

2.

CR-134011

STATISTICAL EVALUATION
 OF TIME SERIES
 ANALYSIS TECHNIQUES

Prepared for:
 NASA JOHNSON
 SPACECRAFT CENTER
 Under Contract
 NAS 912902

(NASA-CR-134011) STATISTICAL EVALUATION
 OF TIME SERIES ANALYSIS TECHNIQUES
 (Behavioral Technology Consultants,
 Silver) 47 p HC \$4.50 CSCL 06P

N73-30995

G3/04 Unclas
 14602

Prepared by:
 Vernon A. Benignus

BEHAVIORAL TECHNOLOGY CONSULTANTS



TABLE OF CONTENTS

	Page
INTRODUCTION	1
PURPOSE	2
PROCEDURES	2
General Spectrum Model	2
Multiple Variable Prediction	4
A Realization of the Multiple Method	9
Monte Carlo Runs	10
Period and Frequency Domains	12
Criteria of Performance	13
RESULTS	15
Criterion For Acceptance of a Peak	15
A Method of Detecting Some Forms of Program Failure	18
Too Few or Many Peaks	23
Multiple Correlation as a Criterion of Fit	24
Other Attempted Improvements	30
Period vs Frequency Domain	34
Other Performance Characteristics	36
Comparison to FFT	36
Amplitude Estimates	38
A New Program	38
CONCLUSION	43

VOLUME II

APPENDICES

1. Listing of Program and Documentation of Use	44
2. Three Examples of Use of SPECT	45

LIST OF FIGURES

FIGURE		PAGE
I	Approximate Main Lobe Shape of a Spectrum Window	5
II	Normalized Plots of Probability of Periods	17
III	Probability Distributions of Period Lengths (corrected and uncorrected)	20
IV	Probability Distributions of Period Lengths	22
V	Period Discrimination at Various Levels of SNR	26
VI	Multiple Squared Correlation as a Function of SNR and Number of Periods	29
VII	Average Windows for the FFT and the NASA Multivariate Program	37
VIII	Distribution of Peak Amplitudes for Various Levels of SNR .	39

INTRODUCTION

It is frequently desirable to detect small changes or shifts of frequency in circadian biological rhythms, especially where there has been some alteration in extrinsic factors which might influence such rhythms. One of the more useful methods employed to analyze biological data for the detection and quantification of circadian rhythms is some form of spectrum analysis (Frazier, Rummel, and Lipscomb, 1968). In standard forms of spectrum analysis, it is possible to resolve or discriminate between two sinusoidal frequencies separated by Δf where

$$\Delta f = \frac{1}{T}, \quad [1]$$

and T is the length of time of the time series record being analyzed (Bendat and Piersol, 1966). Among the many problems in biological data acquisition, one of them is that of obtaining records of long duration. This implies that for most circadian rhythm work Δf , the resolution of the analysis program, will be quite large, due to short time series records. This report is an evaluation of and an improvement upon a spectrum analysis model which achieves finer resolution than $\Delta f = 1/T$ by the use of multiple least squares prediction models.

PURPOSE

The specific purposes of this study were: (1) To further evaluate a model for spectrum analysis, using a multiple regression model, developed by Rummel (1971) and previously tested in a preliminary way (Benignus, 1972); (2) To attempt various modifications and improvements upon this model; and (3) To write an improved program for this model. These objectives have been met and the performance of the analysis model has been documented and compared to the more standard FFT (Fast Fourier Transform) by use of Monte Carlo techniques.

PROCEDURES

General Spectrum Model

The general model for a time series as expressed in the frequency domain is

$$f(t) = K + \sum_{j=1}^J [a_j \sin(\omega_j t) + b_j \cos(\omega_j t)], \quad [2]$$

(Bendat and Piersol, 1966), where $\omega = 2\pi f$, j is the frequency index, $0 < j \leq J$ and k is the D. C. component of the mean of the data. The usual approach to the spectrum analysis of $f(t)$ is analogous to

discrete Fourier analysis where the coefficients in 2 are estimated by

$$\hat{a}_j = \frac{T}{\sum_{t=0}^T} [f(t)\sin(\omega_j t)] \quad [3]$$

$$\hat{b}_j = \frac{T}{\sum_{t=0}^T} [f(t)\cos(\omega_j t)] \quad , \quad [4]$$

where the carat indicates an estimate. It may be shown that equations 3 and 4 are univariate least squares estimates derived from standard regression theory. These estimates of "real" and "imaginary" amplitude (a_j and b_j , respectively) are usually combined to yield

$$\hat{P}_j = \hat{a}_j^2 + \hat{b}_j^2 \quad , \quad [5]$$

the estimated power in frequency power in frequency band j or

$$\hat{A}_j = \sqrt{\hat{a}_j^2 + \hat{b}_j^2} \quad , \quad [6]$$

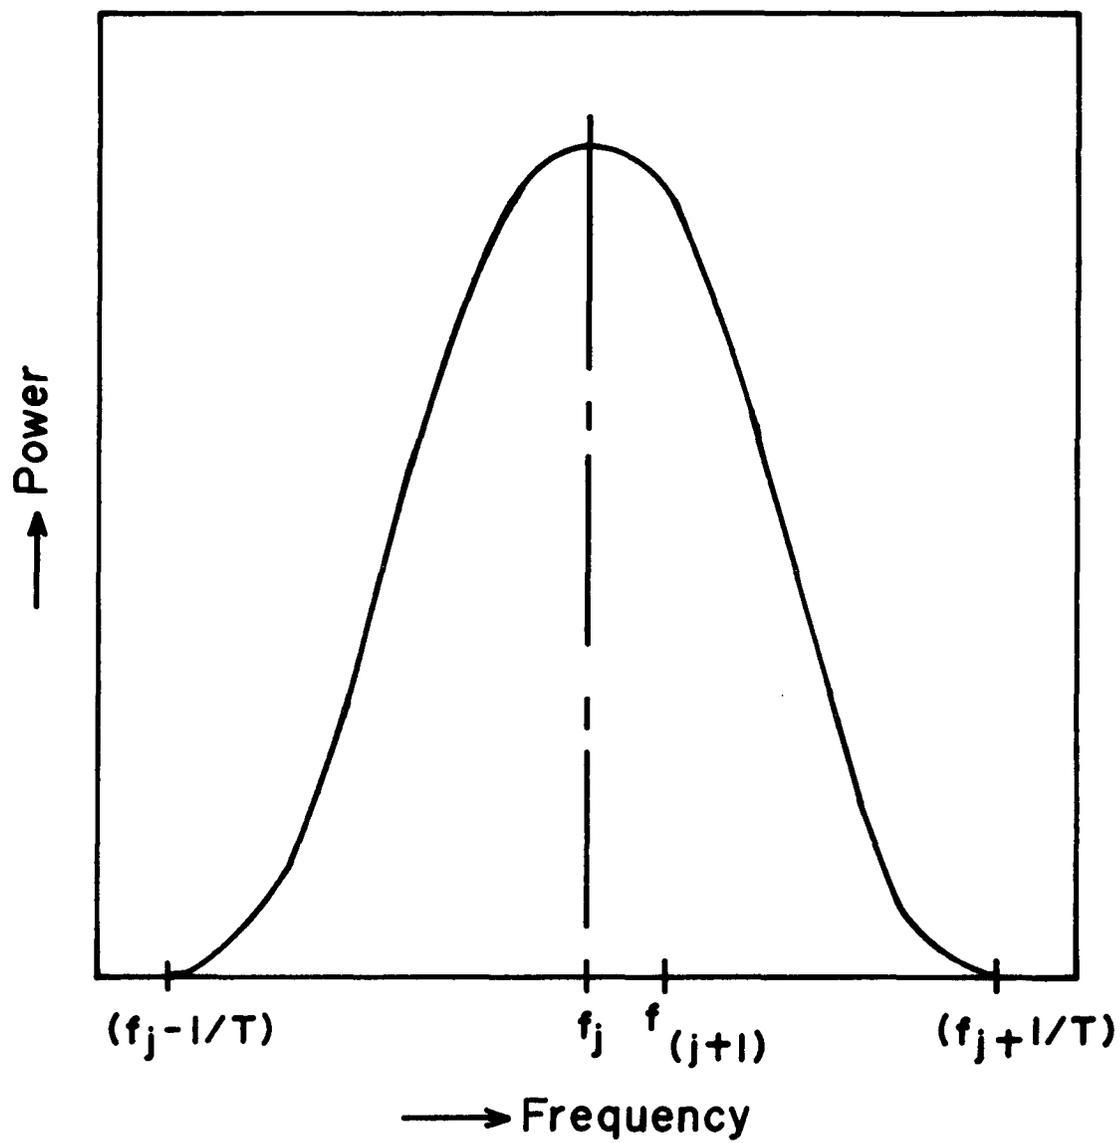
the estimated amplitude in frequency band j (Bendat and Piersol, 1966).

It may be shown that two estimates, \hat{P}_j and $\hat{P}_{(j+1)}$ are orthogonal (uncorrelated) if their corresponding frequencies ω_j and

$\omega_{(j+1)}$, or of course f_j and $f_{(j+1)}$, are spaced such that $\Delta f = 1/T$. If, in the time series $f(t)$ there exist two signals separated in frequency by much less than Δf , then power estimates at those two frequencies \hat{P}_j and $\hat{P}_{(j+1)}$ can be expressed as a continuous function called a frequency domain or spectrum "window", the main lobe of which is shown in Figure 1. The window function shows that for any estimate, \hat{P}_j , if $f(t)$ contains a signal the frequency of which can take on values of $f_j \pm 1/T$, then the value of \hat{P}_j is a function of the true signal frequency. Similarly, if two estimates, \hat{P}_j and $\hat{P}_{(j+1)}$ were made at f_j and $f_{(j+1)}$ as shown in Figure 1, and there were two frequencies present in $f(t)$ at f_j and $f_{(j+1)}$, then the two estimates would be nearly equal because data from the signal at f_j are included in $\hat{P}_{(j+1)}$ and vice versa.

Multiple Variable Prediction

Equations [3] and [4] are univariate prediction equations. They are combined in (5) and (6) as a two variable prediction scheme where the two variables are orthogonal. In usual multiple regression, least squares prediction schemes it is possible to use several predictors simultaneously to estimate the dependent variable. In these cases the several predictors may or may not



APPROXIMATE MAIN LOBE SHAPE OF A SPECTRUM WINDOW

FIGURE 1

be correlated. However, when the predictors are highly inter-correlated, estimates of each predictor's contribution, based on separate univariate estimates, are very inaccurate (Draper and Smith, 1966). When k nonindependent predictors are used simultaneously to estimate a dependent variable, the contribution of each is called the "partial regression weight." This regression coefficient is a least squares estimate of the contribution of a given predictor k with the effects of all the other $k-1$ predictors "accounted for" or "statistically held constant" (Guilford, 1950).

Instead of using univariate predictors such as [3] and [4] to estimate the contribution of a sine/cosine pair of frequency j to $f(t)$, a multiple prediction scheme might be used. In a multiple prediction scheme for estimation of a_j one would not only use a sine wave of frequency j but would include sine waves of several different frequencies in a simultaneous prediction equation. For example if a two predictor scheme were used, then a normalized form of a_j would be estimated using

$$\hat{a} = \frac{R_{dj} - R_{dk}R_{jk}}{1 - R_{jk}} \quad [7]$$

where the R quantities are Pearson product moment correlation coefficients between the variables indicated in the subscripts and the three variables involved are (1) the dependent variable, $f(t)$ which is identified by the subscript "d"; (2) the sine wave of frequency j , the one whose contribution is estimated as \tilde{a}_j and (3) the sine wave of frequency k , the other simultaneous predictor, the effect of which is to be "controlled" or "accounted" for. Examination of [7] reveals that not only is the relation of a sine wave of frequency j to $f(t)$ considered, R_{dj} , but the relations of the other predictor wave to $f(t)$ and the interpredictor relations are also considered. Thus if $R_{jk} \neq 0$, as it will not be if $\Delta f < 1/T$, then this "overlap" will be considered in estimating a_j . This multiple prediction scheme is shown to have a higher resolution than equations [3] and [4]. A similar procedure would be used for estimating b_j , the cosine components. When more than two predictors are used in a simultaneous prediction scheme, matrix methods for estimating the contributions of each predictor must be used as shown in [8] and [9]. The expression for the standardized regression weights in this model is

$$\beta = R_{jj}^{-1} R_{1j} \quad [8]$$

where β is the standardized regression weight vector; R_{jj} is the matrix of correlations among predictors; and R_j^1 is the vector of correlations between the predictors and the dependent variable, $f(t)$. The predictors in this model consist of sin and cos pairs at each of K frequencies. Thus, if there are K frequencies, the order of R_{jj} is $J = 2K$. The matrix of predictors consists of a J by N matrix where N is the number of observations in $f(t)$. The amplitude estimates are computed from the J length vector β by first converting each β weight to a deviation score weight using

$$B_j = \beta_j \frac{S_1}{S_j} \quad [9]$$

where S_1 is the standard deviation of $f(t)$ and S_j is the standard deviation of predictor j , (.707). Then the sin/cos components are combined according to

$$A_k = \sqrt{B_j^2 + B_{j+1}^2} \quad [10]$$

where A_k is the amplitude in frequency band K , B_j is the sin component, and B_{j+1} is the cos component. Here $k = j/2$ where $j = 2, 4, 6, \dots, j$.

A Realization of the Multiple Model

The particular program in this study was designed along the lines of a multiple predictor least squares theme as outlined above.

The procedure of the program is as follows:

- (1) compute a spectrum using one frequency at a time as in equations [3] and [4] ;
- (2) examine this spectrum to locate peaks which exceed a statistical criterion of significance;
- (3) compute a new spectrum where each frequency's contribution, A_j , was evaluated with the contribution of all other significant peaks held constant by the use of multiple least squares prediction as shown in equation [8] ;
- (4) return to step (2) and continue to loop through the procedure until no new peaks are found.

In addition to the above procedure, each time step (3) is executed, the frequency value for the significant peaks is moved up and down around the original value with the spectrum being recomputed for each new frequency component. By moving the significant frequency bands up and/or down around their original

values, it is possible to attempt to optimize the fit of the predictors to the empirical wave $f(t)$. Optimization was carried out by moving a frequency band, testing the spectrum result using the new band for improved fit to $f(t)$, continuing in the same direction if improvement resulted or trying the other direction if the fit was poorer. In this way each band in the multiple prediction model was shifted around in frequency to optimize the model's fit.

A computer program has been written to realize the multiple regression model discussed above. The program consists of a mainline calling IBM Scientific Subroutines (SS) computing multiple regression.

Monte Carlo Runs

In order to evaluate the performance of the multiple predictor spectrum analysis program it is necessary to analyze data which approximate that on which the program will be used. Biological signals which are subject to circadian or other periodic variation can be modeled using a "source of variance" model such as

$$V(\text{Total}) = V(\text{Periodic}) + V(\text{Unaccounted}) \quad [11]$$

where $V(\text{Total})$ is the total variance (or power) in the wave,

V(Periodic) is that portion or component of the wave which is due or correlated with such periodic components as diurnal cycling and V(Unaccounted) are such sources of variation as short term fluctuations due to stress, homeostatic fluctuation, and in general, any source of variation not related to periodic cycles. In this discussion V(Unaccounted) will be referred to as either noise or error variance. The general effect of noise in the biological signal is to "mask" the periodic component both with respect to amplitude and frequency. This results in unreliable estimates of variance in the power spectrum since, as with most transformations, the Fourier transformation has as much variance in the resultant as in the original data.

In this paper data were constructed using [11] as a model. The periodic component, V(Periodic), was simulated by generating a sine wave of a particular frequency. The noise, V(Unaccounted), was simulated by a white Gaussian noise generated by sampling from a random number table (Rand Corp., 1955) which was punched onto cards and loaded into a disk file. It is fully realized that biological noise might not have a white spectrum or have a Gaussian distribution. It is true, however, that the assumption of white, Gaussian noise is usually made and it was felt that

the program should be evaluated on "fair" theoretical grounds.

When random noise is involved in data to be analyzed, it is the long-range, average results which are of interest as well as the variation around these averages. The variation around average results is sometimes expressed as variance, error, confidence intervals, failure rates, etc. In order to assess the program's average performance and variance about these averages, a series of records were generated according to [11]. For each of 100 records the noise was obtained by sampling from a unique section of random number table. Each record had a length of 100 observations; the sinusoids which were used as the signals (sine waves) were at various frequencies. Several SNR (signal to noise ratio) levels were used.

Period and Frequency Domains

The theoretical relations worked out for spectrum analysis are usually expressed in terms of the frequency of sinusoids. It often becomes convenient for one purpose or another to express rate of oscillation in terms of the length of one cycle (period length). Spectra where estimates are spaced along equal increments of frequencies, Δf , will be referred to as frequency domain

data. Spectra where estimates are spaced along equal increments of periods, ΔP , will be called period domain data. Most of the results of this study will be given in the period domain but frequent comparisons and references will be made to the frequency domain.

Criteria of Performance

Several aspects of the performance of the spectrum analysis program were used as criteria as follows: (1) finding the correct frequency of the periodic components, (2) finding the correct amplitudes of the periodic components and (3) program failure to find too few or too many peaks. The average performance as well as the variability about the averages is described for (1) and (2) above and failure probabilities are given for (3).

On any given wave the program generates a spectrum which shows significant amplitude peak(s). Due to the noise component there is some error in the frequency of a peak. This error (variability) is described in terms of the relative number of times that the program made various degrees of error. There is a probability of error for each particular degree of error and

and a graph of error probability vs degrees of error constitutes the probability distribution of frequency errors. Optimally one would want this distribution to be peaked around a mean of zero (high probability of zero error) and to have a narrow width (lower probability of error the greater the degree of error). This probability distribution of frequency errors is analogous to a frequency domain window except that it refers to probability of errors in finding narrow peaks rather than amplitudes.

On any given analysis, the amplitude of any peak hopefully approximates the correct amplitude of the signal but will frequently be greater or less due to noise. In order to evaluate the variability around the correct amplitude the probability of an amplitude estimate falling into a certain amplitude range or category can be computed. Here a distribution of probabilities can be graphed and it would be desirable for this distribution to be peaked around the correct amplitudes (high probability of finding the correct amplitude) and have a narrow width (lower probabilities of finding amplitudes, the farther the value deviates from the correct amplitude). This is essentially the sampling distribution of amplitudes from which confidence limits would be computed.

RESULTS

Criterion For Acceptance of a Peak

In the Rummel program an initial test is made on each peak found in the univariate spectrum to decide whether or not a peak is sufficiently large to be considered statistically significant and hence retained as a predictor. Similarly, each time the program attempts to optimize the period value of a particular peak and a new optimized multiple predictor spectrum is computed, each empirical peak found is again checked for significance. The criterion used for acceptance or rejection of a peak as significant, is the value of Student's t for that peak. Since the peak is a regression coefficient, the significance of a regression coefficient can be evaluated using a form of Student's t (Hays, 1963, p. 521). The particular form of t -test used here is equivalent to

$$t = \sqrt{t_s^2 + t_c^2} \quad [12]$$

where t_s is the t -value for the sine wave predictor at a particular peak's period and t_c is the corresponding cosine predictor's value. The criterion used for the univariate spectrum was

called CHEK and the criterion for the multivariate spectrum was called CHEX.

The effect of various levels of CHEK and CHEX was evaluated. The methods for evaluating these criterion levels were essentially the same as in the previous work. Monte Carlo methods were used with sine waves at two period lengths, 23 and 27, mixed with Gaussian random noise. One hundred points were used in each time series and 100 time series were used for the evaluation of the program's performance. ΔP was set at 0.5. Performance was evaluated at SNR (signal to noise ratios) of 1/.5 and 1/1.

For SNR 1/.5, there were very few differences in probability distribution for periods and none for probability or failure across various values of CHEK and CHEX.

For a signal/noise ratio of 1/1 however, the differences for various levels of CHEK, CHEX were more noticeable, but only very extreme values were meaningfully different. Figure 2 shows the graphs for the four levels of CHEK, CHEX. Values for CHEK, CHEX of (2.56, 2.56) yield a probability distribution which indicated considerably poorer resolution than other values.

RELATIVE PROBABILITY (NORMALIZED FOR PEAKS TO FACILITATE COMPARISONS)

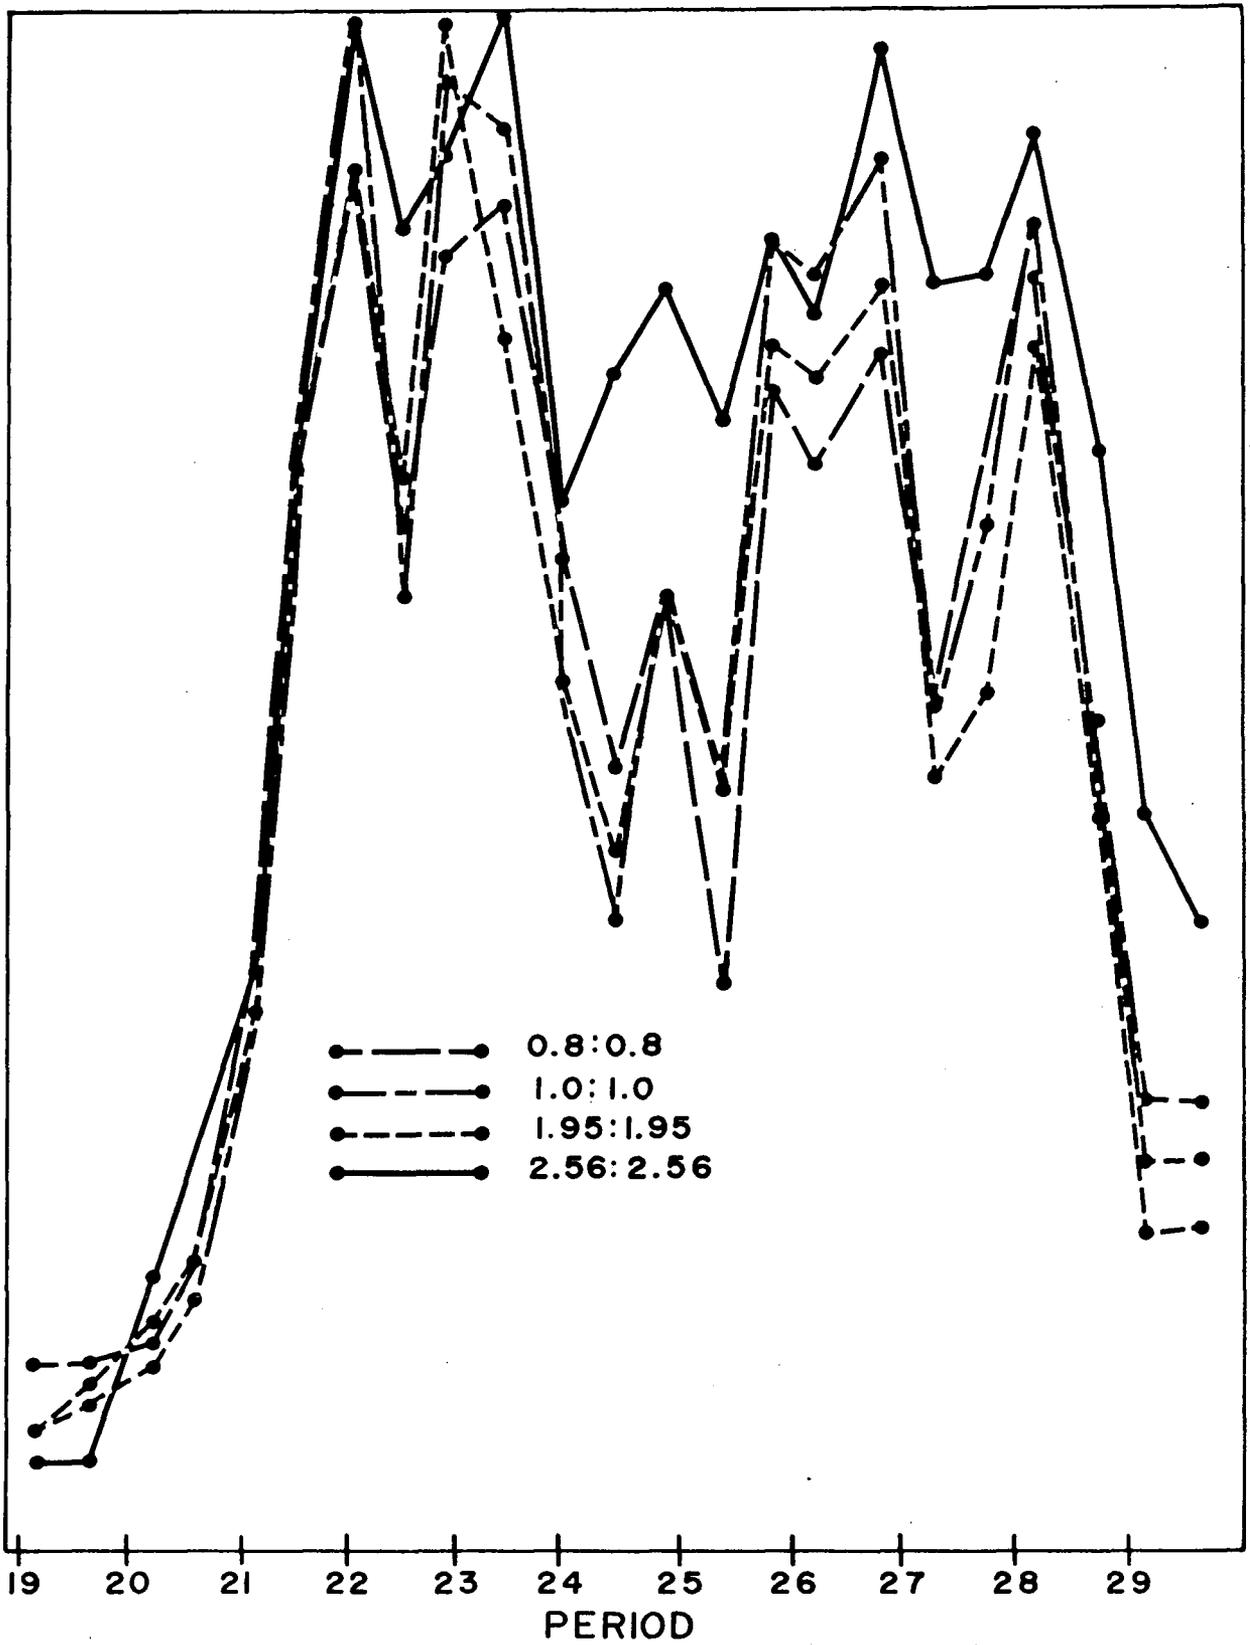


Fig. 2 Normalized Plots of Probability of Periods

SNR = 1.0 : 1.0

The next poorest result is probably that yielded by CHEK, CHEX values of (0.8, 0.8). While the other curves are only marginally different, it would appear that the values of CHEK, CHEX of (1.95, 1.95) are best. Table 1 shows the probabilities of failure for the various values.

Table 4

CHEX	0.80	1.00	1.95	2.56
0.80	0.18			
1.00		0.18		0.37
1.95			0.12	

Examination of this table verifies the above conclusions suggesting that values of (1.95, 1.95) yield the lowest probability of program failure.

A Method of Detecting Some Forms of Program Failure

It was observed in many of the Monte Carlo runs that the program sometimes yielded spectrum results in which the sum of the significant amplitudes was considerably greater than the total amplitude of the original time function which had been analyzed.

As usual, the time series were composed of the sum of the two sine waves and white Gaussian random noise. In order to determine what feature of the signals might have triggered this peculiar failure, the signals of several of the successful runs and the signals of several of the "failure" runs were plotted for inspection. From inspection of such records it was not possible to deduce the causes of the failures.

Since the root mean squared (RMS) amplitude of the time series wave is known, it is possible to compare the total of the significant peaks against this value and reject erroneous results even when the true nature of the time series data is unknown. Thus, while this kind of failure occurs unpredictably, its occurrence is at least detectable and the results can be disregarded. It was therefore decided to recompute the operating characteristics of the program by eliminating all spectra in which the total of the significant peaks was 20% or more greater than the RMS amplitude of the time series. The selection of a 20% cut point was entirely arbitrary.

After having eliminated cases of detectable failure, a new period discrimination window was computed. Figure 3 shows the

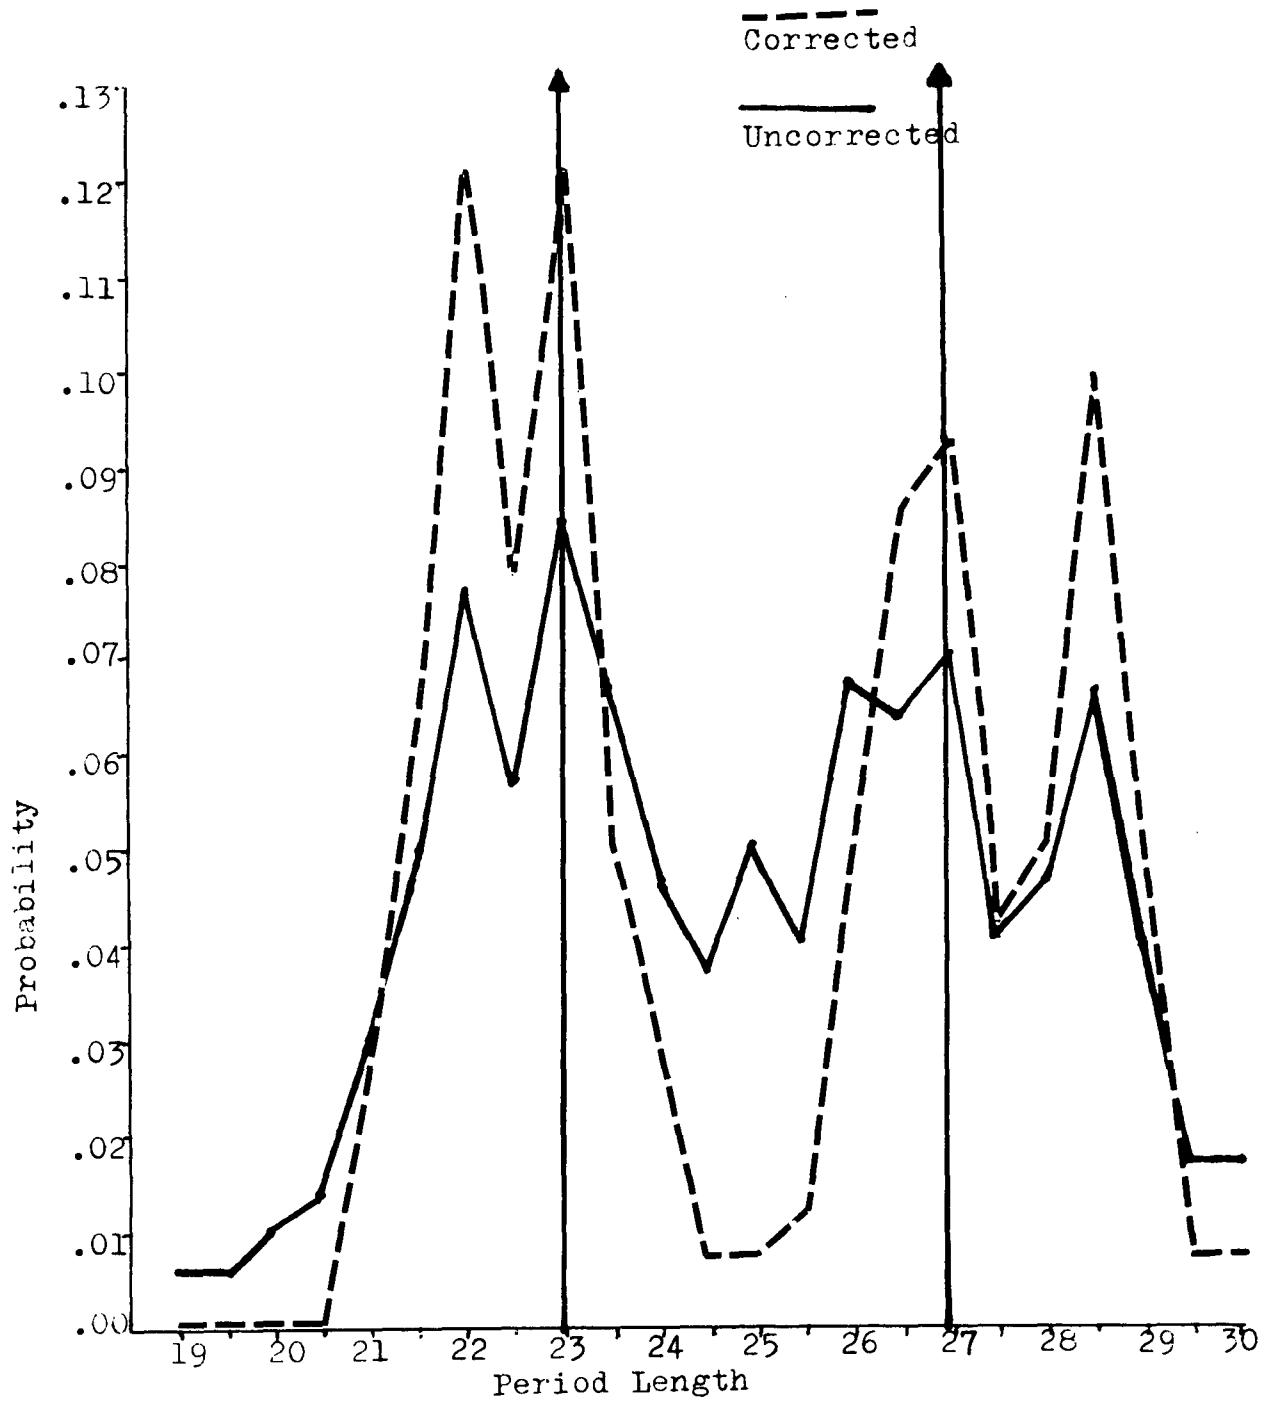


Figure 3. Probability distributions of period lengths, corrected and uncorrected.

period discrimination window for $SNR = 1/1$ as reported in an earlier interim report, as compared to the window obtained for the same SNR , but with the deviant cases removed. It is readily seen that there is better separation between the two periods, 23 and 27, with the deviant data removed. Apparently, those cases were also deviant in that wrong periods were found. Of course, the latter condition would not have been independently detectable. Thus, by eliminating deviant cases, considerable improvement in period discrimination has been attained.

In order to explore the operating characteristics of the program further, period discrimination windows were computed for $SNR = 1/0.5$, $1/1$ and $1/1.5$. The results are displayed in Figure 4. As can be seen, when $SNR = 1/0.5$, the discrimination is excellent and it remains reasonably good for $SNR = 1/1$. The discrimination for $SNR = 1/1.5$, while still better than theoretically expected, is not impressively sharp. It therefore appears that $SNR = 1/1.5$ is about as high an SNR time series as ought to be analyzed. Subsequently presented data will bolster that conclusion.

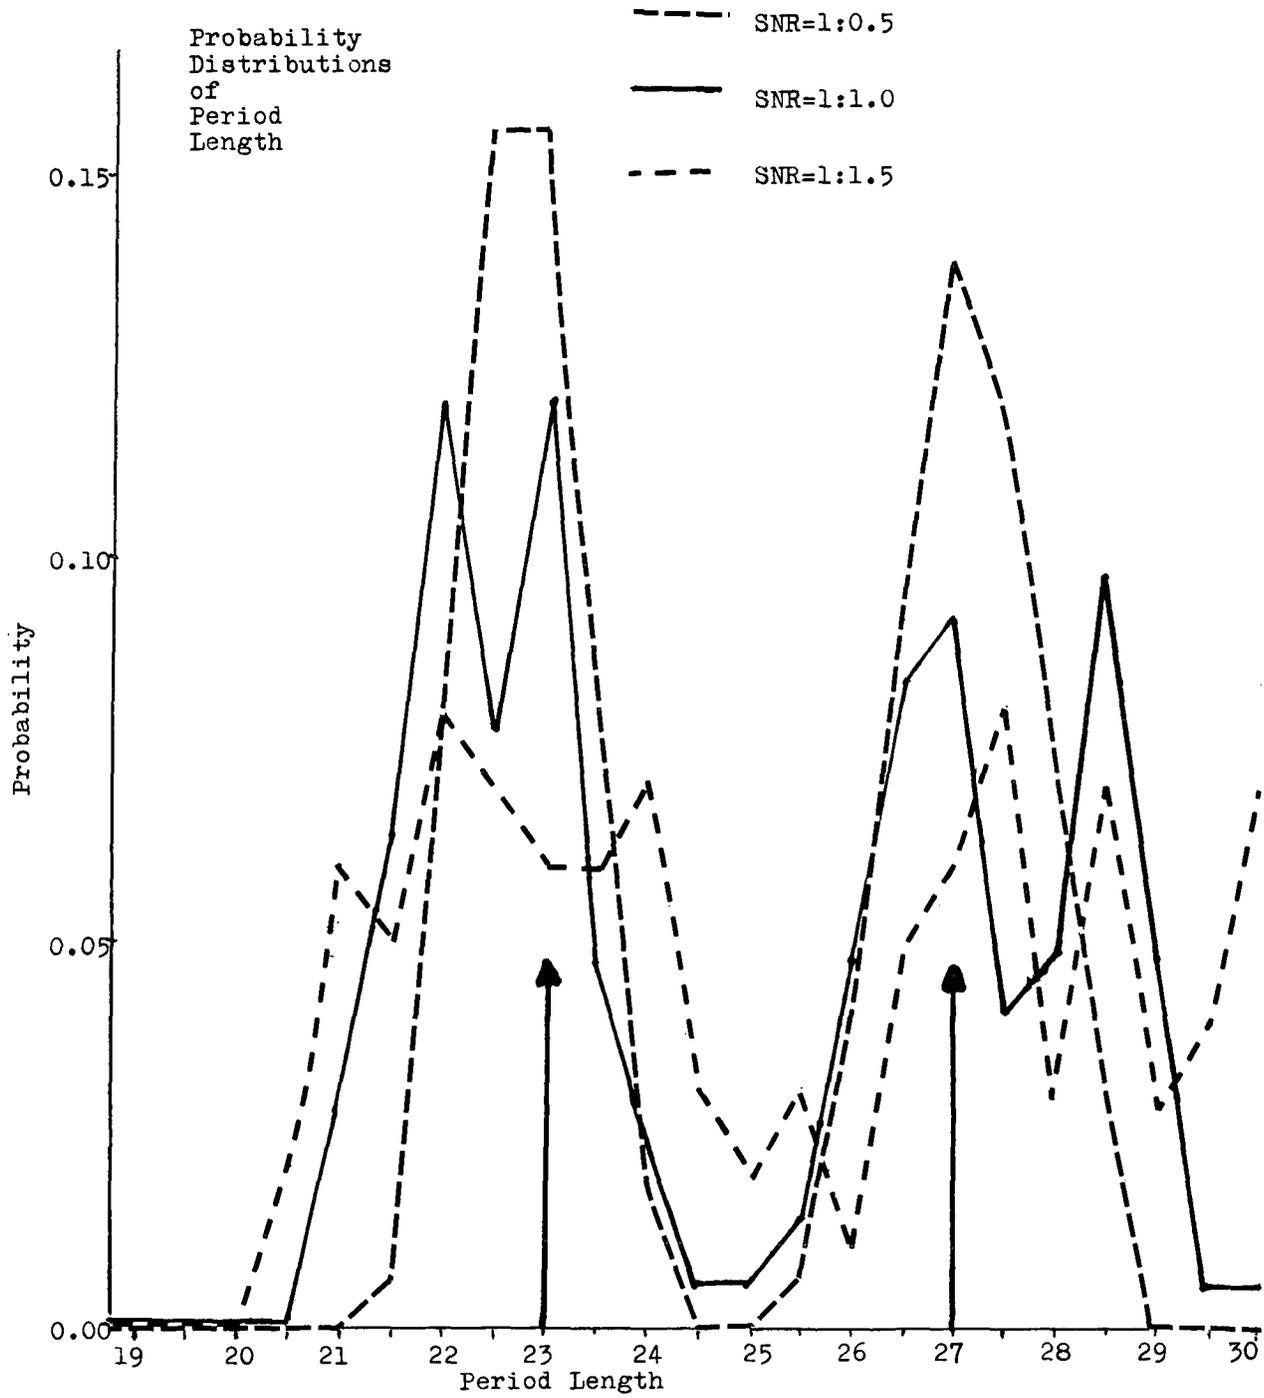


Figure 4.
22

Too Few or Many Peaks

As discussed in previous reports, the program also sometimes finds too many or too few peaks in the spectrum. This kind of failure is not detectable, however, in "real world" situations where the time series population is not known. Therefore, all such failures, while detectable in a Monte Carlo study, must still be included in the results.

Table 2 shows the probability of finding the wrong number of peaks in the spectrum for three levels of SNR.

SNR	p (failure)
1/0.5	0.00
1/1.0	0.05
1/1.5	0.61

Table 2. Probability of program failure by finding wrong number of peaks.

Comparison of Table 2 with results given in earlier reports shows that substantial improvement has been made in the probability of failures to detect the wrong number of peaks. This finding supports the notion that the deviant cases which were eliminated

were deviant on other (but undetectable) grounds than that of finding excessive total amplitude. While this is true for $SNR = 1/0.5$ and for $SNR = 1/1$, one cannot fail to be alarmed at the high (0.61) probability of program failure for $SNR = 1/1.5$. Inspection of the failures reveals that usually the program found one significant period only. Modifications discussed in the next section however, reduce these probabilities of failure even further.

Multiple Correlation as a Criterion of Fit

The original Rummel program not only uses a t-test criterion to decide whether a peak is significant, it also attempts to optimize the multivariate model by optimizing the value of t for the particular peak being moved up or down in period. The program was modified to optimize multiple correlation, the amount of variance accounted for in $f(t)$. The rationale for this change was that possibly t-values for a given period band could increase while the overall goodness-of-fit could decrease. Certainly it is possible for one t-value to increase at the expense of another. At least it was not clear that optimizing individual t-values was the best approach.

Period discrimination optimizing R^2 rather than t was explored through several Monte Carlo test runs. Figure 5 shows the results of three such runs of one hundred segments each at SNR levels of 1/0.5, 1/1, and 1/1.5. The two sinusoids had period lengths of 23 and 27 respectively. These graphs should be compared to the results for optimizing t , shown in Figure 4. It can be seen that the curves for optimized R^2 are generally smoother, but not otherwise dramatically different in shape from those for optimized t . R^2 does achieve higher levels of probability for detection of the correct peaks. This is due to a considerably reduced rate of detectable program failure for optimized R^2 which is discussed below.

One of the more prominent kinds of program failures for the optimized t method was that of excessive amplitude estimates. As for previous runs, it was arbitrarily decided that when the total of the amplitude estimates for all significant peaks in any given spectrum was 20% greater than the amplitude of $f(t)$, the result would be declared a program failure and the data not included in further performance characteristics. Table 3 shows the probabilities for this kind of failure (which is detectable in real data cases) for three levels of SNR. Data for both methods are given.

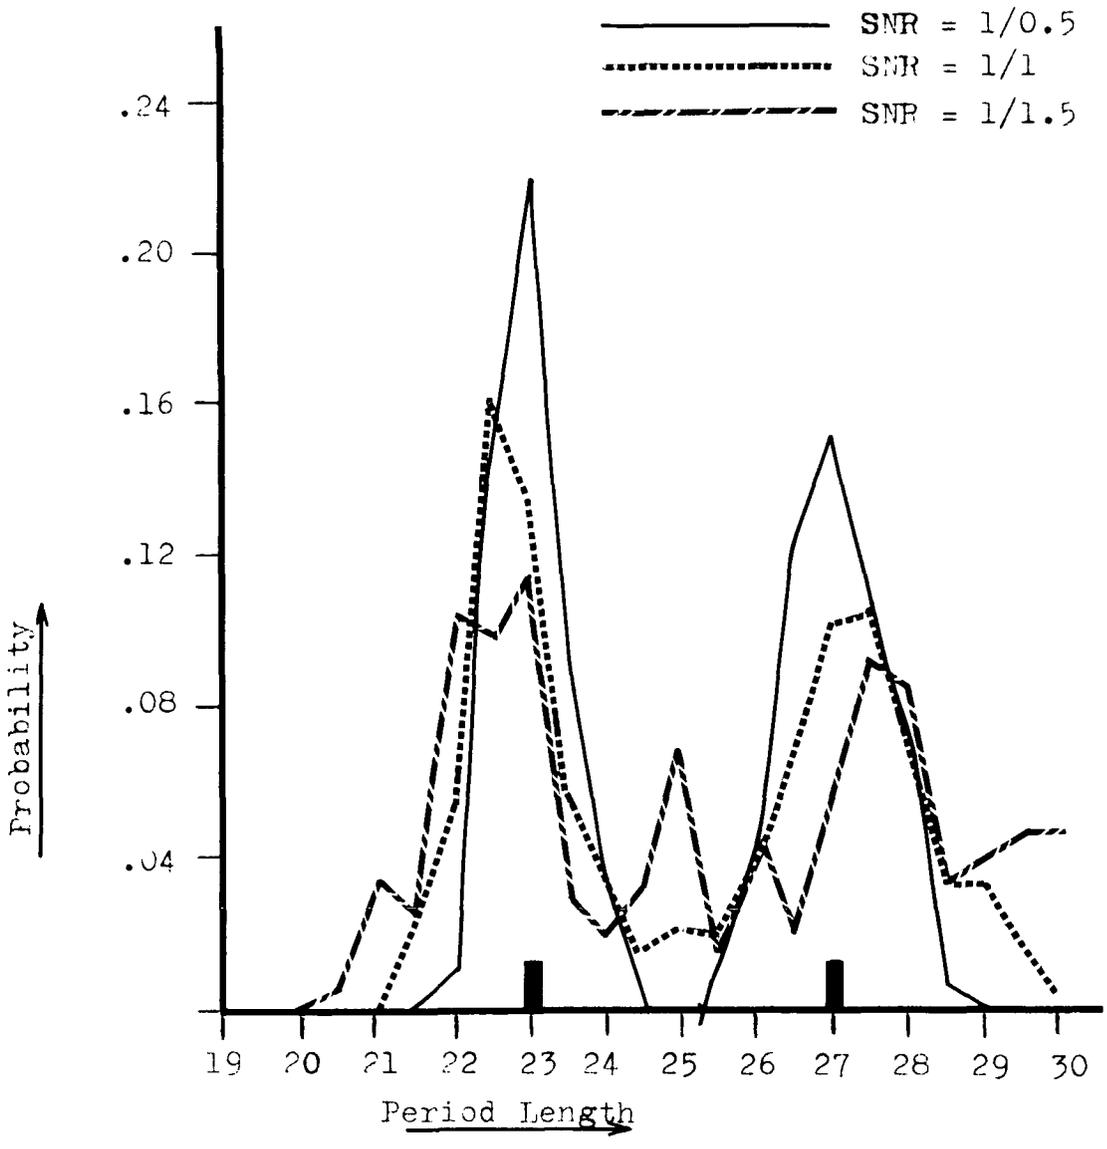


Figure 5. Period discrimination at various levels of SNR.

SNR	optimized t	optimized R^2
1/0.5	.17	.05
1/1.0	.27	.08
1/1.5	.05	.17

Table 3. Probability of program failure by finding excessive total significant peak amplitude.

It is noteworthy that optimized R^2 produces lower probability of failure for all SNR levels than optimized t, except for SNR = 1/1.5.

Another kind of program failure, which is not detectable in a real data case, is that of finding too many or too few significant peaks. Table 4 shows the probabilities of this kind of failure for the two methods at three levels of SNR. Here, it is seen that optimized R^2 is higher at SNR = 1/0.5. Quite possibly, when SNR is known to be low, t-values ought to be optimized rather than R^2 . Even so, optimized R^2 exhibits a relatively low probability of failure, even for SNR = 1/0.5.

SNR	optimized t	optimized R^2
1/0.5	0.00	0.03
1/1.0	0.05	0.02
1/1.5	0.61	0.11

Table 4. Probability of program failure by finding the wrong number of peaks.

Estimation of SNR from the Data. Multiple squared correlation, R^2 , is obviously related to SNR and can be computed from the data. Figure 6 shows a plot of mean R^2 values obtained for Monte Carlo runs using four levels of SNR. It is quite apparent that R^2 can be used to estimate SNR. When only one sinusoid was used in a Monte Carlo run with SNR = 1/1, R^2 was as shown by the asterisk in the plot. For the two sinusoid case, the RMS value of the signal component was higher than for the single sinusoid case and so the effective SNR would be lower. When this is considered, all points would lie close to the plotted line.

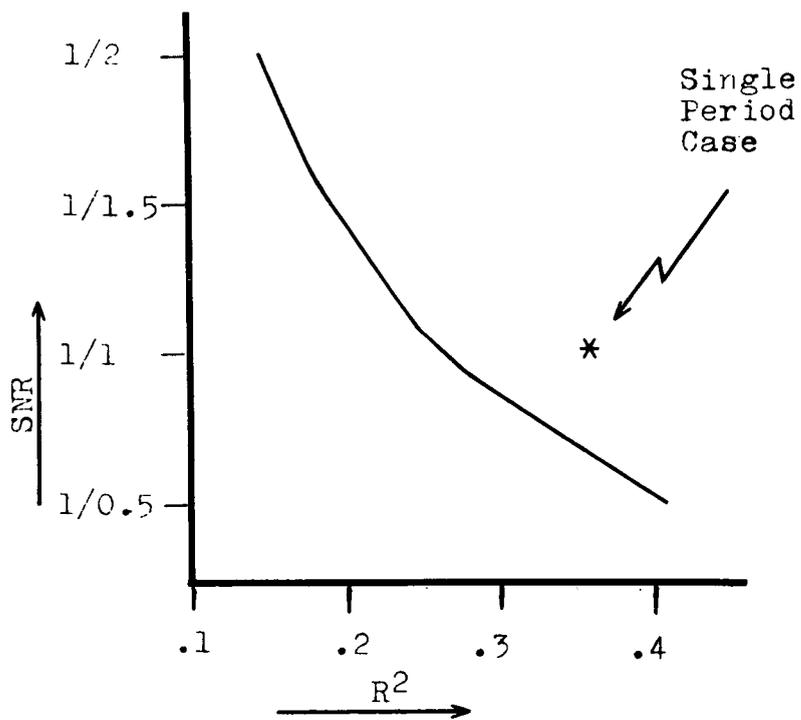


Figure 6. Multiple squared correlation as a function of SNR and number of periods.

Other Attempted Improvements

Several other attempts were made to improve the search routine for the best fit. None of these proved to be clearly superior across the board. Some of the attempts will be described below.

Rather than moving each peak up and/or down only once before computing the whole spectrum anew, using a multi-prediction scheme, it was decided to run through the entire sequence of moving each peak several times. The rationale was that it was possible to affect the optimal value of one peak by moving another since these predictors interact. Looping through the peak-shifting algorithm a number of times did indeed alter specific results, sometimes for the better, sometimes for the worse. The overall long run result showed no improvement however.

Rather than shifting a peak only one increment of ΔP it was decided to optimize the model in two or more stages before computing the whole spectrum. The first stage was as usual, shifting peaks by one increment of ΔP . The second stage began at the shortest period and shifted each peak by two increments

up and/or down in an attempt to optimize R^2 . The third stage shifted by three increments of ΔP , etc. This technique showed some small improvements for a two stage process, but the advantages were very small and tenuous.

Several combinations of the above procedures were attempted with generally similar results. It was finally decided to use a simple one-pass search algorithm as used in the original Rummel program except to offer the alternative of optimizing R^2 instead of t-values.

One method of analysis which was attempted yielded ambiguous results. Rather than entering only significant peaks into the spectrum analysis model, it was attempted to enter all frequencies in the spectrum simultaneously. The rationale for this procedure was that in this way all interpredictor correlations would be accounted for simultaneously and the optimization routine might be avoided. In terms of equation (8), there would be $J = 2k$ estimators or k frequency bands in the whole spectrum, all k of them entered simultaneously.

There were immediate and serious problems with this technique

even when $f(t)$ consisted of two sinusoids without noise. When Δf was set at $\frac{1}{10T}$ and there were as few as 10 frequency bands in the spectrum, the determinant of R_{jj} was zero. With fewer than 10 bands, the determinant was very small, approaching zero as the number of bands approached a limit of 10. Similar results were obtained with other values of Δf .

It was initially expected that the determinants would be small, due to the large correlations among predictors. It was not anticipated and is not presently clear why the determinant should ever equal zero. Apparently as the determinant approaches zero, the machine accuracy eventually truncates the value to zero. This possibility was substantiated by using a double precision program, where it was possible to enter a few more variables than with single precision.

If this problem had been the only difficulty, it might still have been possible to test this model, at least for narrow spectrum ranges. However, even where the spectrum range was restricted so as to obtain non-zero determinants of R_{jj} , some insurmountable problems remained. As long as R_{jj} was not singular and as long as $f(t)$ was not noisy, the program provided accurate spectra. When even a slight amount of noise was added

to $f(t)$, however, the spectrum estimates were exceedingly large and bore little relation to expected results. It was not entirely clear why this result should occur, but the hypothesis was made that the model was grossly "overpredicting" $f(t)$. When the number of predictors becomes large and when these predictors are highly correlated regression weights may become quite large. The scalar expression for the two predictor standard regression coefficient illustrates this:

$$\beta_1 = \frac{r_{0,1} - r_{0,2}r_{1,2}}{1 - r_{1,2}^2} \quad [13]$$

where variable zero is the dependent variable and variables 1 and 2 are predictors. Now

$$\lim_{r_{1,2} \rightarrow 1} \beta_1 = \infty \quad [14]$$

This conclusion may be phrased more analogically and intuitively. The program is attempting to predict the random noise in $f(t)$ by adding all the sinusoids at its disposal. Sometimes the resulting β values are very large.

The above explanation is still speculative because no proofs were performed, but it is considered plausible. The warning

should also be clear. When Δf becomes small, too many predictors can result in nonsensical results. It is not clear how many estimators may be used for any given Δf . It is apparently reasonably safe to use two estimators when Δf is as small as $\frac{1}{5T}$ because these were values used in earlier reports. At this point it was decided to abandon this particular effort in favor of other approaches described below. The effort was obstructed because of mathematical limitations on the model. The above results were obtained when estimators were spaced at equal ΔP or equal Δf intervals.'

Period vs Frequency Domain

Theoretical treatment of spectrum analysis is most often expressed in the frequency domain. An attempt was made to document the performance of the multiple predictor model of spectrum analysis in the frequency rather than the period domain. Again, non-noisy sinusoids were used to test the program. The results were dismayingly close to correct, but the program could not be made to find the correct frequency bands. The results would invariably be a spectrum of two peaks with non-zero estimates around the peaks. Furthermore, the peaks were invariably

exactly two Δf increments above or below the correct value. This was no fault of the regression computations because when the correct values were forced and the program was not allowed to search, the amplitude estimates were exactly correct; all other estimates were zero; and the multiple correlation coefficient was one.

After testing an exhaustingly long list of possible explanations for this error, it was decided to space the predictor waves unequally in the frequency domain, that is $\Delta f_1 < \Delta f_2 < \Delta f_3 \dots\dots\dots$. Under these circumstances the program worked perfectly for the no-noise case. It became obvious that the period domain would provide superior results because of the unequal spacings which it provides in the frequency domain. Inspection of the program's search attempts, when equal Δf s were used, revealed that as the selection procedure began with a frequency reasonably far removed from the correct value and approached the correct band, the criterion for optimum value (the t-value) would appropriately increase. However, superimposed upon the increasing t-value trend was an oscillation such that some steps made more improvement than others. If f_0 was the correct band and the search was approaching from f_4 downward, $t_4 < t_3 < t_2$ as would be expected.

However, $t_2 > t_1$ and although $t_0 < t_1$, the program stopped the search because a maximum value of t had passed. At this point it was decided to continue working in the period domain when $\Delta f < \frac{1}{T}$.

When using R^2 as an optimization criterion, rather than t , as discussed above, the results were quite different. The program found the correct frequencies without problem with approximately the same performance as documented previously. This finding reinforces the use of R^2 rather than t .

Other Performance Characteristics

Using the final version of the program, which optimized R^2 and rejected obviously erroneous results, further operating characteristics were documented.

Comparison to FFT

Performance of the multiple prediction program was compared to that of the more classical FFT and the results were as expected. Figure 7 shows the performance of the two programs when there was one frequency in $f(t)$ and no noise. Quite clearly, while the FFT meets theoretical expectations, the multivariate program is superior.

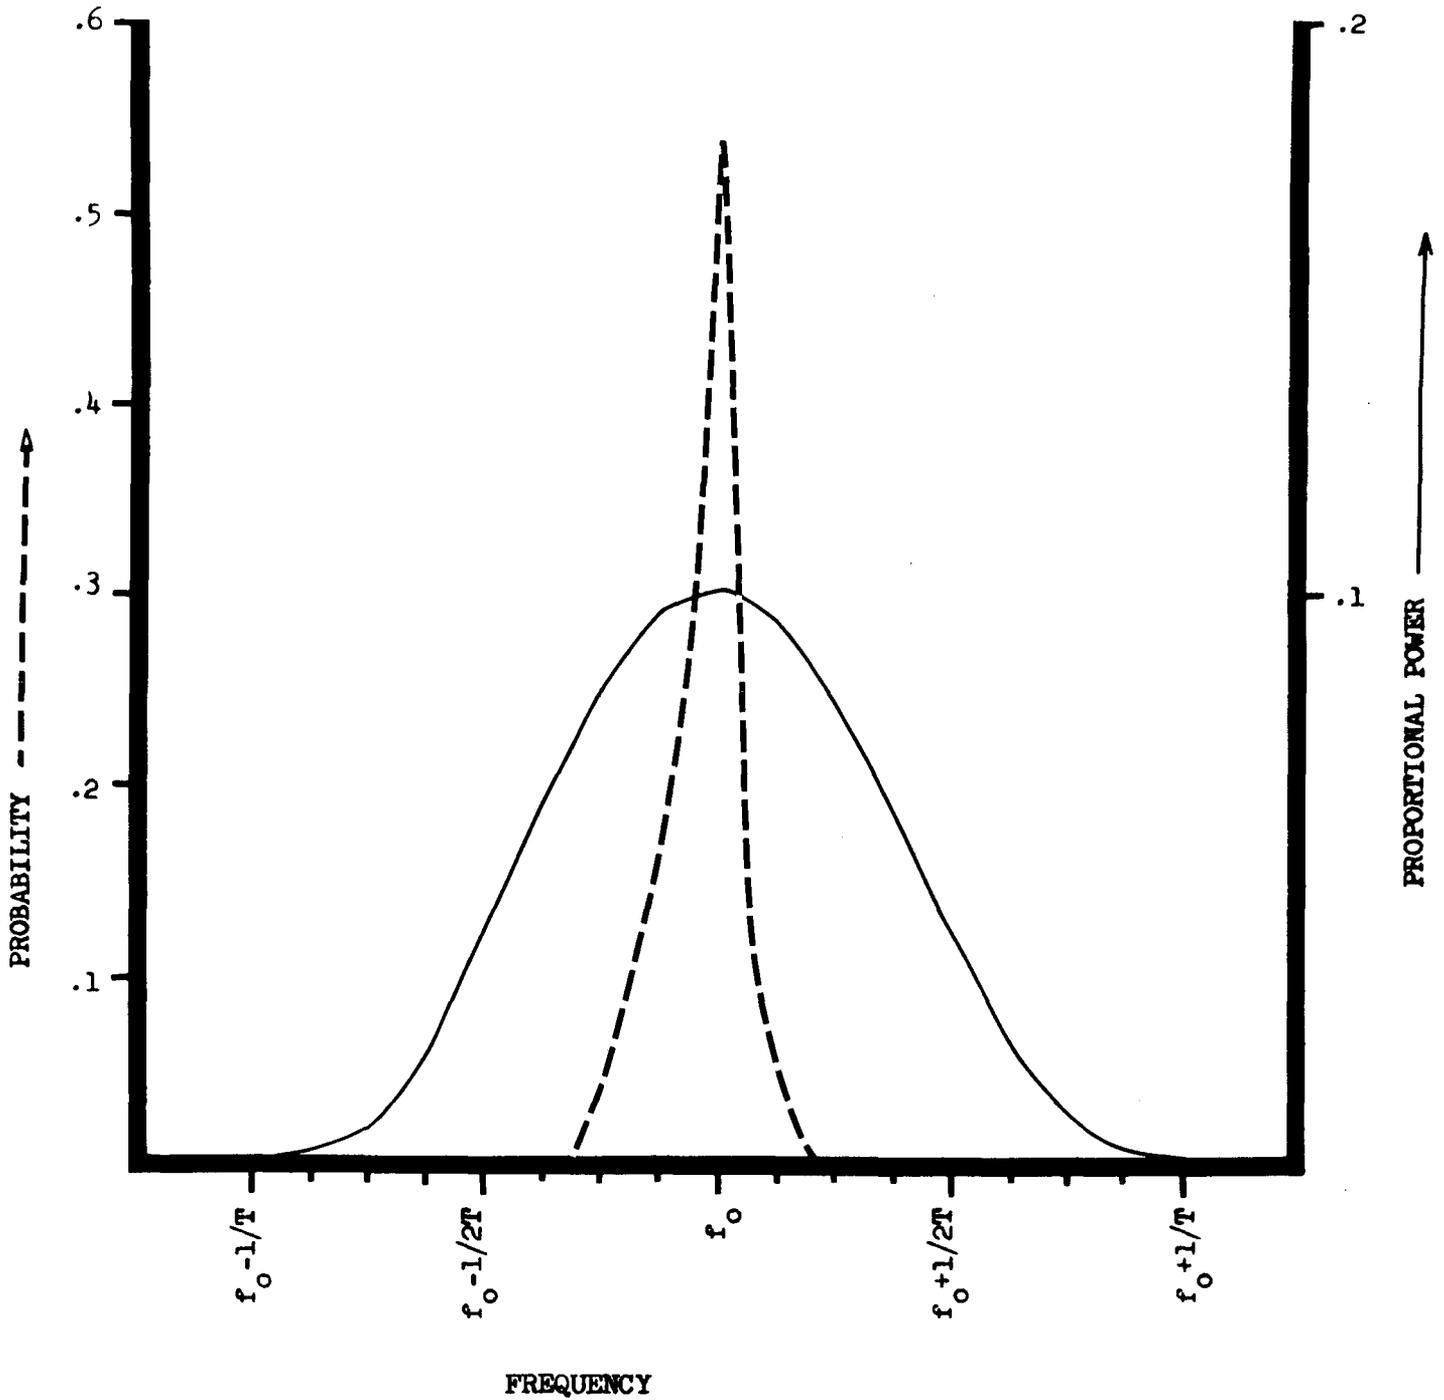


Fig. 7 Average windows for the FFT and the NASA multivariate program. It should be pointed out that the NASA window demonstrates the probability of line-spectrum results falling within a certain frequency range, whereas the FFT window is simply the average of many windows. This is because the FFT never yields a line-spectrum result when estimates are spaced at $1/8T$.

Amplitude Estimates

Figure 8 shows the probability distribution for three levels of SNR. It is noticed that all of the distribution have maxima near the correct value of the signal (10). Thus the program appears to provide unbiased estimates. Also the probability of correct estimates increases as an inverse function of noise.

A New Program

Using the improvements which were devised, a new multivariate program was written. The program, called SPECT, was written in Fortran IV, using only the most generally available features of the language. Thus while some aspects of the program are somewhat awkward, such as if statements rather than logical statements, the program is less machine dependent. The program makes use of extensive comment cards for documentation.

SPECT is written in modular form using IBM scientific subroutines for many of the modules. These subroutines are well "de-bugged" and reduce the probability of programming errors. Other modules were written especially for SPECT.

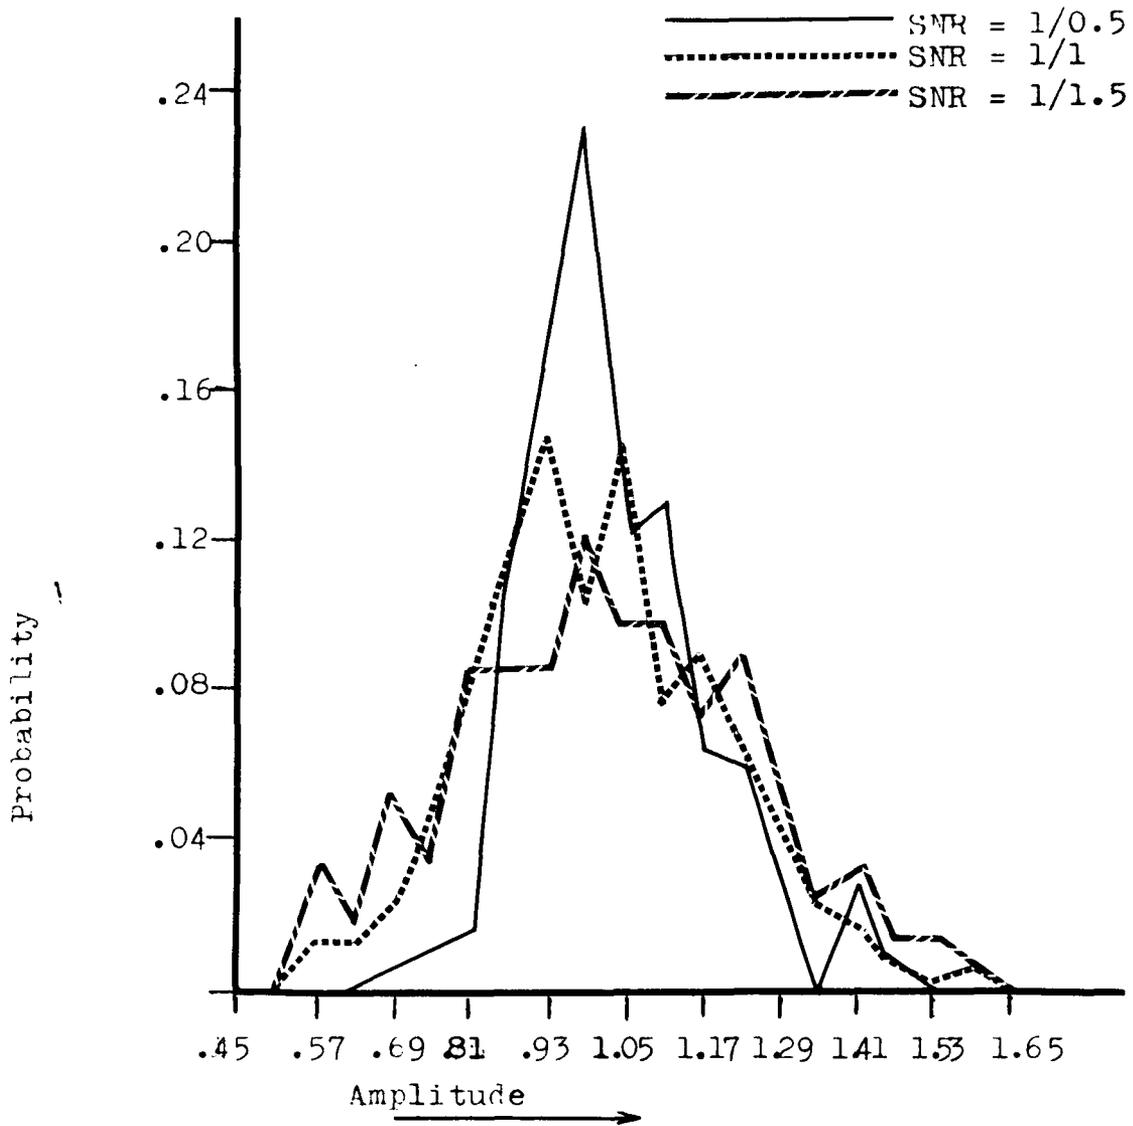


Figure 8. Distribution of peak amplitudes for various levels of SNR.

SPECT can be used in any one of three modes and either in the frequency or time domain. By simply selecting control card parameters, SPECT will either (1) compute an optimized spectrum using either R^2 or t as a criterion, (2) compute a multiple estimator special model where specific period frequency bands are entered by the user, or (3) compute a simple univariate amplitude spectrum. The above three options can be performed in either the Frequency or Period domain.

There are two ways that the user can enter the spectrum limits and values. He can read in the longest period value, ΔP and the number of bands in the spectrum and let the program generate the appropriate values, or he can choose to read in the value of each period individually. The latter provides the ability to perform specially spaced spectra. Of course, if the program is performing a frequency domain analysis, then the frequency band values are read in.

Documentation on the output (printer) is automatically appropriate to the domain of the analysis. Output printing of numerical results is controlled by a three-level control digit so as to either (1) print out only final results; (2) print out

final and intermediate search results; and (3) print out final results, intermediate results plus all matrices, correlations, determinants, etc. All numeric output is inhibited by a fourth level. "Dot plot" graphs are also available under a two-level control, the levels of which correspond to levels 1 and 2 of the numeric output control. All plots may be inhibited by a third level.

The program examines the final results as well as intermediate results, for all of several possible program failures and unreasonable results. Frequently, a suggestion will be printed along with the warning that the result might be unreasonable. Some types of program failure are documented, but no output is permitted.

A complete listing of the program SPECT is provided in Appendix I. This listing gives ample documentation of the program's operation and complete instructions as to its use. About 300K of core requirements can be greatly reduced by reducing the dimensions of the job. Presently the program can handle an $f(t)$ length of 300 observations and can compute a

spectrum with 50 period/frequency bands.

Appendix II gives example runs for each of the several options available. These include execution times.

CONCLUSION

The report contains documentation on the performance of an improved version of the NASA multivariate spectrum analysis program which was written by John A. Rummel. It is shown that the general method of multivariate spectrum analysis is superior to the standard FFT in resolution.

Also included in this report are various other modifications and efforts which were for some reason rejected. The appendices contain listings of a new, highly flexible modular program written by the contractor, as well as several examples of its use.

REFERENCES

- Benignus, Vernon A. Statistical Assessment of a Unique Time Series Analysis Technique. Behavioral Technology Consultants, Inc. Final Report (NASA contract NAS 9-11784), 1972.
- Bendat, J.S. and Piersol, A.G., Measurement and analysis of random data. New York: John Wiley, 1966.
- Draper, N.R. and Smith, H. Applied Regression Analysis. New York: John Wiley, 1966.
- Frazier, T.W., Rummel, J. and Lipscomb, H., Circadian variability in vigilance performance, Aerospace Medicine, 1968, 39, 383-395.
- Guilford, J.P. Fundamental Statistics in Psychology and Education. New York: McGraw-Hill, 1950.
- Hays, William L., Statistics. New York: Holt, Rinehart and Winston, Inc., 1963.
- RAND Corp., A Million Random Digits with 100,000 Normal Deviates. Glencoe, III: Glencoe Free Press, 1955.
- Rummel, J.A. Rhythmic variation in heart rate and respiration rate during space flight - Apollo 15. Proceedings of the International Symposium on Chronobiology. Little Rock, Ark., Oct., 1971.