The RFI 2022 Workshop at ECMWF, 14 – 18 February 2022
2022 URSI International Union of Radio Science

Simulating Spectral Kurtosis Mitigation against Realistic RFI signals

E.Smith*[(1)(2)], Ryan S. Lynch[(3)], D.J. Pisano[(1)(2)(4)]
(1) Department of Physics and Astronomy
West Virginia University
White Hall, Box 6315, Morgantown, WV 26506
(2) Center for Gravitational Waves and Cosmology
West Virginia University
Chestnut Ridge Research Building, Morgantown, WV 26505
(3) Green Bank Observatory (GBO)
155 Observatory Road, Green Bank, WV 24944
(4) Department of Astronomy
University of Cape Town
Private Bag X3
Rondebosch 7701, South Africa

Abstract

We explore the statistical radio frequency interference (RFI) mitigation technique spectral kurtosis ($\widehat{SK}$) in the context of simulated realistic RFI signals. $\widehat{SK}$ is a per-channel RFI detection metric that estimates the kurtosis of a collection of M power values in a single channel to discern between human-made RFI and incoherent astronomical signals of interest. We briefly test the ability of $\widehat{SK}$ to flag signals with various representative modulation types, data rates, and duty cycles, as well as accumulation lengths M and multi-scale $\widehat{SK}$ bin shapes. Multi-scale $\widehat{SK}$ uses a rolling window to combine information from adjacent time-frequency pixels to mitigate weaknesses in single-scale $\widehat{SK}$. High data rate RFI signals with significant sidelobe emission are harder to flag, as well as signals with a 50% effective duty cycle. Multi-scale $\widehat{SK}$ using at least one extra channel can detect both the center channel and side-band interference, flagging most of the signal at the expense of larger false positive rates.

1 Introduction

As the problem of human-made radio frequency interference (RFI) grows and datasets are increasingly large, there exists a need for new techniques of automated online RFI mitigation that works inside the spectrometer. It is also extremely important that these RFI mitigation techniques can detect and remove human-made signals without negatively impacting scientific data. Many modern RFI mitigation techniques happen in the post-averaging phase after the data leaves the spectrometer, which can have their own set of drawbacks [12]. For example, if a low duty cycle interferer affects a longer averaged scan or a very narrow signal saturates a much wider frequency channel, there may be a significant amount of clean data that is thrown out along with the RFI, as is documented for the Canadian Hydrogen Intensity Mapping Experiment (CHIME; [7]). Thus, it is important to characterize online RFI mitigation techniques that can work on un-averaged data inside the spectrometer to increase their effectiveness [2] while ensuring the extra computation involved is optimized enough to keep up with the data rate and not bottleneck the spectrometer.

In these proceedings we explore the ability of the statistics-based spectral kurtosis estimator ($\widehat{SK}$) [9, 10] to detect and flag RFI signals while preserving astronomical signals. The algorithm determines the Gaussianity of channelized voltages by assessing the distribution of their associated squared power values per spectral channel over time, and it is a computationally simple method that has been proven to mitigate RFI in real-time inside the spectrometer [4, 11]. The statistical qualities of different RFI signals can vary in significant ways, and a typical radio observatory can have a large collection of incoming human-made radio signals corrupting data. We aim to determine how these qualities determine how well each is detected by $\widehat{SK}$. It is particularly important to ascertain the optimal setup for an unknown RFI environment that will detect the most interference possible, as well as provide insight into how to detect a specific RFI source with known qualities.

In Section 2, we describe how flagging is applied and evaluated, as well as the process of creating simulated datasets. Results are shown via flagging charts in Section 3, and these are contextualized in Section 4.

2 Methods & Data

The SK estimator of a set of data is calculated with the formula

$$\widehat{SK} = \frac{M+1}{M-1}\Big(\frac{M \cdot S_2}{S_1^2} - 1\Big) \tag{1}$$

with M as the accumulation length of power spectral density (PSD) values. [4, 9, 10, 8, 11]. S_1 and S_2 are defined per spectral channel k as

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

where P_j are the power spectral density values. The statistical distribution of $\widehat{SK}$ values from RFI-free active white Gaussian noise is modelled by a Pearson type III curve centered at 1, so we find thresholds for clean data by numerically solving the related cumulative distribution function at the 3σ level. $\widehat{SK}$ is run individually on each channel on M spectra at a time. The RFI environment and spectrometer characteristics can influence a suitable choice of M, which should be large enough to avoid errors from small statistics without making the time resolution too coarse. The RFI environment and spectrometer characteristics can influence a suitable choice of M.

$\widehat{SK}$ is weak to signals with $\sim 50\%$ duty cycle [9], which has led to the development of multi-scale $\widehat{SK}$ that averages rolling windows of adjacent time-frequency pixels in the arrays of S_1 and S_2 and allows us to bypass this weakness to 50% duty cycle signals [4]. For a multi-scale $\widehat{SK}$ bin width of m channels and n $\widehat{SK}$ time bins, S_1 and S_2 become

$$\begin{aligned} S_{1,ik}^{mn} &= \frac{1}{(m+1)(n+1)} \sum_{j=0}^{m}\sum_{l=0}^{n} S_{1,i+j,k+l} \\ S_{2,ik}^{mn} &= \frac{1}{(m+1)(n+1)} \sum_{j=0}^{m}\sum_{l=0}^{n} S_{2,i+j,k+l} \end{aligned} \tag{3}$$

which changes the effective duty cycle seen by each $\widehat{SK}$ bin and allows for more complete flagging of 50% duty cycle signals. $\widehat{SK}$ is applied with the new $S_{1,ik}$ and $S_{2,ik}$, and if the multi-scale bin is found to have RFI, all of the original single-scale time-frequency pixels within the bin are flagged.

To simulate and flag RFI, we start by summing together the time samples of an RFI signal with additive white gaussian noise. We ramp up the RFI amplitude from zero over the course of the scan to help discern how bright the signal has to be before $\widehat{SK}$ detection. The data stream is then channelized via a poly-phase filterbank routine[1], which uses 24 taps and a Hanning window to smooth the filter shape, as this mimicks the VEGAS spectrometer at the Robert C.Byrd Green Bank Telescope (GBT; [14]). The resulting data is reshaped and squared, resulting in a 2-dimensional array of un-averaged pseudo-power spectra.

The data is split up into chunks of M spectra to apply single-scale and multi-scale $\widehat{SK}$. When multi-scale $\widehat{SK}$ is used, we apply both multi-scale and single-scale and use the union of the two flagging masks. To calculate true positive rates (TPR) and false positive rates (FPR), a comparison mask is generated by channelizing the RFI independently from the noise and convert into decibels using the noise as reference. We consider RFI to be present when the signal is greater than -10dB. These percentages are then recorded for each simulation run. To contextualize the false positive rates, we run the simulations with no RFI present. For $M > 1024$, we approach the Gaussian 3σ level of 0.3%.

As mentioned above, the simulation process is done for a wide range of different RFI characteristics. Table 1 summarizes the various possible values for the different parameters.In order to generate flagging plots as shown in Section 3, a script was written to simulate every combination of two parameters in Table 1 one hundred times each to generate a statistically significant population of flagging results from which means and variances can be calculated.

The general form of each RFI generation function is modeled after examples found at `http://cyclostationary.blog`. Phase-shift keying (PSK) changes the phase of the signal, amplitude-shift keying (ASK) changes the amplitude, and frequency-shift keying (FSK) changes the frequency while holding the phase constant. After the full signal is generated, apply a duty cycle ranging from from 5% to 100% by setting a period of 1 millisecond and setting some fraction of the data to 0 for however many duty cycle periods there are. The signal generation process is described more rigorously in Smith et. al (2022, submitted), which is a full paper outlining results with more signal characteristics and more fully clarifying reasons behind why certain signals are detected or non-detected.

3 Results

We summarize the most key flagging results from comparing signal data rate to either duty cycle or $\widehat{SK}\,M$, as well as how multi-scale $\widehat{SK}$ increases flagging.

[1]The code is adapted from `https://github.com/telegraphic/pfb_introduction`.

Data Rate (ksps)	Duty Cycle (%)	$\widehat{SK}$ M (# spectra)	MS shape
1	5	128	Single-scale
4	10	256	MS-12
20	15	512	MS-21
100	...	1024	MS-22
200		2048	MS-24
	...	4096	MS-42
	95		
	100		

Table 1. Values for the different RFI and flagging parameters.

3.1 Signal Data Rate to Duty Cycle

We start by testing the $\widehat{SK}$ algorithm's known weakness to 50% effective duty cycle. These simulations were run with a duty cycle period of 1ms, which is shorter than the time length of a single $M = 512$ $\widehat{SK}$ bin by a factor of 2.62144. This can cause the effective duty cycle as the $\widehat{SK}$ algorithm sees it to be slightly different from the given intrinsic duty cycle, but we keep this distinction as there is no reason a transmitter duty cycle has to operate at a similar or fortuitous rate (e.g. 3/4) to the radio spectrometer.

Figure 1a shows that single-scale $\widehat{SK}$ minimally detects low data rate BPSK signals with about a 50% duty cycle. The 100 and 200 ksps signals are not sampled entirely accurately, however, which leads to somewhat different flagging responses. For these data rates, the symbols change at a comparable rate to the spectra dump rate, which leads to the minimum flagging occuring at duty cycles higher than 50%. For the 100 ksps signal, the center channel is more completely flagged as the duty cycle rises above 60%, with a drop in the sidelobe flagging effectiveness, leading to the local maximum at 80%. These effects can be eliminated by changing the simulation parameters such that the sampling rate is increased by a factor of 10 and the duty cycle period reduced by the same factor.

Due to their shifts in signal strength, amplitude-shift keyed signals have an effective duty cycle different from the intrinsic duty cycle given during signal generation. Figure 1c shows how this manifests for $\widehat{SK}$ detection. Lower data rates remain largely undetected at around 85% duty cycle. Since the amplitude changes over time the signal appears at full strength for a shorter amount of time than when it's actually on, causing the effective duty cycle of the signal to become larger than the intrinsic duty cycle. We define the intrinsic duty cycle as the percentage of time where the signal is physically on, and the effective duty cycle is when the signal appears to be on at the point of RFI detection. This is explored more in Section 4.2. It is therefore important to realize the difference between effective and intrinsic duty cycle in the context of signals that change amplitude.

We run the same experiment as above with a MS-21 $\widehat{SK}$, which was shown to be the smallest multi-scale $\widehat{SK}$ bin shape shape that significantly improves flagging. Adding one extra channel in the multi-scale shape is indeed enough to detect the center channel when it gets missed by single-scale $\widehat{SK}$ due to 50% duty cycle, as shown in Figure 1b. QPSK and ASK signals are also completely flagged under MS-21 $\widehat{SK}$.

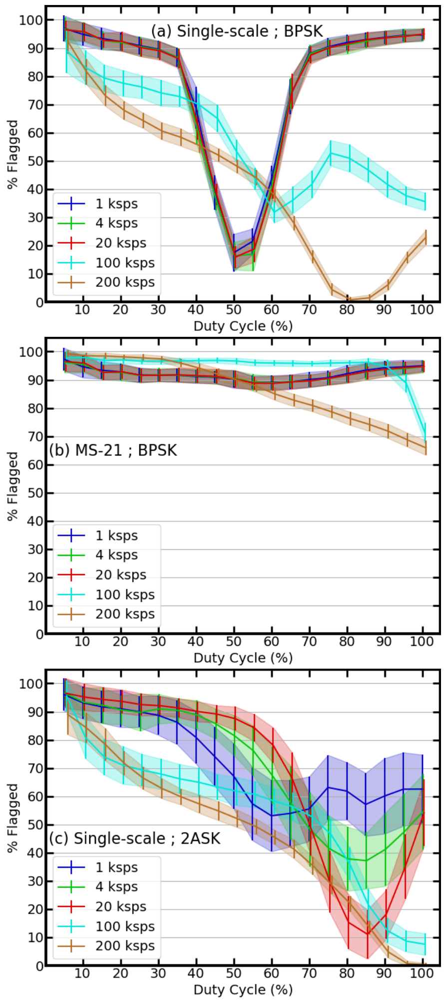

Figure 1. Flagging percentages for various data rates compared to the duty cycle of the RFI signal. (a) Single-scale $\widehat{SK}$ on a BPSK signal. (b) MS-21 $\widehat{SK}$ on the same signal. (c) Single-scale $\widehat{SK}$ on a binary ASK signal.

3.2 Signal Data Rate to $\widehat{SK}$ M

As shown in Figure 2, we see that higher data rates lead to less flagging, due to the fact that sidelobe emission in adjacent channels is not detected as well. There are larger sidelobes at higher data rates because the symbol sequence that modulates the carrier wave acts like a square wave multiplied along the data sequence. This square wave has a smaller and smaller period for higher data rate signals, which translates to a wider emission bandwidth after channelization. As the sidelobes get larger, less of the RFI is flagged. On the other hand, the flat flagging curve at lower data rates in Figure 2a are caused by the fact that there is no sidelobe emission beyond the center channel until the data rates rise high enough. Quadrature phase-shift keyed (QPSK) signals are flagged along slightly lower rates than the BPSK signals, although at slightly lower rates due to slightly larger sidelobes.

Figure panels 2b and 2c show that ASK signals are also flagged better with higher M and at slower data rates, but go almost completely undetected at 200ksps. 4-bit ASK signals are flagged slightly less, but only at higher M. As M rises, the disparity between the two bit rates rises. As the data rate rises, the sidelobes start to get flagged more reliably while the center channel gets less flagged, leading to a local maximum. The $\widehat{SK}$ algorithm is likely affected by how the effective duty cycle of an undersampled ASK signal appears, and we is explored more in Section 4.2. However, we do see that a higher M does lead to higher TPR across the board. A larger value of M is able to detect weaker signals and sidelobe emission somewhat better.

Figure 3 shows the effect of MS-12, MS-21, and MS-42 multi-scale $\widehat{SK}$ on different M and data rates. It is apparent that MS-12 did not increase flagging, while MS-21 and MS-42 $\widehat{SK}$ perform much better than single-scale $\widehat{SK}$, especially at higher M. So, the weakness to sidelobe emission can be mitigated by including at least one extra channel in the multi-scale $\widehat{SK}$ window allows the algorithm to detect some of the sidelobe emission. MS-42 $\widehat{SK}$ was found to heavily overflag clean data, with false positive rates of up to 2.9%, more than 2 percentage points higher than the baseline noise FPR of 0.3%. In these cases, MS-$\widehat{SK}$ flags extra clean channels on either side of the center channel, even when the RFI doesn't actually have any sidelobes. This is apparent in Figure 4, which shows an unmitigated signal, the comparison mask, and the single-scale and MS-42 $\widehat{SK}$ masks. While the sidelobe emission gets missed in the single-scale mask, the MS-42 mask is much wider than the comparison mask, leading to the overflagging. Larger multi-scale shapes have the potential for even more overflagging on clean data, depending on the frequency width of the RFI.

4 Discussion

We summarize the results above and outline some reasons for why certain RFI signal characteristics had different $\widehat{SK}$ detection rates. Again, Smith et. al (2022, submitted) provide a more rigorous insight into these causes.

4.1 Data Rate

Higher data rates were generally harder for single-scale $\widehat{SK}$ to detect, due to the wider sidelobe emission. As mentioned in Section 3.2, the modulation of the carrier wave acts like a square wave in time, with a period related to the data rate. After channelization, a shorter period square wave results in a signal with wider emission bandwidth. This sidelobe emission does not get flagged optimally, leading to the result that the larger the data rate is, the less likely the full signal will be flagged. It is unclear at this point how the statistical properties of sidelobe emission cause it to be missed by single-scale $\widehat{SK}$. However, we do see that using multi-scale $\widehat{SK}$ with enough frequency width to cover the sidelobe emission can detect the entire signal, comparing panels c and d of Figure 4. Despite not knowing why single-scale $\widehat{SK}$ does not detect this emission, we can see that multi-scale $\widehat{SK}$ does nullify this weakness even in the case that the center channel doesn't get flagged.

4.2 Duty Cycle

To expand on some of the differences in flagging across duty cycle and for different modulation types, we state the equation for the spectral variability V_k [8]:

$$V_k = \frac{\sigma_k^2}{\mu_k^2} \quad (4)$$

with μ_k as the mean and σ_k^2 as the variance of power values in frequency channel k. It was found that whenever these two values were very similar, single-scale $\widehat{SK}$ underperformed in comparison to multi-scale $\widehat{SK}$, and equation 4 does approach unity for the channelized voltages of a 50% duty cycle BPSK signal. Different modulation types and data rates affected how μ^2 and σ^2 changed over different intrinsic duty cycles, but this general rule applied as a good marker of flagging effectiveness, and explains the minima in flagging rates in Figure panels 1a and 1c, as well as other tested simulation setups. Again, we restate that multi-scale $\widehat{SK}$ with at least one extra channel in bin size was able to cover these weaknesses effectively. In practice, then, one could detect RFI with a known 50% effective duty cycle by making use of multi-scale $\widehat{SK}$.

5 Conclusions

We simulated realistic RFI signal data sets with varying encoding schemes, data rates, and duty cycles and miti-

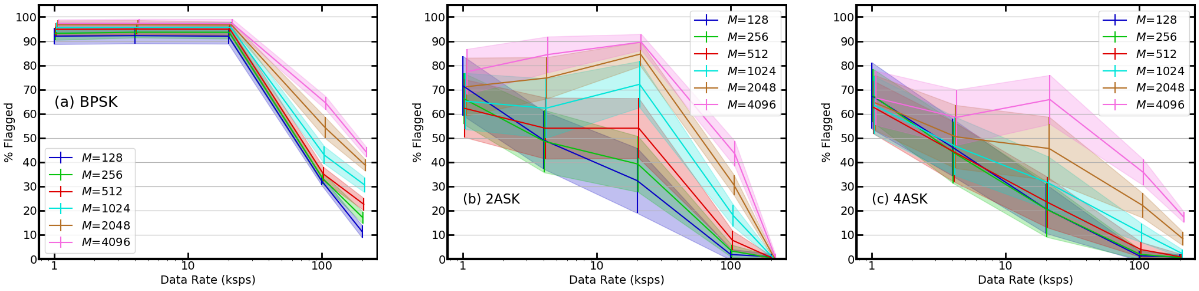

Figure 2. Single-scale $\widehat{SK}$ flagging percentages for various values of M compared to the data rate of the RFI signal. a) BPSK modulation. b) 2-bit ASK modulation. c) 4-bit ASK modulation

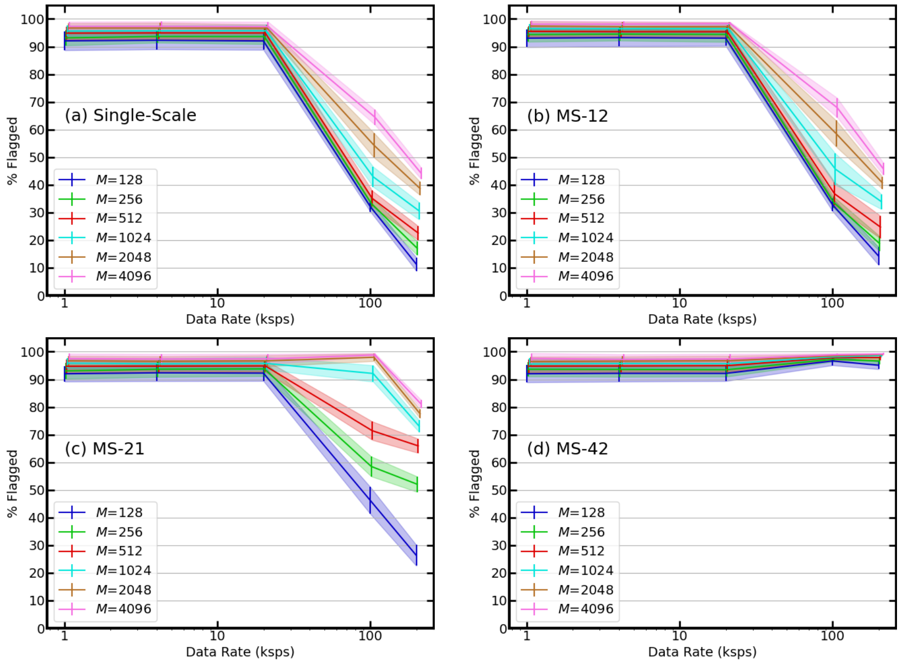

Figure 3. Flagging percentages comparing data rate to M for several multiscale $\widehat{SK}$ bin shapes. a) Single-scale $\widehat{SK}$, b) MS-12, c) MS-21, d) MS-42.

gated them with single-scale and multi-scale $\widehat{SK}$ to test the strengths and weaknesses of this method. We found that signals with higher data rates are typically harder to flag, especially those with significant sidelobe emission. Signals with an effective duty cycle of $\sim$50% also went undetected. Multi-scale $\widehat{SK}$ with $m \geq 2$ can subvert these weaknesses and detect the entire signal, but with the possible expense of significant FPR if the signal's width did not exceed m channels.

For future work, we want to include more complicated modulation types, as the list of possible modern encoding schemes contains lots of possible alternatives, such as minimum-shift keying and spread-spectrum modulation. The basic modulation types we explored in this paper can also be scaled up to larger amounts of bits per symbol such as 32-PSK, which can transmit 32 bits per symbol. Smith et al. (2022, submitted) expand upon more simulation parameters than are tested in these proceedings, which include different carrier frequencies relative to the PFB channel centers and FIR response shape, as well as testing against a realistic stationary and incoherent astronomical signal.

6 Acknowledgements

This work is supported by the National Science Foundation through Advanced Technologies and Instrumentation grant #1910302 and ETS and DJP thank the West Virginia University Eberly College Dean's Office for partial support of this project. This material is based upon work supported by the Green Bank Observatory which is a major facility funded by the National Science Foundation operated by Associated Universities, Inc. DJP is supported through the South African Research Chairs Initiative of the Department of Science and Technology and National Research Foundation. We also acknowledge advice from conversations with David McMahon and Jason Ray, as well as invaluable guidance on realistic signal generation from Dr. Chad Spooner.

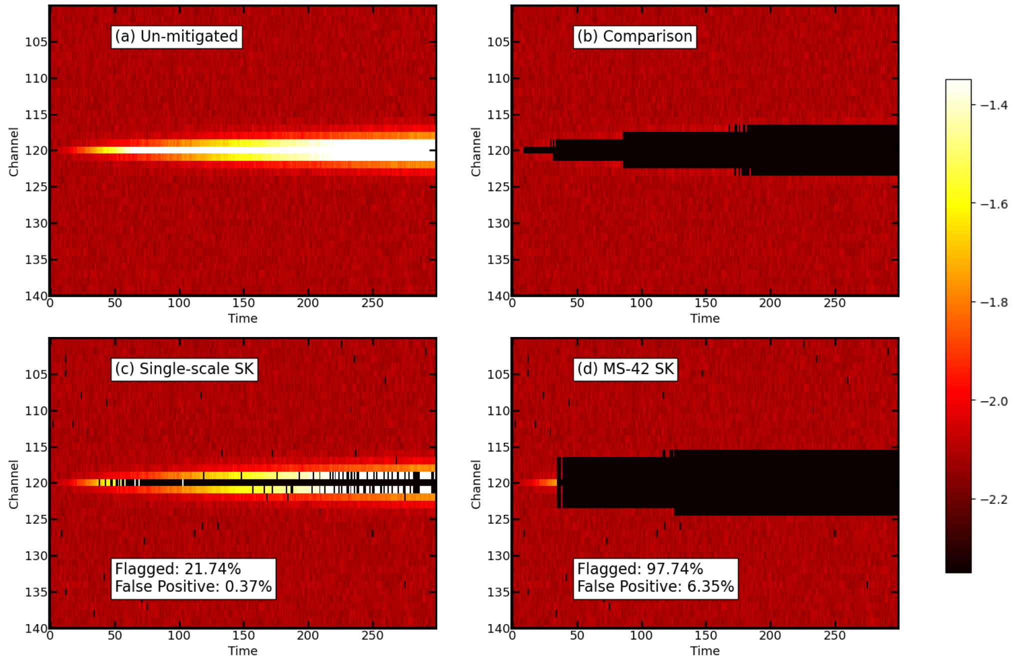

Figure 4. Examples of flagging a 200 ksps BPSK signal. (a) The unmitigated signal. (b) Comparison mask. (c) Single-scale $\widehat{SK}$. (d) MS-42 $\widehat{SK}$.

References

[1] Andrews, Mark J., et al. "Properties of the 500–2000-MHz RFI Environment Observed in High-Latitude Airborne Radiometer Measurements." *IEEE Transactions on Geoscience and Remote Sensing* 60 (2021): 1-11.

[2] Baan, Willem A. "RFI mitigation in radio astronomy." 2011 *XXXth URSI General Assembly and Scientific Symposium*. IEEE, 2011.

[3] Fisher, J. R., et al. "Mitigation of pulsed interference to redshifted HI and OH observations between 960 and 1215 megahertz." *The Astronomical Journal* 129.6 (2005): 2940.

[4] Gary, Dale E., Zhiwei Liu, and Gelu M. Nita. "A wideband spectrometer with RFI detection." *Publications of the Astronomical Society of the Pacific* 122.891 (2010): 560.

[5] Hess, Kelley M., et al. "CHILES: H i morphology and galaxy environment at z= 0.12 and z= 0.17." *Monthly Notices of the Royal Astronomical Society* 484.2 (2019): 2234-2256.

[6] Hunt, Lucas R., D. J. Pisano, and S. Edel. "The Search for HI Emission at z≈0.4 in Gravitationally Lensed Galaxies with the Green Bank Telescope." *The Astronomical Journal* 152.2 (2016): 30.

[7] Mirhosseini, Arash. "High cadence kurtosis based RFI excision for CHIME." *Ph. D. Thesis* (2020).

[8] Nita, Gelu M., et al. "Radio frequency interference excision using spectral-domain statistics." *Publications of the Astronomical Society of the Pacific* 119.857 (2007): 805.

[9] Nita, Gelu M., and Dale E. Gary. "The generalized spectral kurtosis estimator." *Monthly Notices of the Royal Astronomical Society: Letters* 406.1 (2010): L60-L64.

[10] Nita, Gelu M., and Dale E. Gary. "Statistics of the spectral kurtosis estimator." *Publications of the Astronomical Society of the Pacific* 122.891 (2010): 595.

[11] Nita, Gelu M., et al. "EOVSA implementation of a spectral kurtosis correlator for transient detection and classification." *Journal of Astronomical Instrumentation* 5.04 (2016): 1641009.

[12] Offringa, A. R., et al. "The LOFAR radio environment." *Astronomy Astrophysics* 549 (2013): A11.

[13] Offringa, A. R., et al. "The low-frequency environment of the Murchison Widefield Array: radio-frequency interference analysis and mitigation." *Publications of the Astronomical Society of Australia* 32 (2015).

[14] Prestage, Richard M., et al. "The versatile gbt astronomical spectrometer (vegas): Current status and future plans." 2015 *USNC-URSI Radio Science Meeting* (Joint with AP-S Symposium). IEEE, 2015.