Adapting PRESTO Pulsar Validation Techniques for Amateurs PW East

Abstract

PRESTO, the professional pulsar search and analysis software, is now the package of choice for most amateurs detecting and studying pulsars. PRESTO analysis tools examine pulsar data, first reducing the effects of radio frequency interference (RFI) and, after further processing, outputs an image containing a number of sub-plots. These sub-plots analyze various pulsar-specific identifying characteristics which together can provide convincing evidence of a wanted pulsar presence. This paper investigates validation issues for the weaker pulsar intercepts with an integrated/folded signal-to-noise ratio (SNR) below 10:1. These are not clearly highlighted by the efficient statistical processes in PRESTO, that were essentially designed⁽¹⁾ for large dishes to speedily detect strong new pulsars of unknown parameters in a low RFI environment. Here it is shown that, for keen amateurs prepared to examine data in more detail, the recognition features in the PRESTO *prepfold* output plot can be made more discriminatory by replacing the Chi-square amplitude statistic with an SNR measure. By this means, positive intercept can be confirmed much more confidently for the weaker, known pulsar data records; this is illustrated with a PRESTO-inspired image, but using the SNR statistic on a B0329+54 data file with an integrated SNR of less than 5:1.

Introduction

The PRESTO *prepfold* software tool rapidly carries out a number of processes on a data file and ideally outputs a graphical plot as shown in Figure 1 [1,2]. The various processes normally combine to provide sufficient information for an operator to confidently confirm the acquisition of a real pulsar.



Figure 1. Typical 35:1 SNR Example PRESTO prepfold plot using the Chi-square amplitude statistic.

The sub-plots and recognition features based on expected pulsar parameters are summarized below,

1. Pulse profile plot (located top-left of PRESTO plot)

Pulse profile and amplitude on folded noise base - two periods shown to improve visibility at phase extremities. Clear narrow pulse indicated.

- 2. Waterfall and associated Reduced-Chi-square cumulative value plots (located left side). For large SNRs, vertical dark lines or vestiges directly below the pulse profile plot (as is visible here) with possibly scintillation dropouts indicate the intercept of a regular pulse train. Not always clear for even quite large SNRs. The Reduced Chi-square section attached to the right of the waterfall should indicate a continuous, mainly increasing, amplitude trend from start to finish, beginning with a value equal to one.
- 3. Frequency sub-band plot (located in center). Continuous vertical lines consistent with the pulse profile phase position (again, clearly visible here due to the reduced number of band sections) indicating frequency components in all sub-bands (4 sub-bands in this case) as expected from the broad-band nature of pulsar signals.
- 4. Dispersion Measure (DM) sub-plot (located center bottom).A peak is expected at the known DM of the pulsar source. Signifying the signal is extra-terrestrial, typical

of traveling a distance through interstellar space.

5, 6, and 7, P-dot search, period search and correlated p-dot/period search plots respectively (right side). Peaks expected at zero error on the separate period and P-dot search plots; signifying expected accurate pulsar period with negligible rotation frequency (P-dot) drift. The correlated P-dot/Period search plot should indicate a maximum at the center zero-point with no other peak ambiguities and a feature with a slope of -2/N, where N is the number of pulsar periods in the data record. This test confirms the presence of an accurate and stable pulse train throughout the data record.

Positive evidence in all seven sub-plots builds the required recognition confidence.

With a strong signal with an SNR = 35:1 as in Figure 1, the sub-plot amplitude discrimination is significant, so recognition confidence is very high, but with lower level signals the plot quality is not so clear as is obvious in Figure 2 with a lower integrated SNR of 15:1.



Figure 2. Typical 15:1 SNR Example PRESTO prepfold plot using the Chi-square amplitude statistic.

The pulse profile looks good but the time and frequency waterfall pulsar phase lines are just visible. The DM search is valid but the period and P-dot searches are very noisy but showing evidence of peaking at zero search offset. It is noted that the Chi-square discrimination (maximum / minimum) has dropped to about 1.5:1 (from about 5:1 for the Figure 1 example.



Figure 3. Typical 5:1 SNR Example PRESTO prepfold plot using the Chi-square amplitude statistic.

Figure 3 *prepfold* example shows a typical output for a 5:1 SNR acquisition. Here, apart from a visible pulse profile, none of the sub-plots inspire confidence and so the intercept would likely be rejected. But is this

assessment correct?

Once their receiving system has been proven, amateurs have been satisfied with just a large SNR result and in some cases just a positive period and DM search. It is known however, that in modest to low SNR cases, that this is insufficient and that other methods may need to be invoked to improve recognition confidence. There is now a trend for amateurs to require validation evidence only from professional software. This paper shows that this is not always the best option, but instead, offers a tested alternative developed from an understanding of pulsar, noise and RFI properties. It is shown that by duplicating PRESTO algorithms but applying an alternative measure, pulsar validation confidence can be extended to much lower target SNRs. This article considers two approaches. The first uses PRESTO *prepfold* tool together with an SNR program, to manually examine the PRESTO 'best profile' data, and the second develops sub-programs to mimic the PRESTO sub-plot functions.

Associated with the PRESTO waterfall and search plots is the amplitude measure, Reduced χ^2 , which is a statistical measure of the data quality and is discussed in the next section together with the alternative SNR measure.

Reduced Chi-Square (χ^2) and SNR Algorithms

The reduced χ^2 algorithm is used in PRESTO as a goodness-of-fit measure on folded data assuming the underlying data noise is purely Gaussian/Normal distributed. With just noise, the measure gives a value close to unity but when a large pulsar pulse is present the numerical measure can be very large. It is not dependant upon the pulse phase and is easily implemented, so can be useful in period, P-dot and DM searches to track the pulsar amplitude profile offset.

The reduced χ^2 statistic is given by,

$$\chi_{red}^{2} = \frac{1}{N-1} \sum_{p=0}^{p=N-1} \frac{(x_{p} - \overline{x})^{2}}{\sigma_{p}^{2}}$$
(1)

where, N is equal to the number of fold bins (is a measure of the 'degrees of freedom' = N-1).

 \bar{x} , and σ_p are the mean and standard deviation estimates of the folded noise alone.

If x_p were solely folded samples of Gaussian noise, the right-hand part of Equation 1 equates to one, but if a significant pulsar pulse (or skewed RFI) is present then the result will exceed unity.

The corresponding definition of signal-to-noise ratio (SNR) is,

$$SNR = \frac{x_{pm} - \overline{x}}{\sigma_p}$$
 (2)

where, x_{pm} is the folded data maximum response.

Note: There is a strong similarity between these equations and it is noticed that the reduced χ^2 equation is a power ratio and the SNR equation is a voltage ratio. A good approximation for Equation 1, assuming a Gaussian-shaped pulse and a pulsar duty cycle = W/P (W = half height pulse width; P = pulsar period) is,

$$\chi^2_{red} \approx 1 + SNR^2 \frac{W}{\sqrt{2}P} \quad (3)$$

which illustrates the duty cycle dependency. i.e. improved discrimination with increased pulsar duty cycle.

Although the Chi-square formula looks more complex than SNR, it is in fact, much easier to compute, being a simple sum of squares. To optimize SNR on the other hand, it requires the data bandwidth to be matched to the pulse bandwidth and a data search, first to locate and measure the peak response then to separate it from the base noise to compute the standard deviation.

Chi-square/SNR Discrimination Comparison.

Figure 4 models the SNR (red) and χ^2 algorithms (solid blue for the approximately 1% duty-cycle B0329 pulsar, and dashed blue for a 3% duty cycle pulsar to indicate χ^2 dependence. The χ^2 curves show good pulse amplitude tracking at high SNR levels, but for weak signals with SNRs below 20, the curve slope is much reduced offering much less amplitude discrimination. The curvature is dependent upon pulsar pulse duty cycle so the χ^2 plot's curvature sharpens for pulsars with increased on/off ratio.

The main conclusion drawn from Figure 4 is that due to the 'flatness', hence poor amplitude discrimination,

of the χ^2 statistic at low SNR levels. It can only be expected to indicate good parameter search peaks for intrinsic SNRs greater than about 12:1 (certainly for 1% duty-cycle pulsars such as B0329). Below this level it seems that SNR is a better discriminant.

An interesting conclusion can be drawn from this graph. If there is sufficient amplitude discrimination using the χ^2 measure to validate Dispersion Measure (DM) search, Period search and P-dot search graphs in the *prepfold* output plot for pulses with an SNR of 10:1, then by applying the SNR algorithm, the corresponding validation level could now drop down to 4:1 - closer to the expected residual noise peak line. In practice, there are other validating characteristics and impediments that can affect the confidence of recognition success.



Figure 4. Comparison of Chi-square and SNR measures on a Folded Pulse Train

Manual Search Results of 5.5:1 SNR Data

In each of the search plots reproduced in Figures 5, 6 and 7, the red points are the indicated maximum SNR values obtained from the *prepfold* best profile results using an SNR plotting program on the Figure 3, *rfifind*-masked data. The blue curves are those calculated assuming a continuous pulse train with frequency and pulse characteristics identical to the pulsar considered and folded using the standard algorithm. The three theoretical curves may appear identical shapes, but in fact are slightly different, responding to the particular search parameter. The background theory is explained in Reference 3. Rolling SNR measurements are seriously affected by impulsive RFI unless first mitigated and this is probably another reason why it is not the preferred measure in PRESTO.

Figure 5 shows the SNR DM search results which now indicate a clear peak around the expected B0329 pulsar DM of 27, not evident in the χ^2 version of Figure 3. It is interesting to note that the form and accuracy of Figure 5 SNR result more closely matches the χ^2 version of the 15:1 SNR plot in Figure 2, supporting the improved discrimination claim for SNR profiling.



This is sufficient to confirm that the source is broad band and dispersed as expected and the peak DM matches that of the expected source allowing for some distortion due to the underlying noise. The deviation from theory with increasing DM is anticipated as the source signal gets spread out in time, reducing its amplitude below that of the folded noise peaks.

Figure 6 shows the results for SNR period search with the period change ranging from -5 ppm to +5 ppm. There is a clear peak evident now in the SNR plot matching very closely to the expected value (TEMPO

predicted) and the roll-off theory over the range -2 ppm to +2 ppm [3].



Outside this range, random noise peaks appear to take over. The theory assumes a perfect pulse train matching the pulsar parameters implying that an accurate pulsar pulse train is largely present throughout the data file. When the fold period is changed away from the correct value, for example, a true pulsar pulse position in the final period relative to the first period will have shifted. This has three effects after folding and they are, 1. the final summed pulse is broadened, 2. it is reduced in amplitude and 3. the pulse center will shift in relative time/bin. For a given data set these three effects are predictable for any substantially complete pulse train as described in Reference 3. Similar effects are predicted for DM and P-dot searching.



Figure 7. P-dot Search

The period-rate search, P-dot plot in Figure 7, again shows good conformity with theory over the range -2 x 10^{-10} s/s to 2 x 10^{-10} s/s, a peak at zero P-dot and the symmetry proving a high degree of frequency stability of the target (<< 1 µs over the 3 hour observation period) as well as again demonstrating the presence of a continuous pulse train.

Note that these searches report the <u>maximum</u> SNR value in the data fold. An alternative policy is to calculate the pulsar peak bin shift and report the data amplitude of this bin, adding confidence and possibly reducing the outlier values [3].

Practical Results using the SNR Statistic with 5:1 SNR Data File [5]

C-programs were written to duplicate the *prepfold* plots. For the Period/P-dot graphic, the folded peak SNR for each *p* and p^d combination was calculated [4]. This cycled through period values around the pulsar topocentric period of -15ppm to +15ppm and P-dot values -30 x 10⁻¹⁰ s/s to +30 x 10⁻¹⁰ s/s in ppm increments of 0.5ppm and P-dot increments of 1 x 10⁻¹⁰ s/s. For each of the 3600 fold runs the folded data peak SNR (PSNR) was recorded and used to plot a simulant of the *prepfold* graphic considered.

The result together with a 3-D version and amplitude code is shown in Figure 8.

It is observed that there are no ambiguities either on the zero period error or the zero period rate error ordinates.



Figure 8 5:1 SNR '*prepfold*' graphic and 3-D Version based on the PSNR amplitude value.

A number of conclusions can be drawn from the Figure 8 presentation,

1. A clear central peak at coordinates (0,0) is evident confirming period match and nominally zero spindown.

2. A continuous pulse train matching the pulsar properties occurs throughout the data sample.

3. The peak signal slope matches the prediction. N = 10100, p = -15 ppm,

and, $p^d/p = -2/N = 28 \times 10^{-10}/-15 \times 10^{-6} = -1.8 \times 10^{-4}$.

4. Significant strong response occurs over a period variation exceeding 8ppm (compared to less than 4ppm for $p^d = 0$.

5. There appears a slight offset from symmetry about coordinate (0,0). Possibly due to distortion caused by the underlying noise at this modest SNR level.

6. Observing the 3-D version, it appears that there is a extended amplitude response along the diagonal including the peak response. This is possible since folding with an offset positive period value can be partially compensated by a negative P-dot value as illustrated in Figure 9.



Figure 9. Fold Period and P-dot Timing offset Diagram

Normally the optimum fold strategy is to have zero period and P-dot offsets and folding all data record N periods. In the case illustrated, the period value is offset by +p ppm and compensated by a negative P-dot value of p^d such that it crosses the zero line at N/2 periods (integration along the zero line gives maximum response of course). The ellipse identifies the data region offering the best periods to produce a maximum positive folded result. The optimum p^d slope is such that $p/(N/2) = p^d$; confirming the result noted in Equation 3. It can also be inferred that partial compensation for increased period offsets is available for all values of opposite polarity p^d providing the P-dot increment slope crosses the zero line within the N data record periods.

This plot result is only possible if there is a precise, very low drift pulse train present with moderate scintillation, throughout the data record.

The Waterfall Plot

Figure 10 shows a 'waterfall' plot (flowing upwards). For this plot, the nominally 5:1 SNR data file used for the following plots was divided into four equal sections and folded separately before combining in the contour graphic (red strong positive, blue strong negative). The central pulsar signal fills each section as expected but, as discussed above, some noise peaks remain present for two or more sections. When dividing into many more sections, the central pulsar feature reduces into noise and so to invisibility.

In this example the data file pulsar peak SNR is reduced by one half ($1/\sqrt{4}$ for 4 sections).



Figure 10. 5:1 SNR B0329+54 Example Matched-Period 4-Section Fold of 10,000 Periods.

It is this reduction in component SNR with increased waterfall sections that is responsible for the poor visibility of pulsar presence in the Figure 3 waterfall plot. In general, for a waterfall plot, a reasonable goal is to limit the number of sections 'S' so that the section expected SNR does not drop much below 2.5:1; with this proposed constraint the integrated fold SNR should be greater than $2.5\sqrt{S}$.

SNR Running Value Plot

This section investigates using the running SNR measure instead of the *prepfold* running Chi-square statistic. For this plot, the data is again divided into a number of equal sections (in this case, 1000), each section is folded, but now the sections are accumulated linearly and the SNRs calculated for each accumulated section sum and plotted. In this way, the growth of SNR can be observed; the final sum giving the SNR of the complete data fold.



Figure 11. 5:1 SNR B0329+54 Example Accumulated Block SNR

Figure 11 is an interesting result as it shows that the central pulsar candidate's SNR (red) increases (roughly as the square root of the number of sections) as predicted. The variations may also be expected due to source scintillation and possibly, random effects of residual RFI features along the data record. The brown line accumulates the peak SNR of all the data in each section (without knowledge of the pulsar position) and clearly shows that the pulsar takes over from random noise peaks as the accumulated SNR exceeds about 3:1. This is a powerful tool for not only differentiating pulsar candidates and noise, but to recognize weak ranges for integration, due to scintillation or RFI presence (ignoring/blanking these sections can improve the final pulsar SNR as discussed later). For information, the performance of the Chi-square cumulative statistic on this data is shown in blue and confirms that it is virtually insensitive to this weak signal.

Modified Waterfall Plot

Instead of the time waterfall displaying each divided section fold results, a more instructive PRESTO emulation approach is to plot this accumulated peak SNR as a waterfall showing all the data as in Figure 12

below. In this method, section folds are sequentially added together such that by the final section, the sum fold of all the data is presented.

Figure 12 shows much more clearly the suppression of random noise components and the final dominance of this low SNR pulsar response. The vertical lower-level features (yellow/green) show that the final noise peaks can indeed be present throughout the record but tend not to integrate as does the true pulse train. This is a regular feature due to the folding algorithm's selectivity.



Figure 12. 5:1 SNR B0329+54 Example Accumulated SNR Waterfall

The central, red, high signal line towards the end part (top) of the integration extent show the pulsar peak rising above ambient noise. The noise peak SNR remains reasonably constant throughout the integration path and only the continuous pulse train increases in SNR producing the fine red line towards the end of section integration. Other short-lived peaks can of course occur. The pulsar line width in keeping with its pulse width/period ratio.

Scintillation Screening

Once the data has been divided into synchronized sections or discrete periods it is possible the scan the data by means of a moving average window to investigate pulse group amplitude/scintillation variations along the data record.



Figure 13. 5:1 SNR B0329+54 Cumulative SNR Plot with Scintillation Indicator (blue)

The blue plot shows the result of a running/moving average of ± 50 blocks (10% of the data) about the block center (Note: the displayed amplitude is reduced by a factor of four for clarity). The mean SNR is expected to be $5/\sqrt{10} = 1.58:1$ if the pulse train is largely present throughout the record; agreeing with the plot. This means that for the cumulative plot to increase, the measured SNR should exceed this figure. The blue plot also shows the regions where the received pulsar signal is the strongest. Conversely, when the average amplitude falls below 1.58 due to either scintillation or increase in RFI, then the cumulative SNR reduces. Blanking these low SNR regions tends to improve the final pulsar SNR.

Final Combined Data Plot

Combining the information discussed above with the other sub-plot results, it is now possible to emulate the Figure 3 PRESTO plot and compare the results of using SNR rather than Chi-square to sense the pulsar presence.

Figure 14 shows the SNR comparison plot. An extra inclusion is the frequency spectrum (center) plot; the data used here was recorded using a three-band RTL receiver, each of 2 MHz bandwidth centered on 609, 611, and 613 MHz [5]. The DM, Period and P-dot sub-plot data were obtained from a MathCad program operating on the same set of recorded data (Appendix 3 Reference 5).



Figure 14. 5:1 SNR B0329+54 Example PRESTO prepfold Comparison Plot

Comparing Figure 14 with the 5:1 SNR *prepfold* plot of Figure 3, it is clear that pulsar recognition features are now very much evident, just by using the maximum SNR rather than the Chi-square statistic. It should be admitted that at present there is a large time penalty to obtain the detail of Figure 14; this was assembled manually for this exercise but would seriously benefit from the power of a Python GUI!

Conclusions

The PRESTO software tools were built for speed and efficiency in finding unknown pulsars in data acquired by large dishes on a continuous search basis. Although not designed for amateurs to validate detections of

known pulsars, it is useful for pulsars intercepted with an SNR greater than about $\sqrt{\sqrt{2}P_W}$. However, by

emulating PRESTO functions and employing peak SNR rather than the in-built Chi-square statistic, detection and recognition can be confidently extended to SNRs well below this PRESTO limit.

The PRESTO analysis algorithms have been emulated for this exercise, except that the Chi-square statistical measure has been replaced by a peak SNR measure in the various parameter searches. The folded data bandwidth has been matched to optimally pass the pulsar pulse width/shape. By this means, it has been demonstrated that the B0329+54 pulsar is clearly recognized with an SNR = 5:1 (equivalent PRESTO limit ~12:1). This technique is not only useful for amateurs sporting small aperture antennas, but could prove useful for amateurs with larger dishes wanting to explore weaker pulsars near their normal system threshold limit. The cumulative SNR plot is not only useful for indicating and separating likely candidates at lower SNRs than the example here, but blanking low amplitude scintillating pulse sections can usefully improve the final folded SNR prior to facilitating the search processes.

PW East June 2021

References

- [1] SM Ransom. Searching for Pulsars with PRESTO https://www.cv.nrao.edu/~sransom/PRESTO_search_tutorial.pdf
- [2] PRESTO Home, https://www.cv.nrao.edu/~sransom/presto/
- [3] PW East. An Analytical Method of Recognizing Pulsars at Moderate SNR., Journal of the Society of Amateur Radio Astronomers. November-December 2018.
- [4] PW. East, Getting the Best out of PRESTO Part 2 The PRESTO Period/P-Dot Search Graphic., Journal of the Society of Amateur Radio Astronomers. March-April 2021.
- [5] PW East, A Minimal Pulsar Detection System. Journal of the Society of Amateur Radio Astronomers. January-February 2018. p36. http://www.y1pwe.co.uk/RAProgs/MiniPulsarRx.pdf