SETI: A MORE ECLECTIC APPROACH

Bernard M. Oliver
Ames Research Center
National Aeronautics and Space Administration

ABSTRACT: Three groups, one at Ames Research Center, one at the Jet Propulsion Laboratory, and one at Stanford University are currently engaged in the research and development phase of the NASA SETI program Other papers in this session will describe in greater detail some of the work going on at each of these centers. It is my intent here to give an overall picture of the program so that the other papers may be viewed in context.

1. INTRODUCTION

SETI related activities within NASA began in 1970 with the presentation at Ames Research Center of a summer lecture series on interstellar communication. This was followed in 1971 by the summer faculty fellowship program, Project Cyclops (1), which was a design study of a system to detect ETI signals. The Cyclops team attacked the problem of expanding the antenna collecting area if early searches showed this to be necessary. It was found to be technically feasible to use large phased arrays of conventional radio-telescopes to obtain collecting areas of tens of square miles. The study also examined the signal processing and detection problems and presented an optical method for performing the spectral analysis, and an analog scanning system involving the storage of 100 or more spectra to detect drifting CW signals.

During the 70's and early 80's SETI was kept alive at Ames on a shoestring budget that allowed only further paper study and occasional workshop meetings and conferences. The first significant annual funding began in FY 1983 at a $1.5 million dollar per year level. The present R and D phase is concerned with the design of the hardware and software needed for an effective search. This phase is expected to last until FY 1988. Although field tests of the system will be made earlier, the actual full scale observing program will not start until FY 1988.

M. D. Papagiannis (ed.), The Search for Extraterrestrial Life: Recent Developments, 351–360.

2. FREQUENCY RANGE

Since the publication of Cocconi and Morrison's seminal paper in 1959 almost all searches for ETI signals have taken place in the microwave window from 1 to 10 GHz with most attention being given to a narrow frequency range near the hydrogen line at 1.42 GHz. The present NASA program is no exception. The search will cover as much of the microwave window as remains clear of radio frequency interference (RFI) plus perhaps a few spot bands below 1 GHz and above 10 GHz. The search will cover most intensively the "water hole" between the hydrogen and hydroxyl lines, thought by some (1) to be naturally marked for the meeting of intelligent species.

The microwave window still appears to be the best part of the spectrum since the background noise there is lowest and the required transmitter power is thus reduced. Other regions, such as the infrared, come into serious consideration only if some naturally powered source is discovered that makes much higher powers very cheap.

3. RECEIVING ANTENNAS

There is no plan at present to build a dedicated SETI receiving facility. Instead the plan is to use existing radiotelescopes such as at Arecibo, NRAO at Greenbank, Ohio State University, and perhaps others as time becomes available. Use will also be made of antennas in the NASA Deep Space Network at Goldstone, California and at Tidbinbilla, Australia. A dedicated (and expandable) facility would of course be very desirable but is simply not possible at the present funding level.

Because SETI observations at most sites will be interrupted by other telescope users, it may be desirable to replicate the receiving and signal processing hardware to avoid frequent transport of the equipment.

4. SEARCH STRATEGY

Almost all searches to date have used large directive antennas pointed at known solar type stars. The NASA program combines two strategies: a targeted search of late F, G, and early K stars in the solar neighborhood and an all-sky survey. In addition to about 800 solar type stars (and many more if the catalogs can be extended) the targeted search will look at a number of spectrally peculiar stars and may also scan nearby galaxies. With only approximately 1000 targets, the targeted search can afford to spend considerable time on each. If 1000 seconds is devoted to each, a complete search at one frequency band takes only about 12 days. If each frequency band is 10 MHz wide one can cover the water hole in a year, assuming 100% telescope availability.

To cover the entire sky, the sky survey must point the antenna in a number of directions roughly equal to its gain. If the gain is 10^6 and the time per direction (or the time to scan the full width half power beamwidth) is one second, the total time required is the same for both strategies. Because the sky survey can afford much less time per direction the sensitivity is perforce lower. On the other hand, if powerful sources exist that are not in the direction of known solar type stars, the targeted search might miss them and the sky survey might not. In addition the sky survey, being more sensitive and having higher frequency resolution than previous radio astronomy all sky surveys, may well discover objects of astronomical interest.

5. THE MULTI-CHANNEL SPECTRUM ANALYZER

At the sensitivities obtainable with existing radio astronomy antennas we have virtually no hope of detecting broad band modulation such as audio or video signals. We will be lucky if we can detect continuous wave (CW) signals or pulses whose spectra are as narrow as their duration and shape permit. Such signals might be radiated as beacons to attract our attention or they might be radiated for the use of the other society, even as we ourselves are doing today.

To detect a monochromatic (CW) signal one wishes a receiver with a bandwidth as narrow as possible. But since the signal can be anywhere in a band thousands of megahertz wide it would be hopelessly slow to tune a single channel narrow band receiver across the wide spectrum. What we need is a high resolution radio frequency spectroscope: a device that will permit resolving some tens of megahertz of spectrum simultaneously into millions of individual channels on the order of 1Hz wide. This is the function of the multi-channel spectrum analyzer (MCSA).

This instrument will be described in detail in a later paper in this session by Peterson, et al. Suffice it to say here that the present design resolves about 8 MHz of spectrum into about 8 million channels 1 Hz wide and into about 8000 channels 1024 Hz wide or it can provide an intermediate number of channels at intermediate bandwidths. The plan has been to provide one intermediate bandwidth of 32 Hz.

By adding 2 or 4 adjacent channels, or "bins", one can, with relatively little loss, obtain "pseudo-bins" twice or 4 times the width of the true bins. Or by adding 2 or 4 successive samples from a single bin one can, again with only a small penalty, obtain pseudobins 1/2 or 1/4 the width of the true bin. Thus with true bins having widths of 1, 32, and 1024 Hz we can have effective bin widths from 1/4 to 4096 Hz wide in octave steps. This allows one to provide an approximately matched filter for pulses having durations ranging from 250 μ seconds to 4 seconds.

Since the detectability of a pulse depends only on its total energy a pulse 250 μ seconds long would have to have 16,000 times the power of one 4 seconds long to be equally detectable. Thus, it may be more realistic to provide true bins with 1, 8, and 64 Hz bandwidths and to go only by a factor of 2 in each direction by pseudobinning.

The first step in the detection of a signal occurs when the output of a bin exceeds some threshold set at a level such that the probability of noise alone exceeding it (and thus causing a false alarm) is tolerably low. This thresholding operation is done on-board the MCSA and only the hits exceeding threshold are reported out.

Although the MCSA performs a complex Fourier transform and delivers for each bin the real and imaginary amplitudes, these are squared and added to give the power and the thresholding is performed on this power. The performance is then identical to a sin x/x filter followed by a square law detector.

6. CW DETECTION

The signal-to-noise ratio needed to exceed threshold for a CW signal can be reduced by adding or averaging several successive time samples, i.e. by integration of the output of each bin. Figure 1 shows the input signal to noise ratio (SNR) needed to have a 50-50 chance of detecting or missing the signal when the threshold is set to give a probability of a false alarm per bin, p_{fa}, equal to the value shown for each curve. The abscissa is the number, n, of samples that are accumulated.

If we had a billion bins in our accumulator we might want $p_{fa} = 10^{-10}$. This would give a false alarm in one out of 10 accumulations. We see that for a single sample (n=1) the SNR must be 13.53 dB while if n = 100 the SNR need be only -1.09 dB so the system sensitivity has been increased by 14.62 dB. If there were no problem of frequency drift, simple accumulation of samples would be a good method of detecting CW signals.

Good frequency standards have stabilities on the order of a part in 10^{15} over a 1000 second interval. Thus a transmitter and receiver controlled by frequency standards might have an rms combined drift at 1.5 GHz of about 0.002 Hz in a 1000 second observation time.

A much more important source of drift is Doppler drift due to acceleration of the transmitter and receiver along the line of sight. For the diurnal rotation of the Earth $\dot{\nu}/\nu$ can be as much as 10^{-10}/sec which, at 1.5 GHz would produce a drift of 150 Hz in 1000 seconds. Other, more rapidly rotating planets could produce several times this amount.

We can, of course, compensate for Doppler drift in the direction we are observing by appropriately sweeping our own local oscillator. If the signal we receive is a beacon pointed at us we can assume that the ETI will have corrected for their Doppler. In this case there is no problem. But if the signal comes from an omni directional beacon, or is merely intended for their own use, there may be appreciable Doppler drift.

The Cyclops system proposed to detect drifting CW signals by storing each time sample of the spectrum during the entire observation period of n samples, then adding those spectra with 2n+1 different amounts of skew, thus anticipating any drift rate between -1 and +1 bins/sample. This is a very good method but with n = 100 or, even worse, 1000, it requires a tremendous amount of memory.

We are currently evaluating a method in which the spectrum is accumulated into m registers with occasional shifts left or right to match m equally spaced drift rates that cover the range. The maximum drift matching error in n examples is thus n/2m bins. We then examine the spectra for energy "mounds" up to n/2m bins wide by seeing if running sums up to this length exceed appropriate thresholds giving false alarm probabilities on the order of 0.01.

As a next step we zero in on the drift rate for each of these mounds by assigning portions of our m memories to each mound and drifting the storage process for each mound with a new set of drift rates clustered about the drift rate found in the first phase. Again p_{fa} may be set on the order of 0.01. Any mound that had a signal and survived phase 1 will survive phase 2, but 99% of those that were due to noise alone will be rejected by phase 2.

Phase 2 may be repeated with further refined drift rates. There would then be a probability of 10^{-6} that a surviving peak was due to noise alone.

Finally, with the drift rate(s) now known exactly, the surviving portion(s) of the spectrum may be accumulated with the drift removed and the threshold set to give a $p_{fa} = 10^{-4}$ for an overall $p_{fa} = 10^{-10}$.

Preliminary calculations seem to indicate that the Cyclops system performance can be duplicated with $m \approx 9$ rather than 100. The advantage comes from accepting a lot of early false alarms and then letting the computer weed them out. This seems smarter than erecting a high barrier against all but the strongest events, and is applicable to pulse detection as well. In fact the concept began there.

Drifting CW detection is important only for the targeted search. For the proposed sky survey the binwidths are wide enough and the observation time per beamwidth short enough that drift is no problem. On the other hand a successful sky survey requires that CW signals be present. It would be very unlikely to detect a low duty cycle pulse.

7. PULSE DETECTION

Pulses offer a fundamental advantage over CW signals in that, all other things being equal, they can be detected over greater distances with the same average power, or over the same distance with less average power. From Figure 1 with $p_{fa} = 10^{-10}$ we saw that integrating 100 samples reduced the required SNR by 14.62 dB, from 13.53 dB for $n = 1$ down to -1.09 dB for $n = 100$. But in 100 samples we have received 100 times the energy received in any one sample. If all this energy had been radiated as a pulse during one sample the received SNR would have been $20 - 1.09 = 18.91$ dB. This is 5.38 dB above the SNR required for single sample ($n=1$) detection.

The reason for this is the poor performance of square law detection at low signal-to-noise ratios. With a synchronous detector the required SNR drops 3 db with each doubling of n. Square law detectors approach this same improvement at high SNR's but at low SNR's they approach an improvement that is only half as great, or 1.5 dB for each doubling of n. Thus in going from $n = 1$ to $n = 100$ with a square law detector we did not get 20 db improvement, but only 14.62 dB.

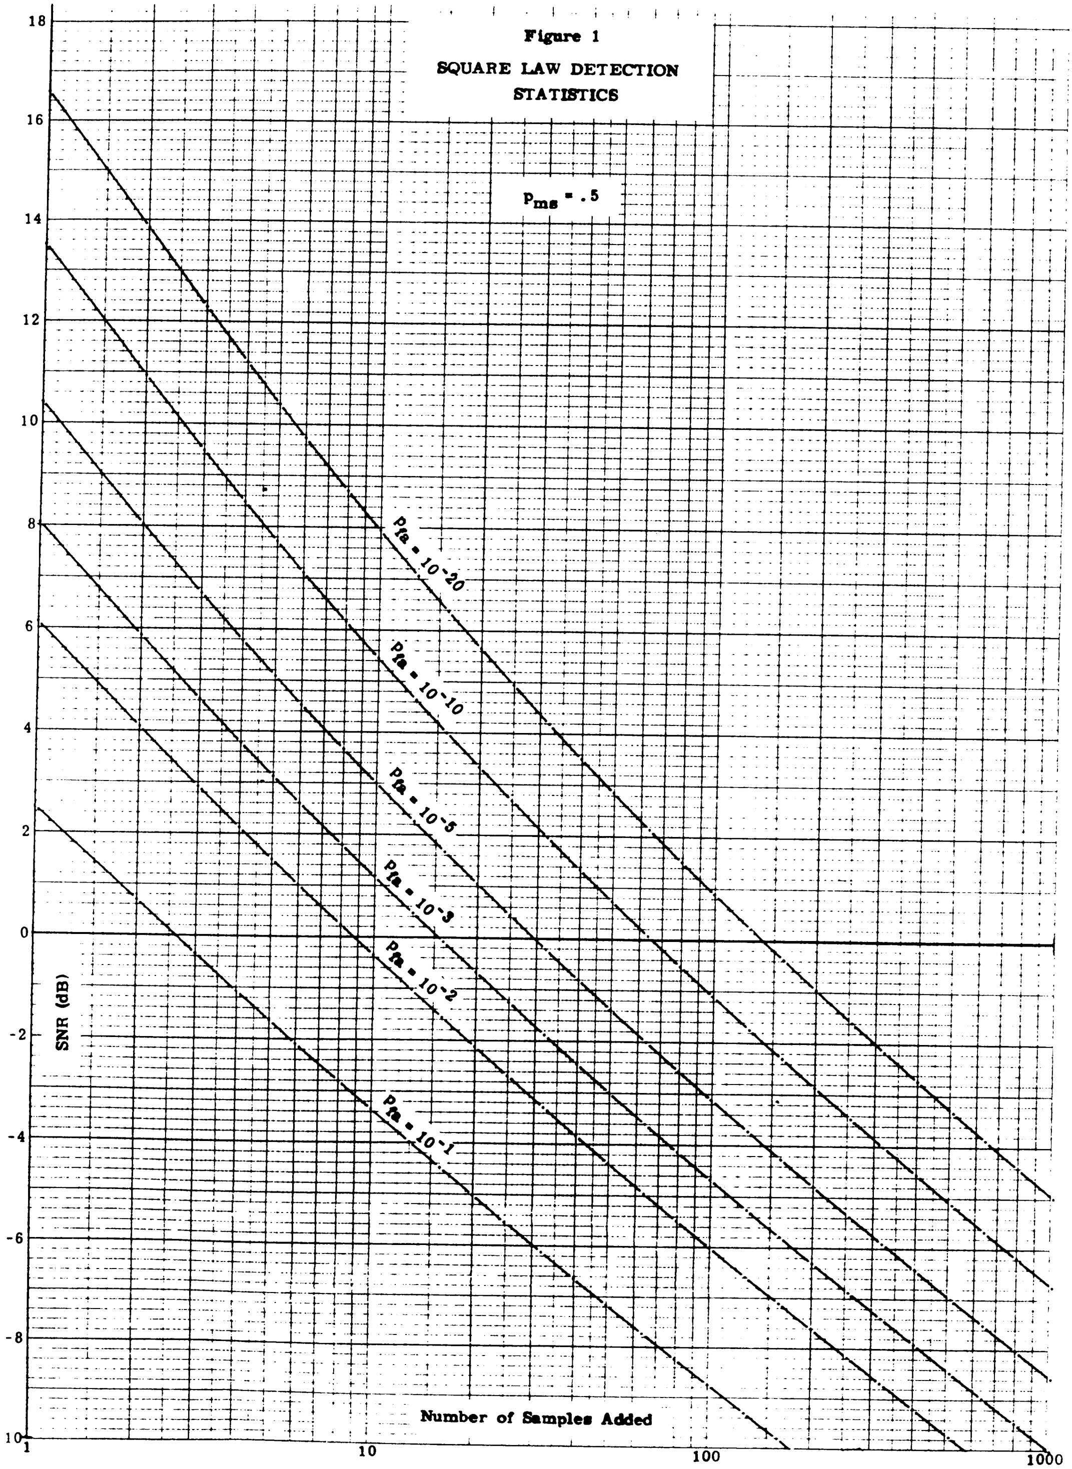
Figure 1
SQUARE LAW DETECTION
STATISTICS
$P_{ms} = .5$
$P_{fa} = 10^{-20}$
$P_{fa} = 10^{-10}$
$P_{fa} = 10^{-5}$
$P_{fa} = 10^{-3}$
$P_{fa} = 10^{-2}$
$P_{fa} = 10^{-1}$
SNR (dB)
18
16
14
12
10
8
6
4
2
0
-2
-4
-6
-8
10
Number of Samples Added
1
10
100
1000

Let us consider rectangular pulses having a duration equal to one sampling time and a repetition period equal to n sampling times. Then the fractional on time, or duty cycle, $\delta = 1/n$. The improvement over a CW signal is then

$$\Delta \text{ SNR} = 10 \log n - \text{SNR}(p'_{fa},1) + \text{SNR}(p_{fa},n) \qquad (1)$$

Here p_{fa} is the false alarm probability per bin per sample that we can afford in CW detection and p'_{fa} is the same quantity in pulse detection.

Assume the number of bins per spectrum is N. Then if we are doing drifting CW detection a la Cyclops to a maximum drift rate of ± 1/2 bin per spectrum, the number of cells in our final accumulations will be ≈ nN. This is also the number of cells in which a false alarm can occur in the pulse case. Thus we have $p'_{fa} \approx p_{fa}$. The improvement given by (1) is then as shown in Figure 2.

On the other hand, if we are not doing drifting CW detection but are merely accumulating successive samples into a simple register we then have that the number of cells for the CW case is N but for the pulse case it is nN. This means that now we should set $p'_{fa} = p_{fa}/n$. The improvement is then less as shown in Figure 3. Note that if $\delta >$ 1/5 pulses are actually slightly worse for the case $p_{fa} = 0.1$.

The detection of a single pulse, however strong, would hardly convince us that we had an artificial signal. We would expect a regularly spaced train of pulses, perhaps alternating irregularly between left and right handed circular polarization to provide an information bearing binary channel. This means that we should look for regularly spaced strings of pulses in the combined pattern of hits exceeding threshold from both the LHC and RHC receivers.

We do not require that the probability of each hit being due to noise alone equal our overall false alarm probability, but only that the joint probability of the entire string being due to noise alone be this low. This means we can lower the threshold to that corresponding to a p_{fa} on the order of 10^{-3} and let the computer reject hits that do not form strings. This means additional sensitivity improvement of around 5 1/2 dB for a total gain of 10 dB over CW detection a la Cyclops with n = 100. The actual threshold reduction and sensitivity improvement will depend upon how rich a field of hits the computer can sift through in real time. Special purpose hardware will certainly be required.

This pulse detection approach, based on a suggestion by John H. Wolfe, was developed by Kent Cullers and is described by him in a companion paper.

8. CONCLUSION

The proposed NASA observing program for SETI is a multiple approach that places fewer restrictions on the nature of the signal and its point of origin than any known previous or concurrent search. In one phase of the program the entire sky will be scanned for CW signals at all frequencies in the terrestrial microwave window from 1 to 10 GHz.

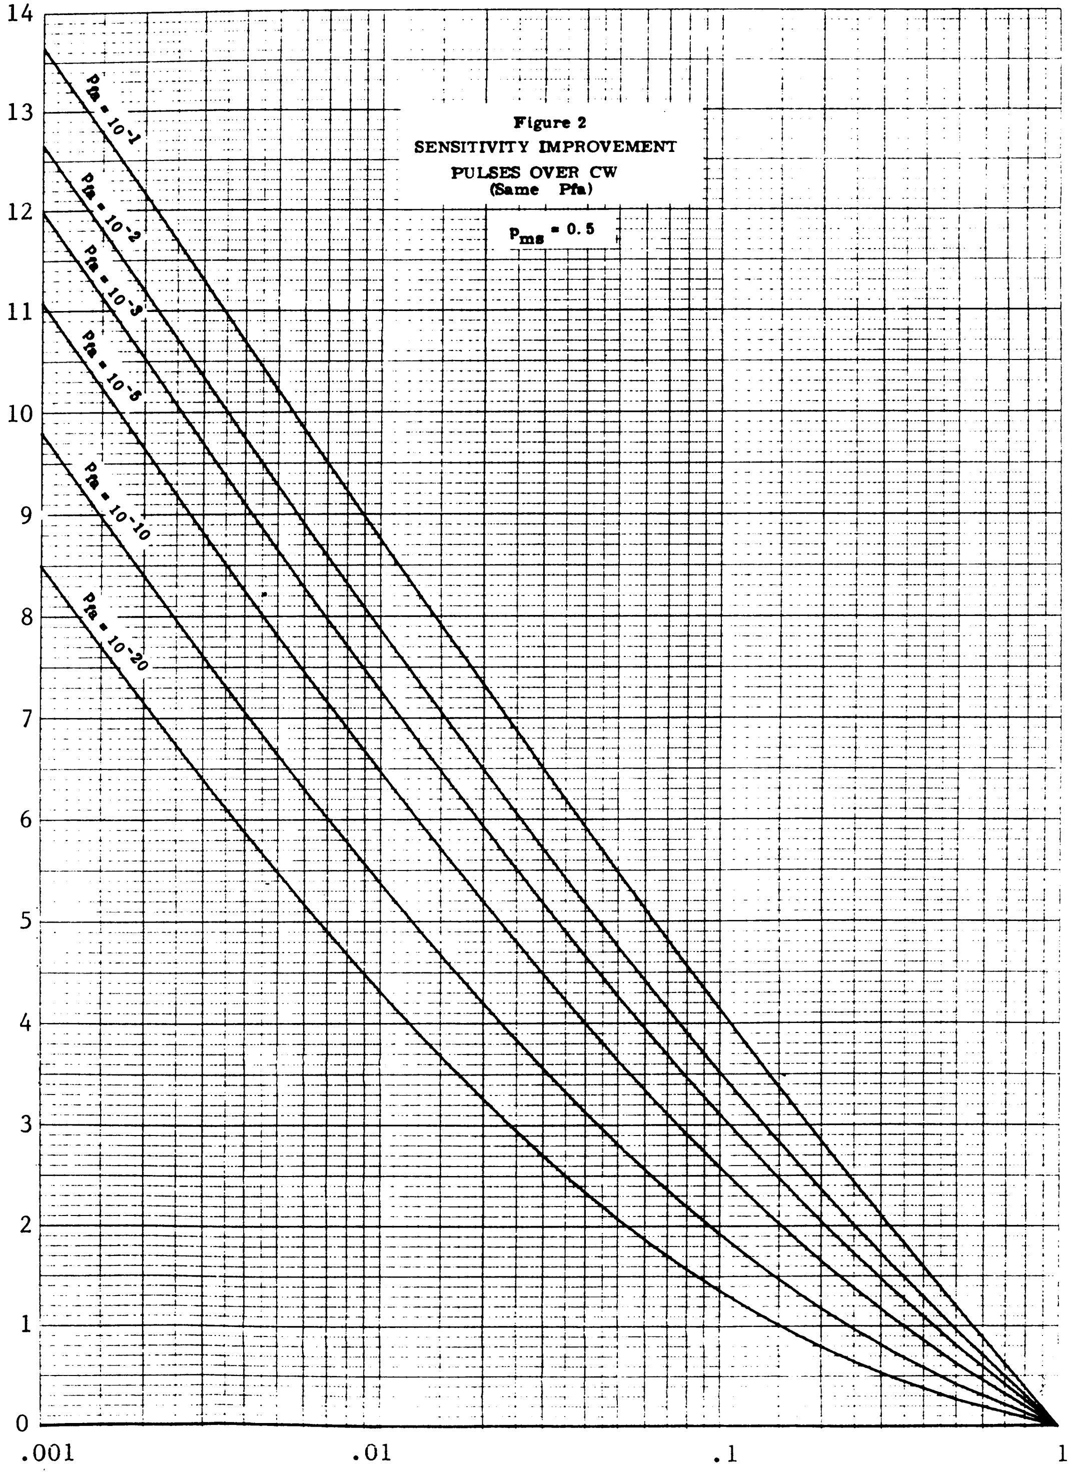

Figure 2
SENSITIVITY IMPROVEMENT
PULSES OVER CW
(Same Pfa)

$P_{ms} = 0.5$

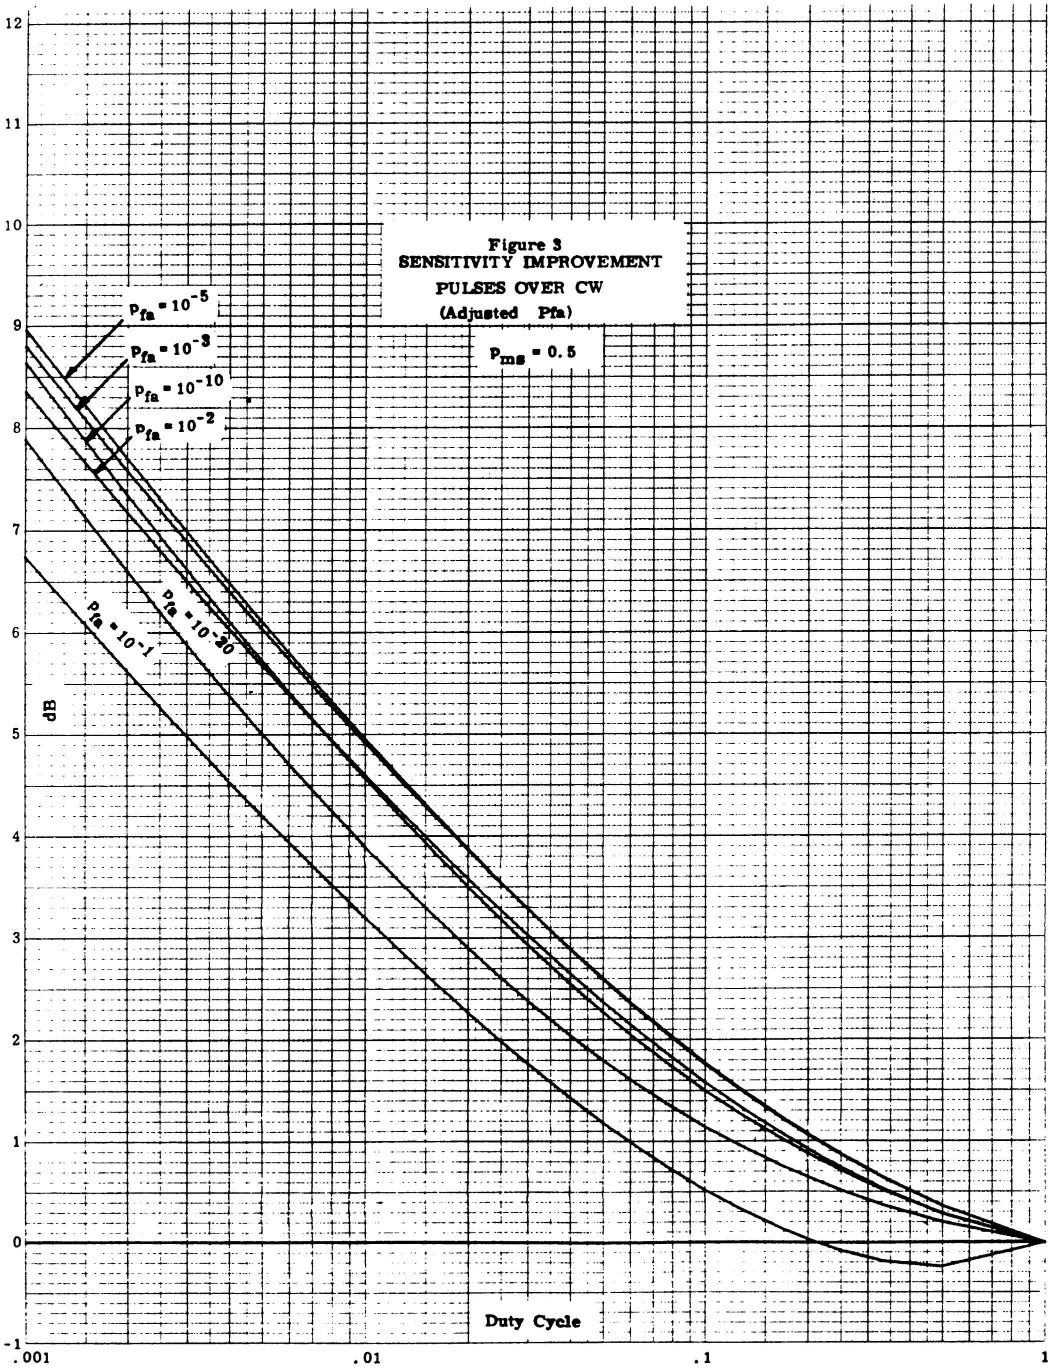

Figure 3
SENSITIVITY IMPROVEMENT
PULSES OVER CW
(Adjusted Pfa)

In the other phase of the program all solar type stars in any catalog (as well as many peculiar objects) will be searched not only for drifting and non-drifting CW signals but also for drifting or non-drifting pulses covering a wide range of durations and repetition rates. The targeted search will concentrate on the low frequency end of the microwave window including the water hole but will cover other "magic" frequencies and as much additional spectrum as time allows. Both right and left handed circular polarizations will be observed simultaneously.

Algorithms are being developed that, for the first time in SETI history, will allow the computer following the spectrum analyzer to weed out (in real time) false alarms both in CW and in pulse detection. This allows a reduction in the initial threshold and a consequent increase in detection sensitivity.

REFERENCE

1. OLIVER, B.M., and BILLINGHAM, J. (Eds.), Project Cyclops, NASA CR11445, Revised Edition (1973).