



Experimental RFI Mitigation of Pulsar and Spectral Line Data

Evan Smith, Ryan Lynch

Date: July 28, 2025
Keys: RFI Mitigation

Contents

1	Introduction	2
2	Methods	2
2.1	Spectral Kurtosis	2
2.2	RFI Mitigation Using Cyclostationary Signal Processing	3
2.3	Testing Methodology	4
3	Data	6
3.1	Spectral Line Targets	6
3.2	Pulsar Targets	8
4	Results	8
4.1	Spectral Kurtosis	8
4.1.1	OH Megamasers	8
4.1.2	Neutral Hydrogen galaxies	10
4.1.3	Pulsars	12
4.2	Cyclostationary Processing	14
5	Conclusions	15
6	Acknowledgements	16

Changelog

- 319.1** Smith et al. (2024-04-28) — Initial publishable version.
- 319.2** Smith et al. (2025-05-28) — Reorganization and edits.
- 319.2** Smith et al. (2025-07-25) — Final Draft for publication.

1 Introduction

In the current environment of increasing amounts of radio frequency interference (RFI), it is important to develop cutting-edge real-time RFI mitigation methods to preserve radio astronomical science. In order to deal with a more crowded spectrum, telecommunication companies are developing more complex ways to encode signals that present differently in the statistics of the data. In this memo we detail two experimental statistical radio frequency interference (RFI) mitigation techniques, the Spectral Kurtosis estimator (\widehat{SK}) and cyclostationary signal processing (CSP). Spectral Kurtosis detects RFI by estimating the Gaussianity of the pre-square law detected data and flagging when it is not white gaussian noise. Cyclostationary signal processing detects the periodic modulation of human-made signals. Both of these RFI detection methods work on baseband data that hasn't yet been accumulated or averaged over time, so that we can test what their efficacy would be in real-time without actually requiring the optimization necessary for such usage. Our data set includes observations of two HI galaxies, two OH megamasers, and two pulsars. We compare the science-ready data products (SRDPs) between the mitigated and unmitigated versions.

2 Methods

2.1 Spectral Kurtosis

Spectral Kurtosis is a computationally simple statistical technique to determine the presence of RFI. It is briefly described here, with more details found in (Gary et al., 2010; Nita, 2016; Nita and Gary, 2010b; Nita and Gary, 2010a; Nita, Gary, et al., 2007; Nita, Hickish, et al., 2016).

The \widehat{SK} estimator is applied independently in each frequency channel k and across M power spectral density (PSD) values, and is defined as

$$\widehat{SK} = \frac{MNd + 1}{M - 1} \left(\frac{M \cdot S_2}{S_1^2} - 1 \right) \quad (1)$$

with N as the number of pre-disk write accumulations in the spectrometer, and d as an empirically determined shape factor. N is essentially is the amount of spectra that are averaged together before the data is written to disk or we apply \widehat{SK} RFI detection. S_1 and S_2 are defined per spectral channel k as

$$\begin{aligned} S_1(k) &= \sum_{i=1}^M \left(\sum_{j=1}^N P_j \right)_i \\ S_2(k) &= \sum_{i=1}^M \left(\sum_{j=1}^N P_j \right)_i^2 \end{aligned} \quad (2)$$

Or for $N=1$:

$$\begin{aligned} S_1(k) &= \sum_{i=1}^M (P_i) \\ S_2(k) &= \sum_{i=1}^M (P_i)^2 \end{aligned} \quad (3)$$

where P_i, P_j are the power spectral density values. \widehat{SK} is a well-defined and bias independent metric for estimating how closely the raw time-series voltages follow a normal distribution. When data is affected by RFI,

the presence of a large non-Gaussian carrier or modulated signal can shift the time-series voltages away from a normal distribution. This change is the metric by which \widehat{SK} detects RFI. The statistical distribution of \widehat{SK} can be computationally modeled by a Pearson type III curve centered at 1, with a spread determined almost entirely by the accumulation length M . The formulas are given in Python form in the pip-installable package used for this work¹. The default option is to take a 3-sigma threshold level, which corresponds to a probability of false alarm (PFA) value of 0.13499%. The accumulation length should be large enough to avoid errors from small statistics without being too large that RFI is flagged on unacceptably long time scales. A reasonable choice of M will depend on the spectrometer characteristics and RFI environment.

Many spectrometers do not have the ability to output raw, unaccumulated spectra (Nita and Gary, 2010b). In the case that spectra are averaged before writing to disk, then $N > 1$ and each of the power values input to the algorithm is the inside sum of the unaccumulated power values. With a large enough N , the PDF of the \widehat{SK} estimator becomes that of a Pearson type VI curve. However, as described in Section 3, we used the Versatile Green Bank Astronomical Spectrometer (VEGAS) in its baseband recording modes, which output the unaccumulated channelized voltages and allows us to set $N = 1$. Any real-time system would work on the same channelized voltage data, so by these raw recording voltages to disk, we can simulate real-time operation and focus on comparing different algorithms. A future stage of the project is to optimize the algorithms so that they could run in real-time and would be part of a second-stage effort.

2.2 RFI Mitigation Using Cyclostationary Signal Processing

We also tested an RFI mitigation algorithm that makes use of cyclostationary signal processing (CSP) as alternative to SK. Cyclostationary processes have statistical moments that vary periodically. Such signals are common in telecommunication and arise naturally from data encoding schemes that modify the phase, amplitude, or frequency of a carrier wave. However, most astronomical signals are wide-sense stationary, i.e. their statistical moments are time-invariant. This makes CSP a potentially powerful means of distinguishing RFI from astronomical signals of interest. One important caveat is that pulsars are cyclostationary, since they are periodically amplitude-modulated. They thus represent a challenging but important use-case for CSP-based RFI mitigation.

Here we present the results of applying our CSP algorithm to baseband data collected on the millisecond pulsar J1713+0747. The algorithm is described in detail in Lynch et al. (in prep), but we provide a brief overview.

In order to use CSP for blind RFI mitigation, one must be able to detect cyclostationary signals with modulation frequencies that are not known a priori. We use the strip spectral correlation analyzer (SSCA; Roberts et al. 1991), which we implemented in Python. Specifically, we calculate the spectral coherence function at discrete spectral and modulation frequencies using both the non-conjugate and conjugate configurations. The spectral coherence is desirable because it is a normalized measure of the cyclostationary content of a signal, making it easier to set detection thresholds. Different modulation schemes have different signatures in the non-conjugate and conjugate configurations, making it important to calculate both.

The maximum value of the spectral coherence function for normally-distributed random values follows a generalized extreme value distribution. We flag data as containing RFI if the observed maximum value of the spectral coherence exceeds the 0.99-quantile for an extreme value distribution with shape parameters that we determined empirically for different SSCA parameters.

Our algorithm has two parameters in addition to the flagging threshold: M controls the spectral frequency resolution and N controls the total amount of data analyzed, which also sets the modulation frequency resolution. In **lsh24** we found that larger values of N tend to produce better results. However, when analyzing pulsar data it is important that the value of N does not exceed the number of samples corresponding to one rotation of the pulsar, lest the algorithm detect the cyclostationary nature of the pulsar and mistakenly flag it as RFI. The exact number of samples corresponding to a rotational period will depend on the sampling rate and the pulsar being observed.

As mentioned above, we analyzed data taken on the bright MSP J1713+0747. We used a VEGAS mode with 800 MHz of sampled bandwidth and 512 coarse frequency channels, corresponding to a sampling rate of 1.5625 MHz (or a sampling time of 640 ns). The rotational period of J1713+0747 is 4.57 ms, corresponding to about 7140

¹<https://github.com/etsmit/nettingi>

samples. Since our algorithm is currently only implemented for N that are powers-of-two, we used $M = 32$ and $N = 4096$.

2.3 Testing Methodology

We implemented the algorithms in Python, in a pip-installable package¹. As expanded on in later sections, we then run RFI detection and flagging on each block, replacing any flagged data, and write that replaced block back into the a new file, preserving the headers and metadata. The flagging result is a boolean-type mask of 0's and 1's. We assumed that any RFI in one polarization will also be in the other polarization, so we applied the union of the flags for both polarizations to each independent polarization. After flagging arrays have been generated, we used two different methods to replace flagged data, depending on the type of RFI flagged and the scientific objective of the data.

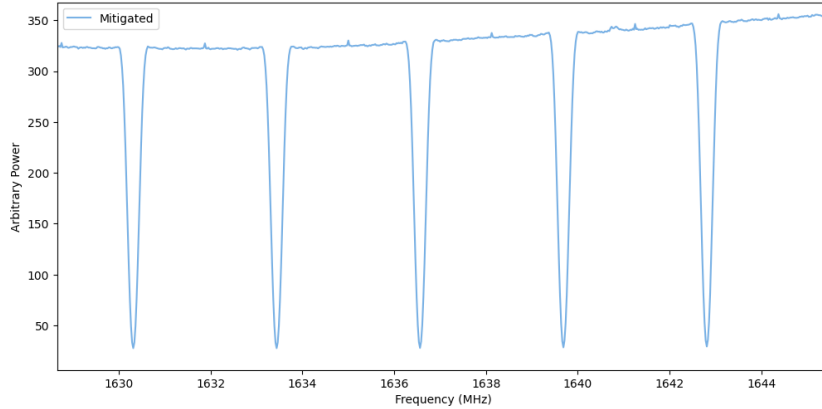


Figure 1: A few coarse channels with 3.125 MHz width that have been fine-channelized to 24 kHz. The scalloping is caused by the poly-phase filterbank sensitivity dropping off between channels to minimize leakage.

For the pulsar data, we replaced flagged data with zero-kurtosis white Gaussian noise that has been modulated to mimic the frequency response of the polyphase filterbank channelization used in the VEGAS backend. The frequency response in VEGAS data is tailored to reduce leakage between coarse channels. An example of fine-channelized data is shown in Figure 1. The mean of the noise is set to zero and the standard deviation is taken from unflagged bins in the same and adjacent frequency channels across the whole data block, since the bandpass does not change on the timescale of a single data block, which about one-third of a second. Data blocks are described in more detail in Section 3. A variation of this method only takes unflagged bins in nearby frequency channels only from the same time index. This avoids “averaging over” the time variable emission of interest for pulsars, especially in the case of millisecond pulsars.

Another alteration to this method involved placing a physical absorber over the feed horn, with a short scan using each pertinent VEGAS baseband mode. This absorber data is representative of the frequency response of the receiver only, and does not contain RFI or atmospheric emission. For each frequency channel, we determine the variance of the complex voltages in the absorber spectrum, perform a simple scalar fit to the variance of the unmasked data, and use those scaled variances to generate zero-kurtosis Gaussian noise. This way, the remaining RFI in partially flagged channels will not contaminate the noise generation. However, we found that the absorber caused a frequency-dependent impedance mismatch with the receiver that caused small but meaningful changes in the total bandwidth instrument response compared to on-sky observations, which made the simple scalar fit inaccurate. An example is shown in Figure 2, with various absorber configurations. These configurations include covering the receiver feed horn radome, hovering the absorber about 6 inches above radome, and angling the absorber at about 30 degrees off of the feed horn². More work is needed to find ways of using these absorber data.

²These different configurations were chosen to test mitigating a certain harmonic or self-reflection that appeared in the “Completely Covered” case, and were secondary to the purpose of using the absorber model.

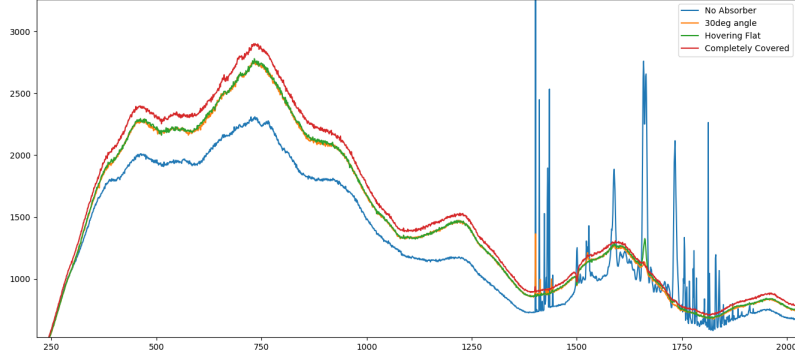


Figure 2: Comparisons of L-Band frequency response with a few different absorber configurations.

For the spectral line data, we replaced flagged data by NaNs at the coarse channel level and proceeded to fine channelize to a 24 kHz resolution, as described below. Given the nature of the polyphase filterbank routine used and `NUMPY.FFT.FFT`, we needed to extend the flags slightly to avoid inputting NaN values into the channelization. This increased the flagging percentages by a few tenths of a percent at most, but also allowed for the recovery of the OH signal buried in the Iridium RFI.

Other options for data replacement that were developed but not used included replacing flagged data with zeros and copying previously unflagged data from the same frequency channel. Replacing with zeros would result in heavy artifacts during pulsar analysis, and replacing with previously unflagged data was incredibly unoptimized due to the corner cases at the time and frequency edges of each block, as well as when an entire channel was flagged at 100%.

After data replacement, we ran identical data reduction techniques on both the original and mitigated datasets and used a series of metrics to judge if the RFI mitigation improved the scientific results. For the spectral line targets, the data was fine-channelized to 24.4 kHz, averaged over the whole data file, and inspected by eye to make sure the scientific signal of interest was not flagged or reduced in any way. We also inspected the quality of the spectrum and the removal of RFI.

For the pulsar targets, we ran the PRESTO `rfifind` procedure and compared the masks for both data sets. The metric for success is that if our algorithms improved the data, `rfifind` should detect less RFI in the mitigated data set. Next, we searched the data for pulsar candidates in 4 different variations: with or without the `rfifind` mask, and with or without RFI mitigation. For improved data, we expect to see a smaller amount of spurious candidates in the mitigated data regardless of the application of the `rfifind` mask, as well as seeing a smaller amount of spurious candidates when only our techniques are applied, compared to when only `rfifind` is applied. An improved dataset will also feature a higher SNR of the pulsar detection in the search results. The next check we made was to take a Fourier transform (FFT) of the time series data. The pulsar and its harmonics will clearly show at the correct frequencies, as will any RFI. A well-mitigated data set will have less spurious non-pulsar signals. After that, we compared the folded and integrated pulse profiles. For these, a simple SNR was calculated from the maximum height of the pulse and the standard deviation of the off-pulse noise. We expect RFI mitigation to at least not diminish the pulse profile strength, and to reduce the amount of off-pulse noise, thus increasing the SNR. Finally, we calculated the pulse time of arrivals (TOAs) and compared their statistics. We measured the RMS of the residuals, which are the TOA values with the model subtracted away. With less RFI, we should see less outliers and the TOAs will be more accurate, leading to smaller RMS values. We then compare the mean of the uncertainties on the TOAs. Again, with sufficient RFI mitigation, the uncertainties will be smaller. We run the previous two tests on the unmitigated and mitigated data, with and without the standard PSRCHIVE RFI `paz` mitigation applied, leading to four different states to compare.

We ran this process for many different algorithm input parameters. For \widehat{SK} , these are the accumulation length M , multiscale shape, and the flagging threshold. For CSP, these are the spectral frequency resolution M , window size N , and flagging threshold. The data replacement scheme chosen is independent of the detection algorithm.

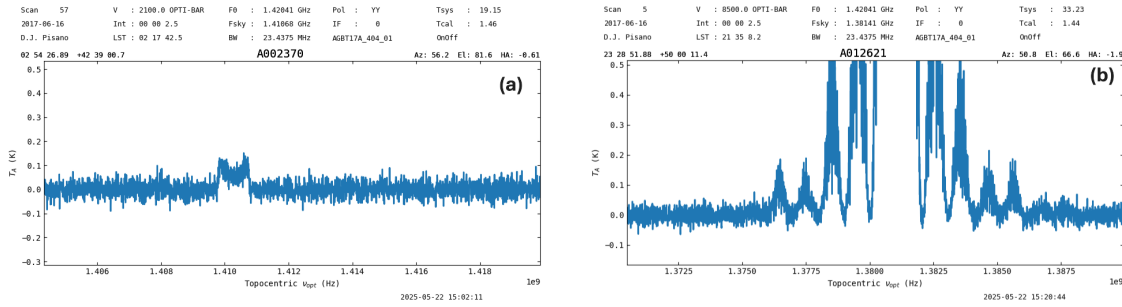


Figure 3: Reference spectra of (a) UGC 2370 and (b) CGCG 536-030 made by the project GBT17A-404. The CGCG 536-030 signal is entirely washed out by the periodic GPS-L3 RFI.

3 Data

To test the ability of our algorithms to mitigate RFI while leaving real astronomical signals untouched, we took data with the GUPPI³ and VEGAS backends on the Green Bank Telescope (GBT) as part of projects TGBT20A-505, TGBT22B-601, and TGBT23B-610. Since we took this data with the intention of testing the viability of various real-time RFI mitigation schemes, we used the pulsar baseband modes, which do not accumulate or average in any way and instead output the non-squared channelized voltages. Using this data, we can test the effectiveness of algorithms as they would be implemented in a real-time application without yet needing to optimize them to run at real-time speeds. Unfortunately, this means that the data rates are up to hundreds of GB every few minutes and we had to use short scans in order to keep our storage requirements reasonable. Fortunately, most of the RFI we captured is persistent and easily detected on such short timescales.

The contents of each data file are organized as a series of headers and data blocks. The headers are simple text and contain metadata for the following data block, such as timestamps, telescope location, and spectrometer information. The data block is binary encoded and contains a 3-dimensional array for frequency, time, and polarization. The exact size of each data block is specific to the sampling rate and number of frequency channels, so in order to avoid mismatches at the end of the block, we were careful to use \widehat{SK} flagging cadences that evenly divided the time length of each block.

The raw data was entirely processed offline, allowing us to simulate real-time processing while changing RFI mitigation schemes as we saw fit. The GUPPI raw files were read into Python using the `GuppiRaw` class of the `BLIMPY` package⁴ courtesy of the SETI group at UC Berkeley. The values in the array were channelized voltages - already Fourier transformed, but needing a magnitude calculation to be considered power values for the \widehat{SK} calculation.

Several scientific targets were chosen such that they had a large enough signal to detect on short time-scales and had a high chance of being corrupted by RFI. These targets are described in the next two subsections.

3.1 Spectral Line Targets

We observed two types of spectral line targets. All targets were chosen to be in or near heavily RFI-affected parts of the spectrum, and bright enough to be easily observable in a few minutes.

We observed the HI spectra of two nearby galaxies, UGC 2370 and CGCG 536-030, that are respectively redshifted to 1410 MHz and 1393 MHz. This places them near the 1381 ± 4 MHz range which is also occupied by the GPS-L3 Nuclear Detonation Detection System, an alert system for surface nuclear detonations. The system sends out periodic tests at 1381.05 MHz, which drown out any HI signal in that frequency range. We chose a

³<https://safe.nrao.edu/wiki/bin/view/CICADA/GUPPISupportGuide>

⁴<https://github.com/UCBerkeleySETI/blimp>

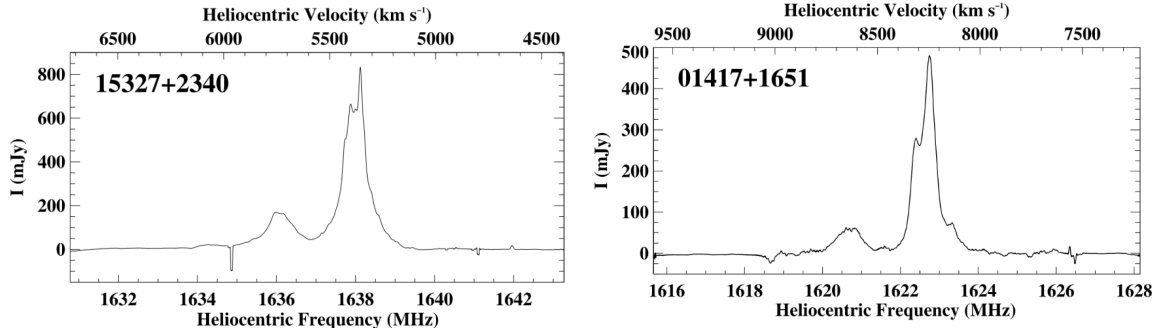


Figure 4: Reference spectra of the OH masers in (a) Arp 220 and (b) III Zwicky 35, taken by (Robishaw et al., 2008) with the Arecibo telescope. The III Zwicky 35 falls entirely in the Iridium band, and was recovered by using the median rather than averaging over time.

200 MHz bandwidth to include many other sources of RFI, including other GPS frequency bands. Figure 3a shows UGC2370, and Figure 3b shows GPS-L3 washing out any potential signal. These targets were picked from archival data courtesy of GBT project AGBT17A-404.

We observed two bright Hydroxyl Masers in the nearby galaxies III Zw 35 and Arp 220 (Robishaw et al., 2008). Masers emit coherent light which, if strong enough, can supposedly be detected by \widehat{SK} . The ability of \widehat{SK} to detect these science sources is an important metric for flagging RFI without removing celestial signals. The maser in III Zw 35 is redshifted to 1638 MHz, which is right outside of the 1620-1626 band that is occupied by the Iridium global satellite phone service. The maser in Arp 220 is redshifted slightly further to 1623 MHz, which means it is almost totally covered by RFI. Arecibo detections of the two masers are shown in Figure 4. A close-up spectrogram of the Iridium RFI is shown in Figure 5.

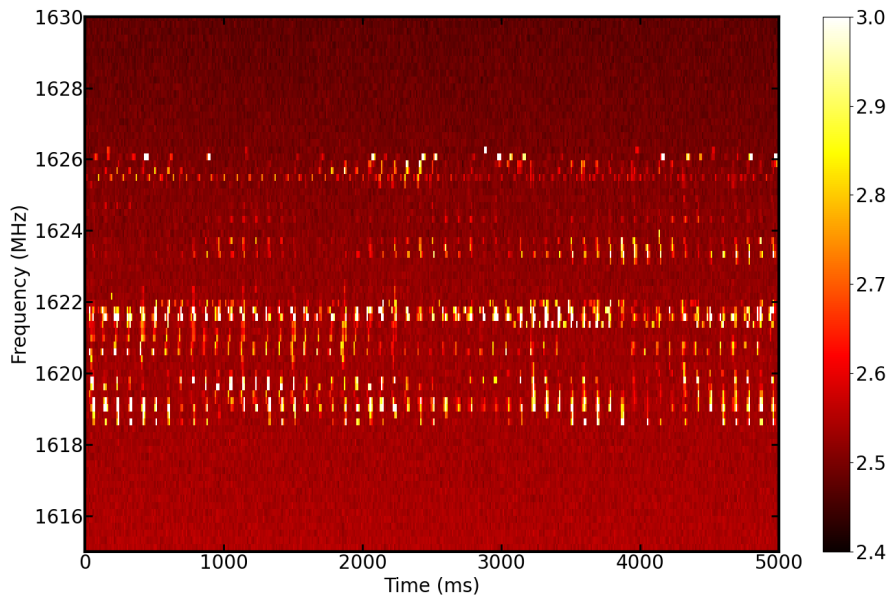


Figure 5: Example of the time and frequency variance of the Iridium satellite RFI. In a single frequency channel, satellites are expected to transmit for 8.28 ms every 90 ms for a 9.2% duty cycle. The modulation scheme is a 50 kbps differentially-encoded quadrature phase-shift keying (QPSK) scheme.

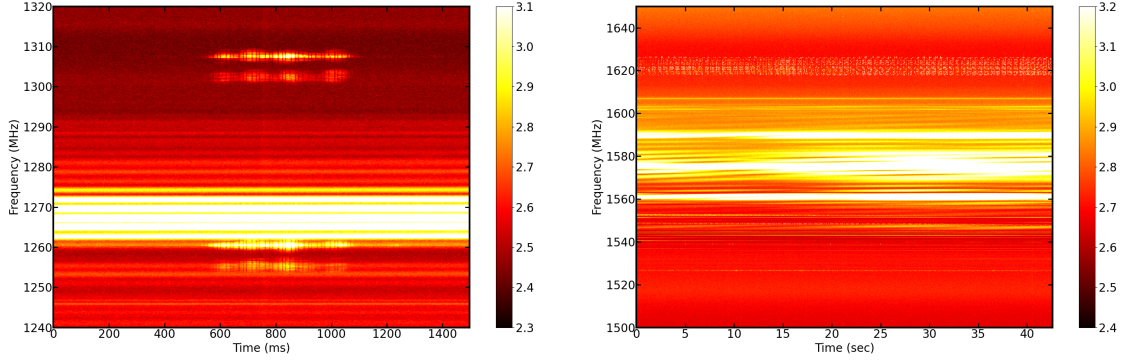


Figure 6: Examples of RFI encountered in our wideband pulsar observations. On the left, the FAA radar located in Bedford, VA is seen every 12 seconds and we detect the dual-frequency emission and response. The wide persistent RFI at 1268 MHz is likely from a satellite system that tracks bird migrations. On the right, there are several features. The large, persistent features at 1560, 1575, and 1590 MHz are associated with radio-navigational systems such as GPS and GLONASS. The more narrow-band RFI between 1540 and 1560 MHz is mostly due to meteorological satellite data downlinks, and the complex RFI in the 1616-1626 MHz band is the Iridium RFI.

3.2 Pulsar Targets

We also observed three pulsars as their emission is both broadband and periodic, different from our spectral line targets. The first pulsar is J1713+0747, a very bright millisecond pulsar with a period of 4.57 milliseconds and a dispersion measure (DM) of 15.992196. The second pulsar is B0329+54, a non-recycled pulsar with a period of 0.715 seconds and a DM of 26.833. The final pulsar is B0335+54, another non-recycled pulsar with a period of 0.15 seconds and a DM of 57.142.

Each pulsar was observed with a center frequency of 1500 MHz, a bandwidth of 800 MHz, and a series of different total channel numbers starting from 32 and up to 4096, in steps of powers of 2. The scan lengths were all at least two minutes, which, given the flux density of these pulsars, is sufficient to derive pulse profiles and acquire TOA residuals. These relatively wideband observations captured the abundance of RFI at around 1500 MHz due to meteorological satellites, as well as radar pulses from the Bedford FAA station at 1255 MHz and 1304 MHz and the GPS L1 signal at 1575 MHz. See Figure 6 for examples.

4 Results

Here we summarize the results from \widehat{SK} and CSP mitigation and data replacement for the targets described above.

4.1 Spectral Kurtosis

4.1.1 OH Megamasers

As mentioned above, OH Megamasers at 1665/1667 MHz can serve as a good spectral line comparison to the HI targets as they result in a similarly shaped detection in a spectrum but they are caused by coherent emission, not thermal emission which results in Gaussian noise.

Arp 220 is not redshifted into the Iridium RFI range, and the results are shown in Figure 7. Much of the persistent RFI at lower frequencies is completely flagged, as evidenced by the lack of a mitigated spectrum plotted in the ranges 1560-1567 MHz, 1570-1582 MHz, and 1585-1595 MHz. The RFI at 1605 has been completely missed by \widehat{SK} , but the Iridium is heavily flagged and reduced in power. The OH detection is marked at 1638 MHz, and

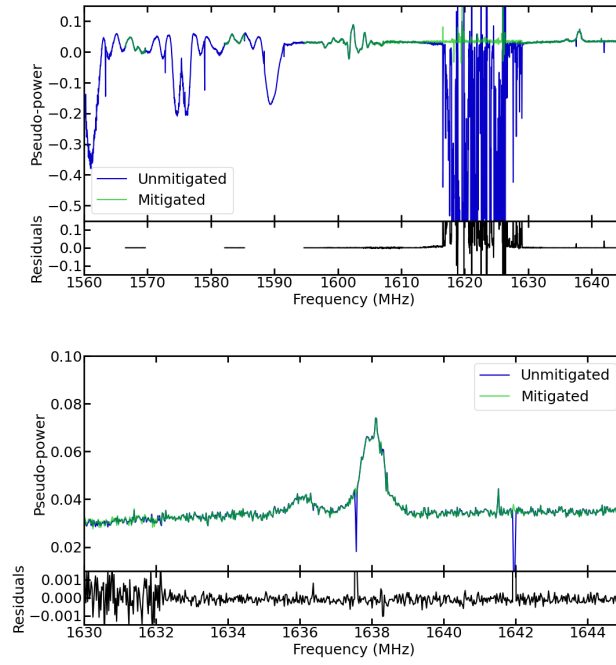


Figure 7: Top: RFI in the vicinity of Arp220, and the detection at 1638 MHz marked by (a) This data was taken with 64 coarse channels across 200 MHz of bandwidth and fine channelized to 24.4 kHz. Bottom: Zoomed in on the detection.

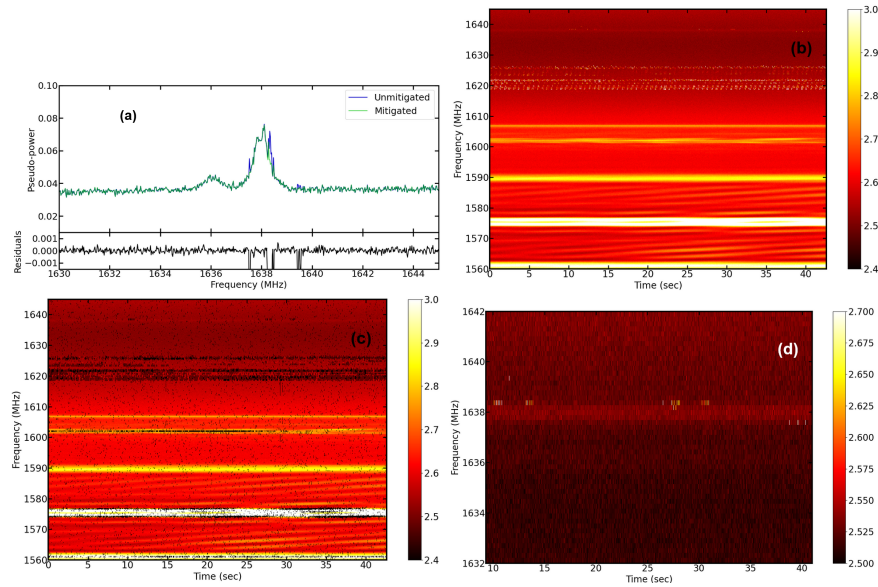


Figure 8: (a) Arp 220 OH detection, with an original number of coarse channels of 1024 fine channelized to 24.4 kHz frequency resolution. (b) Spectrogram of the dataset. (c) Flagged spectrogram of the dataset, with flags denoted as black boxes. (d) A focus on the broad, dim Arp 220 detection at 1638 MHz, with the intermittent RFI showing up at the edges of the signal.

shown more clearly in the bottom of the figure. There are no large features in the residuals except for the large spike at 1637.5 MHz and a slight difference in power at 1638.5 MHz.

When comparing to the Arp 220 detection with 1024 coarse channels in Figure 8, there is a small amount of RFI at 1638 MHz that is flagged. The source of this RFI is extremely intermittent and unknown. Despite the RFI, the Arp 220 OH signal remains similarly undetected, even when it is strong enough on a coarse channel resolution to be detectable.

The OH signal from III Zw 35 was clearly revealed in the mitigated data, despite being completely obscured by Iridium RFI in the unmitigated, as shown in the top panel of Figure 9. While we did not perform flux or antenna calibration during these observations, the strength of the signal is roughly 60% of the strength of the Arp 220 GBT detection, which matches with the relative strengths derived from the reference spectra in Figure 4.

The bottom panel of Figure 9 demonstrates the difference between using the Gaussian noise replacement for flagged data versus blanking data. Coarse channels that have been completely flagged do not show up in the masked spectrum, as evidenced by the gaps in the mitigated spectrum. These channels in the noise-replaced data are also completely flat and devoid of any features, including both RFI and potential spectral lines.

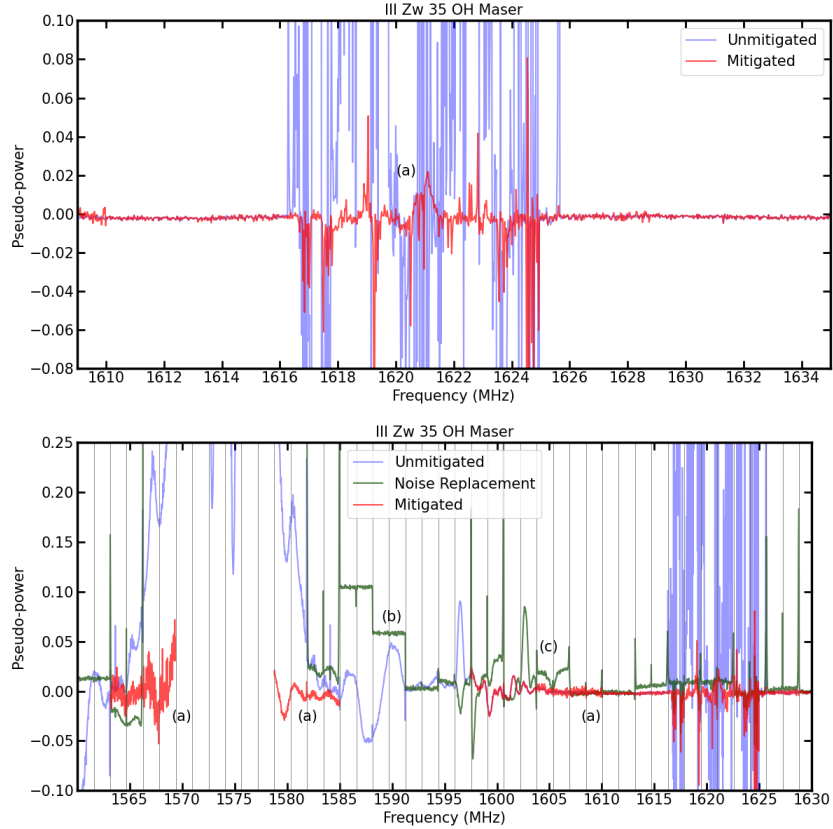


Figure 9: Top: III Zw 35 OH detection marked by (a). Bottom: Comparison of using Gaussian noise replacement vs. masked data during the fine channelization stage. (a) Higher noise segments due to less unflagged integration time in a single coarse channel. (b) Fully flagged channels end up completely flat and devoid of any features. (c) Partially flagged coarse channels result in inverted or different RFI features.

4.1.2 Neutral Hydrogen galaxies

\widehat{SK} mitigation did not flag the H I signal from the bright galaxies we observed. Figure 10 shows the relatively dim H I signal compared to the GPS-L3 RFI signal nearby. However, as we show in Figure 11, the GPS-L3 signal

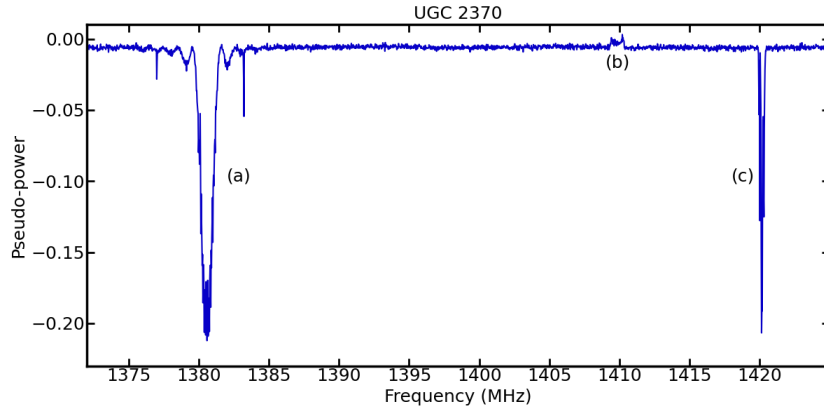


Figure 10: (a) GPS-L3 interference, (b) UGC 2370 double-horned profile, (c) Milky Way HI emission

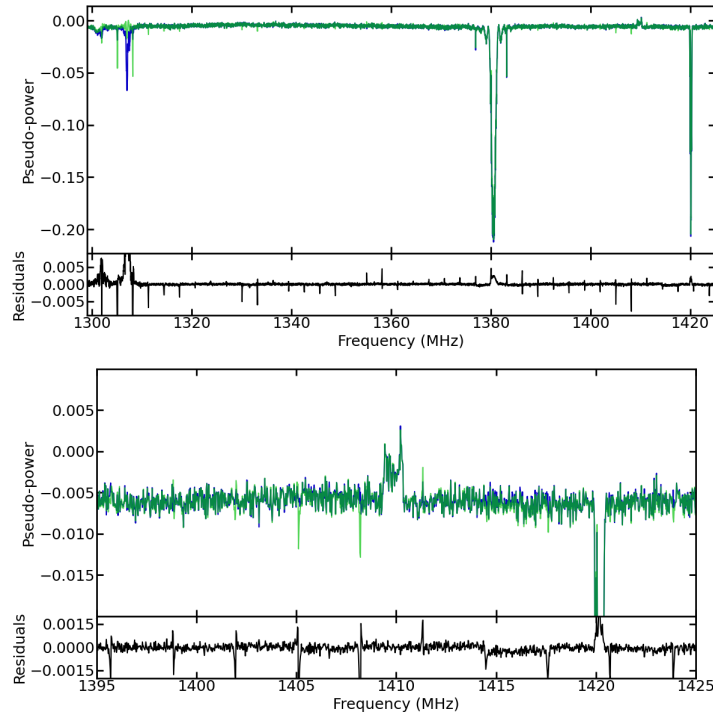


Figure 11: Top: Full spectrum of UGC2370, including the flagged DTUSat system at 1302 MHz, GPS-L3 RFI at 1381 MHz, UGC 2370 at 1410 MHz, and Milky Way. Bottom: Close-up of UGC 2370.

is also not flagged more than 2.5%, at least in this case. No other \widehat{SK} configuration of M or multi-scale kurtosis performed any better. The primary reason for this is that GPS signals are modulated heavily to model Gaussian noise. This holds true when we attempt to flag the other GPS bands as well. Work is ongoing to determine if it is possible for \widehat{SK} to detect GPS signals at all. The bottom of Figure 11 shows that there is no residual signature on the UGC 2370 signal, which means it is invisible to \widehat{SK} mitigation. The 3.125 MHz regularly spaced spikes along the residuals plot on both panels are a holdover from the Gaussian noise replacement scheme described in Section 2.3 and the boundaries of the coarse channels.

In some cases where the Milky Way HI emission was particularly strong and covered a few coarse channels, it

would be partially flagged. While the emission is clearly Gaussian in nature, in these cases the transition from cold sky to bright HI would be significant enough across the width of a single coarse channel that the Gaussian statistics would be a mix of different variances, and thus trigger \widehat{SK} flagging. Hence, care should be taken when using \widehat{SK} and observing Milky Way emission near the plane of the galaxy.

Spectra of CGCG 536-030 are not shown, as the results are similar. Further, the GPS-L3 signal was (un)fortunately not present during any of the scans on this source.

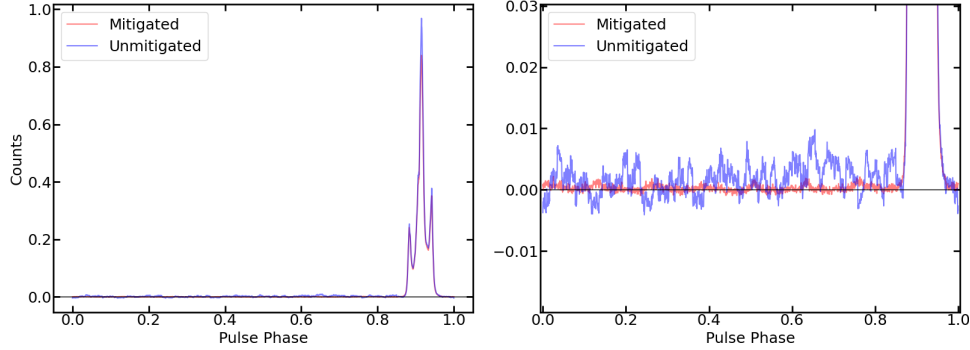


Figure 12: Left: Folded pulse profile with 256 channels, $M = 512$, MSSK shape of (1,2), and a sigma threshold of 3.6. Right: Zooming in on the drastically reduced off-pulse noise levels.

4.1.3 Pulsars

The B0329+54 observations were markedly improved by many (but not all) \widehat{SK} mitigation schemes. The largest benefit, at least for the SNR of the folded pulse profile, came with using an observing setup of 800 MHz and 256 channels, $M = 512$, $m, n = 1, 2$, and $\sigma = 3.6$. With this \widehat{SK} configuration, we flagged 1.26% and 1.37% of data in the X and Y polarizations respectively, leading to a union of 1.68% total data flagged. The folded pulse strength was decreased to 87% of its unmitigated level, but the noise was lowered enough to increase the SNR to 419% of its unmitigated value, as demonstrated in Figure 12. This does unfortunately mean that the pulse emission itself is getting flagged by \widehat{SK} in this case. Many other configurations of frequency resolution, M , and MSSK also resulted in a ratio of 100%-300% between the mitigated and unmitigated SNR values. The results were only worsened when \widehat{SK} flagging rates approached or exceeded roughly 20%, which meant there was a significant amount of overflagging. The best SNR increase with a \widehat{SK} configuration that did not flag pulse emission was with $M = 256$, $m, n = 1, 1$ and $\sigma = 3.0$ with the same frequency resolution as above. The folded profile height was lowered by 2% and the SNR was increased by 302%, with a 1.66% flagging rate.

For the former case (SNR increased by 419%), the `rfifind` routine flagged 41.7% in the unmitigated data and 39.2% in the mitigated data, which is a marginal decrease. For the TOA calculations, the RMS of the residuals dropped from 5235.539 μs to 4075.058 μs in the non-paz flagged data, and dropped from 4118.536 μs to 4074.078 μs in the paz flagged data. The uncertainties dropped from 85.0748 μs to 57.62 μs in the non-paz flagged data, and dropped from 58.6185 μs to 57.6111 μs in the paz flagged data. The non-paz flagged residuals are shown in Figure 13. It was found that significant overflagging could artificially drop these RMS and uncertainty values, so the TOA results in overflagged trials are marked as irrelevant.

Figure 14 shows the `prepfold` results for the unmitigated and mitigated data. The mitigated data shows a much cleaner pulse profile and a 16% higher detection significance. To summarize the other metrics for this best-case mitigation, the low-frequency noise in the FFT was significantly reduced. The number of candidates found in the blind search did drop with RFI mitigation, but the best result came with `rfifind` applied with no \widehat{SK} , which means `rfifind` performed better at reducing spurious signals. However, the significance of our detected pulsar was highest with \widehat{SK} applied and no `rfifind`.

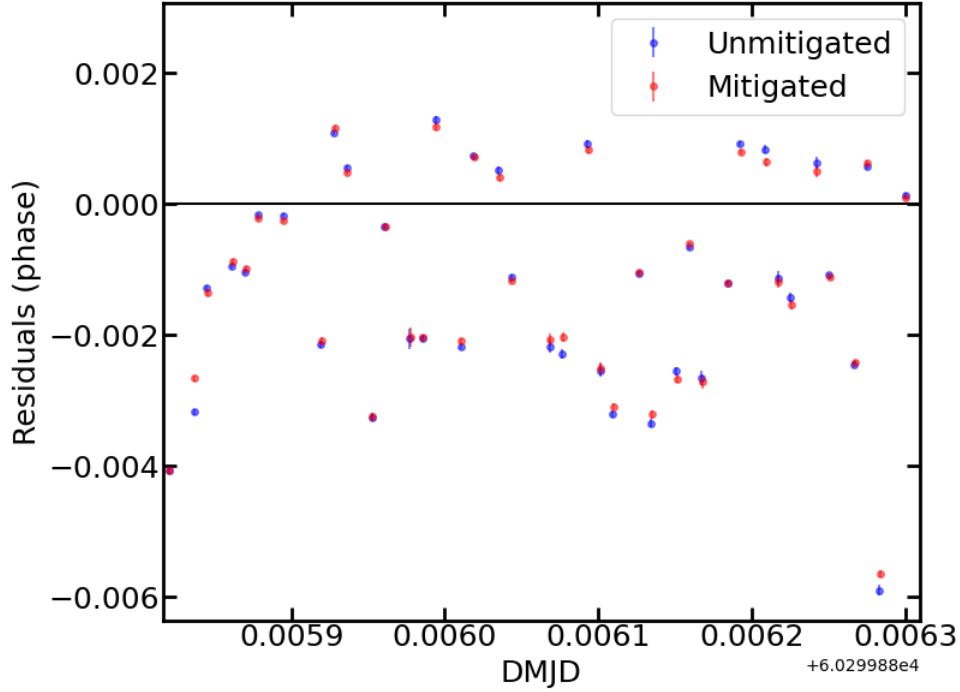


Figure 13: Pulsar TOA residuals for the best-case mitigation of B0329+54. The mitigated data has smaller magnitudes of residuals and smaller uncertainties. The uncertainties on most TOA residuals here are much smaller than their magnitude.

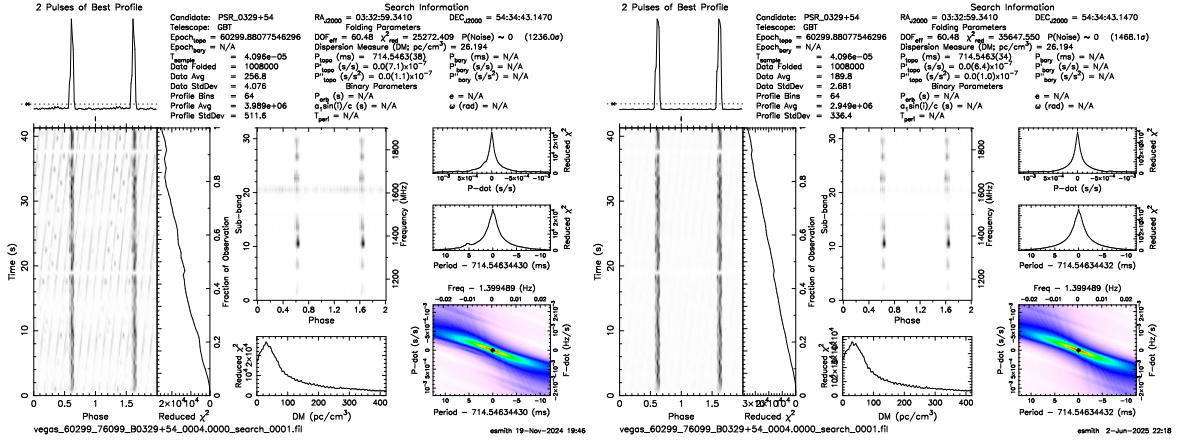


Figure 14: prepfold results of the unmitigated and mitigated data on the left and right, respectively. The mitigated result shows a cleaner pulse profile and higher detection significance.

For the millisecond pulsar J1713+0747 data, we were able to match the SNR values between the mitigated and unmitigated pulse profiles in a variety of cases, but not improve them. We tried using a similar flagging time resolution to our B0329+54 data, scaled by the difference in pulse period, as well as various MSSK configurations and number of channels. In every case, we unfortunately matched or decreased the SNR.

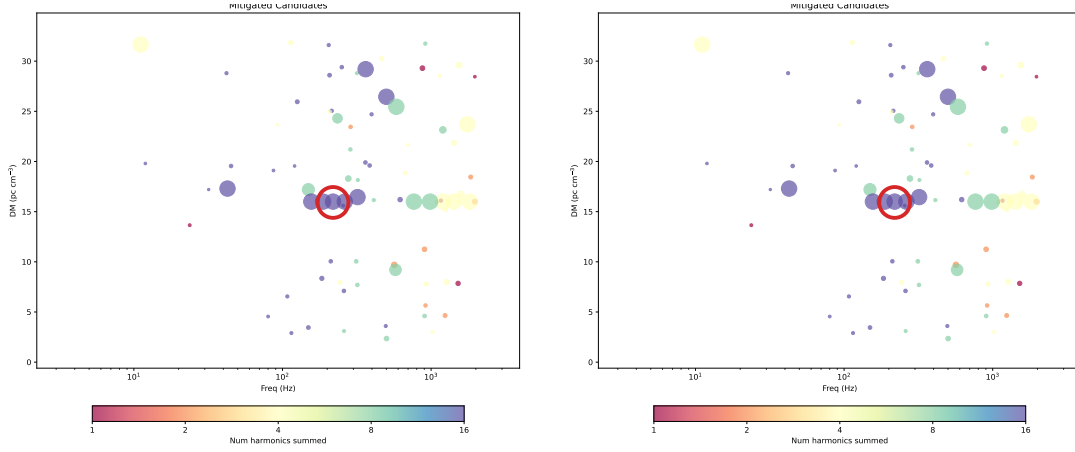


Figure 15: A summary of candidates found using PRESTO. The x-axis is the rotational frequency of the candidate and the y-axis is the dispersion measure. The colorbar indicates the number of harmonics of the signal that were detected, and the size of the symbol indicates the S/N. The red circle marks the true signal from the pulsar, which is also visible at its harmonics. Mitigated data is shown in the left panel and unmitigated data is shown in the right panel. We see far fewer candidates in the mitigated data.

4.2 Cyclostationary Processing

We performed the same data reduction procedures as described in the above sections for data mitigated with \widehat{SK} . We find some improvement in several metrics when comparing data mitigated and unmitigated data products that do not make use of the PRESTO `rfifind` routine. However, when we use `rfifind`, the differences in the mitigated and unmitigated data products are minimal. This suggests that CSP may be a promising way of reducing the need for other RFI mitigation in post-processing. However, when an observer records very high time resolution power spectra, as is the case for pulsar observations, post-processing methods perform similarly to CSP.

First, we examine the number of pulsar candidates detected in a “blind” search of the data. Without using a PRESTO RFI mask, we find 109 candidate pulsars in the unmitigated data and 85 in the mitigated data. Since most of these candidates are spurious and caused by RFI, this represents a 22% improvement. In both datasets, the most significant candidate is the pulsar J1713+0747, which is detected with the same period and dispersion measure. However, the S/N of J1713+0747 is only about 1.4% higher in the mitigated data. Figure 15 is a visualization of the candidates for each data set.

When using the PRESTO RFI mask, we find 79 candidates in the unmitigated data and 88 in the mitigated data, and the S/N is 3.6% lower in the mitigated data. This suggests that the data-replacement method is causing `rfifind` to identify some of the data we replace as being RFI due to slightly incorrect statistics. We are continuing to investigate refinements to the data-replacement strategy.

Next, we examine the main input to the PRESTO search routines — an FFT of the integrated time series. Figure 16 compares the FFT for mitigated and unmitigated data, without and with use of the PRESTO RFI mask, respectively. We see much fewer spurious signals in the mitigated data when not using an RFI mask. With the RFI mask, the results are similar, though there is less power at low frequencies in the mitigated data, suggesting that it is flagging RFI that is varying more slowly. These results are in-line with the number and significance of the candidates found in a blind search.

Next, we consider the quality of data that has been phase-folded with the precise, known rotational parameters of the pulsar. This mimics post-processing for observations of a known pulsar, as opposed to a blind search for candidate pulsars. Figure 17 shows data folded with the PRESTO routine `prepfold` for the mitigated and unmitigated data, respectively. The pulsar is clearly detected in both, but with a 2.8% lower significance in the mitigated data. As with the candidates found in the blind search, this indicates that the data replacement strategy

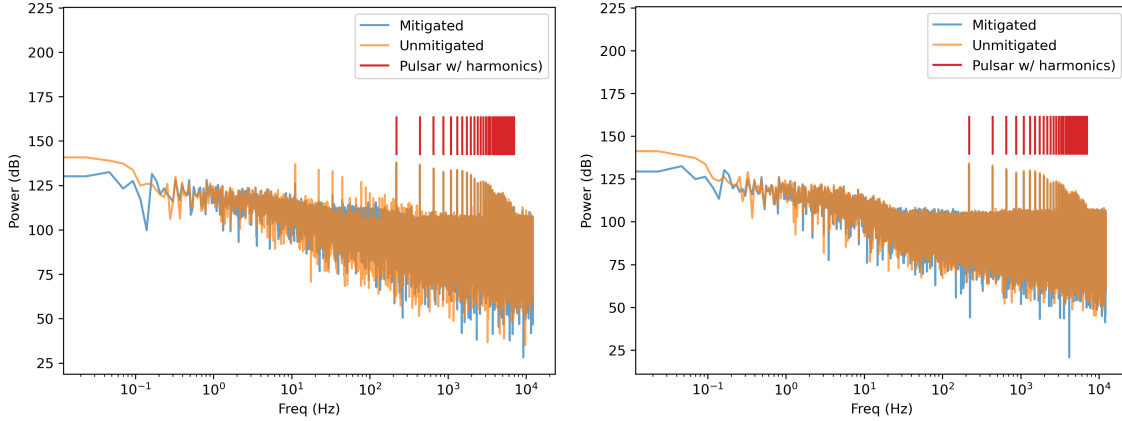


Figure 16: Left: FFT power spectra for mitigated and unmitigated data when not using an RFI mask. We see far fewer spurious signals in the mitigated data set, while the pulsar and its harmonics are still clearly visible. This spectrum was made using data dedispersed at the true DM of the pulsar. Right: FFT power spectra using an RFI mask. Spurious signals are reduced in both cases, but there is still less power at low frequencies in the mitigated data.

needs further fine tuning. However, we note that the pulsar is still detected with very high significance in both cases.

Although we observed a broad-band pulsar rather than a spectral line source, we can still examine the integrated spectrum for the mitigated and unmitigated data. In Figures 18–20 we compare the integrated spectra for mitigated and unmitigated data at various frequencies of interest. First, we show frequencies around the Iridium band. We can see that the algorithm works well at flagging Iridium, similar to the \widehat{SK} approach. Fine tuning of the algorithm may be able to further reduce what little Iridium signal is still present in the data. Next, we show a region of the spectrum from approximately 1530–1565 MHz, which is replete with many complex sources of RFI. Here, results are mixed. Many sources of RFI are greatly reduced in power, but the data replacement has produced artifacts in the form of elevated power levels, due to the difficulty in determining the appropriate variance of RFI-free data for the normally distributed noise. There are some sources of RFI that are not detected by our algorithm. As noted previously, we use a relatively small value of N to avoid identifying the pulsar as cyclostationary, and this may be limiting the effectiveness of the algorithm for non-pulsar signals. Both of these results demonstrate the importance of selecting the appropriate algorithmic approach for different science cases. Finally, we show the two frequencies at which the Bedford, NC FAA radar transmits. Our algorithm does a reasonably good job of reducing the power of this source of RFI. Because the radar repeats approximately every 12 seconds, it is frequently detected in blind pulsar searches. The effectiveness of our algorithm at flagging the radar signal partly explains the reason that there are fewer candidates in the mitigated data when it is searched blindly for pulsar candidates.

5 Conclusions

We used \widehat{SK} to greatly improve the detections of the OH megamaser in III Zw 35 and the pulsar B0329+54. In the other spectral line observations, we matched the mitigated and unmitigated SNRs, which means the celestial signal of interest wasn't inadvertently flagged. We did not see improvement in the results for the millisecond pulsar J1713+0747. \widehat{SK} had issues detecting Gaussian-modulated interference, and will continue to be investigated as a possible real-time RFI mitigation technique for the GBT.

CSP is a promising technique for RFI mitigation. We find moderate improvements in some metrics when examining pulsar data products, especially when CSP is used without other post-processing tools to reduce RFI. Pulsars are a challenging use-case for CSP since they are themselves cyclostationary. Integrated spectra show that our algorithm identifies and removes a variety of sources of RFI at different frequency ranges. We will continue to fine

tune our algorithm and test it on additional astronomical sources, especially sources of molecular line emission and masers.

6 Acknowledgements

This work is supported by the National Science Foundation through Advanced Technologies and Instrumentation grant #1910302.

References

- Gary, Dale E et al. (2010). “A wideband spectrometer with RFI detection”. In: *Publications of the Astronomical Society of the Pacific* 122.891, p. 560.
- Nita, Gelu M (2016). “Spectral Kurtosis statistics of transient signals”. In: *Monthly Notices of the Royal Astronomical Society* 458.3, pp. 2530–2540.
- Nita, Gelu M and Dale E Gary (2010a). “Statistics of the spectral kurtosis estimator”. In: *Publications of the Astronomical Society of the Pacific* 122.891, p. 595.
- (2010b). “The generalized spectral kurtosis estimator”. In: *Monthly Notices of the Royal Astronomical Society: Letters* 406.1, pp. L60–L64.
- Nita, Gelu M, Dale E Gary, et al. (2007). “Radio Frequency Interference Excision Using Spectral-Domain Statistics”. In: *Publications of the Astronomical Society of the Pacific* 119.857, p. 805.
- Nita, Gelu M, Jack Hickish, et al. (2016). “EOVSA Implementation of a Spectral Kurtosis Correlator for Transient Detection and Classification”. In: *Journal of Astronomical Instrumentation* 5.04, p. 1641009.
- Roberts, Randy S. et al. (Apr. 1991). “Computationally efficient algorithms for cyclic spectral analysis”. In: *IEEE Signal Processing Magazine* 8, pp. 38–49. DOI: [10.1109/79.81008](https://doi.org/10.1109/79.81008).
- Robishaw, Timothy et al. (2008). “Extragalactic Zeeman detections in OH megamasers”. In: *The Astrophysical Journal* 680.2, p. 981.

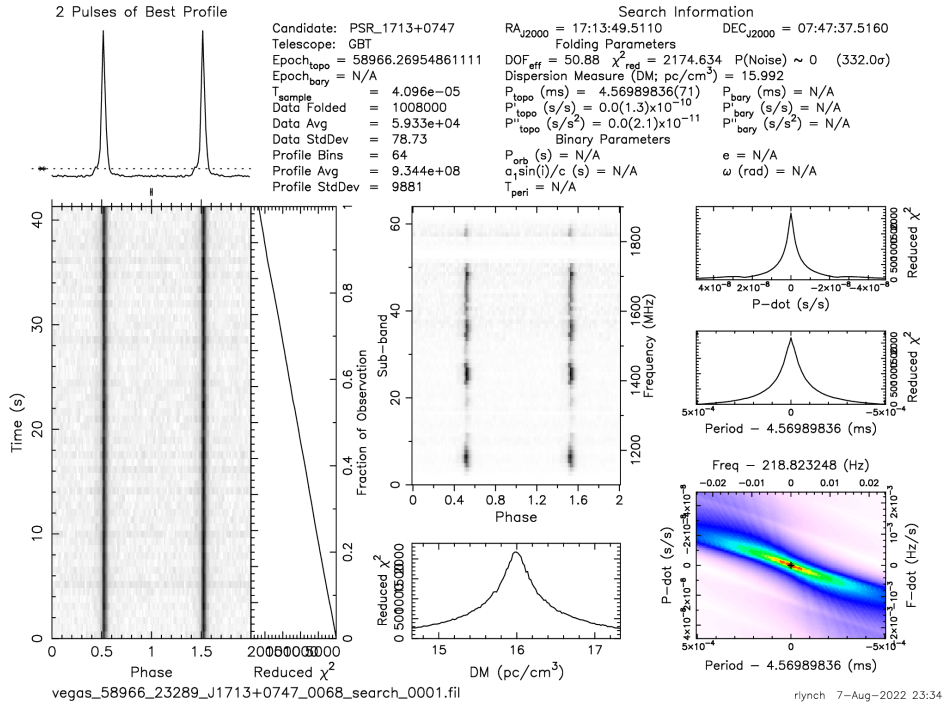
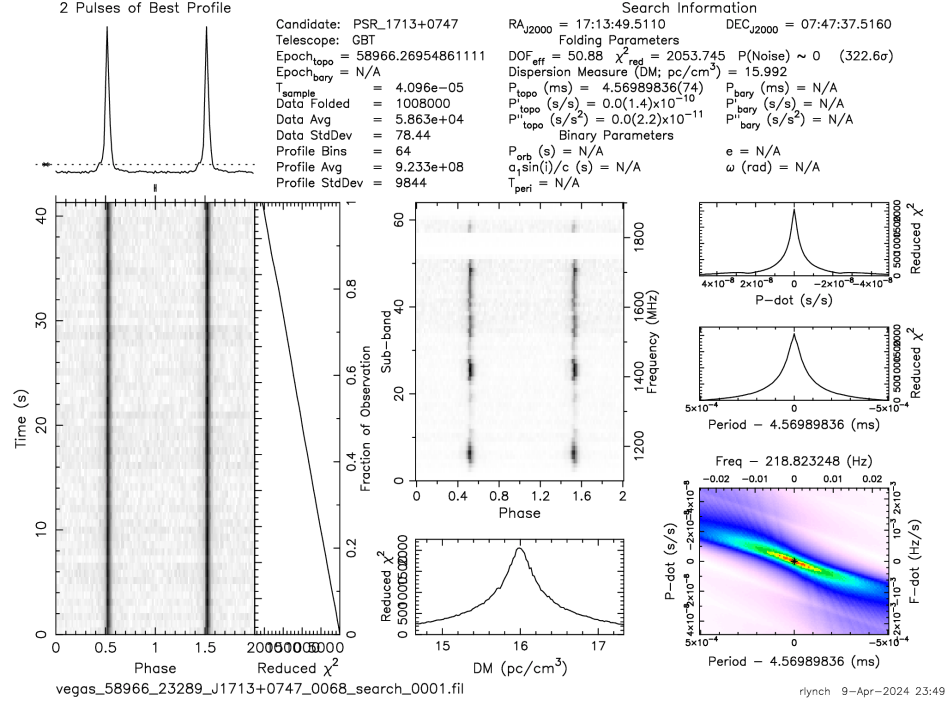


Figure 17: Output from the PRESTO prepfold routine for mitigated data on the top panel, and unmitigated data on the bottom panel.

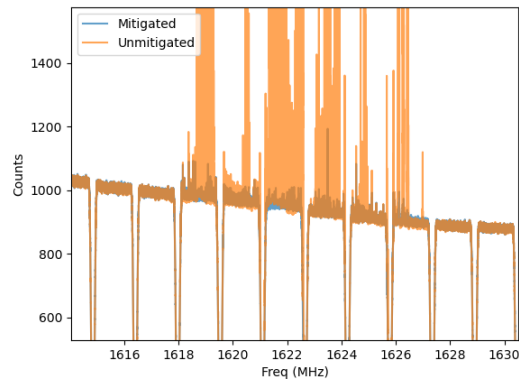


Figure 18: Integrated spectrum for mitigated data (blue) and unmitigated data (orange) in the Iridium band. Our algorithm greatly reduces the power of Iridium transmissions while preserving the overall spectral shape.

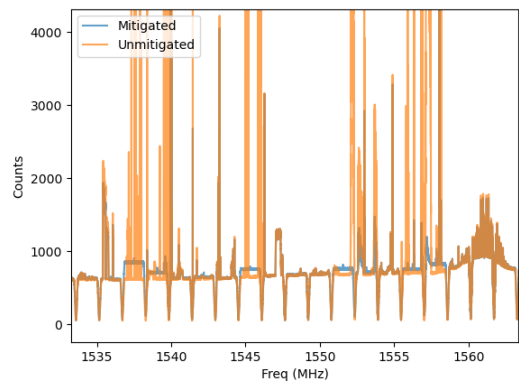


Figure 19: Integrated spectrum for mitigated data (blue) and unmitigated data (orange) from ~1530–1565 MHz. Results here are more mixed, with some sources being reliably detected as RFI, while some are not picked up by our algorithm. The noise-replacement approach has difficulty determining appropriate variance for the normally-distributed noise is difficult at frequencies with persistent RFI, causing elevated power levels in some channels.

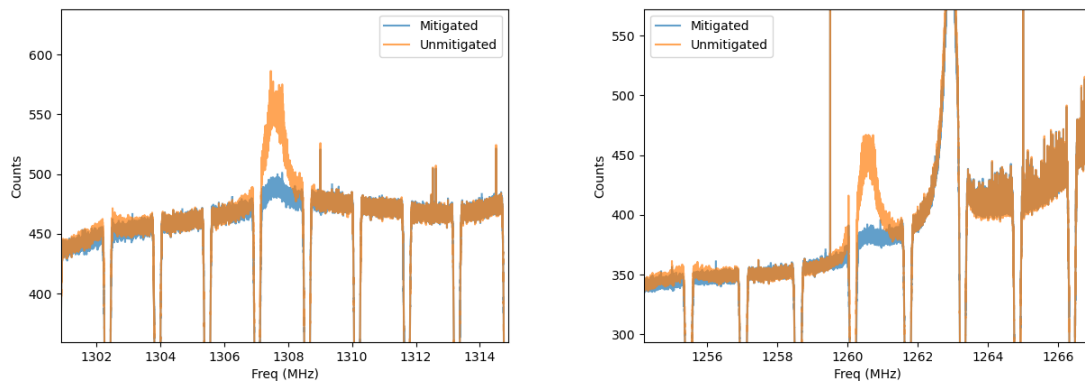


Figure 20: Integrated spectrum for mitigated data (blue) and unmitigated data (orange) for the two frequencies used by the Bedford radar. Our algorithm does a good job of removing this source of RFI.