DFTs that are zero-padded to twice the length of the actual time series in order to avoid wrap-around effects. This allows pulse broadening times up to about 20% of the time series length to be included. The 1/e time of the field PBF is $2\tau_d$ if τ_d is the 1/e width for the intensity. If the field PBF has a long tail extending to $3 \times 2\tau_d$, then without zero-adding, avoiding wrap-around requires $6\tau_d < N_{\text{FFT}}/2$ or $\tau_d/N_{\text{FFT}} < 1/12$.

Method 1 Application

Algorithm: These results suggest that, at least for the stochastic-exponential PBF, that the mean Hilbert-transform time \bar{t}_{HT} can be scaled to the true mean time of the PBF. In the mean-shift regime, this implies that TOA correction is very simple:

1. Calculate the spectrum $S(\nu) = |\tilde{\varepsilon}(\nu)|^2$ averaged synchronously over as many pulse periods as allowed by the DISS time Δt_{ISS} . For $N_{b,\text{max}} \approx \Delta t_{\text{ISS}}/P \gg 1$, self noise in the pulsar signal will cause a fractional error in $S(\nu)$ that scales as $N_{b,\text{max}}^{-1/2}$. Additive noise will scale the same way but the net error depends on the signal-to-noise ratio.
2. Use the Hilbert transform and the spectrum to obtain the minimum-delay solution for the PBF and the resulting mean time, \bar{t}_{HT} .
3. Use the ratio $R = \bar{t}_{\text{HT}}/\bar{t}$ to scale the HT mean time to an estimate for the TOA correction, $\Delta\text{TOA} = R^{-1}\bar{t}_{\text{HT}}$. The ratio R depends on the data-block length and the pulse-broadening time as discussed above, so defining the ratio relies on additional information (such as an estimate of the scintillation bandwidth) or can be gotten iteratively.
4. Calculate the TOA from the pulse profile for the same set of data using standard methods.
5. Subtract ΔTOA from the nominal TOA.
6. The accuracy of the correction depends on the number of scintles in the band and, of course, on the number of blocks averaged and on the S/N of the pulsar signal.

Notes on Applying the Hilbert Transform

Phase retrieval using the HT on the magnitude of $\tilde{h}(\nu)$ typically underestimates the pulse broadening time, as shown above.

The underlying reason is that the true phase can undergo one or more wraps that are not captured by the HT solution.

Wraps can be added to the HT solution by applying one or more all-pass (AP) filters. A single-stage AP can add up to 2π of phase change across the full frequency band and most of the change can be concentrated in a narrow frequency range through suitable choice of the amplitude of the pole. If the amplitude of the pole is zero, the impulse response is simply a Kronecker delta yielding a delay of one time step.

Properties of All-pass Filters

A single-stage AP filter is parameterized by a single complex number $a = |a|e^{i\theta}$, which is the pole of $e^{ik/N}$. It produces up to 2π radians of phase rotation across the frequency band with the amount and bandwidth of largest phase gradient determined by $|a|$. The location of the largest phase gradient is determined by θ .

Figure 19 shows an AP filter that has a pole at $e^{i\omega} = 0$, i.e. $|a| = 0$, which yields a linear phase change of 2π with a negative slope, corresponding to a time delay of one sample. This filter simply shifts the input by one sample. Concatenating β stages corresponds to the single-stage filter being taken to the power β , which yields a total phase change of $2\pi\beta$ and a shift by β samples.

Figure 20 shows a family of filters for fixed $\theta = 180^\circ$ for different magnitudes of a .

Figure 21 shows a family of AP filters for fixed $|a| = 0.8$ for different θ .

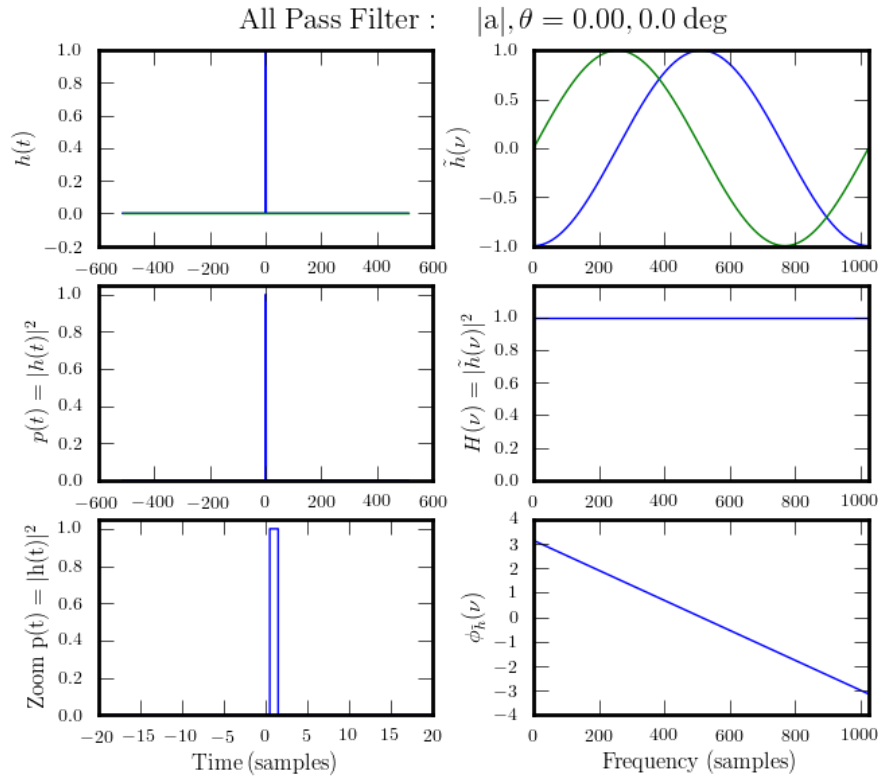


Fig. 19.— Plots showing the properties of an all-pass filter with a pole at zero ($a = 0$). The left-hand column of plots shows the field PBF $h(t)$ and the intensity PBF $p(t)$, including a zoomed version bottom left. The right-hand column shows frequency-domain quantities: FT of the field PBF (top), the squared magnitude of the FT (middle) that is by definition unity, and the phase, which is linear with negative slope and rotates over 2π .

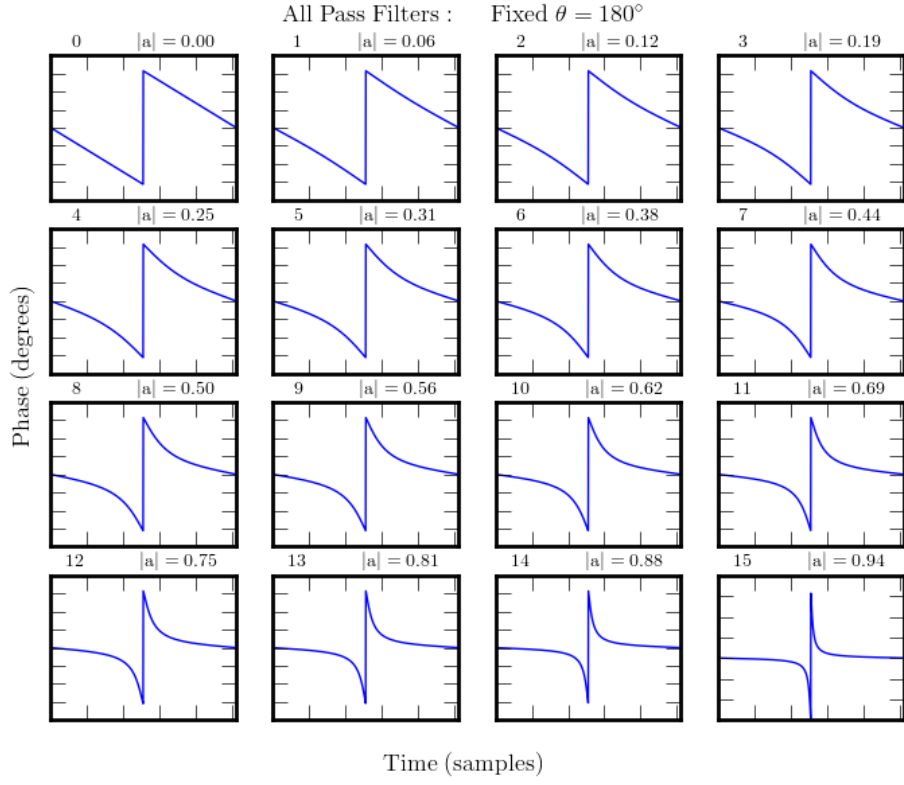


Fig. 20.— Plots showing a sequence of all-pass filters where the pole describing the filter at $a = |a|e^{i\theta}$ is varied in amplitude as the angle is fixed. The angle is related to the frequency index k by $\theta = 360^\circ k/N$ with $N = 1024$.




Fig. 21.— Plots showing a sequence of all-pass filters where the pole describing the filter is at $a = |a|e^{i\theta}$ where the magnitude is $|a| = 0.8$ held fixed and the angle is varied. The angle is related to the frequency index k by $\theta = 360^\circ k/N$ with $N = 1024$.




Fig. 22.— A sequence of all pass filters based on a single-stage filter $\tilde{h}_1(\nu)$ with $|a| = 0.8$ and $\theta = 0^\circ$. The transfer function is $\tilde{h}_1^\beta(\nu)$. The left-hand column shows impulse responses for different β values and the right-hand column shows the corresponding phase variations vs. frequency.

Idealized Cases for Determining $h(t)$ III.

Phase retrieval using the Hilbert transform augmented with all-pass filters: Figure 23 shows another example where the phase is retrieved using only the HT. The middle-left panel shows the true phase of $\tilde{h}(\nu)$ while the panel below it shows the HT phase. In this case the HT solution has less phase wrap than the true solution.

To identify *where* additional phase wraps are needed from AP filters we note the following:

1. As a rule of thumb, a mean time delay of τ_d samples corresponds to $2\pi\tau$ total phase wrap;
2. The overall phase trend in the ACF $\Gamma_{\tilde{h}}$ has negative slope, consistent with the causality of the PBF;
3. Extra phase wraps in the true PBF occur rapidly with changes in frequency;
4. The extra, fast phase wraps also tend to have negative slope (ignoring phase wraps associated with the inverse-tangent function);
5. Phase jumps also occur in the HT solution, $\tilde{h}_{HT}(\nu)$ at the same locations as fast jumps in the true transfer function. However these are *positive* going jumps as compared to the negative-going jumps in \tilde{h} .

These results suggests that a comparison of neighboring frequencies can indicate the locations of fast phase wraps. This is where the 2D frequency ACF comes into play,

$$\Gamma_{\tilde{h}}(\nu, \delta\nu) = \tilde{h}(\nu)\tilde{h}^*(\nu + \delta\nu).$$

For a fixed value of $\delta\nu$, the phase of the ACF is the difference of the phase of \tilde{h} . The locations of fast phase wraps appear as *positive* impulses in the phase of $\Gamma_{\tilde{h}}(\nu, \delta\nu)$. These can be seen in the bottom-right panel of Figure 23 (blue curve). Phase pulses are also seen in the ACF of the HT solution, $\Gamma_{\tilde{h},HT}$ (the green curve in the bottom right panel); however, these are *negative* impulses. At other frequencies the two ACFs ($\Gamma_{\tilde{h}}$ and $\Gamma_{\tilde{h},HT}$) are identical; Therefore the phase of their product, $\Gamma_{\tilde{h}}(\nu)\Gamma_{\tilde{h},HT}^*(\nu)$, is expected to be positive near a phase jump and asymptotic to zero elsewhere in frequency. We use the locations of pulses in the product ACF to define one of the parameters of a corresponding all-pass filter.

In addition to fast phase wraps, simulations show that a uniform phase rotation of an integer multiple β of -2π is sufficient to retrieve the phase in some cases. These rotations can be implemented by using AP filters with coefficients $a = 0$. In z -transform language, these AP filters are simply $z^{-\beta} = e^{-2\pi i\nu\beta}$, the same phase factor expected from application of the shift theorem. Below, we refer to these as “zero-pole” (ZP) filters.

Idealized Cases for Determining $h(t)$ IV.

Algorithm: For the idealized case where the data are equal to $\tilde{h}(\nu)$ (i.e. the pulsar signal is a delta function and there is no additive noise), but we pretend that we can only obtain the spectrum $H(\nu) = |\tilde{h}(\nu)|^2$ and the 2D ACF $\Gamma_{\tilde{h}}(\nu, \delta\nu)$ (both intensity-like quantities), the overall algorithm for phase retrieval is:

1. Calculate $H(\nu) = |\tilde{h}(\nu)|^2$ and $\Gamma_{\tilde{h}}(\nu, \delta\nu)$.
2. Use the Hilbert transform to calculate the phase $\phi_{\text{HT}}(\nu) = \text{HT}\{\frac{1}{2} \ln H(\nu)\}$ and the resulting transfer-function estimate $\tilde{h}_{\text{HT}}(\nu) = H^{1/2}(\nu)e^{i\phi_{\text{HT}}(\nu)}$.
3. Calculate the cost (i.e. χ^2) by comparing ACFs of \tilde{h} and \tilde{h}_{HT} :

$$C = \sum_{\nu} |\Gamma_{\tilde{h}}(\nu, \delta\nu) - \Gamma_{\tilde{h}_{\text{HT}}}(\nu, \delta\nu)|^2.$$

4. Find the zero-pole (ZP) filter that produces the lowest cost. I.e. For a new transfer function $\text{ZP}_{\beta}(\nu)\tilde{h}_{\text{HT}}(\nu)$ with

$$\text{ZP}_{\beta}(\nu) = e^{-2\pi i \beta k/N},$$

find the integer β that minimizes the cost. Note that $\beta = 0$ is possible.

5. Find the locations $\nu_j, j = 1, \dots, N_j$ of phase jumps in the product $\Gamma_{\tilde{h}}(\nu, \delta\nu)\Gamma_{\tilde{h}, \text{HT}}^*(\nu, \delta\nu)$. Here ν_j is in units of frequency bins. So far, jumps are found by finding impulses in $\Gamma_{\tilde{h}}(\nu, \delta\nu)$ that are above a threshold $\phi_{\min} = 0.5$ rad. Then averages are made over the points above threshold that correspond to the same jump.
6. Step through the phase-jump list and calculate the angular part of the complex pole: $\theta_j = 2\pi\nu_j/N$. Then find the magnitude a_j for each jump by minimizing the cost C using the new transfer function $\text{AP}(a_j, \theta_j, \nu)\text{ZP}_{\beta}(\nu)\tilde{h}_{\text{HT}}(\nu)$. Note the jumps can be fitted individually because they do not overlap and because the corresponding filters only change the phase.
7. Calculate the complete phase-retrieved PBF as

$$\tilde{h}_r(\nu) = \text{ZP}_{\beta}(\nu) \times \tilde{h}_{\text{HT}}(\nu) \times \prod_{j=0}^{N_j-1} \text{AP}(a_j, \theta_j, \nu).$$

8. Calculate the impulse response $h_r(t)$ as the inverse FT of $\tilde{h}_r(\nu)$.

Idealized Cases for Determining $h(t)$ V.

Figure 24 shows the results from applying this algorithm to the simulated case. The left-hand panels show (from top to bottom) the true phase of the PBF (not accessible to actual measurement), the phase of the ACF of the PBF (which is accessible to measurement) with the locations of phase jumps identified with red and black symbols, the retrieved phase using the HT, the phase of from the best fit ZP filter (in this case a -2π phase wrap) , the phase of the best-fit AP filters (two filters added together for this case, and the final total retrieved phase.

The right-hand panels show field PBFs in the top three panels (true, HT, and total retrieved PBF) and the corresponding intensity PBFs in the bottom three panels.

For the particular example, the mean pulse-broadening times using Equation 5 are

$$\begin{aligned}\bar{t} &= 2.98 \text{ samples} && \text{true PBF} \\ \bar{t} &= 2.12 \text{ samples} && \text{PBF} = \text{HT solution} \\ \bar{t} &= 3.12 \text{ samples} && \text{PBF} = \text{HT} * \text{ZP} \\ \bar{t} &= 3.21 \text{ samples} && \text{PBF} = \text{HT} * \text{ZP} * \text{AP}.\end{aligned}$$

For this example, the HT solution underestimates the mean PB time while combining it with the ZP solution slightly overestimates the PB time and with the AP solution even more so.

Other simulated examples show that the complete solution scatters around the true PB time with up to 5 to 10% error.

So far only a coarse grid has been used to find the best amplitudes a_j for the AP filters and the thresholding for identifying the jump list has not been optimized, so the error finding the PBF may be reduced.

Questions: One question is whether the phase retrieval *should* be perfect, i.e. can a unique solution be found that is the true solution? For the idealized case where the data are uncontaminated by either self-noise in the pulsar or by additive noise, it should be possible to match the measurable quantities $H(\nu)$ and $\Gamma_{\tilde{h}}(\nu, \delta\nu)$ perfectly using a sufficient number of parameters. However, our algorithm attempts to model the PBF with a minimal number of parameters, so perfect agreement is not expected.

In realistic situations, self-noise and additive noise are relevant and we cannot expect uniqueness. The most important question, however, is how well can arrival times be corrected?

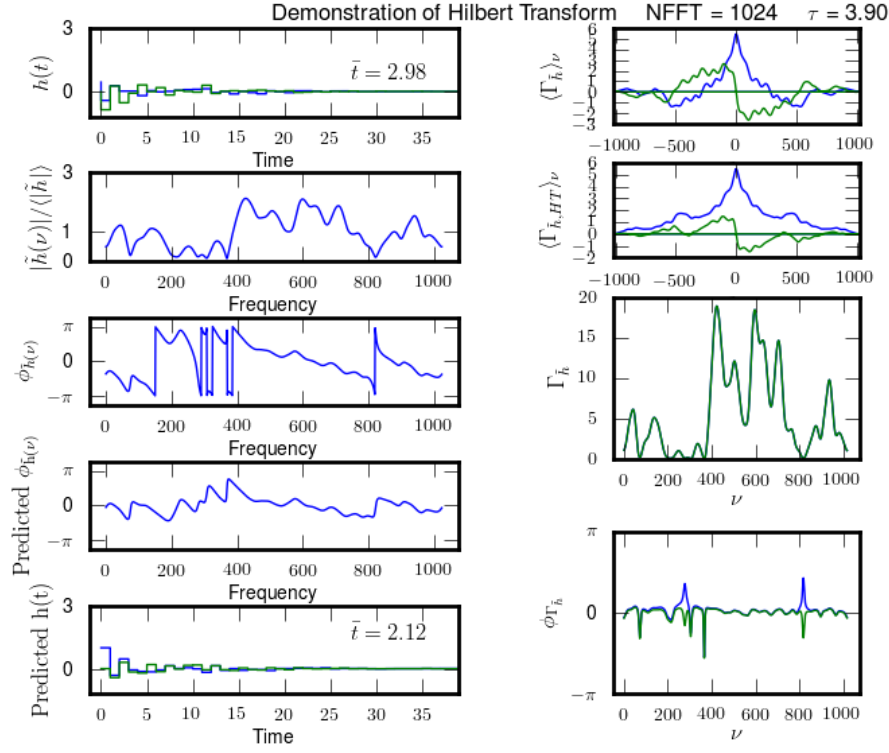


Fig. 23.— Demonstration of full phase retrieval of the PBF. From top to bottom the left-hand panels show the true PBF $h(t)$, the normalized magnitude of the FT $\tilde{h}(\nu)$, the phase of $\tilde{h}(\nu)$, the phase calculated using the HT of the FT magnitude, and the PBF calculated from the HT phase. On the right side, the panels are the ACF of \tilde{h} integrated over frequency, the ACF of the HT solution for \tilde{h} , the magnitude of the ACF of \tilde{h} calculated for a fixed lag and plotted against ν , and the phase of the ACF.

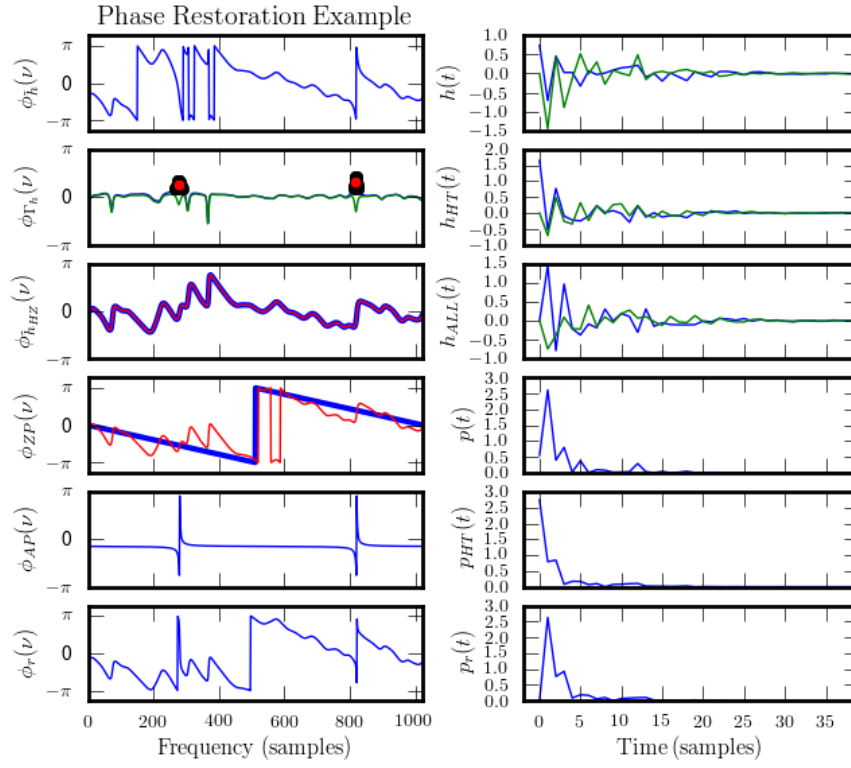


Fig. 24.— Demonstration of full phase retrieval of the PBF (continued). From top to bottom, the left-hand panels show phase vs frequency for: the true PBF, the ACF of the PBF for a fixed lag, the HT phase, the phase for the best-fit zero-pole filter, the phase for the best all-pass filters, and the total retrieved phase $\phi_r(\nu)$. The right hand side shows PBFs: the true field PBF $h(t)$, the HT solution, and the full solution; and intensity PBFs for the same cases.

Realistic Cases for Determining $h(t)$

The real task is to determine $h(t)$ for an actual pulsar signal that has the properties of AMN and where there is additive noise. The algorithm is the same as before, except the data products used for the phase retrieval are the intensity spectrum $S(\nu)$ and the field ACF $\Gamma_{\tilde{\epsilon}}(\nu)$ instead of $H(\nu)$ and $\Gamma_{\tilde{h}}(\nu)$. A new quantity comes into play, the intrinsic intensity pulse shape $A(t) \iff \tilde{A}(\nu)$ that must be used in the cost function.

First things first: before attempting a full retrieval we see how well, with sufficient averaging that the HT analysis applied to the full pulsar signal (including additive noise) does compared to analyzing the pure $h(t)$ signal.

Figures 25 - 32 show examples of the phase retrieval on the full pulsar signal. The specified signal to noise ratio is for a single pulse and is the ratio of the pulse peak amplitude to the *additive* noise. Self noise from the pulsar implies a separate signal-to-noise ratio, $S/N_{\text{self}} = 1/\sqrt{N_r}$ where N_r is the number of realizations averaged. When both S/N values are large, HT retrieval applied to the full pulsar yields the same result as HT retrieval of the pure PBF, to within the errors implied by the S/N.

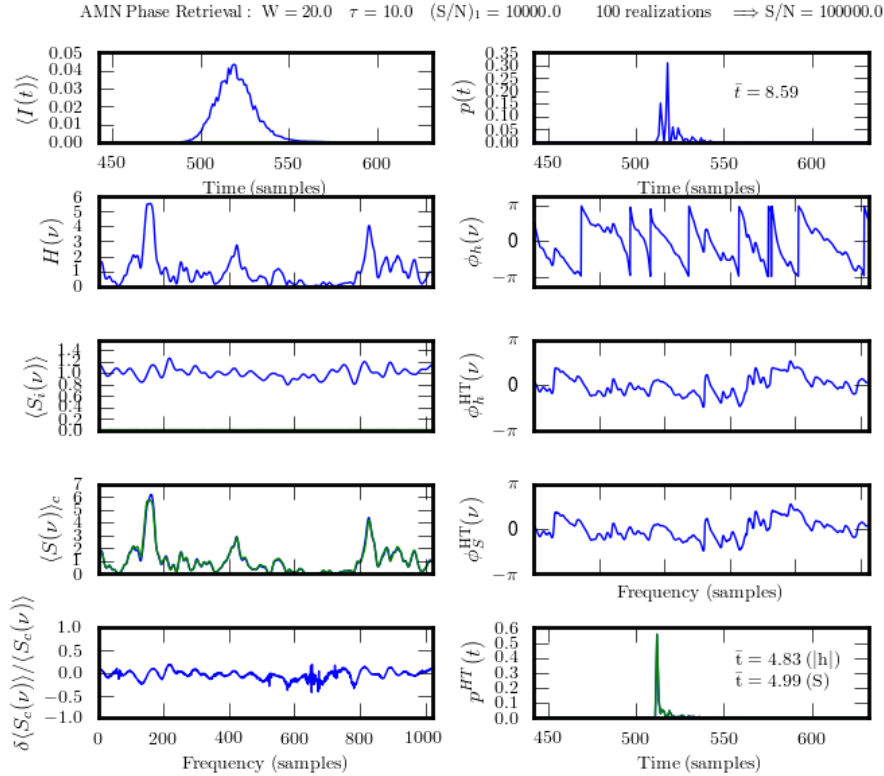


Fig. 25.— Demonstration of HT phase retrieval applied to the full pulsar signal. This case has $W = 20$, $\tau = 10$, net $S/N = 10^5$ (after averaging) and averages 100 realizations.

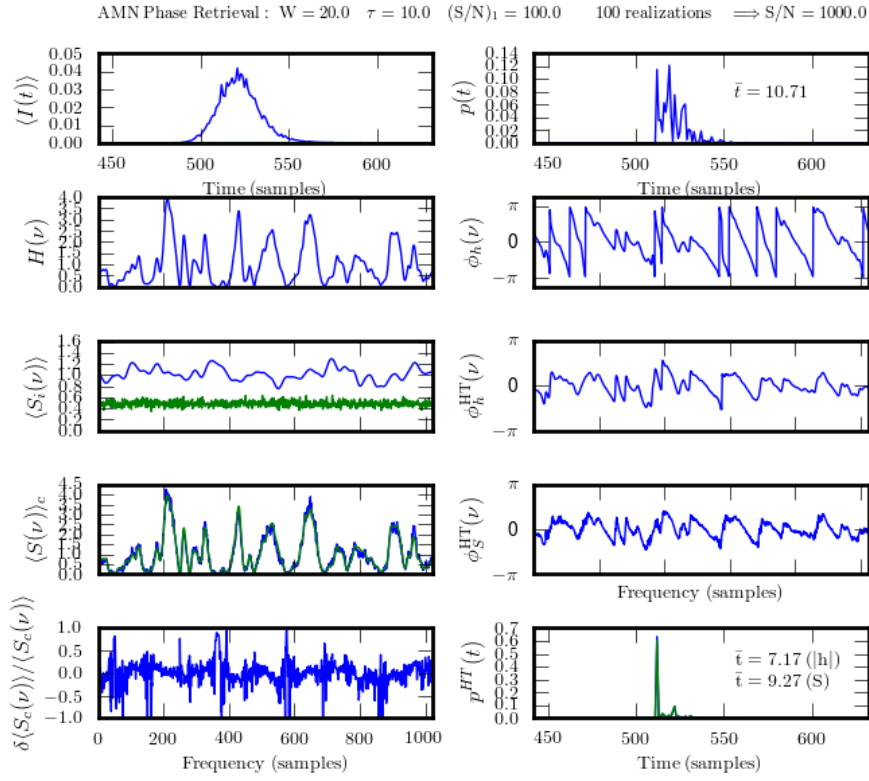


Fig. 26.— Demonstration of HT phase retrieval applied to the full pulsar signal. This case has $W = 20$, $\tau = 10$, $S/N = 10^3$ and averages 100 realizations.




Fig. 27.— Demonstration of HT phase retrieval applied to the full pulsar signal. This case has $W = 20$, $\tau = 10$, $S/N = 3.2 \times 10^3$ and averages 1000 realizations.

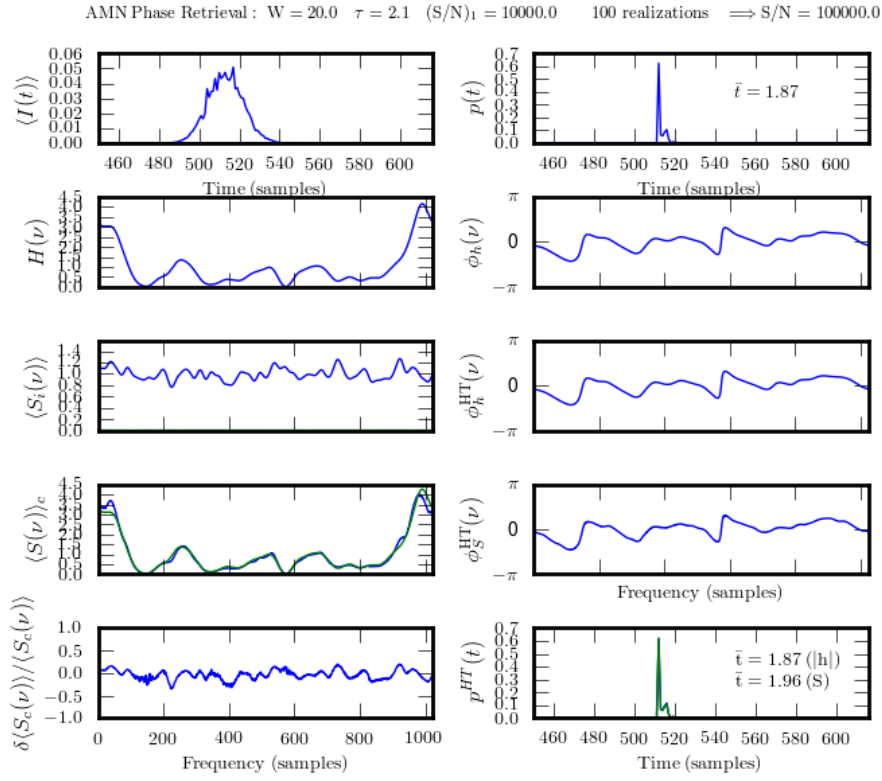


Fig. 28.— Demonstration of HT phase retrieval applied to the full pulsar signal. This case has $W = 20$, $\tau = 2.1$, $S/N = 10^5$ and averages 100 realizations.

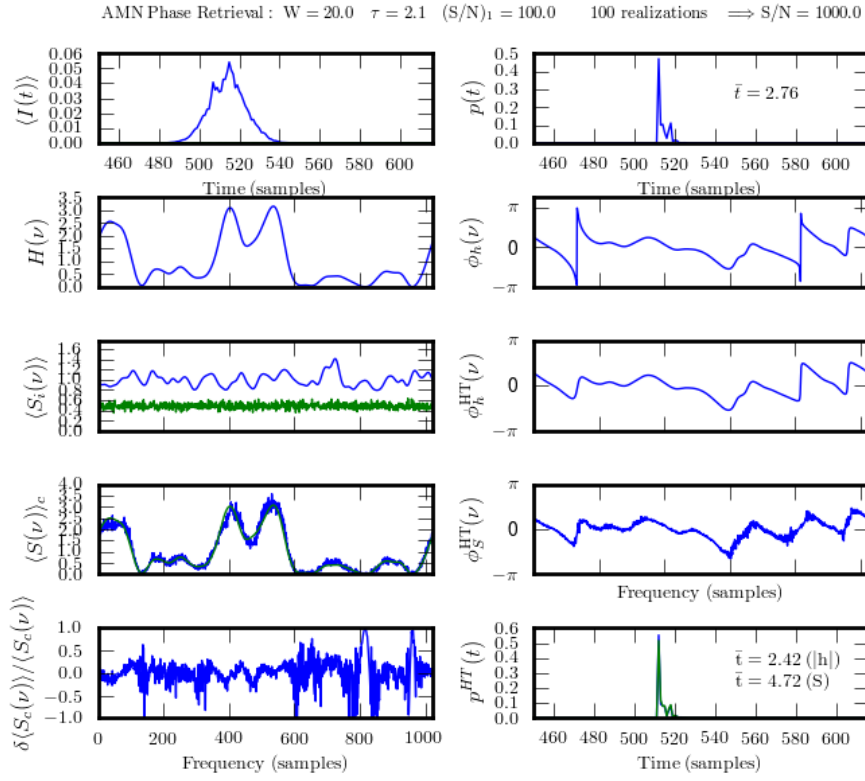


Fig. 29.— Demonstration of HT phase retrieval applied to the full pulsar signal. This case has $W = 20$, $\tau = 2.1$, $S/N = 10^3$ and averages 100 realizations.

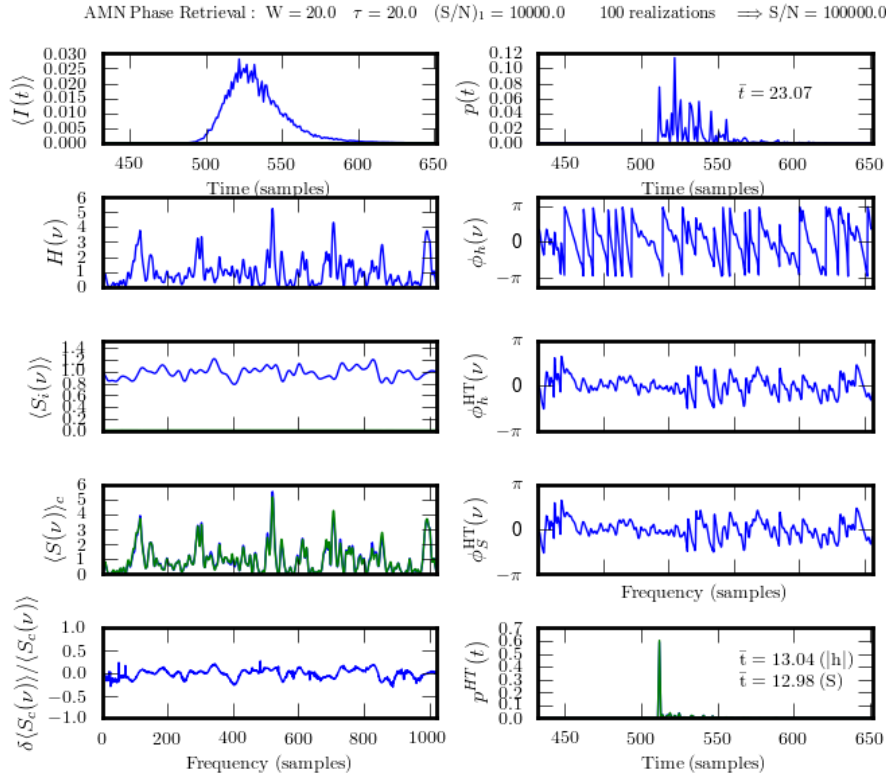


Fig. 30.— Demonstration of HT phase retrieval applied to the full pulsar signal. This case has $W = 20$, $\tau = 20$, $S/N = 10^5$ and averages 100 realizations.




Fig. 31.— Demonstration of HT phase retrieval applied to the full pulsar signal. This case has $W = 20$, $\tau = 40$, $S/N = 10^5$ and averages 100 realizations.

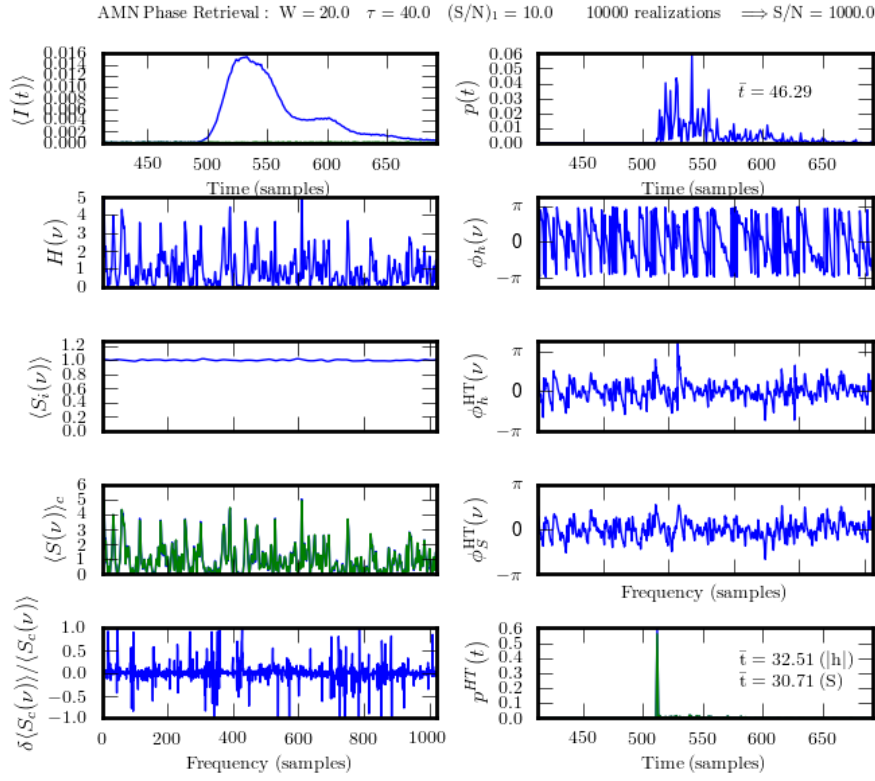


Fig. 32.— Demonstration of HT phase retrieval applied to the full pulsar signal. This case has $W = 20$, $\tau = 40$, $S/N = 10^3$ and averages 10^4 realizations.

Toy Model for Obtaining the PBF

Suppose the pulsar signal is intrinsically a delta function. Observations then provide a direct measure of $h(t)$, or its moments. If combined with additive noise, we need to integrate an intensity-like quantity, such as the frequency-domain ACF,

$$\Gamma_{\tilde{h}}(\nu, \delta\nu) = \tilde{h}(\nu)\tilde{h}^*(\nu - \delta\nu),$$

where we keep the dependence on ν explicit. For the case where the FT has a length of one period, the sample interval is $\Delta\nu = 1/P$ and the lags of the ACF are $\delta\nu = k/P$.

We write the ACF in terms of the magnitude $\tilde{h}_0(\nu_j)$ and phase $\phi(\nu_j)$ of \tilde{h} for frequency samples ν_j , $j = 0, \dots, N_\nu - 1$ and lags $\delta\nu_k$, $k = 0, \dots, N_{\delta\nu}$,

$$\Gamma_{\tilde{h}}(\nu_j, \delta\nu_k) = \tilde{h}_0(\nu_j)\tilde{h}_0(\nu_j - \delta\nu_k)e^{i[\phi(\nu_j) - \phi(\nu_j - \delta\nu_k)]}.$$

We might have a situation where $N_\nu = 1024$ and $N_{\delta\nu} \approx 20$ (the number of harmonics that describe the pulse shape). This yields a large number of equations from which the $2N_\nu$ unknowns can be solved.

Consider the case where scintles are well resolved by the frequency sampling, i.e. $\Delta\nu \ll \Delta\nu_d$ corresponding to $P \gg 2\pi\tau_d$. If we consider only the first lag of the ACF, $\delta\nu_1 = 1/P$, we expect the amplitude to be nearly constant and the phase will vary slowly. For other lags this may also hold. Assuming the spectrum $H(\nu) = |\tilde{h}(\nu)|^2$ is available, we have

$$\tilde{h}_0(\nu_j)\tilde{h}_0(\nu_j - \delta\nu_k) = [H(\nu_j)H(\nu_j - \nu_k)]^{1/2}.$$

We can start with the lowest frequency channel, ν_0 , and assign it zero phase: $\phi_{\nu_0} = 0$. We can define a new quantity

$$G_{jk} = G(\nu_j, \delta\nu_k) = \frac{\Gamma_{\tilde{h}}(\nu_j, \delta\nu_k)}{[H(\nu_j)H(\nu_j - \nu_k)]^{1/2}}.$$

Using this, we can show that

$$\begin{aligned} e^{i\phi(\nu_1)} &= G_{11} \\ e^{i\phi(\nu_2)} &= G_{21}G_{11} \\ &\vdots \\ e^{i\phi(\nu_j)} &= \prod_{\ell=1}^j G_{\ell 1}. \end{aligned}$$

Thus, in principle, the individual phases of \tilde{h} can be determined. The advantage of this approach is that no phase is calculated for $\Gamma_{\tilde{h}}$ and so phase-wraps are avoided. The method depends on any real errors in the ACF being small. Since the solution is from smallest to largest frequency, errors will accrue and the phase curve will random walk away from the true phase.

A possible remedy is to also work from largest to smallest frequency and to combine the resulting phase function with that obtained by working from smallest to largest.

Errors can be reduced by using other lags $\delta\nu_k$ of the ACF. These too must satisfy $k/P \ll \Delta\nu_d$ or $P \gg 2\pi k\tau_d$.

Inverse Problems Using the Hilbert Transform

Generally a function cannot be determined uniquely from its autocorrelation function (e.g. Scargle 1981, ApJS, 45, 1). For example, the time reverse of a function has the same ACF and so does any delayed version of the function.

If a function is causal, there is a fundamental constraint on its Fourier transform (FT): the real and imaginary parts are related by a Hilbert transform (HT). In physics, causality appears in the Kramers-Kronig relations for dispersion relations in dielectrics.

For a causal, minimum-delay function, the HT can be used to calculate the FT from the magnitude of the FT. In particular, if $|\tilde{h}|$ is known the phase of \tilde{h} can be calculated as

$$\phi_{\tilde{h}}(\nu) = \frac{1}{\pi} \mathcal{P} \int_{-\infty}^{\infty} d\nu' \frac{\ln |\tilde{h}(\nu)|}{(\nu - \nu')},$$

where \mathcal{P} denotes Cauchy principal component.

Using the phase, the complete FT can be calculated.

For a function that does not have the minimum-delay property, the minimum-delay phase still has a role. Let $\tilde{h}_{\text{HT}}(\nu)$ be the FT obtained using the HT to calculate the minimum-delay solution. It can be shown (e.g. Oppenheim & Shafer 1989) that the true FT can be written as

$$\tilde{h}(\nu) = \tilde{h}_{\text{HT}}(\nu) Z(\nu),$$

where $Z(\nu)$ is an “all-pass” filter with unit amplitude that comprises some number of zeros of $\exp(-2\pi i\nu)$ in the complex plane and that changes only the phase.

The all-pass filter can be written as the product of Q individual factors involving the zeros s_j^{-1} where $\exp(2\pi i\nu) = 0$,

$$Z(\nu) = \prod_{j=0}^{Q-1} \left(\frac{e^{-2\pi i\nu} - s_j}{1 - s_j^* e^{-2\pi i\nu}} \right),$$

and where each factor and thus Z has unity magnitude. The numerical task is to determine the values for s_j ($2Q$ unknowns) and the value of Q itself. Determination is aided by the fact that the zeros s_j^{-1} are outside the unit circle and by the possibility that Q is much smaller than the overall length of a time series.

Figure of merit: least squares + positivity constraint on the deconvolution. Could use entropy to do that.

Hilbert Transform Example




Fig. 33.— Example phase retrieval for a one-sided exponential that initiates at the time origin, $h(t) = e^{-t/\tau}U(t)$ where U is the unit step function.

REFERENCES

- Bhat, N. D. R., Cordes, J. M., Camilo, F., Nice, D. J., & Lorimer, D. R. 2004, *ApJ*, 605, 759
- Klibanov, M. V., Sacks, P. E., Tikhonravov, A. V. 1995, *Inverse Problems*, 11, 1
- Kuz'min, A. D., & Izvekova, V. A. 1993, *MNRAS*, 260, 724
- Oppenheim, A. V. & Shafer, R. W. 1974, *Digital Signal Processing*, Prentice-Hall, Englewood Cliffs, NJ.
- Anholm, M., Ballmer, S., Creighton, J. D. E., Price, L. R., & Siemens, X. 2009, *Phys. Rev. D*, 79, 084030
- Armstrong, J. W. 1984, *Nature*, 307, 527
- Armstrong, J. W., Rickett, B. J., & Spangler, S. R. 1995, *ApJ*, 443, 209
- Ahuja, A. L., Mitra, D., & Gupta, Y. 2007, *MNRAS*, 377, 677
- Backer, D. C. et al. 1982, *Nature*, 300, 615
- Backer, D. C., & Hellings, R. W. 1986, *ARA&A*, 24, 537
- Backer, D. C. et al. 1993, *ApJ*, 404, 636
- Bailes, M. 2007, *arXiv:astro-ph/0702698*
- Barnard, J. J., & Arons, J. 1986, *ApJ*, 302, 138
- Bertotti, B., Carr, B. J., & Rees, M. J. 1983, *MNRAS*, 203, 945
- Bhat, N. D. R., Cordes, J. M., & Chatterjee, S. 2003, *ApJ*, 584, 782
- Bhat, N. D. R. et al. 2004, *ApJ*, 605, 759
- Briskin, W. F. et al. 2010, *ApJ*, 708, 232
- Britton, M. C. 2000, *ApJ*, 532, 1240
- Cognard, I., & Lestrade, J.-F. 1997, *A&A*, 323, 211
- Coles, W. A., Rickett, B. J., Gao, J. J., Hobbs, G., & Verbiest, J. P. W. 2010, *ApJ*, 717, 1206
- Cordes, J. M. 1993, *Planets Around Pulsars*, 36, 43
- Cordes, J. M., Weisberg, J. M., & Boriakoff, V. 1985, *ApJ*, 288, 221
- Cordes, J. M. & Downs, G. S. 1985, *ApJS*, 59, 343
- Cordes, J. M., & Stinebring, D. R. 1984, *ApJ*, 277, L53

- Cordes, J. M., Pidwerbetsky, A., & Lovelace, R. V. E. 1986, *ApJ*, 310, 737
- Cordes, J. M. et al. 1990, *ApJ*, 349, 245
- Cordes, J. M., & Lazio, T. J. 1991, *ApJ*, 376, 123
- Cordes, J. M. & Rickett, B. J. *ApJ*, 507, 846
- Cordes, J. M. & Lazio, T.J.W. 2001, *ApJ*, 549, 997
- Cordes, J. M., & Lazio, T. J. W. 2002, *arXiv:astro-ph/0207156*
- Cordes, J. M. et al. 2004, *ApJ*, 612, 375
- Cordes, J. M. et al. 2004, *New Astronomy Review*, 48, 1413
- Cordes, J. M. et al. 2006, *ApJ*, 637, 346
- Cortés Medellín, G. 2007, SKA Memo 95, “Antenna Noise Temperature Calculation,”
<http://www.skatelescope.org/>
- Craft, H. D., Jr. 1970, Ph.D. Thesis
- Damour, T., & Vilenkin, A. 2005, *Phys. Rev. D*, 71, 063510
- Deneva, J. S., Cordes, J. M., & Lazio, T. J. W. 2009, *ApJ*, 702, L177
- D’Alessandro, F., McCulloch, P. M., King, E. A., Hamilton, P. A., & McConnell, D. 1993, *MNRAS*, 261, 883
- Demorest, P. B. 2007, Ph.D. Thesis, U.C. Berkeley.
- Demorest, P., et al. 2009, *astro2010: The Astronomy and Astrophysics Decadal Survey*, 2010, 64
- Desai, K. M., & Fey, A. L. 2001, *ApJS*, 133, 395
- Detweiler, S. 1979, *ApJ*, 234, 1100
- Downs, G. S., & Reichley, P. E. 1983, *ApJS*, 53, 169
- Edwards, R. T. & Stappers, B. W. 2003, *å*, 407, 273
- Edwards, R. T., Hobbs, G. B., & Manchester, R. N. 2006, *MNRAS*, 372, 1549
- Finn, L. S., & Lommen, A. N. 2010, *arXiv:1004.3499*
- Foster, R. S., & Backer, D. C. 1990, *ApJ*, 361, 300
- Foster, R. S., & Cordes, J. M. 1990, *ApJ*, 364, 123
- Freire, P. C. et al. 2009, *arXiv:0902.2891*

- Haslam, C. G. T. et al. 1982, A&AS, 47, 1
- Helfand, D. J., Manchester, R. N., & Taylor, J. H. 1975, ApJ, 198, 661
- Hellings, R. W., & Downs, G. S. 1983, ApJ, 265, L39
- Hemberger, D. A., & Stinebring, D. R. 2008, ApJ, 674, L37
- Hill, A. S. et al. 2003, ApJ, 599, 457
- Hill, A. S. et al. 2005, ApJ, 619, L171
- Hobbs, G. B., et al. 2009, Publications of the Astronomical Society of Australia, 26, 103
- Hobbs, G., et al. 2010, Classical and Quantum Gravity, 27, 084013
- Hotan, A. W., Bailes, M., & Ord, S. M. 2006, MNRAS, 369, 1502
- Hu, W., Stinebring, D. R., & Romani, R. W. 1991, ApJ, 366, L33
- Jacoby, B. A. 2005, ApJ, 629, L113
- Jaffe, A. H., & Backer, D. C. 2003, ApJ, 583, 616
- Jenet, F. A. et al. 1998, ApJ, 498, 365
- Jenet, F. A., Anderson, S. B., & Prince, T. A. 2001, ApJ, 546, 394
- Jenet, F. A., & Gil, J. 2004, ApJ, 602, L89
- Jenet, F. A. et al. 2004, ApJ, 606, 799
- Jenet, F. A. et al. 2005, Binary Radio Pulsars, 328, 399
- Jenet, F. A. et al. 2005, ApJ, 625, L123
- Jenet, F. A. et al. 2006, ApJ, 653, 1571
- Jenet, F. A. et al. 2009, arXiv:0909.1058
- Kaspi, V. M., Taylor, J. H., & Ryba, M. F. 1994, ApJ, 428, 713
- Kinkhabwala, A., & Thorsett, S. E. 2000, ApJ, 535, 365
- Kramer, M. et al. 2004, New Astronomy Review, 48, 993
- Kramer, M., et al. 2006, Science, 314, 97
- Kuz'min, A. D., & Izvekova, V. A. 1993, MNRAS, 260, 724
- Lambert, H. C. & Rickett, B. J. 1999, ApJ, 517, 299

- Lambert, H. C. & Rickett, B. J. 2000, *ApJ*, 531, 883
- Lazio, T. J. W. 2004, *ApJ*, 613, 1023
- Lee, K. J., Jenet, F. A., & Price, R. H. 2008, *ApJ*, 685, 1304
- Löhmer, O. et al. 2001, *ApJ*, 562, L157
- Löhmer, O. et al. 2004, *A&A*, 425, 569
- Lorimer, D. R. et al. 1995, *MNRAS*, 273, 411
- Lundgren, S. C. et al. 1995, *ApJ*, 453, 433
- Lyne, A. G. 2010, *Science*, 329, 408
- Manchester, R. N., Hobbs, G. B., Teoh, A. & Hobbs, M., 2005, *AJ*, 129, 1993
- McLaughlin, M. A. et al. 2002, *ApJ*, 564, 333
- Mitra, D., & Rankin, J. M. 2002, *ApJ*, 577, 322
- Narayan, R., & Goodman, J. 1989, *MNRAS*, 238, 963
- Phillips, J. A., & Wolszczan, A. 1991, *ApJ*, 382, L27
- Pshirkov, M. S., & Tuntsov, A. V. 2010, *Phys. Rev. D*, 81, 083519
- Ramachandran, R. et al. 2006, *ApJ*, 645, 303
- Rathnasree, N., & Rankin, J. M. 1995, *ApJ*, 452, 814
- Rickett, B. J. 1990, *ARA&A*, 28, 56
- Rickett, B. et al. 2009, *MNRAS*, 395, 1391
- Romani, R. W., & Taylor, J. H. 1983, *ApJ*, 265, L35
- Romani, R. W., Narayan, R., & Blandford, R. 1986, *MNRAS*, 220, 19
- Ruderman, M. & Sutherland, P. 1975, *ApJ*, 196, 51
- Sallmen, S. et al. 1999, *ApJ*, 517, 460
- Scargle, J. D. 1981, *ApJS*, 45, 1
- Sesana, A., & Vecchio, A. 2010, *Classical and Quantum Gravity*, 27, 084016
- Seto, N., & Cooray, A. 2007, *ApJ*, 659, L33
- Shannon, R. & Cordes, J. M. 2010, *ApJ*, submitted

- Spangler, S. R., & Gwinn, C. R. 1990, *ApJ*, 353, L29
- Stinebring, D. R. et al. 1990, *Physical Review Letters*, 65, 285
- Stinebring, D. R. et al. 2001, *ApJ*, 549, L97
- Tanenbaum, B. S., Zeissig, G. A., & Drake, F. D. 1968, *Science*, 160, 760
- Taylor, J. H., & Weisberg, J. M. 1982, *ApJ*, 253, 908
- van Haasteren, R., Levin, Y., McDonald, P., & Lu, T. 2009, *MNRAS*, 395, 1005
- van Straten, W. 2009, *ApJ*, 694, 1413
- Verbiest, J. P. W., et al. 2009, *MNRAS*, 400, 951
- Walker, M. A. et al. 2008, *MNRAS*, 388, 1214
- You, X. P., et al. 2007, *MNRAS*, 378, 493

A. Role of Pulse Jitter

Most if not all pulsars show phase variations of individual pulses within the on-pulse window defined by the average profile. Typically the rms phase variation is approximately equal to the width of a single pulse.

The amplitude modulation defined previously was strictly periodic. We can introduce jitter by including a term $\phi_j P$ to the argument of each pulse:

$$a(t) = \sum_{j=0}^{N_p} a_1(t - jP - \phi_j P).$$

The FT is now

$$\tilde{A}(\nu) = \tilde{A}_1(\nu) \sum_{j=0}^{N_p} e^{-2\pi i j \nu P} e^{-2\pi i \phi_j \nu P}$$

If ϕ_j is a Gaussian random variable with zero mean and variance σ_ϕ^2 , the ensemble average of the complex exponential is

$$\langle e^{-2\pi i \phi_j \nu P} \rangle = e^{-(2\pi \sigma_\phi \nu P)^2}.$$

The average FT of the periodic signal + jitter is then

$$\tilde{A}(\nu) = P^{-1} \left[\tilde{A}_1(\nu) e^{-(2\pi \sigma_\phi \nu P)^2} \right] \delta(\nu - k/P).$$

Interpretation: jitter is automatically included in the average pulse profile. For an ensemble average, the jitter factor is identically the characteristic function of the pulse-phase jitter distribution. (In reality, of course, there are also pulse amplitude variations that would be included with jitter in a 2D joint PDF.) For a 1D Gaussian PDF for the jitter, the characteristic function is another Gaussian, as above. For an FT from a finite-length data set, the contribution of jitter to the pulse shape will vary from realization to realization.

B. Notes on Signal Levels, etc.

Normalizations etc. are given for the FFT used in python. Unless stated, all expressions are for the no-smoothing case.

Statistics for single realizations:

White, Gaussian, complex noise:

$$n(t) \longrightarrow \tilde{n}(\nu) \longrightarrow N(\nu) = \langle |\tilde{n}(\nu)|^2 \rangle = \sigma_n^2 N_{\text{FFT}} \equiv N_0 N_{\text{FFT}}.$$

The mean power spectrum level is $N = N_0 N_{\text{FFT}}$ and the rms value is equal to the mean, $\sigma_N = N$.

Amplitude modulated noise:

$$\varepsilon_i(t) \equiv a(t)m(t) \longrightarrow \tilde{\varepsilon}_i(\nu) \longrightarrow S_i(\nu) = \langle |\tilde{\varepsilon}_i(\nu)|^2 \rangle = \sigma_m^2 \sum_j a^2(t_j) \equiv M_0 N_{\text{FFT}} \bar{A} = M_0 A_{\text{max}} W_A,$$

where $\bar{A} = N_{\text{FFT}}^{-1} \sum_t A(t)$ with $A(t) = a^2(t)$. We can write $\bar{A} \equiv A_{\text{max}} W_A / N_{\text{FFT}}$ where W_A is an equivalent width in units of samples.

The mean power spectrum is $S_i = M_0 A_{\text{max}} W_A$ and the rms value is equal to the mean, $\sigma_{S_i} = S_i$.

The net mean spectrum (excluding DISS) is $S(\nu) = S_i(\nu) + N(\nu)$ with a ratio of spectral levels

$$R_\nu = \frac{S_i}{N} = \frac{M_0 N_{\text{FFT}} \bar{A}}{N_0 N_{\text{FFT}}} = \frac{M_0 A_{\text{max}} W_A}{N_0 N_{\text{FFT}}}.$$

We define the signal-to-noise ratio of the signal ϵ_i to the additive noise n in the time domain as the ratio of pulse peak to rms noise in the intensity:

$$R_t = S/N \equiv \frac{A_{\text{max}}}{\sigma_n^2} = \frac{A_{\text{max}}}{N_0},$$

which gives

$$R_\nu = \frac{M_0 W_A R_t}{N_{\text{FFT}}} \longrightarrow \frac{W_A R_t}{N_{\text{FFT}}},$$

where the rightmost expression is for $M_0 = 1$, as used in the code.

The mean spectra of both the noise and the signal are flat and the rms values of each component are equal to the respective means.

Scaling: We can scale both spectral components so that the mean level of the signal part (S_i) is unity. The appropriate scale factor of $\varepsilon_i(t) + n(t)$ to achieve this is $s = 1/\sqrt{M_0 A_{\text{max}} W_A}$. To get a specified S/N in the time domain (i.e. R_t), we also multiply $n(t)$ by $1/\sqrt{R_t}$.

Thus to have $S_i = 1$ we have:

$$\begin{aligned} \varepsilon_i(t) &\longrightarrow \varepsilon_i(t) / \sqrt{M_0 A_{\text{max}} W_A}. \\ n(t) &\longrightarrow n(t) / \sqrt{M_0 A_{\text{max}} W_A R_t}. \end{aligned}$$

In the time domain this yields a pulse maximum

$$I_{\max} = \frac{1}{W_A}$$

and mean and rms off-pulse values

$$\begin{aligned}\langle N(t) \rangle &= \frac{1}{W_A R_t} \\ \sigma_N &= \frac{1}{W_A R_t \sqrt{N_r}},\end{aligned}$$

where N_r is the number of realizations averaged.

For the signal spectrum set to unit mean, the noise spectral level is

$$N(\nu) = \frac{N_0 N_{\text{FFT}}}{(M_0 A_{\max} W_A R_t)} \longrightarrow \frac{N_{\text{FFT}}}{W_A R_t},$$

where the rightmost expression is for $N_0 = M_0 = A_{\max} = 1$, as implemented in the code.

The rms values of each component are still equal to the respective means.

Averaging over N_r realizations:

When N_r realizations are averaged, intensity-like quantities in either domain have the same mean values but their rms values decrease by $N_r^{-1/2}$.

For the scaled signal (so that $S_i = 1$), we have

$$\begin{aligned}\sigma_{N(t)} &= \frac{N_0}{W_A R_t \sqrt{N_r}} \\ \sigma_{S_i} &= \frac{1}{\sqrt{N_r}} \\ \sigma_{N(\nu)} &= \frac{N_0 N_{\text{FFT}}}{W_A R_t \sqrt{N_r}} = \sigma_{N(t)} N_{\text{FFT}}. \\ R_t &= \frac{A_{\max}}{N_0 \sqrt{R_t}}.\end{aligned}$$

c. Pulse Broadening Function

We define the PBF in terms of moments of the wavefield. Let the scalar wavefield at the observer's location at frequency ν be $\varepsilon(\nu)$. In practice, a narrowband portion of the field centered on frequency ν_0 is processed, which we define as

$$\varepsilon_B(\nu) = b(\nu - \nu_0)\varepsilon(\nu), \quad (\text{C1})$$

where $b(\nu)$ is a low-pass function with bandwidth B . We define $b(\nu)$ to have unit area. We will need the Fourier transform of $b(t)$,

$$\tilde{b}(t) = \int d\nu b(\nu) e^{2\pi i \nu t}, \quad (\text{C2})$$

which is a pulse-like function with characteristic width $\sim B^{-1}$. An impulse from the pulsar propagating through a vacuum would yield a measured pulse of this shape.

An estimator (denoted with a caret) for the autocorrelation function (ACF) vs. frequency lag $\delta\nu$ of the narrowband field is

$$\hat{\Gamma}_\varepsilon(\delta\nu) = \int d\nu b(\nu - \nu_0) b^*(\nu + \delta\nu - \nu_0) \varepsilon(\nu) \varepsilon^*(\nu + \delta\nu). \quad (\text{C3})$$

Its ensemble average is

$$\langle \hat{\Gamma}_\varepsilon(\delta\nu) \rangle = \Gamma_\varepsilon(\delta\nu) R_b(\delta\nu), \quad (\text{C4})$$

where the ensemble average of the true ACF is

$$\Gamma_\varepsilon(\delta\nu) = \langle \varepsilon(\nu) \varepsilon^*(\nu + \delta\nu) \rangle. \quad (\text{C5})$$

and the ACF of the bandpass function is

$$R_b(\delta\nu) = \int d\nu b(\nu - \nu_0) b^*(\nu + \delta\nu - \nu_0). \quad (\text{C6})$$

An impulse $\delta(t)$ at the source produces a response — the PBF — at the observer's location given by the Fourier transform of the wavefield ACF

$$p_d(t) = \int d\delta\nu \hat{\Gamma}_\varepsilon(\delta\nu) e^{2\pi i \delta\nu t}. \quad (\text{C7})$$

The PBF is a real function because $\Gamma_\varepsilon(\delta\nu)$ is Hermitian and it has positive area, $\int dt p_d(t) = \hat{\Gamma}_\varepsilon(0) \geq 0$. The PBF is also a function of epoch because the geometry changes owing to motion of pulsar, screen and observer, so it should also be considered a function of epochal time.

The ensemble mean PBF is

$$\bar{p}_d(t) = \int d\delta\nu \langle \hat{\Gamma}_\varepsilon(\delta\nu) \rangle e^{2\pi i \delta\nu t}, \quad (\text{C8})$$

whose Fourier transform is the convolution of the true field correlation function Γ_ε and the instrumental function R_b .

We introduce time dependence in order to define the dynamic spectrum,

$$I(t, \nu) = |\varepsilon(t, \nu)|^2. \quad (\text{C9})$$

The two-dimensional ACF of the dynamic spectrum is

$$\Gamma_I(\delta t, \delta \nu) = \langle I(t, \nu) I(t + \delta t, \nu + \delta \nu) \rangle = \langle |\varepsilon(t, \nu)|^2 |\varepsilon(t + \delta t, \nu + \delta \nu)|^2 \rangle = \langle I \rangle^2 + |\Gamma_\varepsilon(\delta t, \delta \nu)|^2, \quad (\text{C10})$$

where the last equality follows for strong scattering when the scattered wavefield has Gaussian statistics (e.g. Rickett 1990). We are interested in the zero time-lag slice of the 2D ACF and we define the second term as the intensity autocovariance function (ACV),

$$\Gamma_{\delta I}(\delta \nu) = \langle \delta I(t, \nu) \delta I(t, \nu + \delta \nu) \rangle = |\Gamma_\varepsilon(\delta \nu)|^2 = \Gamma_I(0, \delta \nu). \quad (\text{C11})$$

The scintillation bandwidth, $\Delta \nu_d$, is defined as the HWHM of this function. This definition along with the choice of $1/e$ width of the PBF is motivated by the case where the PBF is a one-sided exponential, for which $2\pi\Delta\nu_d\tau_d = 1$. In practice, of course, we operate on the narrowband field and intensity, so the above integrals will include factors involving $b(\nu)$.

The secondary spectrum is the power spectrum of the dynamic spectrum, i.e. a fourth moment of the field:

$$S_2(f_t, f_\nu) = \left| \tilde{I}(f_t, f_\nu) \right|^2. \quad (\text{C12})$$

The intensity ACV defined above, $\Gamma_{\delta I}(\delta \nu)$, is the Fourier transform of the integrated secondary spectrum,

$$\Gamma_{\delta I}(\delta \nu) \Longleftrightarrow \int df_t S_2(f_t, f_\nu). \quad (\text{C13})$$

The integral of the secondary spectrum over f_t is the ACF of the pulse broadening function,

$$\Gamma_{p_d}(\tau) \equiv \int dt p_d(t) p_d(t + \tau) \propto \int df_t S_2(f_t, \tau) \propto \Gamma_{p_d}(\tau). \quad (\text{C14})$$

Formally f_ν is a variable conjugate to $\delta \nu$ that has units of time and often is referred to as a “delay.” However, it is really the time lag in the ACF of the PBF and is not directly related to the delay of a pulse imposed by ISS.

AMN Components, Example 2: Single Pulse

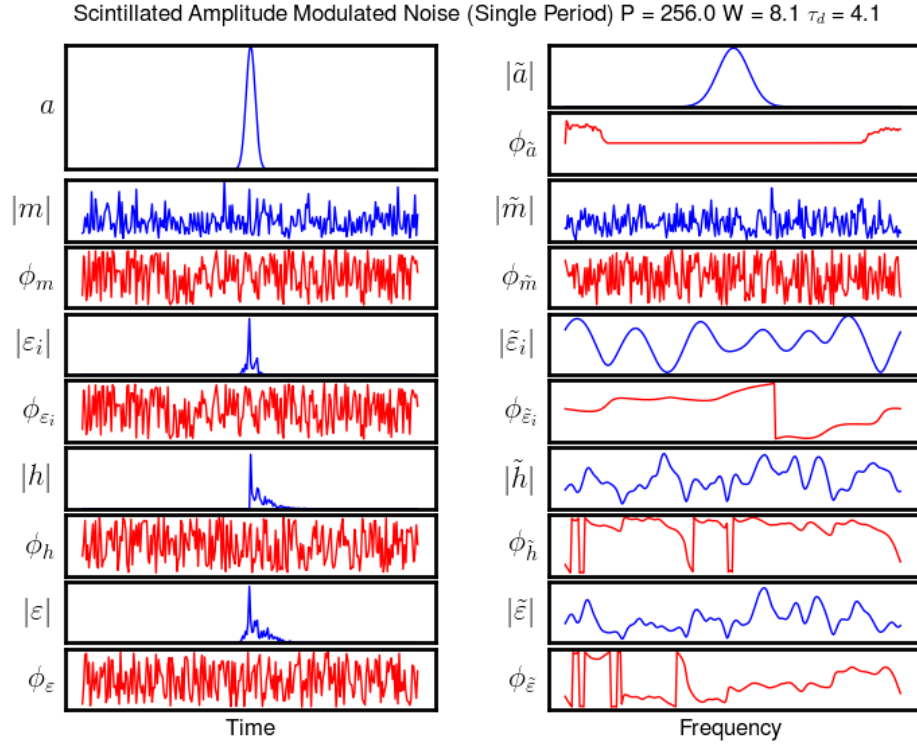


Fig. 34.— Components of scintillated amplitude modulated noise in the time and frequency domain. Amplitudes are plotted in blue, phases in red.

AMN Components, Example 3: Single Pulse

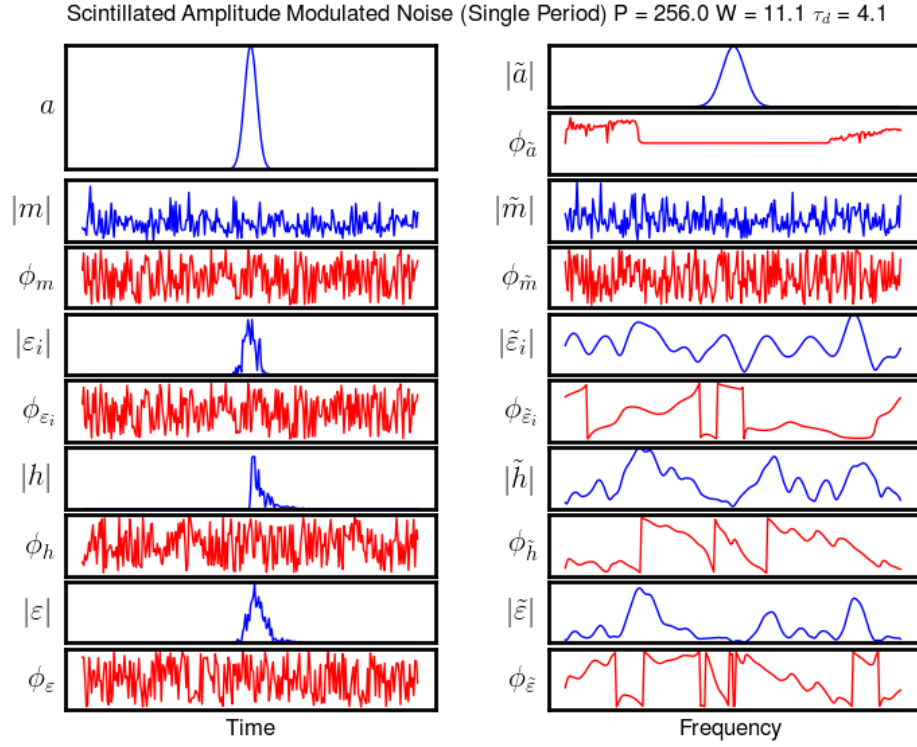


Fig. 35.— Components of scintillated amplitude modulated noise in the time and frequency domain for a single pulse. Amplitudes are plotted in blue, phases in red.

AMN Components, Example 4: Periodic Case

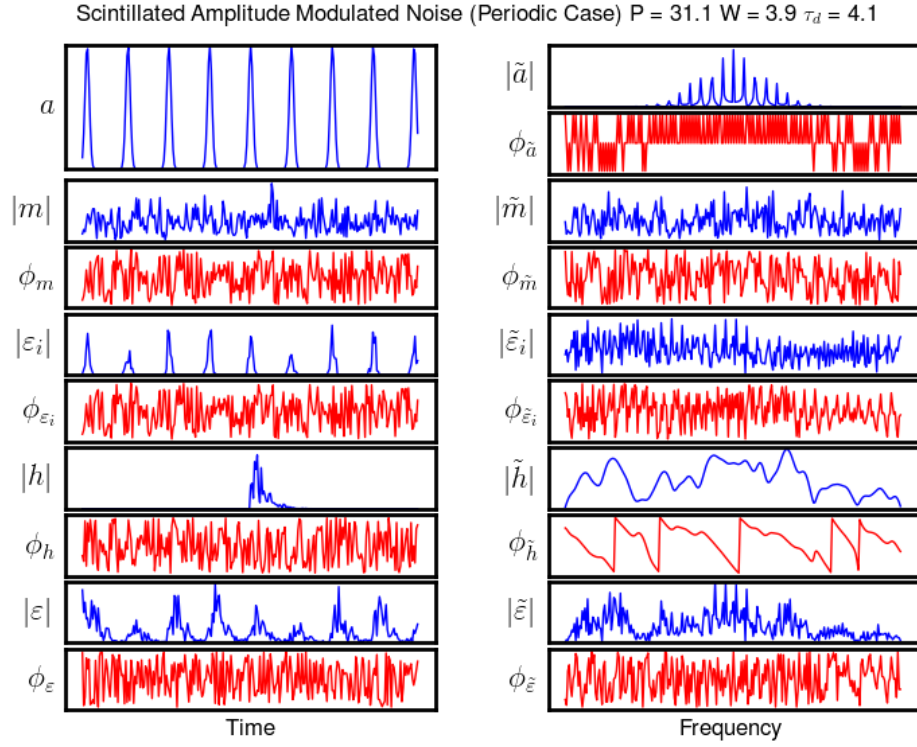


Fig. 36.— Components of scintillated amplitude modulated noise in the time and frequency domain for a periodic train of pulses. Amplitudes are plotted in blue, phases in red.

AMN Components, Example 5: Periodic Case

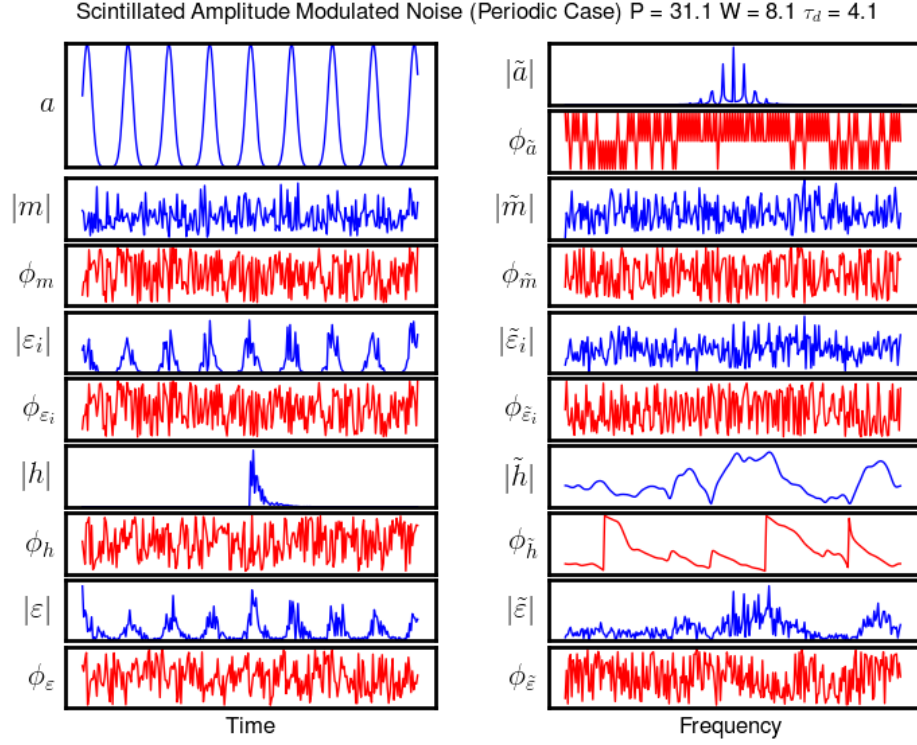


Fig. 37.— Components of scintillated amplitude modulated noise in the time and frequency domain for a periodic train of pulses. Amplitudes are plotted in blue, phases in red.

AMN Components, Example 6: Periodic Case

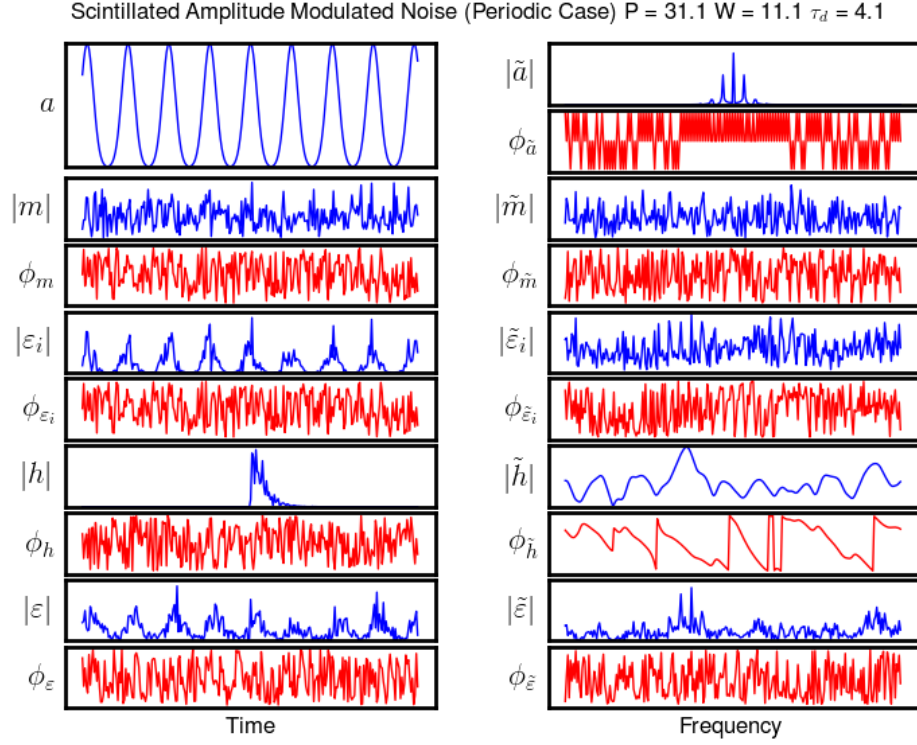


Fig. 38.— Components of scintillated amplitude modulated noise in the time and frequency domain for a periodic train of pulses. Amplitudes are plotted in blue, phases in red.

SAMN Intensity Statistics Example 2




Fig. 39.— Time domain intensity-like quantities are shown on the left. Frequency-domain quantities are on the right. The dotted lines in the plots for $S_i(\nu)$ and $S(\nu)$ are averages over 10 realizations. The ISM transfer function $H(\nu)$ is kept fixed for all realizations, so the average of $S(\nu)$ tends toward the shape of $H(\nu)$. The ACFs $\Gamma_{\tilde{\varepsilon}_i}$ and $\Gamma_{\tilde{\varepsilon}}$ are averaged over 10 realizations, as indicated.

SAMN Intensity Statistics Example 3

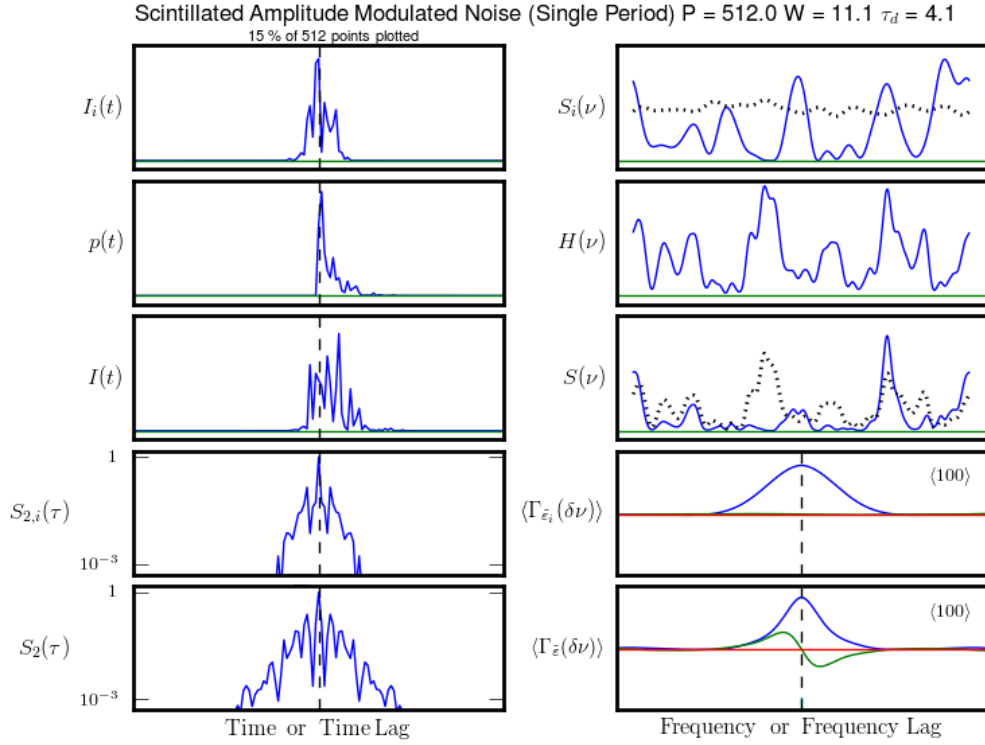


Fig. 40.— Time domain intensity-like quantities are shown on the left. Frequency-domain quantities are on the right. The dotted lines in the plots for $S_i(\nu)$ and $S(\nu)$ are averages over 100 realizations. The ISM transfer function $H(\nu)$ is kept fixed for all realizations, so the average of $S(\nu)$ tends toward the shape of $H(\nu)$. The ACFs $\Gamma_{\tilde{\varepsilon}_i}$ and $\Gamma_{\tilde{\varepsilon}}$ are averaged over 100 realizations, as indicated.

SAMN Intensity Statistics Example 4




Fig. 41.— Time domain intensity-like quantities are shown on the left. Frequency-domain quantities are on the right. The dotted lines in the plots for $S_i(\nu)$ and $S(\nu)$ are averages over 10^4 realizations. The ISM transfer function $H(\nu)$ is kept fixed for all realizations, so the average of $S(\nu)$ tends toward the shape of $H(\nu)$. The ACFs $\Gamma_{\tilde{\varepsilon}_i}$ and $\Gamma_{\tilde{\varepsilon}}$ are averaged over 10^4 realizations, as indicated.

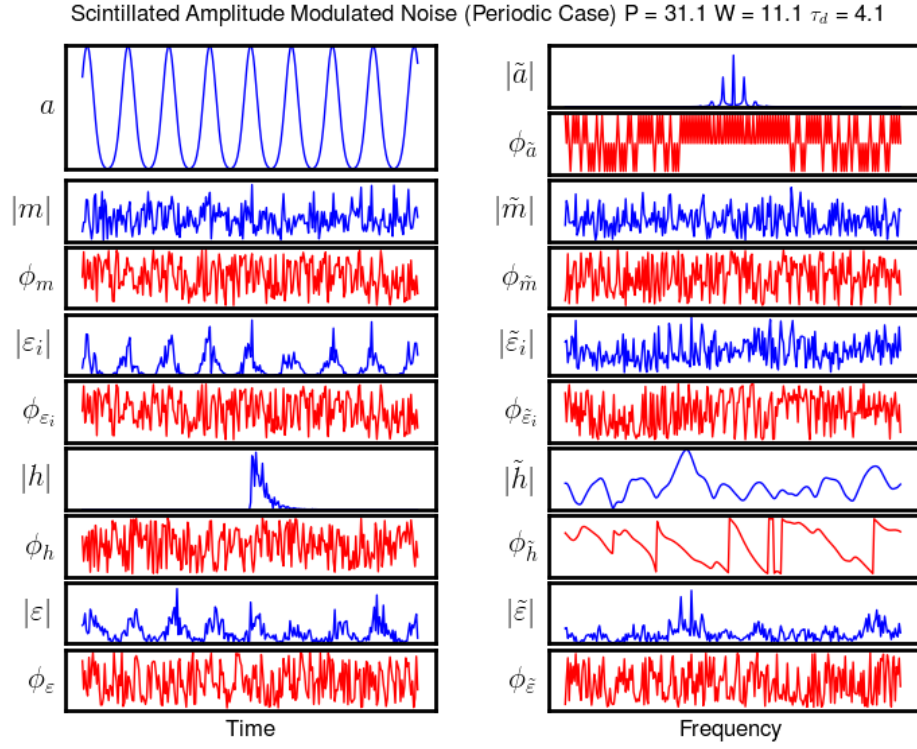


Fig. 42.— Components of scintillated amplitude modulated noise in the time and frequency domain for a periodic train of pulses. Amplitudes are plotted in blue, phases in red.