Commissioning Report for the GBT Cyclic Spectroscopy Backend System

Ryan Lynch, Paul Marganian, Jacob Turner, Tyler Hise, Wolfgang Baudler, Dane Sizemore, Ross Jennings, Joe Brandt, Chris Clark, and Kirstin Morin

Date: June 25, 2026
Keywords: cyclic spectroscopy, pulsars, interstellar medium

Abstract

Cyclic Spectroscopy (CS) is a signal processing technique that allows pulsars to be observed with both high pulse phase and radio frequency resolution. CS can be used to study the ionized interstellar medium (IISM) and its effects on pulsar signals, leading to a better understanding of the structure of the ISM and better high-precision pulsar timing. Green Bank Observatory (GBO) has built an expert-user CS backend (CSB) for the Green Bank Telescope (GBT). The CSB produces data products known as periodic spectra, and it works in parallel with the Versatile Green Bank Astronomical Spectrometer (VEGAS) to simultaneously produce traditional pulsar data products. The CSB can process up to 4.5 GHz of dual-polarization instantaneous bandwidth (3.75 GHz of which is usable), and thus can be used over the full bandwidth of the 0.7 – 4 GHz ultrawideband receiver (UWBR), as well as lower bandwidth receivers. The CSB processes data offline using custom GPU-accelerated software known as Cyclid, working automatically immediately after a scan ends to produce periodic spectra in ≤ 2x the scan length. The CSB can operate continuously as long as individual scans are ≤ 1 hour in duration. The CSB is configured using the standard GBT observing control software (Astrid), and the backend is managed using a web-based interface called Cyclops.

We performed commissioning observations for the CSB between August and November 2025 using the PF800, L-Band, and UWB receivers. We observed a bright millisecond pulsar (MSP) using the most computationally demanding CSB observing modes and compared the calibrated flux density, polarization properties, and pulsar times of arrival to those obtained with VEGAS, finding good agreement. We also observed two additional MSPs and one long-period pulsar, demonstrating the

unique ability of the CSB to study the IISM and to characterize interstellar scatting delays in ways that cannot be accomplished using VEGAS. We also verified the integration of the CSB with Astrid, tested the functionality of Cyclops, and verified the ability of the CSB to process data with the requisite speed for supported observing modes, to operate continuously, and to automatically process data and verify its quality using basic quality assurance checks.

This commissioning report provides a scientific and technical overview of the CSB and presents the results of commissioning observations.

Table of Contents

Change Log

1.0 Lynch et al. (2026-06-25) – Initial published version

Introduction

Pulsars are powerful tools for studying a wide range of phenomena. The most relevant in the context of CS is the IISM and its impact on high-precision pulsar timing. Density variations in the IISM diffract pulsar signals at radio frequencies, scattering rays into an observer's line of sight that would not otherwise intersect with the Earth. The differing path lengths upon which these signals travel causes them to arrive at Earth with different phases, leading to constructive and destructive interference known as scintillation. The resulting interference pattern varies with time and frequency leading to increases in flux density that are localized over a characteristic scintillation bandwidth, $\Delta\nu_{\rm scint}$, and timescale, $t_{\rm scint}$. Multi-path propagation also causes pulsed emission to arrive at the Earth with varying time delays that are functionally described as a one-sided exponential that decays by 1/*e* over a scattering timescale, τ_s. The relationship between the scattering timescale and scintillation bandwidth is given by

$$\Delta\nu_{\rm scint}\,\tau_s = 2\pi C_1$$

where C_1 is a constant of proportionality related to the geometry of density variations in the IISM. For a uniform electron density model with a Kolmogorov turbulence spectrum $C_1 = 1.16$, while for a thin-screen density model and Kolmogorov spectrum $C_1 = 0.957$.

Within pulsar astronomy a measurement of pulse flux density as a function of time and frequency (i.e. an observation of the diffractive interference pattern) is known as a dynamic spectrum, and such a measurement encodes a wealth of information about the IISM. Specifically, measurement of the scintillation bandwidth can be used to estimate the scattering delay up to the value of C_1. Furthermore, a 2-D Fourier transform of the dynamic spectrum, known as a secondary spectrum, often reveals parabolic arcs and inverted parabolic arclets which are related to the distance to the scattering screens. The secondary spectrum is thus a valuable probe of the geometry of the IISM.

While scattering is useful for studying the IISM, it serves as a nuisance term for high-precision pulsar timing. Scattering delays bias measurements of pulse times of arrival (TOAs), and epoch-to-epoch changes in τ_s are a source of stochastic red noise in pulsar timing models. The influence of nanohertz-frequency gravitational waves (GWs) also manifests as a stochastic red noise process, so unmodeled scattering delays decrease pulsar timing array sensitivity to GWs.

Scattering delays can be estimated by measuring $\Delta\nu_{\rm scint}$ from dynamic spectra (which, as previously mentioned, are also useful for studying the IISM), but this is not always possible using traditional signal processing techniques. Autocorrelation and polyphase filterbank spectrometers rely on the Fourier transform to measure radio frequency power spectra. Because time and frequency are conjugate variables, they are subject to the uncertainty principle. Thus, obtaining fine frequency resolution necessarily implies coarse time resolution, and vice versa. Pulsars, in particular MSPs, frequently exhibit temporal structure on the order of 0.1 – 0.01 milliseconds, and resolving this temporal structure is critical for high-precision pulsar timing. As such, pulsar observers typically prioritize high time resolution at the expense of frequency resolution. A common observing setup obtains frequency resolution of about 1 MHz and time resolution of about 1 μs. This resolution is not always sufficient to resolve scintles in highly scattered pulsars, making it difficult to reliably characterize scattering delays and to directly account for the influence of scattering in high-precision pulsar timing.

The Fourier transform uncertainty principle specifically applies to wide-sense stationary processes. However, the periodic nature of pulsars makes them cyclostationary, i.e. their statistical properties are periodic in nature. This property allows one to use advanced signal processing techniques to break the degeneracy between time and frequency resolution and simultaneously achieve very high radio frequency and pulse phase resolution. The resulting cyclic spectrum measures amplitude and phase as a function of radio frequency (ν) and cycle frequency (α), with a frequency resolution limited only by $1/P$, where P is the pulsar period.

This has several benefits. First, it may be possible to measure $\Delta\nu_{scint}$, and thus estimate τ_s, while maintaining adequate pulse phase resolution. Second, measuring the average phase slope of the cyclic spectrum provides another measure of the total time delay that a pulsar signal experiences. Finally, under certain conditions it may be possible to measure the impulse response function (IRF) of the IISM, and to thus measure τ_s directly. However, CS is computationally demanding, and has therefore only been implemented in custom-built backends operating over relatively narrow bandwidths, or in limited cases when baseband voltage data can be recorded and processed offline.

GBO received funding from the National Science Foundation to build a CS observing system for the GBT that can operate over the full bandwidth offered by UWBR. This will be the first observatory-supported CS observing system. The CSB operates simultaneously with VEGAS to record baseband voltages to a dedicated storage system, which are then processed using custom GPU-accelerated software to produce periodic spectra (which are related to the cyclic spectrum by a Fourier transform) shortly after an observation ends. While officially an expert-user system, the CSB is controlled via the usual GBT observing interface, processes data automatically, and is monitored using a simple browser-based interface.

We have performed scientific commissioning observations of the CSB using the PF800, L-Band, and UWB receivers. We observed several pulsars, producing dynamic, secondary spectra and cyclic spectra. We also compared flux and polarization calibrated integrated pulse profiles and TOAs to those obtained with VEGAS. In the following section we present the results of these commissioning observations and certify that the CSB is operating correctly and ready for use by GBT observers.

Technical Overview

Before reporting on commissioning observations we will provide a technical overview of the operation of the CSB. The CSB uses existing GBT receivers, the existing IF system, and existing VEGAS ROACH2 boards to coarsely channelize the data via PFBs. The CSB diverges from VEGAS after the data have been formatted into UDP packets and transmitted from the ROACH2s. When the CSB is not being used the data are uni-cast to the VEGAS HPCs via 10 GbE. However, when the CSB is being used the data are multi-cast to both the VEGAS HPCs (where they are processed using traditional methods) and to the CS processing cluster. The detailed signal path and data processing steps are as follows.

1) Analog baseband data is routed from the Converter Modules to the VEGAS ROACH2 boards.

2) The ROACH2 boards use PFBs to coarsely channelize the baseband data. Each PFB channel is a narrowband, Nyquist-sampled complex voltage time series that together cover the full sampled bandwidth.

3) The ROACH2 boards can sample 800 or 1500 MHz (1250 MHz of which is usable) of bandwidth. When using UWBR three ROACH2 boards are used to sample 1500 MHz of bandwidth each so as to cover the full UWBR frequency range, as well as 3375 MHz of bandwidth for other receivers.

4) The full bandwidth sampled by each ROACH2 is broken into eight sub-bands, each containing 1/8th of the total number of PFB channels per ROACH2. Depending on the recorded bandwidth, these are then multi-cast via 10 Gb Ethernet to one, two, eight, or 24 VEGAS HPCS, and to the CS processing cluster. We will refer to each of these sub-bands as a data stream.

5) Data that are multi-cast to VEGAS are acquired and processed using the guppi_daq software. guppi_daq obtains data and associated meta-data from Shared Memory on each VEGAS HPC. Data are processed on GPUs (one for each HPC) and written to disk in the PSRFITS format. This is all managed by the VEGAS Managers, and various processes are controlled via TaskMaster.

6) Data multi-cast to the CS processing cluster (located in the Data Center) are initially captured by a separate instance of guppi_daq running on each CS host. guppi_daq acquires the data and associated meta-data from Shared Memory on each CS node, and Shared Memory meta-data is mirrored from the associated VEGAS HPC. Rather than further process the data in real time, the guppi_daq instance on the CS nodes is used to write baseband data in a specialized format (GUPPI RAW). These baseband data are essentially the PFB output.

7) After a scan ends it enters a queue to be processed. Once that scan is reached in the processing queue, GPU-accelerated software known as Cyclid processes the data to produce the periodic spectra in the PSRFITS format.

8) After processing is complete, the PSRFITS data undergo automated quality assurance tests. If they pass, the GUPPI RAW data are deleted.

9) Data acquisition and processing on the CS cluster is managed via Cyclops, which has a browser-based interface. A number of processes are controlled via TaskMaster. There is no M&C Manager or CLEO application for CS data processing.

Each HPC in the CS processing cluster receives up to four data streams, and each HPC is equipped with two GPUs per data stream (eight in total). Furthemore, each HPC can store up to 16 TB of baseband data per data stream, which equates to 6 hours of data when using 1500 MHz bandwidth VEGAS modes. However, to maintain optimal I/O performance disk usage should not exceed 14 TB.

Observers are restricted to observing setups that can be processed in ≤ 2x the scan length, and they are further restricted to a maximum scan length of one hour. Thus, the longest a scan will take to process is two hours. Since there are two GPUs per data stream, processing of one scan can begin at the same time as a second scan starts. Under normal operations the first scan will be processed and the baseband data will be automatically deleted by the the time a third scan ends, thus there will never be more than three hours worth of data on disk at any given time. This is equivalent to 8.1 TB, which is significantly less than the 14 TB maximum effective capacity of a disk, with sufficient overhead to accommodate unforeseen processing errors without impacting observing. Thus, the CSB can operate continuously without exceeding the capacity of the storage disks.

Figure 1 is a schematic representation of the CSB and its relationship to VEGAS (note that only one ROACH2 sub-band is shown).

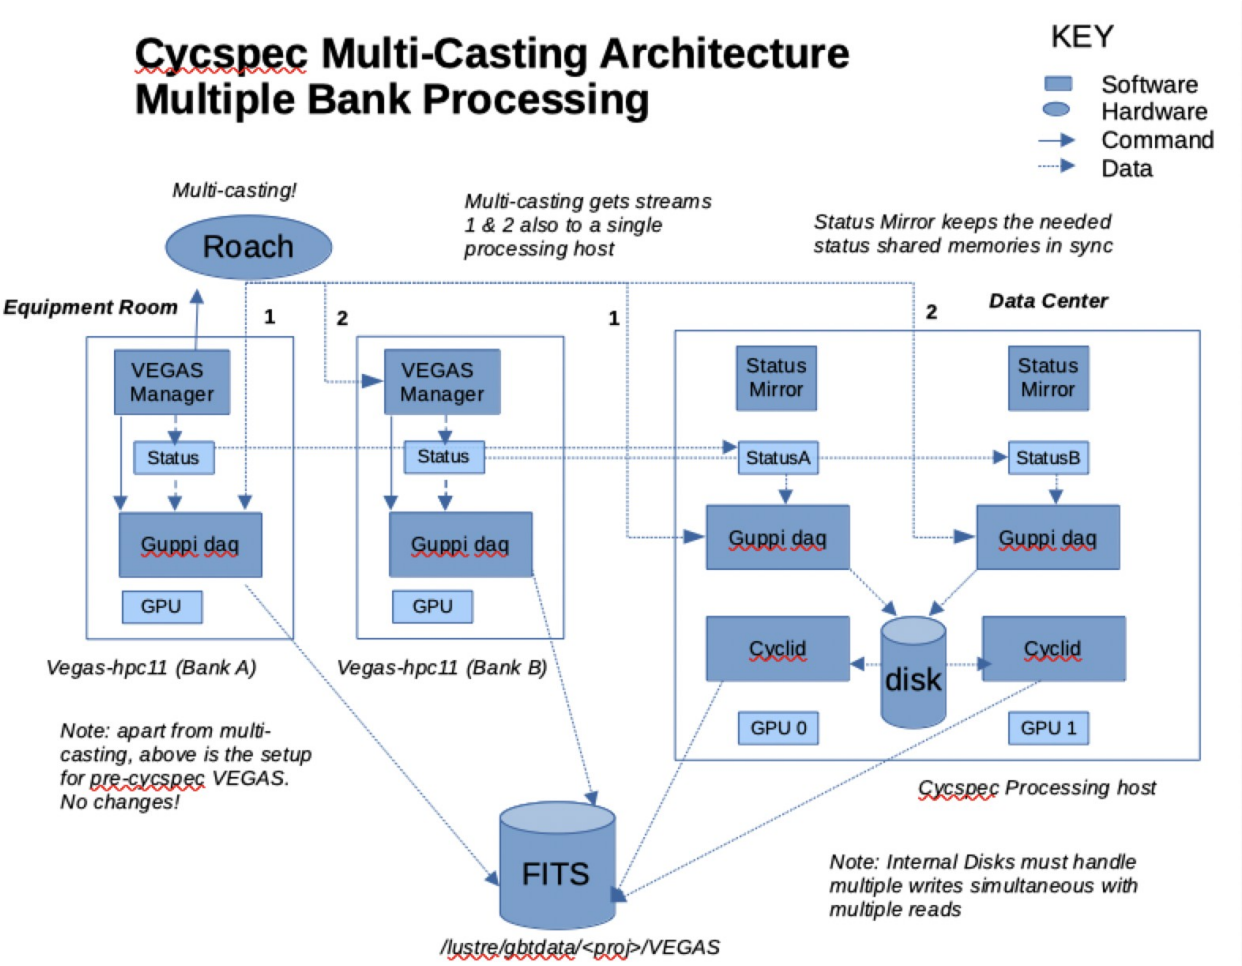

Figure 1: A schematic representation of the CSB and its relationship with VEGAS and other M&C and computing infrastructure.

Cyclid

Cyclid is a custom GPU processing package for producing periodic spectra from GUPPI RAW baseband data. Cyclid takes the following as input:

- A pulsar ephemeris, or "par file". This is a simple text file that contains rotational, astrometric, dispersive delay, and (when necessary) binary orbital parameters. An ephemeris file is used with a third-party pulsar software package known as tempo to predict the observed rotational phase of a pulsar as a function of time. These predictions are used by Cyclid to phase-fold the periodic spectrum.
- The number of pulse phase bins. This is the number of discrete samples that the rotational phase of the pulsar is divided into during phase-folding.
- The number of "cyclic channels". This represents the additional frequency resolution that can be obtained using CS. The frequency resolution of the periodic spectrum is given by $\Delta\nu = \mathrm{BW}/(n_{\mathrm{pfb}} \times n_{\mathrm{cyc}})$.

- The sub-integration time. This is the time over which periodic spectra are accumulated before being written to disk. Ten seconds is typical.

In addition to computing the periodic spectrum and phase-folding the data, Cyclid also coherently dedisperses the data. In place of a par file, Cyclid can fold calibration scans of a receiver's switched noise diode. Following usual GBT pulsar observing conventions, the noise diode must be switched at 0.04 seconds (25 Hz).

All of the inputs options for Cyclid are specified by observers using Astrid configuration keywords.

Cyclid writes data to project directories in /stor/gbtdata. Within each project directory, data is organized into sub-directories of the form CYCSPEC/<bank>, where <bank> is a one-letter code specifying the data stream from which the periodic spectra were produced (this is equivalent to the VEGAS Bank ID). The output from each data stream is then "merged" into PSRFITS files written to /stor/pulsar/gbtdata which cover the full observing bandwidth. Data are merged by a cron job that runs every hour.

Cyclops

As mentioned previously, the CSB does not have an independent Manager. Instead, operation of the CSB is managed by a tool known as Cyclops. When Cyclops recognizes that an observer has configured for a CS scan it tracks the status of data acquisition. Cyclops also produces simple diagnostic plots that can be used by an observer to assess the quality of the baseband data. Once a scan ends, Cyclops automatically starts an instance of Cyclid to process the baseband data acquired as part of that scan. If Cyclid finishes without errors then Cyclops next starts a quality assurance script, which verifies that the periodic spectra produced by Cyclid were processed with the correct parameters and produced files with the expected format and duration. If these quality assurance checks pass, then Cyclops automatically deletes the baseband data, the diagnostic plots, and marks the bank's scan as 'completed'. If Cyclid or the quality assurance checks fail, the bank's scan is marked as 'Failed' (similarly, if Cyclid is externally interrupted, the bank's scan is marked as 'Aborted'). Reprocessing of data requires human intervention: GBO staff can mark a bank's scan as 'not started', and Cyclops will attempt to reprocess the data. If these further attempts at reprocessing fail, GBO staff can mark a bank's scan as 'terminated' and Cyclops will ignore these in the future.

Cyclops has a browser-based interface that can be accessed at https://cyclops.gb.nrao.edu. Users must log in with their my.nrao.edu credentials, and access is restricted to projects that an observer is associated with the in the Dynamic Scheduling System. GBO staff has special privileges that allow them to view all projects and to manually mark scans for deletion when needed.

Data and Processing Scripts

All data presented in this commissioning report were recorded under GBT project TGBT25B_605. Observing logs can be found in /users/rlynch/CycSpec/Commissioning/2025_August and /users/rlynch/CycSpec/Commissioning/2025_November. GBT data can be found in /stor/gbtdata ("unmerged" data produced automatically by cyclops and VEGAS) and /stor/pulsar/gbtdata (merged

data covering the full observing bandwidth, as well as data produced manually to test additional CS parameter combinations). A description of each session follows.

- TGBT25B_605_04 – Observations of PSR B1937+21 used to compare the CSB to VEGAS 200 MHz bandwidth modes and to create dynamic, secondary, and cyclic spectra
- TGBT25B_605_05 – Observations PSR B0523+11 used to create dynamic and secondary spectra of a long-period pulsar using 200 MHz bandwidth modes
- TGBT25B_605_06 – Observations of PSRs J1600-3053 and J1643-1224 used to create dynamic and secondary spectra of MSPs using 200 MHz bandwidth modes
- TGBT25B_605_07 and TGBT25B_605_10 – Observations of PSR B1937+21 used to compare the CSB to VEGAS 800 MHz bandwidth modes and to create dynamic, secondary, and cyclic spectra
- TGBT25B_605_08 – Observations of PSR B0523+11 used to create dynamic and secondary spectra of a long-period pulsar using 800 MHz bandwidth modes
- TGBT25B_605_09 – Observations of PSRs J1600-3053 used to create dynamic and secondary of an MSP using 800 MHz bandwidth modes
- TGBT25B_605_11 – Observations of PSRs J1643-1224 used to create dynamic and secondary spectra of an MSP using 800 MHz bandwidth modes
- TGBT25B_605_12 and TGBT25B_605_13 – Maintenance tests
- TGBT25B_605_14, TGBT25B_605_18, TGBT25B_605_19, and TGBT25B_605_20: Observations of B1937+21 used to compare the CSB to VEGAS using UWBR

Shell scripts used to manually run Cyclid to test many parameter combinations can be found in subdirectories in /stor/pulsar/gbtdata/TGBT25B_605_04/c0200*, /stor/pulsar/gbtdata/TGBT25B_605_07/c0800*, /stor/pulsar/gbtdata/TGBT25B_605_10/c0800*, and /stor/pulsar/gbtdata/TGBT25B_605*/c1500*. Scripts to merge these data can also be found in the same sub-directories. Scripts used to calibrate the data and obtain TOAs can be found in sub-directories in /home/scratch/rlynch/CycSpec/Commissioning/Data/2025_August/ and /home/scratch/rlynch/CycSpec/Commissioning/Data/2025_November. These subdirectories also contain Python scripts used to compare CSB data to VEGAS data. Diagnostic plots can be found in /home/scratch/rlynch/CycSpec/Commissioning/Data/2025_August and /home/scratch/rlynch/CycSpec/Commissioning/Data/2025_November/.

Results

Allowed Observing Modes

As noted above, observers are restricted to observing parameters that result in a Cyclid processing time which is ≤ 2x the scan length. We determined which parameter combinations meet this criteria by analyzing drift-scan data and processing it with Cyclid for all possible parameter combinations. The most important parameters are

1) The sampled bandwidth
2) The number of PFB channels
3) The number of cyclic channels
4) The number of pulse phase bins

There is also some dependence on the observing frequency and the pulsar's dispersion measure (DM), as this influences the computational speed of the coherent dedispersion step and, more subtly, the way that GUPPI RAW files are formatted. Specifically, GUPPI RAW files are organized into data blocks, each of which overlaps slightly with the previous in order to facilitate coherent dedispersion. While the total data volume is determined only by the sampled bandwidth, the size of an individual data block depends in part on the pulsar's DM. Because Cyclid operates on the scale of individual data blocks, this in turn has some effect on Cyclid's processing speed.

The aforementioned drift scans were 60 seconds in duration and the data were processed on the CS processing cluster. We took data for DMs of 10, 100, and 1000 pc/cm^3 and at the low- and high-frequencies covered by UWBR (700 MHz and 4 GHz, respectively). We then determined the parameter combinations that met the 2x processing time factor criterion. For any given parameter combination, we used the largest processing time factor for the combinations of observing frequency and DM to determine the allowed modes. Table 1 shows the maximum value of n_{cyc} for different values of sampled bandwidth, n_{pfb}, and n_{bin}. There are several things to note:

1) Many combinations are superseded in either frequency or pulse phase resolution, i.e. the same frequency resolution can be obtained while also obtaining a higher pulse phase resolution. While other parameter combinations are possible, those highlighted in green are the ones that offer the best pulse phase resolution for a given frequency resolution, and these are the modes we expect to be used by observers.
2) The 100, 200, and 800 MHz bandwidth modes all use the same VEGAS firmware designs and all produce data streams covering 100 MHz. While observers think of them as different modes because the total bandwidth and number of channels differ, from the perspective of VEGAS and the CSB they are in fact identical. However, they are often used at different observing frequencies (100 MHz with PF342, 200 MHz with PF800, and 800 MHz with L-Band).

3) The target use case defined in the original CS proposal uses 1500 MHz bandwidth per ROACH2, 1024 PFB channels, n_{cyc} = 128, and n_{bin} = 512. NANOGrav will use the 1500 MHZ x 1024-channel VEGAS mode with UWBR.

Table 1: Maximum Allowable n_{cyc}

$n_{bin} \rightarrow$ $n_{pfb} \downarrow$	**32**	**64**	**128**	**256**	**512**	**1024**	**2048**
100 MHz Bandwith Modes							
64	512	512	512	256	256	128	64
128	512	512	256	256	128	64	64
256	512	256	256	128	64	64	32
512	256	128	128	64	64	32	N/A
200 MHz Bandwidth Modes							
64	256	256	256	256	128	128	128
128	512	512	512	256	256	128	64
256	512	512	256	256	128	64	64
512	512	256	256	128	64	64	32
1024	256	128	128	64	64	32	N/A
800 MHz Bandwidth Modes							
32	32	32	64	64	32	32	32
64	128	128	128	128	128	128	128
128	128	256	128	128	128	128	128
256	256	256	256	256	256	128	128
512	512	512	512	256	256	128	64
1024	512	512	256	256	128	64	64
2048	512	256	256	128	64	64	32
4096	256	128	128	64	64	32	N/A
1500 MHz Bandwidth Modes							
64	64	16	32	64	16	16	16
128	64	64	64	64	64	64	64
256	128	128	128	128	128	128	64
512	256	256	128	128	128	64	64
1024	256	256	256	128	128[†]	64	16
2048	256	256	128	128	64	32	8
4096	256	128	64	64	16	4	N/A

[†] Target use case defined in original proposal.

Comparison with VEGAS

CS offers better frequency resolution at an equivalent pulse phase resolution, and access to phase information that is not possible with traditional observing. However, when integrated in frequency the periodic spectrum should reduce to an integrated pulse profile that is identical to that obtained with traditional techniques. A fundamental test of the CSB is that it produce data that agree with VEGAS in the following ways:

1) Integrated pulse profile shapes should be the same.
2) Integrated mean flux density and S/N should be the same.
3) Linearly and circularly polarized flux should be the same.
4) Pulsar TOAs should agree to within a constant offset. This constant offset arises from the different signal processing steps involved in CS and is trivial to compensate for once measured (similar instrumental offsets are ubiquitous in pulsar astronomy, and characterizing them is a normal part of combining data from different instruments).

To verify these properties we simultaneously obtained CS and VEGAS data on the bright MSP B1937+21. We also obtained a flux calibration solution using 3C295[1]. VEGAS data were always obtained using 2048 pulse phase bins, and we downsampled in pulse phase to make a fair comparison to the CS data taken with different pulse phase resolution. We used the PSRCHIVE software package to calibrate the total and polarized flux, correct for Faraday rotation, remove RFI, create standard pulsar timing templates, and to obtain TOAs. We used the tempo pulsar timing package to produce timing residuals from those TOAs and to compare the residuals from the CSB and VEGAS. We also used tempo to measure the instrumental offsets.

In all cases CSB data agree with VEGAS to within 3%. The CSB data frequently have higher S/N and/or lower average TOA uncertainty than VEGAS. The differences that we observe can be attributed to RFI – the finer channelization available with the CSB backend allows more precise RFI excision, which in some cases improves S/N. However, since power is spread over more frequency channels, low-level RFI is sometimes not removed as effectively in the CSB data compared to VEGAS. Figures 2 and 3 compare integrated pulse profiles in Stokes I, Q, U, and V, and pulse times of arrival, for a representative case. We empirically measure a timing offset of 2.4 μs in the 200 and 800 MHz-bandwidth modes. The timing offset in 1500-MHz modes (including UWBR modes) is consistent with zero.

1 An operating system error unrelated to CS system but impacting the CS processing cluster forced us to cut some of our UWBR commissioning observations short, and we were not able to acquire flux calibration data in all modes. However, uncalibrated VEGAS and CS data are in good agreement and given the excellent agreement that we see in calibrated data for other modes (including the c1500x1024uwb mode) we recommend releasing UWBR modes for general use.

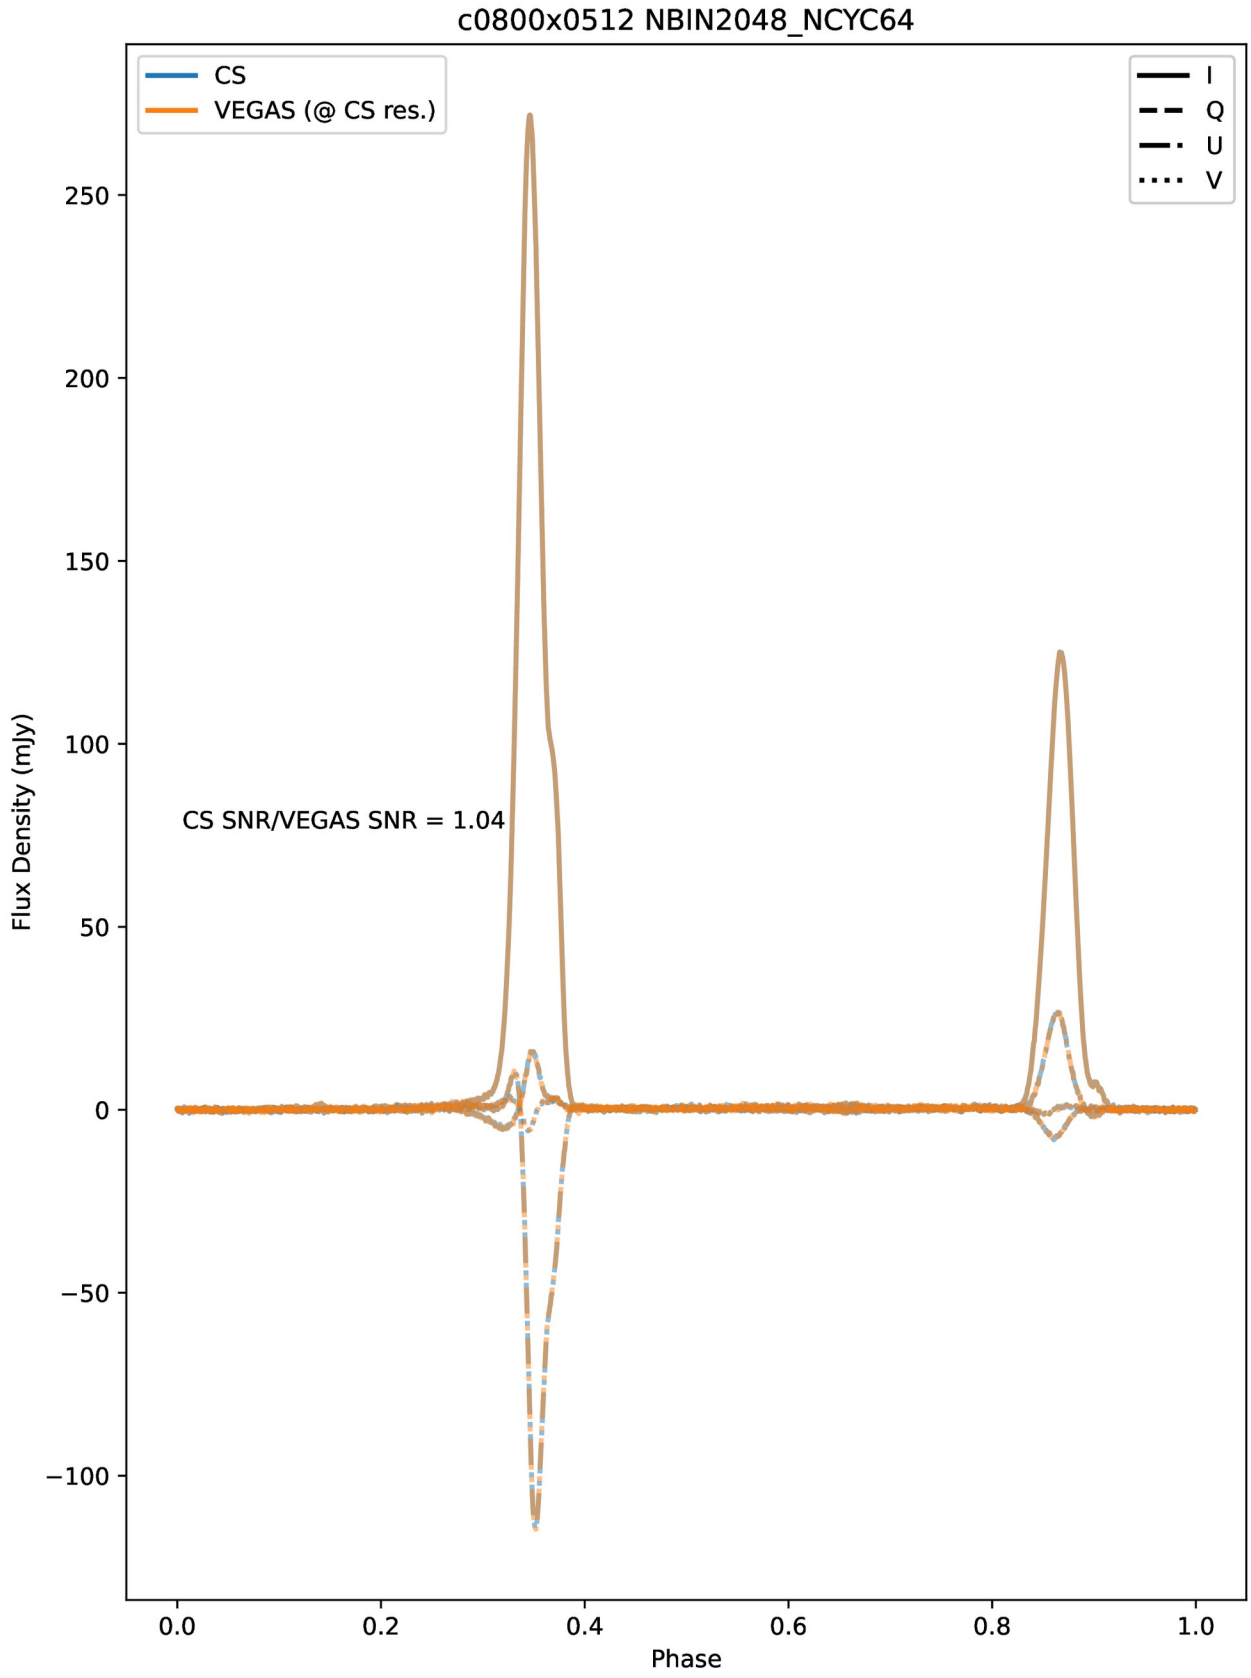

Figure 2: Integrated full Stokes pulse profiles for B1937+21 at L-Band for the CSB and VEGAS. The data have been fully calibrated and RFI has been excised using the psrchive package. The representative example shown here uses the same VEGAS mode as NANOGrav uses at L-Band, and the maximum pulse phase resolution.

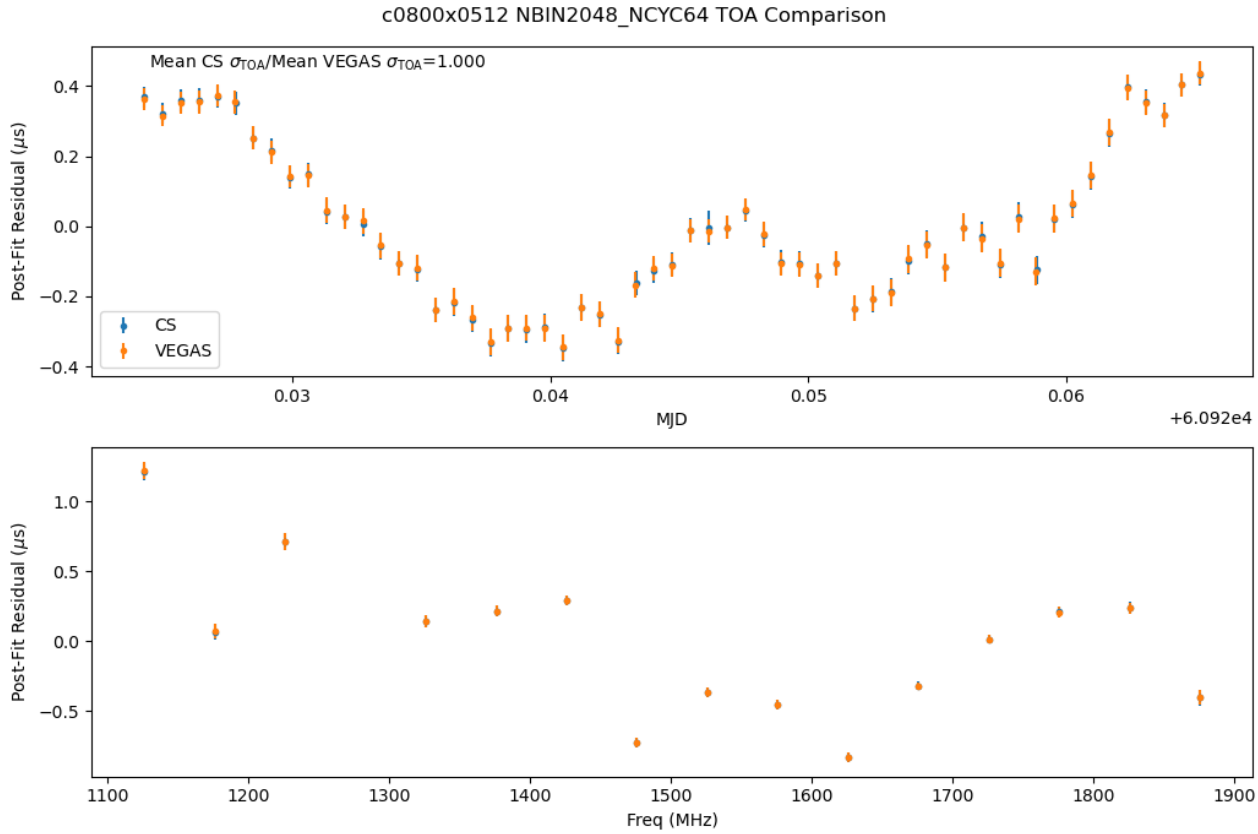

Figure 3: Pulsar timing residuals generated with TEMPO comparing VEGAS and the CSB. TOAs are shown as a function of both time and frequency. We have omitted TOAs at 1250 MHz because they are biased by the notch filter. A constant timing offsets between the CSB and VEGAS has been removed. This representative example uses the same mode and CS processing parameters as the pulse profile shown in Fig 2.

There are two aspects of the CSB and VEGAS that are worth mentioning here. The first is that modes which use a very small number of pulse phase bins are not recommended for pulsar timing analysis. The poor pulse resolution makes it difficult to obtain accurate TOAs. Specifically, modes with n_{bin} = 32 should not be used for pulsar timing at all, and modes with n_{bin} = 64 should only be used with caution, and for pulsars with broad pulse profiles. We emphasize that these limitations are not unique to the CSB. However, there is almost no advantage to using such a low number of pulse phase bins with VEGAS, so these limitations are not encountered in practice. The CSB can provide the best frequency resolution when using a small number of pulse bins, so modes with n_{bin} < 128 may be useful for studies of the IISM that do not require accurate pulsar timing.

The second issue to note is that off-pulse baseline artifacts may be present when using $n_{bin} \leq 128$. These artifacts are related to the removal of dispersive delays, and are not unique to the CSB or VEGAS. Interchannel dispersive delays must be removed to phase-align pulsar emission at different frequencies. Standard practice is to Fourier transform the pulse profile, apply a phase shift, and inverse Fourier transform the pulse profiles at each frequency. This is equivalent to a sinc

interpolation in the time domain, and when the rotation is by fractional phase bins, ringing can be introduced, especially for pulsars with sharp profile features. The amplitude of any artifacts depends on the exact phase rotation and pulse profile shape. Again, we emphasize that this is not unique to the CSB or VEGAS, and is reflects fundamental signal processing challenges.

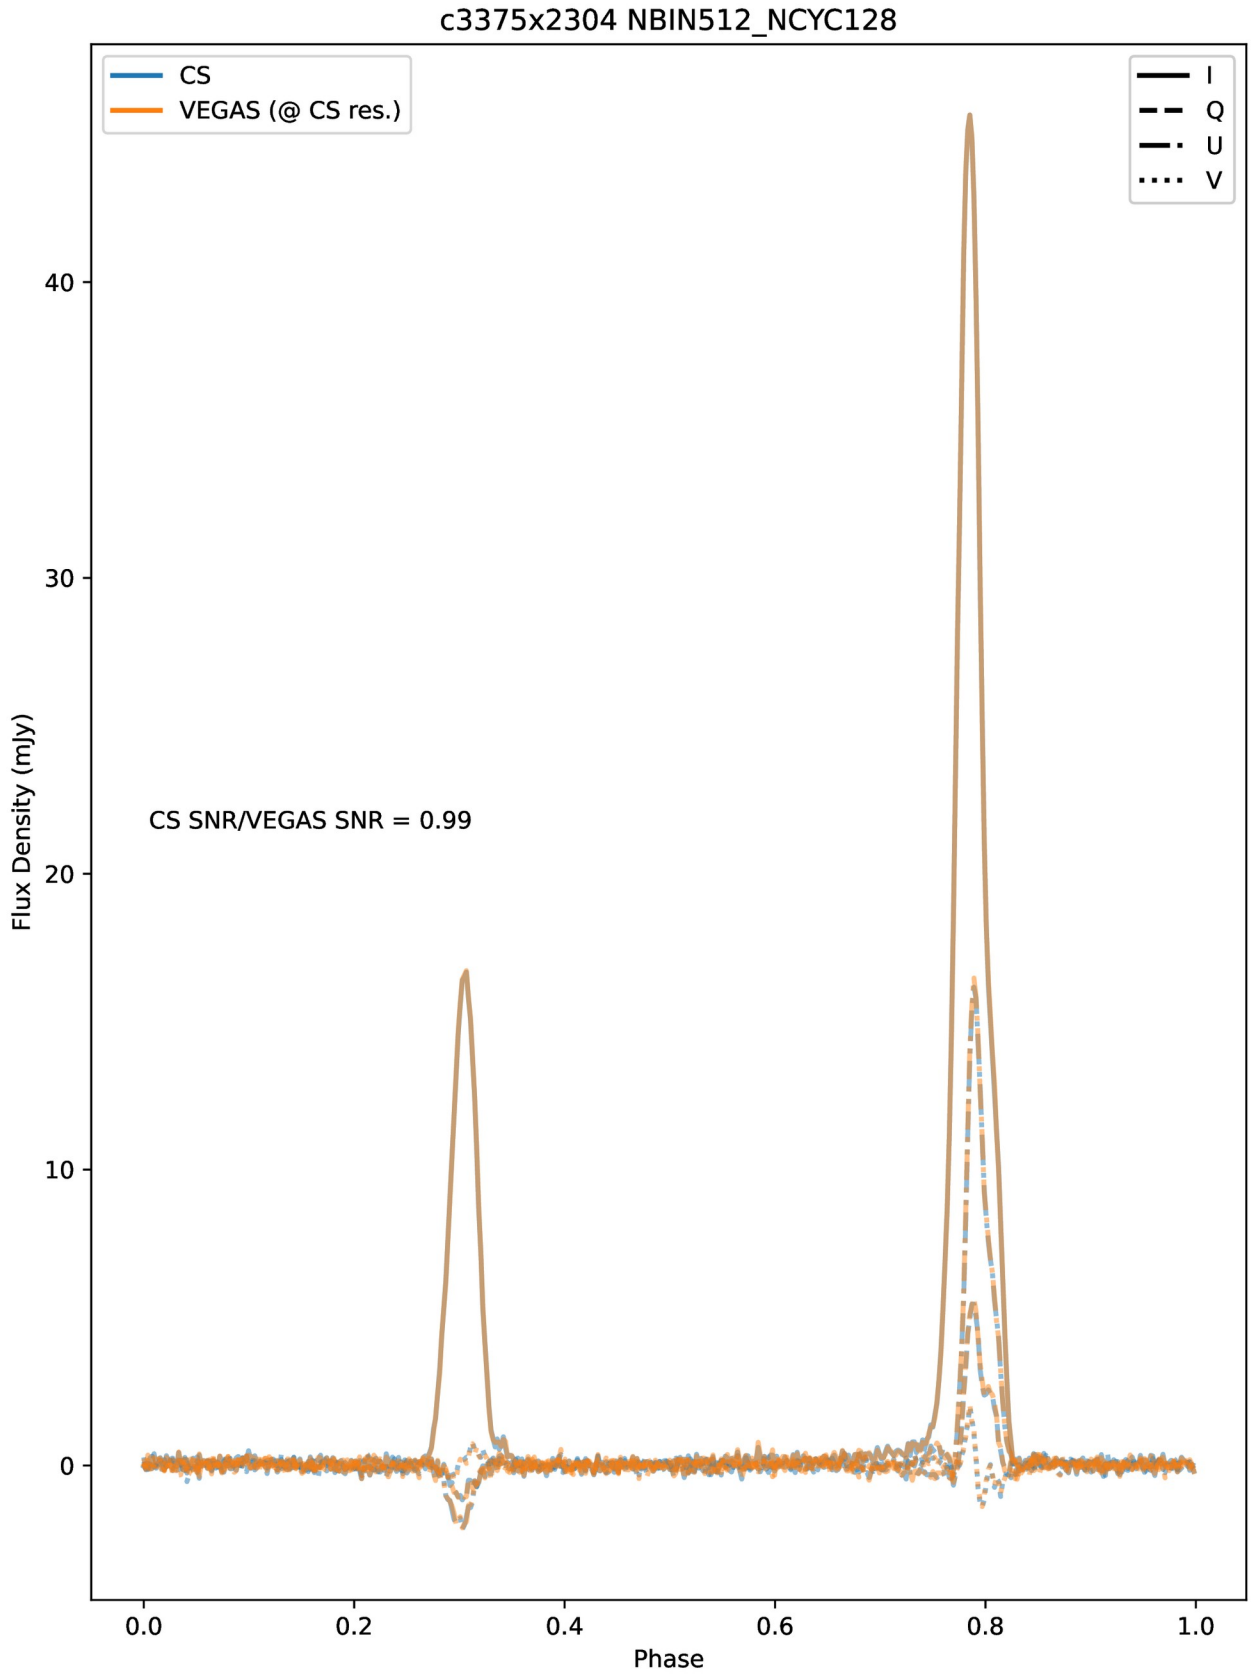

Figure 4: CSB and VEGAS pulse profile comparison for the UWBR c1500x1024uwb mode.

Dynamic Spectra

Dynamic spectra measure a pulsar's integrated on-pulse flux as a function of time and frequency. Variations in flux are the hallmark of scintillation. The better frequency resolution that can be obtained with the CSB compared to VEGAS makes it possible to create higher-quality dynamic spectra. We created dynamic spectra by dedispersing the periodic spectra, summing over both polarizations, and subtracting a constant baseline offset. We then took the difference between the integrated on- and off-pulse components of the profile at each pixel in time-frequency and plotted the result to create the dynamic spectrum. Fig. 5 shows dynamic specta for PSR B1937+21 obtained using the PF800 receiver. The left panel shows data obtained with the CSB using a frequency resolution of 48.8 kHz, while the right panel shows the same data at the best-possible VEGAS frequency resolution that can be achieved while using coherent dedispersion. While scintillation is clearly visible in the right panel, it is also not fully resolved. With the CSB, scintles are fully resolved and the S/N is higher. Additional secondary spectra are shown in the appendix.

The scintillation bandwidth and scintillation timescale can be measured by computing the 2-D autocorrelation function (ACF) of the dynamic spectrum. These quantities can then be used to infer information about the structure and evolution of the interstellar medium along the line of sight to a pulsar, including estimations of pulse time of arrival delays from interstellar scattering, interstellar medium refractive timescales, electron density wavenumber spectrum turbulence, and pulsar transverse velocity. Fig. 6 shows the ACF of the dynamic spectrum shown in Fig. 6 for B1937+21. We recover t_{scint} = 126 s and $\Delta\nu_{scint}$ = 62 kHz, which implies τ_s = 97 μs, which is consistent with literature values for this pulsar. Additional ACFs are shown in the appendix.

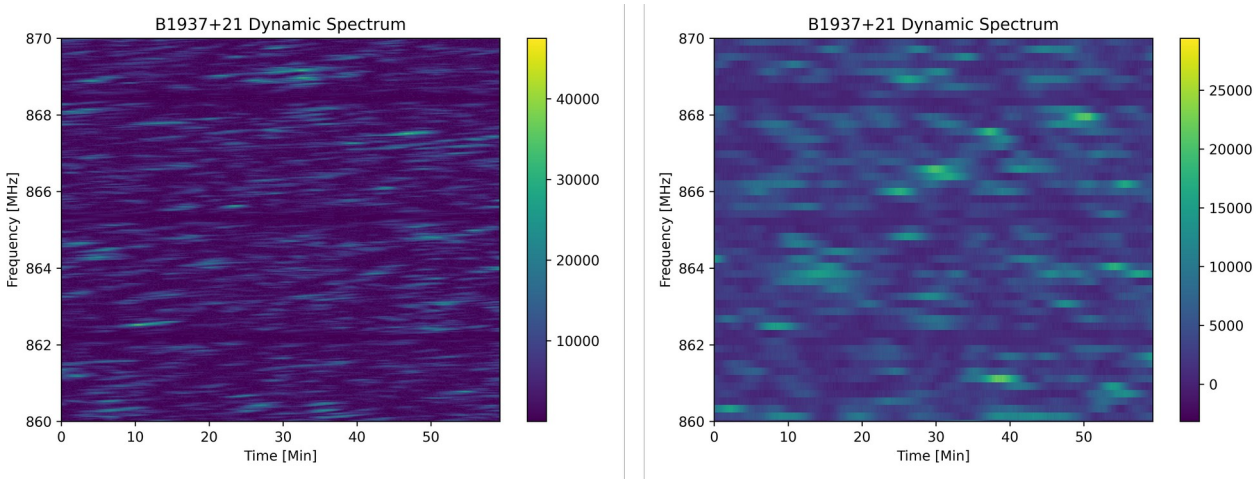

Figure 5: Dynamic spectrum of B1937+21 obtained with the CSB and the PF800 receiver. Note that only 10 MHz out of the full 200 MHz is shown so that narrow-bandwidth scintles can be seen. The right panel shows the same data but at the best possible resolution provided by VEGAS coherent dedispersion modes. Scintles are now unresolved.

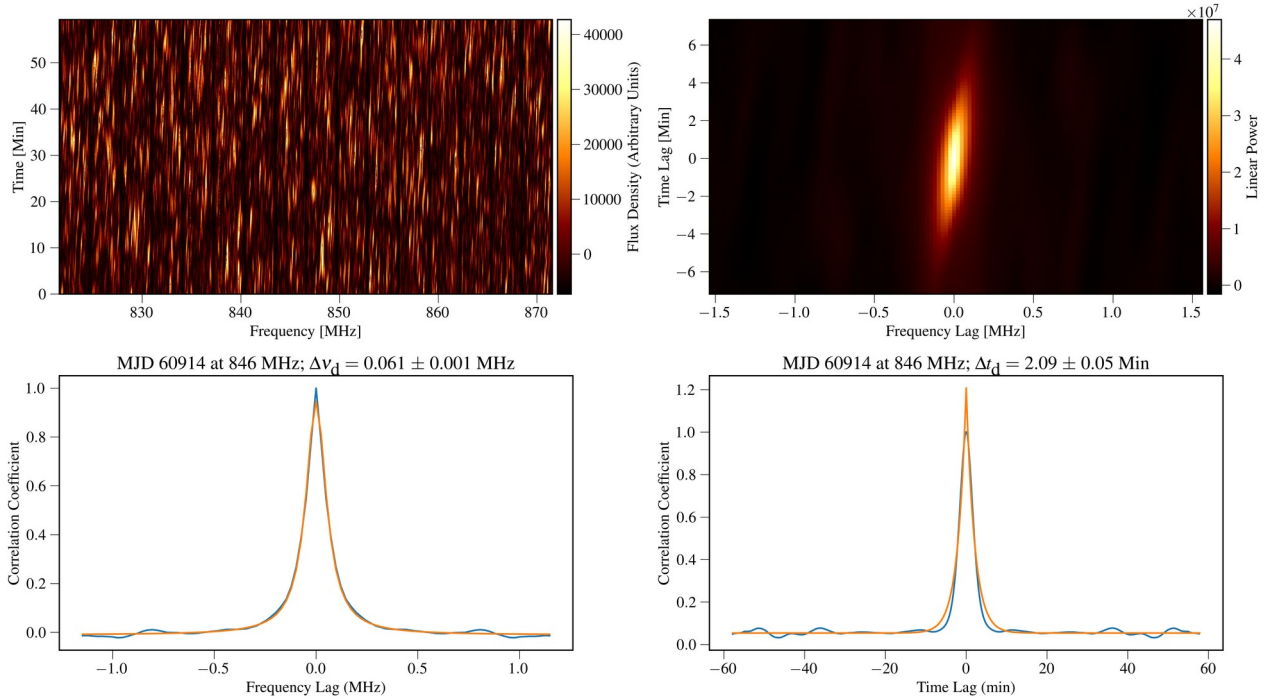

Figure 6: Dynamic spectrum of PSR B1937+21 using the PF800 receiver (upper left) and it's 2-D autocorrelation functoin (upper right). The lower panels show 1-D cuts through the 2-D ACF along frequency (lower left) and time (lower right), as well as fits to the 1-D functions. The characteristic width of the ACF is used to measure $\Delta\nu_s$ and t_{scint}.

Secondary Spectra

The secondary spectrum was created by taking the logarithm of the squared modulus of the dynamic spectrum's 2-D Fourier transform. The secondary spectrum contains structures, known as scintillation arcs, that form due to spatially localized structure in the IISM along the line of sight to the pulsar. These thin structures are responsible for the majority of the scattering observed in any given pulsar. Each image in the spectrum is the result of interference between scattered photons and line of sight photons, with the parabolic shape arising from the quadratic relation between the geometric time delay and the diffractive scattering angle. In contrast, if scattering was highly isotropic rather than localized, we would instead see a diffuse column of power extending along the delay axis.

Fig. 7 shows a secondary spectrum for PSR B1937+21 using the PF800 receiver. The left panel was produced using the same CSB dynamic spectrum as shown in the left panel of Fig. 5, while the right panel was produced with the lower-resolution dynamic spectrum shown in the right panel of Fig. 5. Scintillation arcs are visible in the higher-resolution data but are not visible in the lower resolution data. In Fig. 8 we show a secondary spectrum produced using the CSB at L-Band, in which the arcs are visible with higher S/N.

Scintillation arcs encode information about the geometry of the scattering medium relative to the pulsar and the Earth. In Fig 8 we show a fit to the scintillation arc observed in B1937+21 with the

L-Band receiver. Using the arc curvature of 8315 m^{-1} mHz^{-2}, and assuming a screen orientation that is perpendicular to the line of sight, and Earth velocity of 30 km/s, an ISM velocity of 10 km/s, and a pulsar distance of 2.6 kpc, we estimate that the scattering screen responsible for this arc lies 2.25 kpc from Earth. Previous estimates that used a now-outdated pulsar distance estimate of 3.5 kpc placed the screen at 2.9 kpc from the Earth, but in both cases the relative distance is consistent (83% vs 87% of the source distance). We note that there is a degeneracy between screen distance and orientation (hence our assumption of a perpendicular geometry) – this degeneracy can be broken with observations at multiple epochs.

Finally, the secondary spectrum power profile can be used to estimate the electron density wave number along the line of sight, which is a measure of ISM turbulence. A Kolmogorov medium in which small and large scale structure are statistically homogeneous, isotropic, and independent of each other should have a wave number spectral index of -2.33. In Fig 9 we show a fit to the power spectrum of the secondary spectrum shown below, and find a power law index of -2.46, indicating some deviation from a purely Kolmogorov turbulence profile, which is not uncommon.

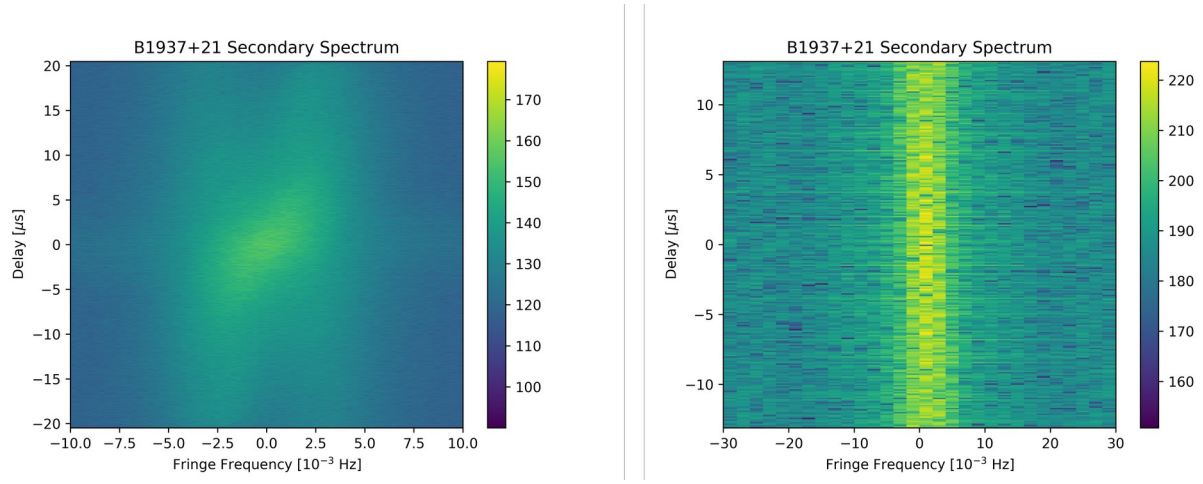

Figure 7: Secondary spectrum of B1937+21 obtained using the CSB at PF800. The panel on the left was produced from the fully resolved dynamic spectrum shown in Fig. 5. The panel on the right is the secondary spectrum that would be obtained using VEGAS with the best possible frequency resolution in coherent dedispersion mode. The scintillation arcs visible in the left cannot be seen on the right.

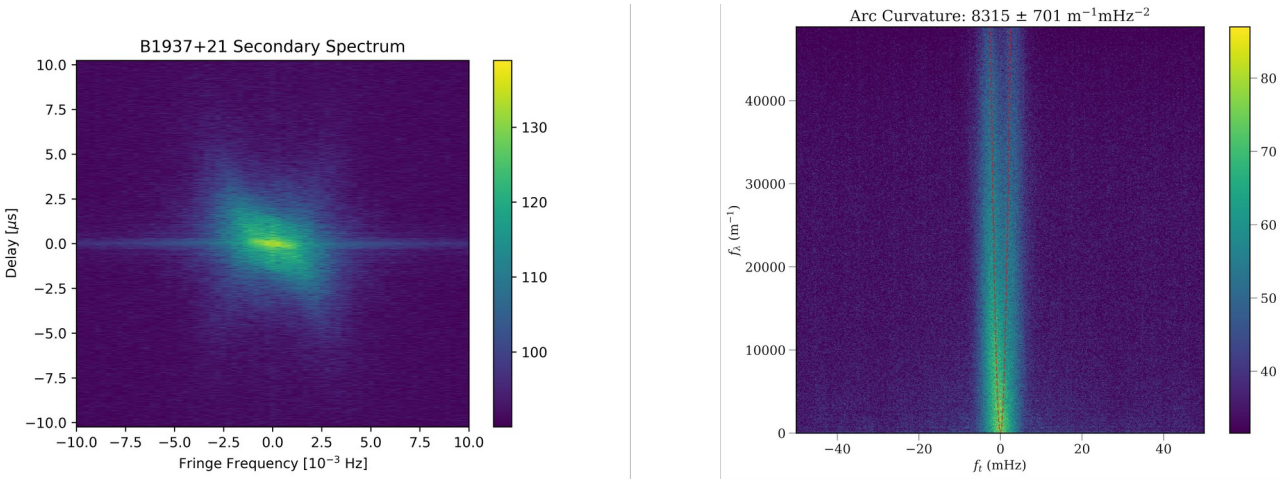

Figure 8: Left: Secondary spectrum of B1937+21 at L-Band produced with the CSB. Scintillation arcs are visible with higher S/N. Right: Fit to the scintillation arc, from which we infer a screen distance of 2.25 kpc.

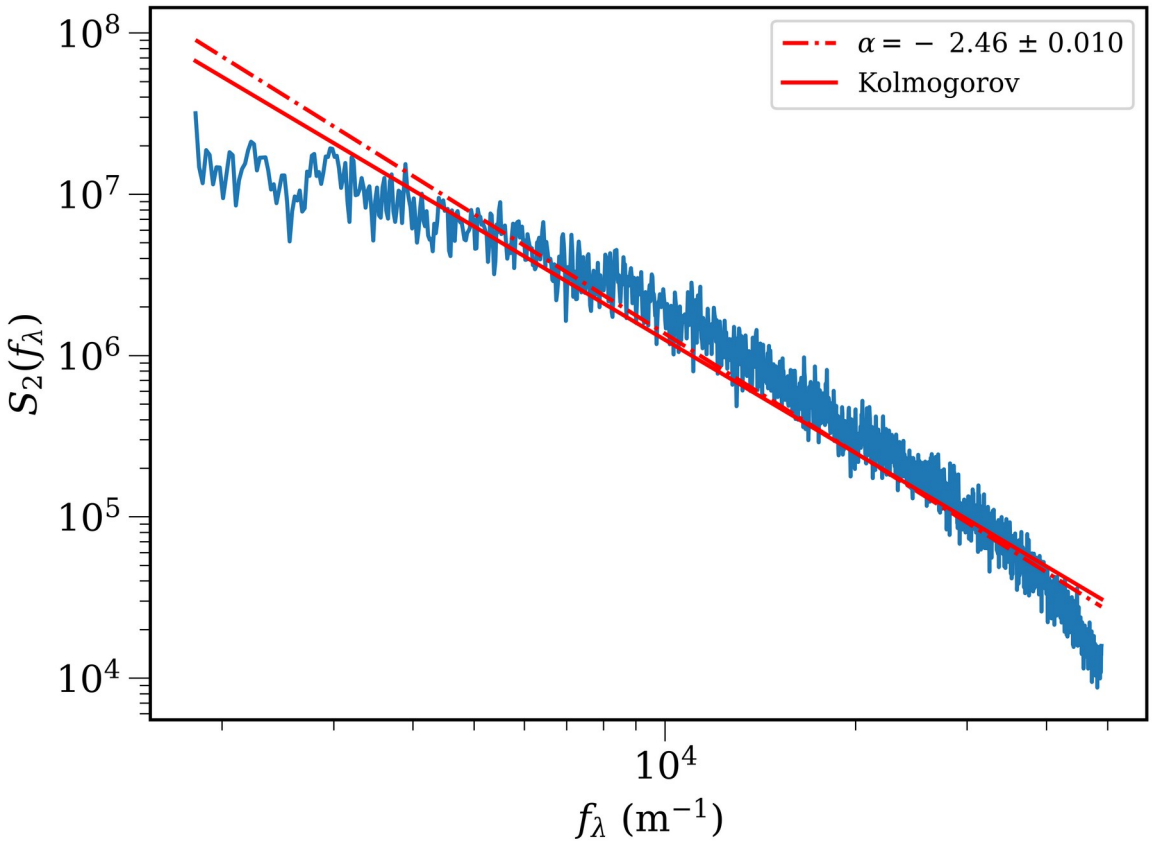

Figure 9: Secondary spectrum power profile (blue) and a fit (dashed line) showing slight deviation from a Kologorov profile (solid line).

Cyclic Spectrum Dynamic Wavefield

While dynamic and secondary spectra can be produced using traditional pulsar observing modes (though perhaps limited in frequency resolution), the cyclic spectrum is a unique data product that can only be obtained using cyclic spectroscopy. The cyclic spectrum presents the information contained within a time-averaged periodic signal as a function of observing frequency and harmonics of the periodic frequency (in this case, harmonics of the pulsar spin frequency). Through direct manipulation of the cyclic spectrum via algorithms such as the iterative deconvolution approach described in Walker, Demorest, & van Straten (2013), and implemented in the pycyc software package, one can deconvolve the complex impulse response of the ISM, known as the pulse broadening function, from the intrinsic pulsar signal. The pulse broadening function provides a direct means through which to measure the scattering delay present in an observation, in contrast with the much more common indirect estimate achieved through measured scintillation bandwidths in pulsar dynamic spectra. In the Fourier domain, the pulse broadening function becomes the transfer function, or frequency response, of the IISM.

Fig 10 shows the cyclic spectrum for B1937+21 obtained using the PF800 receiver. Power is visible in the cyclic spectrum in about the first two dozen harmonics across the full observing band.

By recovering cyclic spectra at intervals less than a scintillation timescale, one can stack recovered transfer functions in time to create an observation's dynamic wavefield, which constitutes the full, complex, phase-reconstructed, time-evolving signal the IISM imparts onto pulsar emission. The dynamic spectrum, which is the most commonly used data product through which pulsar based-studies of the interstellar medium are performed, contains the amplitude of this function only at the 0^{th} harmonic. Fig. 11 shows the dynamic wavefield of B1937+21 obtained using the PF800 receiver. From the cyclic phase we estimate a delay of 629 ns.

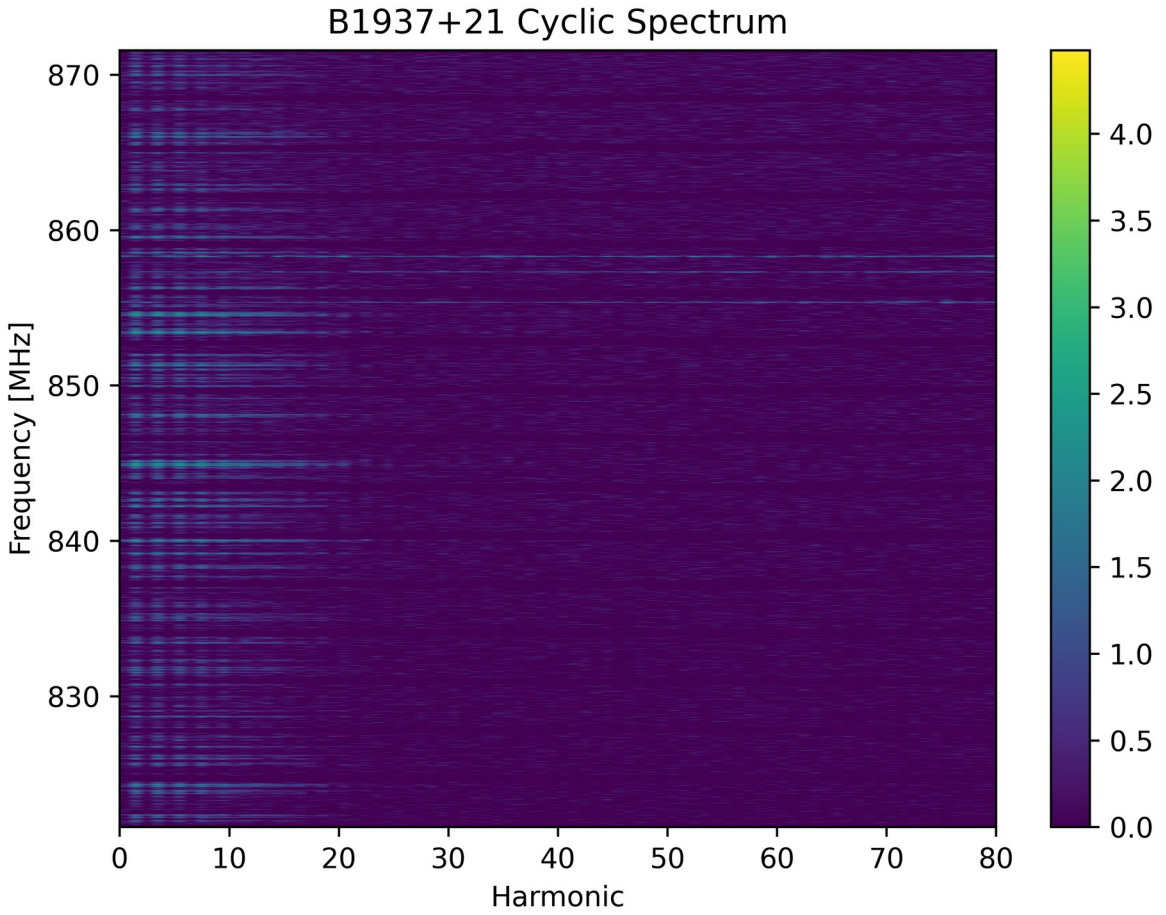

Figure 10: Cyclic spectrum of B1937+21 at 820 MHz. Power is visible in about the first 20-25 harmonics.

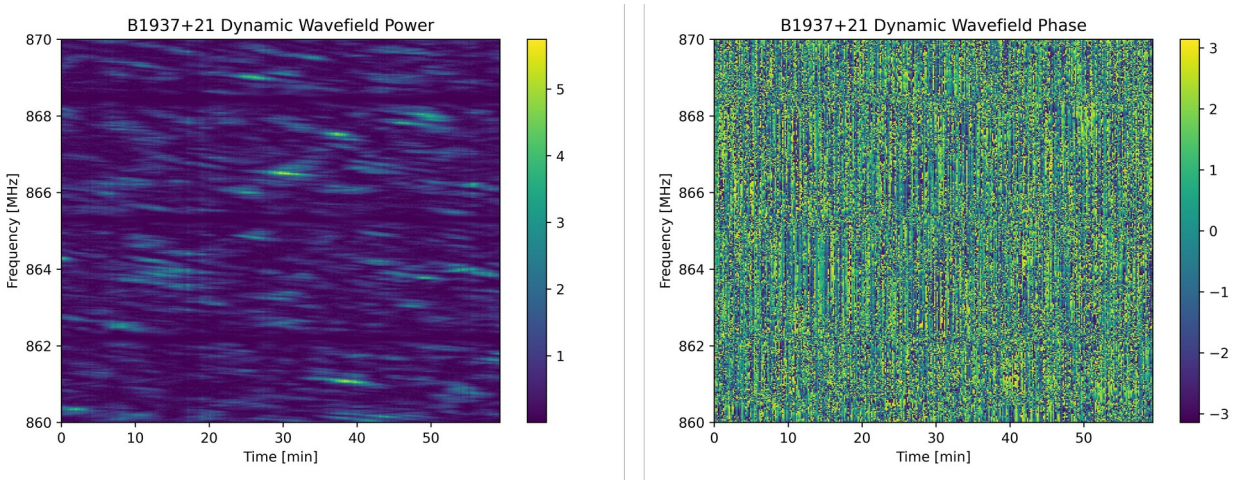

Figure 11: Dynamic wavefield power and phase for B1937+21 obtained using the PF800 receiver.

Operational Considerations

Errors and Aborted Scans

We observed using the CSB for a total of 28.7 hours. The CSB has proven to be very reliable from an operational perspective. During our first observing session and during maintenance testing before our UWBR observations, some processes had to be restarted following some earlier testing. Other than those initial startup issues (which did not reoccur during any other sessions), the M&C system aborted scans nine times, but none were a result of errors with the CSB (seven aborts were caused by VEGAS for reasons that are well understood; two were issued by other systems not related to the CSB). We encountered an error with a VEGAS shared memory parameter and had difficult balancing during the UWBR observing session but these were unrelated to the CSB.

During commissioning observations with UWBR we did encounter an operating system error that interfered with data acquisition. Specifically, a bug in a recently released version of the RHEL8 Linux kernel led to memory leaks and out-of-memory errors on the CS data processing nodes. As a result we were not able to acquire full flux calibration data in all UWBR modes, and we were not able to test all Cyclid processing parameter combinations. However, as noted previously, the fully calibrated data are in good agreement with VEGAS, as are the uncalibrated data. As such, we are confident that the CS system is operating well for UWBR modes, and we recommend releasing UWBR modes for general use. The kernel bug that led the out-of-memory errors has since been fixed, we have confirmed that the CS system now behaves normally in UWBR and non-UWBR modes, and we do not expect this issue to arise again in the future.

Automatic processing and quality assurance checks were all performed without any errors. All quality assurance checks passed with the exception of one scan, which was identified as by GBO staff as being in a failed state. Thus, the offline processing stage has proven to be very robust.

M&C Messages When In Error States

To demonstrate how the CSB responds when there *are* errors, we **intentionally** triggered certain failure modes to test whether M&C messages would appear in the CLEO Messages application. Specifically, we

1) Used Task Master to terminate the cycspecProcHost, cycspecProcess, and cycspecQualityCheck daemons. These processes mirror shared memory from VEGAS to the CSB, run Cyclid, and run the quality check scripts, respectively.

2) Prevented baseband data from being written.

3) Triggered a warning about disk usage by changing a configuration file so that the threshold for issuing a warning was very low (so as to trigger the warning even though the disks were not full).

4) Triggered a warning about dropped packets by changing a configuration file so that the threshold for issue warning was very low (so as to trigger the warning even though dropped packets were within normally acceptable levels).

In all cases M&C messages were issued and appeared in CLEO Messages (see Fig. 12). We thus conclude that errors during data acquisition are being properly caught and reported. We note that errors during data processing must be identified via Cyclops.

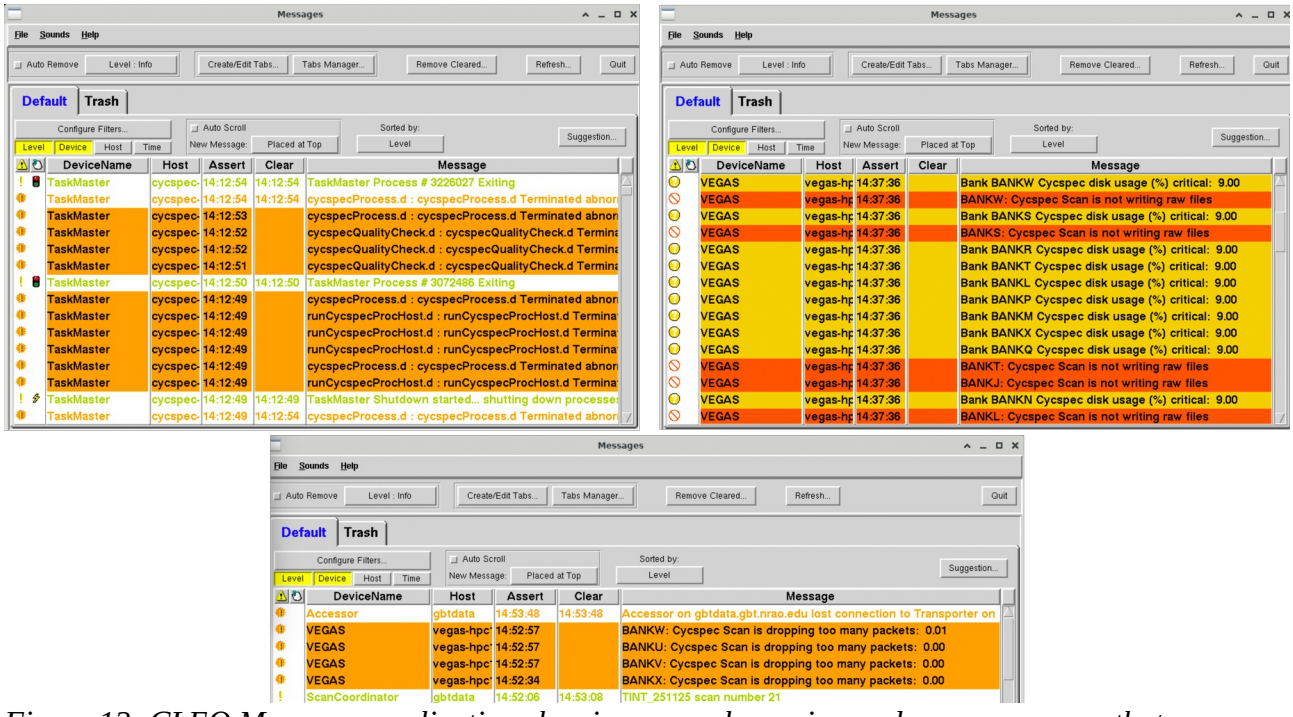

Figure 12: CLEO Messages application showing several warning and error messages that may appear during data acquisition. All the messages shown here were intentionally induced by interfering with normal operation of the system or by lowering the thresholds at which messages are triggered to demonstrate their functionality.

The Cyclops Browser-based Interface

We used Cyclops throughout commissioning observations to monitor data acquisition and offline processing, and we encountered no issues. The system updated in response to the start and end of scans and the start and end of processing in a timely fashion. Baseband data quality plots (see Fig. 13 for an example) were produced in a timely fashion and were useful for checking input power levels and internal VEGAS requantization scale settings.

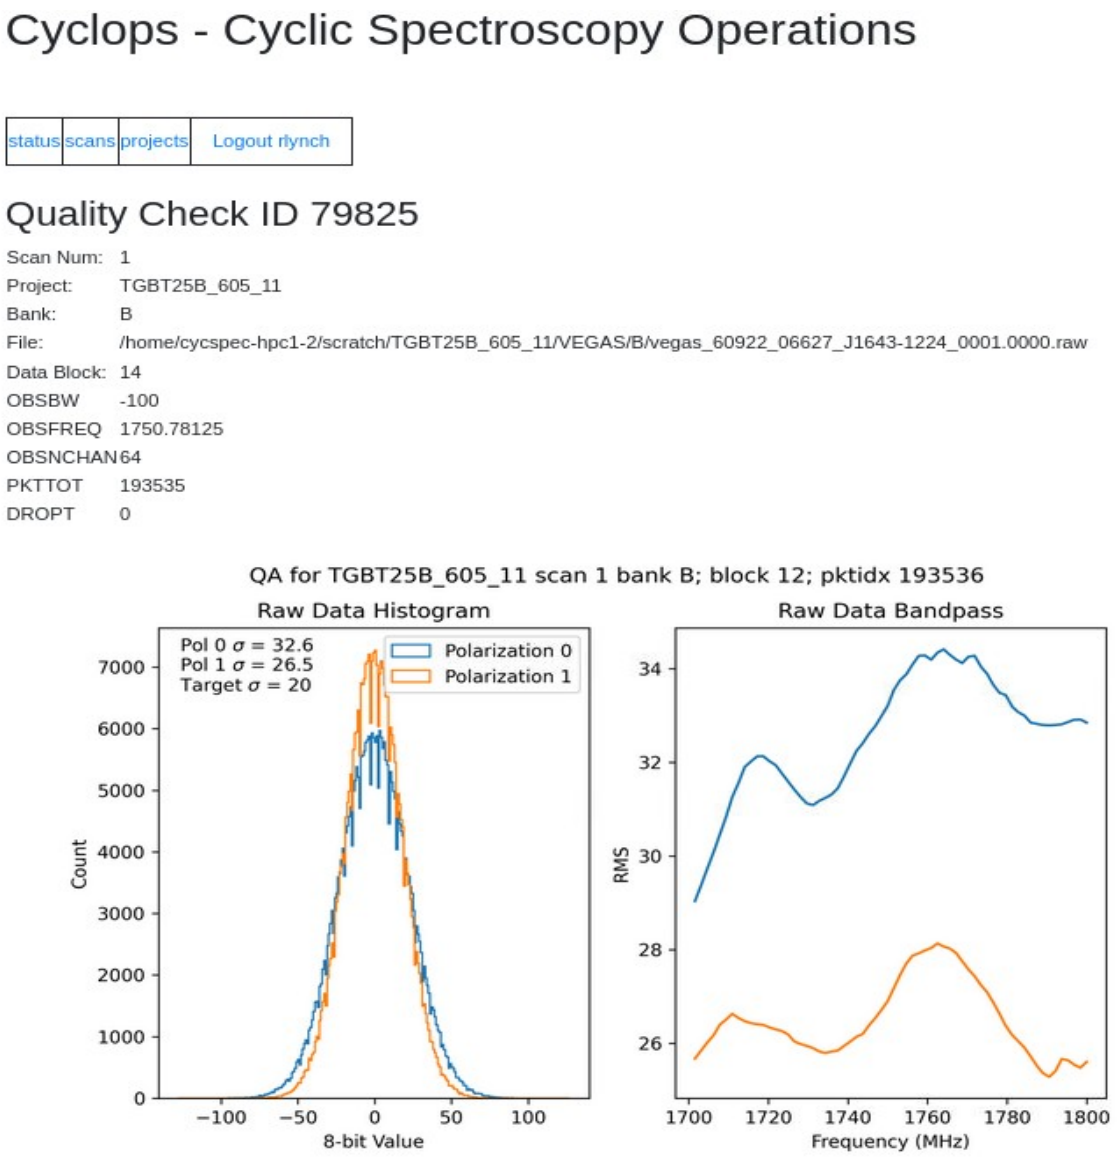

Figure 13: An example of the Cyclops baseband quality check page, showing a histogram of the 8-bit integer baseband data values and the uncalibrated bandpass for this data stream.

Conclusions

The CSB provides a new capability for the GBT and the first observatory-supported CS backend in the world. We have demonstrated that it produces data which agree with VEGAS in integrated

properties, but crucially the CSB provides capabilities that VEGAS does not. It provides better frequency resolution at an equivalent pulse phase resolution than can be achieved with VEGAS, and exceeds the maximum achievable frequency resolution that can be obtained with VEGAS when using coherent dedispersion. The cyclic spectrum is a unique data product that cannot be produced with VEGAS, and which encodes valuable information about the IISM and its effects on pulsar signals. We have demonstrated the ability of the CSB to produce high quality dynamic and secondary spectra, from which we measured scintillation bandwidths and timescales, scattering screen distances, and the turbulent power spectrum of the ISM. We further demonstrated the ability of the CSB to receover cyclic phase, from which we estimated scattering delays.

In this report we present representative data which we can compare with literature values, but we anticipate that NANOGrav will use the CSB for several pulsars in order to better characterize interstellar scattering delays, and we look forward to the CSB's unique capabilities being used for studies of the IISM.

Acknowledgments

This work was funded by the National Science Foundation as a supplement to the GBO Cooperative Agreement (AST-1928936). The National Radio Astronomy Observatory and Green Bank Observatory are facilities of the U.S. National Science Foundation operated under cooperative agreement by Associated Universities, Inc.

Appendix

Dynamic Spectra and ACF Fits

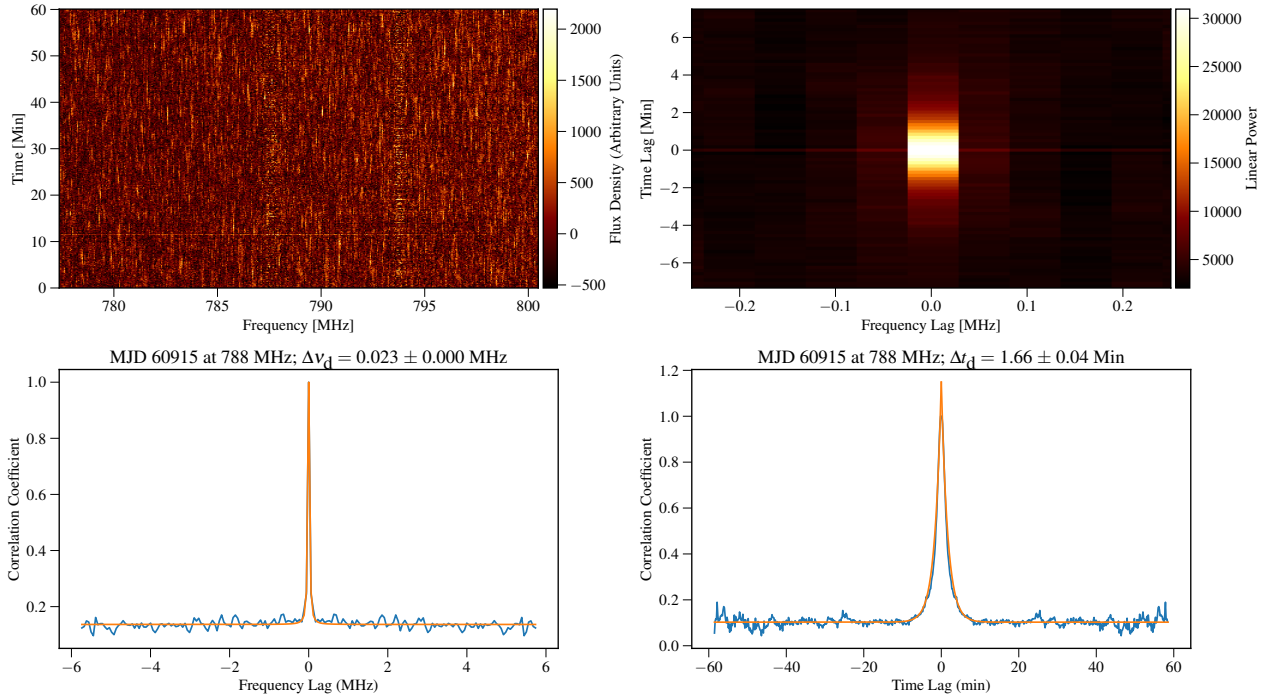

Figure 1: Dynamic spectrum and ACF fit of B0523+11 at 820 MHz.

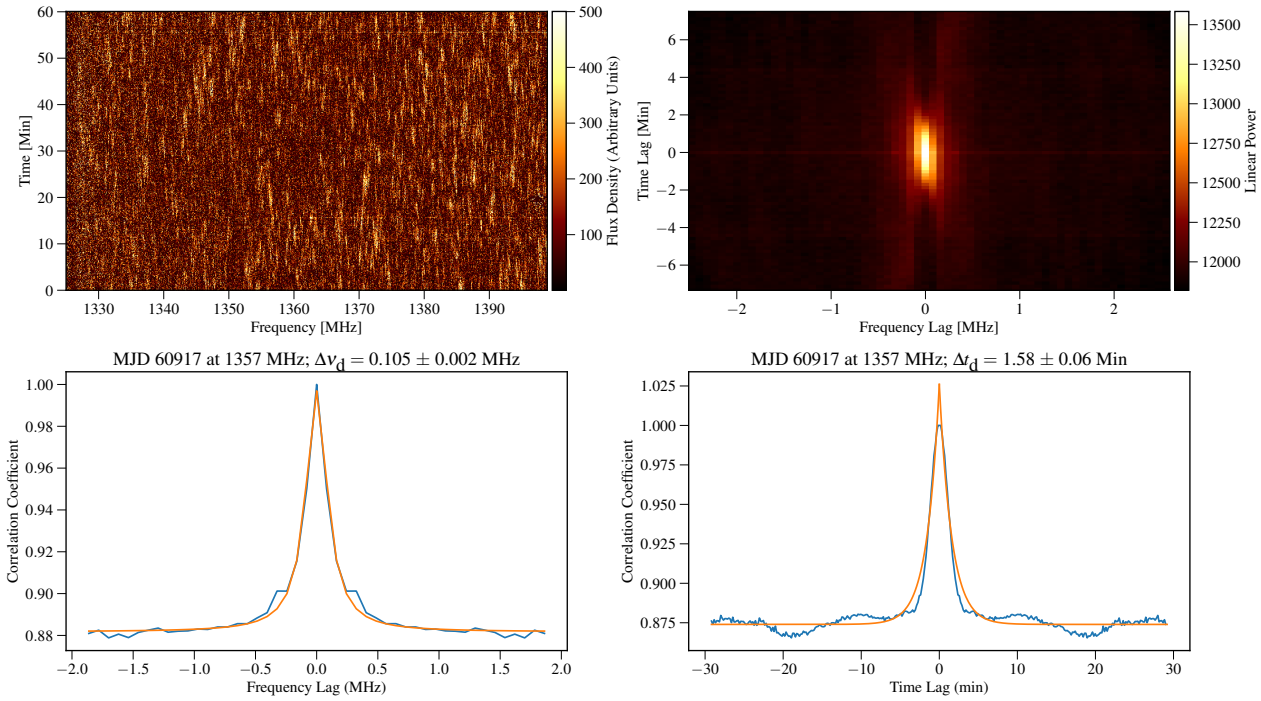

Figure 2: Dynamic spectrum and ACF fit of B0523+11 at L-Band.

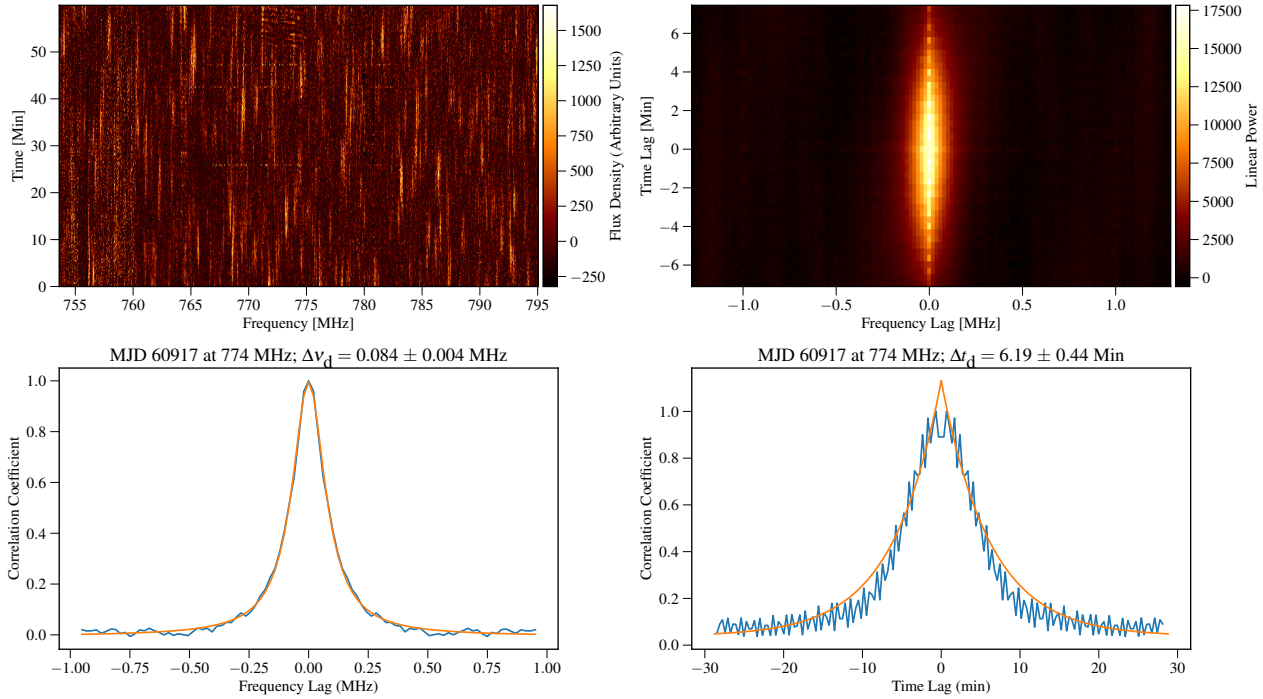

Figure 3: Dynamic spectrum and ACF fit of J1600-3053 at 820 MHz.

Secondary Spectra

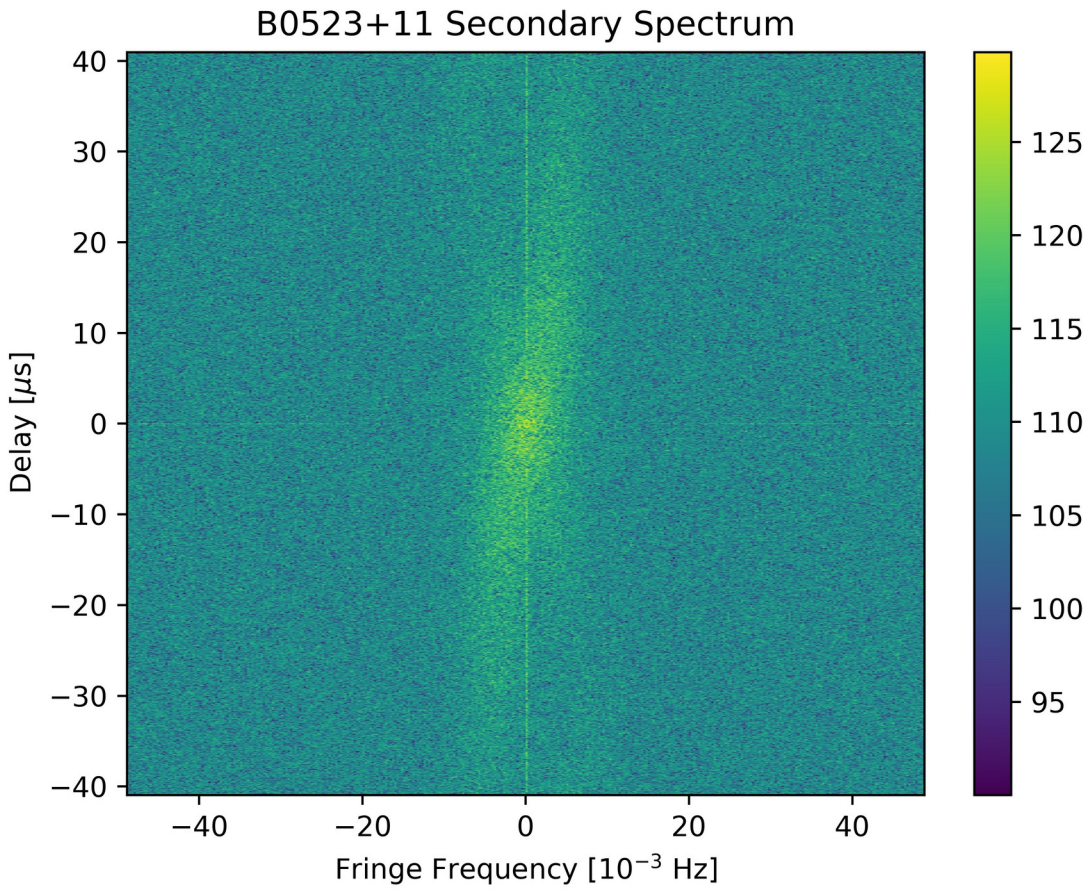

Figure 4: Secondary spectrum of B0523+11 at 820 MHz showing scintillation arc.

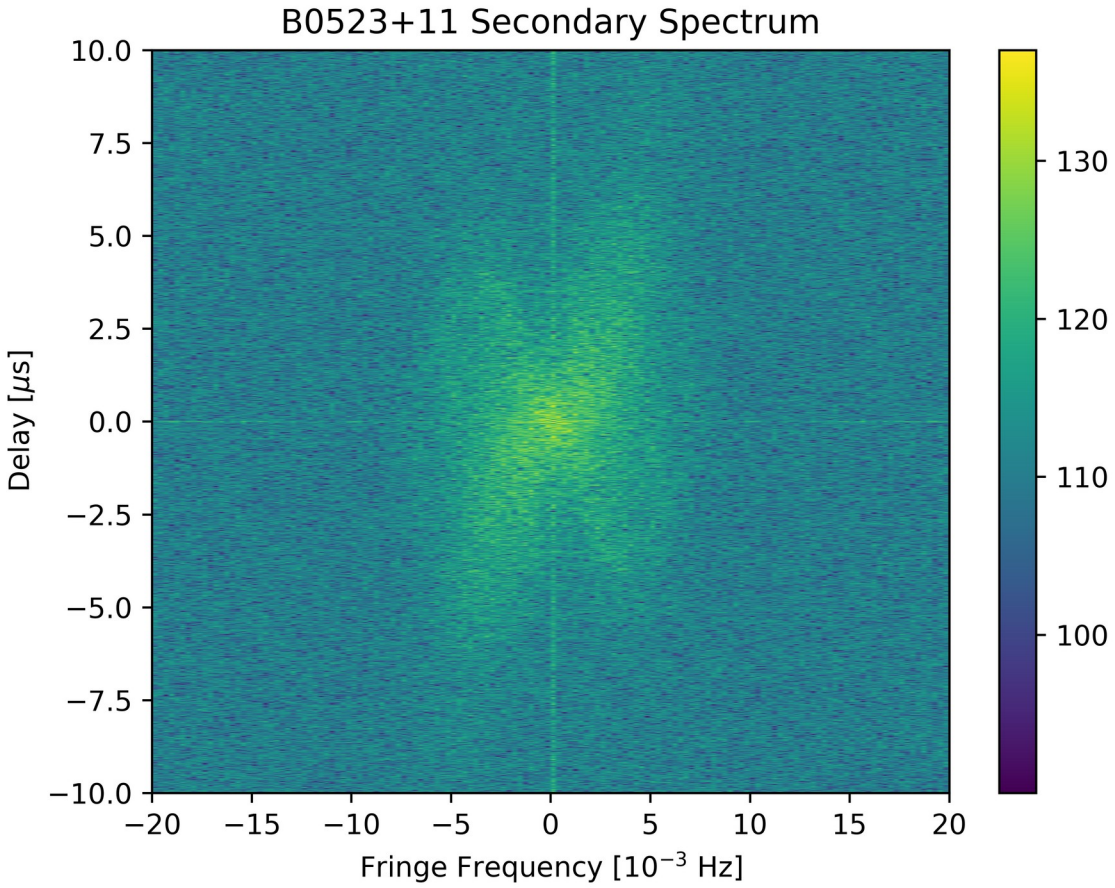

Figure 5: Secondary spectrum of B0523+11 at L-Band, showing scintillation arcs.