SOFTWARE IMPLEMENTATION OF DETECTION ALGORITHMS FOR THE MCSA

D. Kent Cullers
San Francisco State University
Ames Research Center
Mountain View CA 94035

ABSTRACT: The NASA Multichannel Spectrum Analyzer (MCSA) supplies data in several formats and in a wide range of frequency resolutions. To decide if an extraterrestrial signal is present, this data is searched using detection algorithms particularly sensitive to signals concentrated in frequency and/or time. The algorithms, whose ultimate sensitivities are determined by theoretical considerations, are also constrained by available computing power. At present, practical detection of repetitive pulses is possible at an average power about five times lower than that for incoherent CW detection. After a period of field testing, the best detection algorithms, presently implemented in software, will be converted to hardware to increase their speed.

1. INTRODUCTION

Detection of signals in noisy data is a typical communications problem. SETI is unusual because nothing is known a priori about the nature of the incoming signal and because algorithms are required that have very few false positives. Thus, many of the standard signal detection techniques are difficult to apply directly to SETI.

A realization that natural astronomical sources do not produce narrowband signals, coupled with the Cyclops[1)] design study (which described a system capable of quickly searching out such signals over broad frequency bands with 1971 technology), focused attention on narrowband signals as the most detectable interstellar artifacts. B.M. Oliver first pointed out in the late 70's that if a signal is spread between many data samples, it is more difficult to detect using square-law detection and a test on the total power than if all the signal energy is concentrated in a single sample. Essentially, this is because, when only one sample containing all the signal is tested, the noise from other samples is excluded. Thus, it became clear that, on the basis of average power, pulse transmission is more efficient than CW.

John Wolfe pointed out that, though single pulses could be efficiently detected, one large data point would never be convincing scientific evidence for the existence of extraterrestrial intelligence.

M. D. Papagiannis (ed.), The Search for Extraterrestrial Life: Recent Developments, 385–390.

Furthermore, the probability of multiple large events would be much lower, if only noise were present, than that for a single event of the same amplitude. This observation, combined with my own realization that two pulses would determine all subsequent pulse train positions, and the number of possible noise events giving false alarms would be much reduced, led me to an analysis and development of multiple pulse detectors. The analysis[2)] found that, if energy were concentrated in only a few pulses and the signal were large compared to the noise in samples where it was present, multiple pulse detection using a square-law detector and a threshold which each pulse in the train must pass, was almost as efficient as a matched filter. Speaking somewhat more quantitatively, given pulses with peak powers of order 10 times the mean single sample noise, trains of less than 10 pulses, and false alarm rates of order 10^{-10} or less, the sensitivity of multiple pulse detection is within about 1 dB of that of a matched filter.

Signal detection software has been developed for three types of signals: First, narrowband carriers, signals concentrated in frequency. This is consistent with the strategy proposed in the Cyclops report. Second, single pulses, signals concentrated in time. To this end, the MCSA has been designed to supply multiple resolutions and pseudoresolutions matching a wide range of possible pulse lengths. The signal detection algorithms can operate at any resolution, though, at present, not on all simultaneously. Third, the software algorithms search for pulse trains, repetitive events above a threshold set by the desired false alarm rate. All algorithms have been implemented in software so that bad ones can be easily deleted and inspirations easily added. At the moment, the initial algorithms are designed to detect signals which one would expect extraterrestrials to send because they are efficient (concentrated) and easy for us to detect (require square-law detectors and not matched filters).

2. DETECTION SENSITIVITY

The sensitivity of a detection algorithm is determined in a three step process.

1. An acceptable false alarm rate is chosen.

2. A threshold is calculated that yields the correct false alarm rate.

3. Signal strength is varied about the threshold value to yield a curve giving the probability of missing the signal (p_{ms}) as a function of signal-to-noise ratio (SNR).

This curve depends on the characteristics of both the algorithm and the signal being detected. It is a function of such parameters as false alarm rate, number of samples, and type of baseline employed. Though no algorithm has been found that is best for all signal types, one algorithm may characteristically outperform another in important situations.

https://doi.org/10.1017/S0074180900146728 Published online by Cambridge University Press

Figure 1 shows curves for three detection situations with the same false alarm rate. The parameters were chosen as typical of those that will be used in a full sized NASA SETI system. Note that both pulse trains are detected at lower average powers than that of the CW signal. The very simple pulse detector used here is quite fast because it is optimized for pulses synchronized with the sampling interval (frame) of the MCSA. An unsynchronized pulse tends to spread its power among several adjacent samples, none of which pass the detection threshold for synchronized pulses. This problem can be cured by summing pulses in adjacent samples into pseudobins as was discussed by B.M. Oliver earlier in this volume. In fact, this is exactly what is done in the more general pulse detection algorithm now ready for field testing. The object of Figure 1 is to show that even very simple pulse detection schemes can often outperform incoherent CW detection. At their best, as when the synchronized pulse is detected, they approach the sensitivity of a matched filter.

All pulse detectors developed to date look for regularly recurring pulses such as those in Figure 2. In general, the first and last pulse in an allowed train can occur in any spectrum of an observation. If, including the end points, there are r regularly spaced occurrences above threshold and if each of these has a probability p, then

$$p_{fa} \;\; f \; n^2 \; p^r \qquad (1)$$

where p_{fa} is the probability of false alarm in a single channel, and n is the number of spectra in the observation. In words, the false alarm rate is proportional to the number of distinguishable pulse trains times the probability of a particular train being filled. Often not all end points are allowed, which effects the constant of proportionality in (1). This will be discussed further in the next section.

Equation (1) applies to a single channel. If one considers signals drifting in frequency, one must multiply the false alarm rate in (1) by the number of allowed distinguishable drift rates. Since more drifts can be distinguished as the observation gets longer

$$p_{fad} \;\; f \; n^3 \; p^r \qquad (2)$$

where p_{fad} is the false alarm rate for drifting pulses during an observation. This appears to be an unfortunate state of affairs if one wishes to observe for long time intervals. However, since processing time grows linearly with observation time, the number of operations a computer must perform per second increases with the square of n. Though this is not ideal, it is manageable.

3. CURRENT SOFTWARE FEATURES

Though pulse detection has been emphasized throughout this paper, incoherent CW detection is not being neglected in the field test stage. Such detection has been thoroughly analyzed in Cyclops and consists of testing an accumulator against a threshold set for the desired false alarm rate. A program to do this in the nondrifting case is now running on the VAX 11/750 which will be used to process the MCSA output. If, as expected, the MCSA will perform on board accumulation, this program can easily handle detection in real time.

Drifting CW can in principle be incoherently detected using the same algorithm as for nondrifting CW. However, many more accumulators are required to sum samples into synthesized bins in which drifting signals remain during the whole observation. The number of accumulators required is the number of bins in a spectrum times the number of drifts. As the MCSA size and observation length are increased, the memory for these accumulators becomes a large expense. Schemes other than simple incoherent addition are therefore being considered. They will be field tested, but not in real time. Also, matched-filter CW detection has been attempted by using large look-up tables with the object of recovering some of the efficiency shown by the pulse detector. Thus far, these approaches have not succeeded though work is continuing. Over the next few years a method of CW detection that makes the best compromise between high sensitivity, reasonable memory size, and cheap real-time operation must be determined.

Many pulse detection programs exist in software. All of them allow pulses to drift in frequency. Most of them do not require that pulses be synchronized with the sampling frame. They sum powers in adjacent time and frequency samples to pick up as much of the signal as possible. They can operate either on data that has already been organized into pseudobins or do the pseudobinning within the program.

If thresholding and pseudobinning is performed outside the VAX, all detectors can operate in real time on the 74k channels of the prototype MCSA. Early detection schemes operated in a pipeline mode, waiting for completion of an observation before data analysis. These versions determined possible positions by looking at first and last points of pulse trains and assuming increasing numbers of pulses between them. Generally such pulses did not fall on integer positions, so powers near these positions were summed to decide if the criteria for a pulse were met.

More recently, detectors have been developed that determine new pulse positions as data comes in. This allows the detection of strong signals early, but necessitates more complex algorithms for determining pulse train positions. This is because extrapolation from early data is more liable to error than interpolation between end points.

Both detectors are fast, that is, they can perform pulse detection in real time on data from the prototype MCSA. The interpolative schemes are somewhat more computationally intensive than the extrapolative schemes,which require more memory.

All pulse detection programs allow missing pulses. This decreases the probability that a real pulse train will be missed because noise has caused one or two of its samples to fall below threshold. Such schemes are probably not ideal. A better approach is to test the total power in all trains whose individual pulses exceed a very low threshold. Test programs now exist to do this. It is expected that this total power approach will replace schemes that allow one or two missing pulses.

The constant of proportionality in equation (1) can be much reduced if one considers not all regularly spaced pulse configurations but only

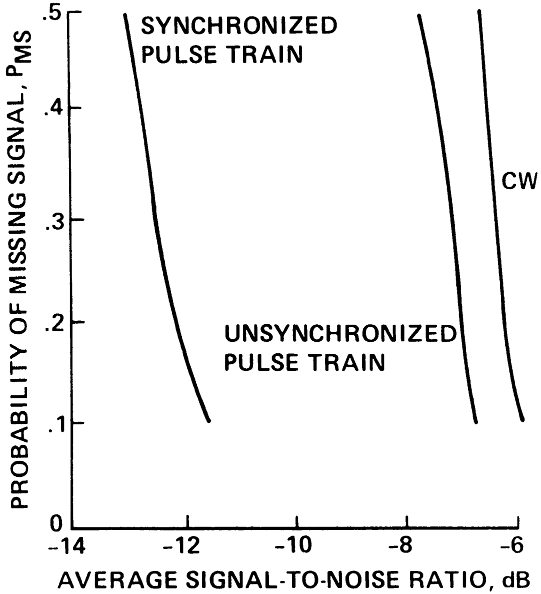

Figure 1:
Comparison, by Monte Carlo simulation, of the performance of a pulse and CW detector for a 10^6 channel MCSA during a 1024 sec observation where the threshold permitted about one false alarm per observation, and detection of drift rates greater than 1 Hz/sec was excluded.

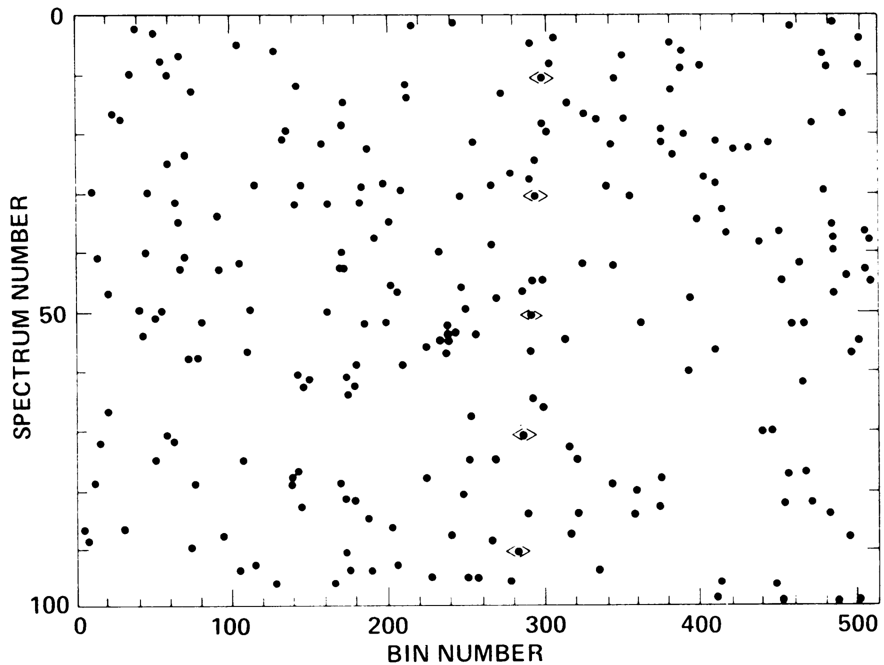

Figure 2:
Regular pulse train in a field of Poisson events. Shown are the pulses above threshold in 500 frequency bins for an observation 100 spectra long.

https://doi.org/10.1017/S0074180900146728 Published online by Cambridge University Press

those that persist throughout the observation. This is all well and good for a rotating beacon that sends an essentially infinite train of pulses to the receiver, but a series of pulses from a low repetition rate radar during some part of an observation would be ruled out by such a detection scheme. To satisfy everyone, we have programs which allow regular but finite pulse trains. They run more slowly than the infinite train detectors because many more possibilities must be investigated.

4. CONCLUSIONS

A large array of detection programs now exists on the VAX ready for field test. Code for detecting pulses sensitively and rapidly exists in software, though much fine tuning still needs to be done. Generally, it is expected that conversion of the pulse detector logic to hardware should be straightforward. However, the scaling of detection complexity with the square of the observation period is an unattractive feature of the algorithms developed thus far. A method of tree pruning which will cause the number of possiblilties to scale as nlog(n) is most desirable but is not yet in sight. It may be impossible without a great loss in sensitivity.

Incoherent CW detection is possible using the brute force Cyclops approach, but it is very memory consuming. So alternative ways of detecting narrowband signals are being investigated. A method that achieves sensitivity for narrowband signals comparable to that for pulses is desirable. To this end, schemes which attempt to quickly synthesize matched filters using a large precomputed table have been tried without success.

The most promising new approach is one proposed by Oliver which selects drifting CW signals by successive tests in a manner analogous to that now used for pulses. Only the relatively small number of starting candidates that pass initial low threshold tests need to be examined in depth, saving both memory and computing time. Development of an optimum drifting CW detection algorithm is one of the biggest challenges of the next few years.

REFERENCES:

1. PROJECT CYCLOPS: A Design Study of a System for Detecting Extraterrestrial Life; Oliver, B.M., and Billingham, J., Eds., rev.ed. 7/73, NASA CR 114445

2, SETI Science Working Group Report; Drake, F., Wolfe, J.H., and Seeger, C.L., Eds.; NASA Technical Paper 2244, 1984, p.49

https://doi.org/10.1017/S0074180900146728 Published online by Cambridge University Press