THE MEASUREMENT OF LINEAR FREQUENCY DRIFT IN OSCILLATORS

James A. Barnes
Austron, Inc., Austin, Texas

ABSTRACT

A linear drift in frequency is an important element in most stochastic models of oscillator performance. Quartz crystal oscillators often have drifts in excess of a part in ten to the tenth power per day. Even commercial cesium beam devices often show drifts of a few parts in ten to the thirteenth per year. There are many ways to estimate the drift rates from data samples (e.g., regress the phase on a quadratic; regress the frequency on a linear; compute the simple mean of the first difference of frequency; use Kalman filters with a drift term as one element in the state vector; and others). Although most of these estimators are unbiased, they vary in efficiency (i.e., confidence intervals). Further, the estimation of confidence intervals using the standard analysis of variance (typically associated with the specific estimation technique) can give amazingly optimistic results. The source of these problems is not an error in, say, the regressions techniques, but rather the problems arise from correlations within the residuals. That is, the oscillator model is often not consistent with constraints on the analysis technique or, in other words, some specific analysis techniques are often inappropriate for the task at hand.

The appropriateness of a specific analysis technique is critically dependent on the oscillator model and can often be checked with a simple "whiteness" test on the residuals. Following a brief review of linear regression techniques, the paper provides guidelines for appropriate drift estimation for various oscillator models, including estimation of realistic confidence intervals for the drift.

I. INTRODUCTION

Almost all oscillators display a superposition of random and deterministic variations in frequency and phase. The most typical model used is[1]:

$$X(t) = a + b \cdot t + Dr \cdot t^2/2 + \phi(t) \tag{1}$$

where $X(t)$ is the time (phase) error of the oscillator (or clock) relative to some standard; a, b, and Dr are constants for the particular clock; and $\phi(t)$ is the random part. $X(t)$ is a random variable by virtue of its dependence on $\phi(t)$.

Even though one cannot predict future values of X(t) exactly, there are often significant autocorrelations within the random parts of the model. These correlations allow forecasts which can significantly reduce clock errors. Errors in each element of the model (Eq. 1) contribute their own uncertainties to the prediction. These time uncertainties depend on the duration of the forecast interval, τ, as shown below in Table 1:

TABLE 1. GROWTH OF TIME ERRORS

MODEL ELEMENT NAME	CLOCK PARAMETER	RMS TIME ERROR
Initial Time Error	a	Constant
Initial Freq Error	b	$\sim \tau$
Frequency Drift	Dr	$\sim \tau^2$
Random Variations	$\phi(t)$	$\sim \tau^{3/2}$*

*The growth of time uncertainties due to the random component, $\phi(t)$, can have various time dependencies. The three-halves power-law shown here is a "worst case" model.[2]

One of the most significant points provided by Table 1 is that eventually, the linear drift term in the model over-powers all other uncertainties for sufficiently long forecast intervals! While one can certainly measure (i.e., estimate) the drift coefficient, Dr, and make corrections, there must always remain some uncertainty in the value used. That is, the effect of a drift correction based on a measurement of Dr, is to reduce (hopefully!) the magnitude of the drift error, but not remove it. Thus, even with drift corrections, the drift term eventually dominates all time uncertainties in the model.

As with any random process, one wants not only the point estimate of a parameter, but one also wants the confidence interval. For example, one might be happy to know that a particular value (e.g., clock time error) can be estimated without bias, he may still want to know how large an error range he should expect. Clearly, an error in the drift estimate (see Eq. 1) leads directly to a time error and hence the drift confidence interval leads directly to a confidence interval for the forecast time.

II. LEAST SQUARES REGRESSION OF PHASE ON A QUADRATIC

A conventional least squares regression of oscillator phase data on a quadratic function reveals a great deal about the general problems. A slight modification of Eq. 1 provides a conventional model used in regression analysis[3]:

$$X(t) = a + b \cdot t + c \cdot t^2 + \phi(t) \tag{2}$$

where c = Dr/2. In regression analysis, it is customary to use the symbol "Y" as the dependent variable and "X" as the independent variable. This is in conflict with usage in time and frequency where "X" and "Y" (time error and frequency error, respectively) are dependent on a coarse measure of time, t, the independent variable. This paper will follow the time and frequency custom even though this may cause some confusion in the use of regression analysis text books.

The model given by Eq. 2 is complete if the random component, $\phi(t)$, is a white noise (i.e., random, uncorrelated, normally distributed, zero mean, and finite variance).

II. EXAMPLE

One must emphasize here that ALL results regarding parameter error magnitudes and their distributions are totally dependent on the adequacy of the model. A primary source of errors is often autocorrelation of the residuals (contrary to the explicit model assumptions). While simple visual inspection of the residuals is often sufficient to recognize the autocorrelation problem, "whiteness tests" can be more objective and precise.

This section analyzes a set of 94 hourly values of the time difference between two oscillators. Figure 1 is a plot of the time difference (measured in microseconds) between the two oscillators. The general curve of the data along with the general expectation of frequency drift in crystal oscillators leads one to try the quadratic behavior (Model #1; models #2 and #3 discussed below). While it is not common to find white phase noise on oscillators at levels indicated on the plot, that assumption will be made temporarily. The results of the regression are summarized in a conventional Analysis of Variance, Table 2.

TABLE 2. ANALYSIS OF VARIANCE QUADRATIC FIT TO PHASE
(Units: seconds squared)

SOURCE	SUM OF SQUARES	d.f.	MEAN SQUARE
Regression	2.32E-9	3	
Residuals	4.26E-12	91	4.68E-14
Total	2.329E-9	94	2.48E-11

Coefficient of simple determination   0.99713

Parameters:

$\hat{a}$ = 1.10795E-5 (seconds)   t-ratio = 161.98

$\hat{b}$ = 1.4034E-10 (sec/sec)   t-ratio = 152.02

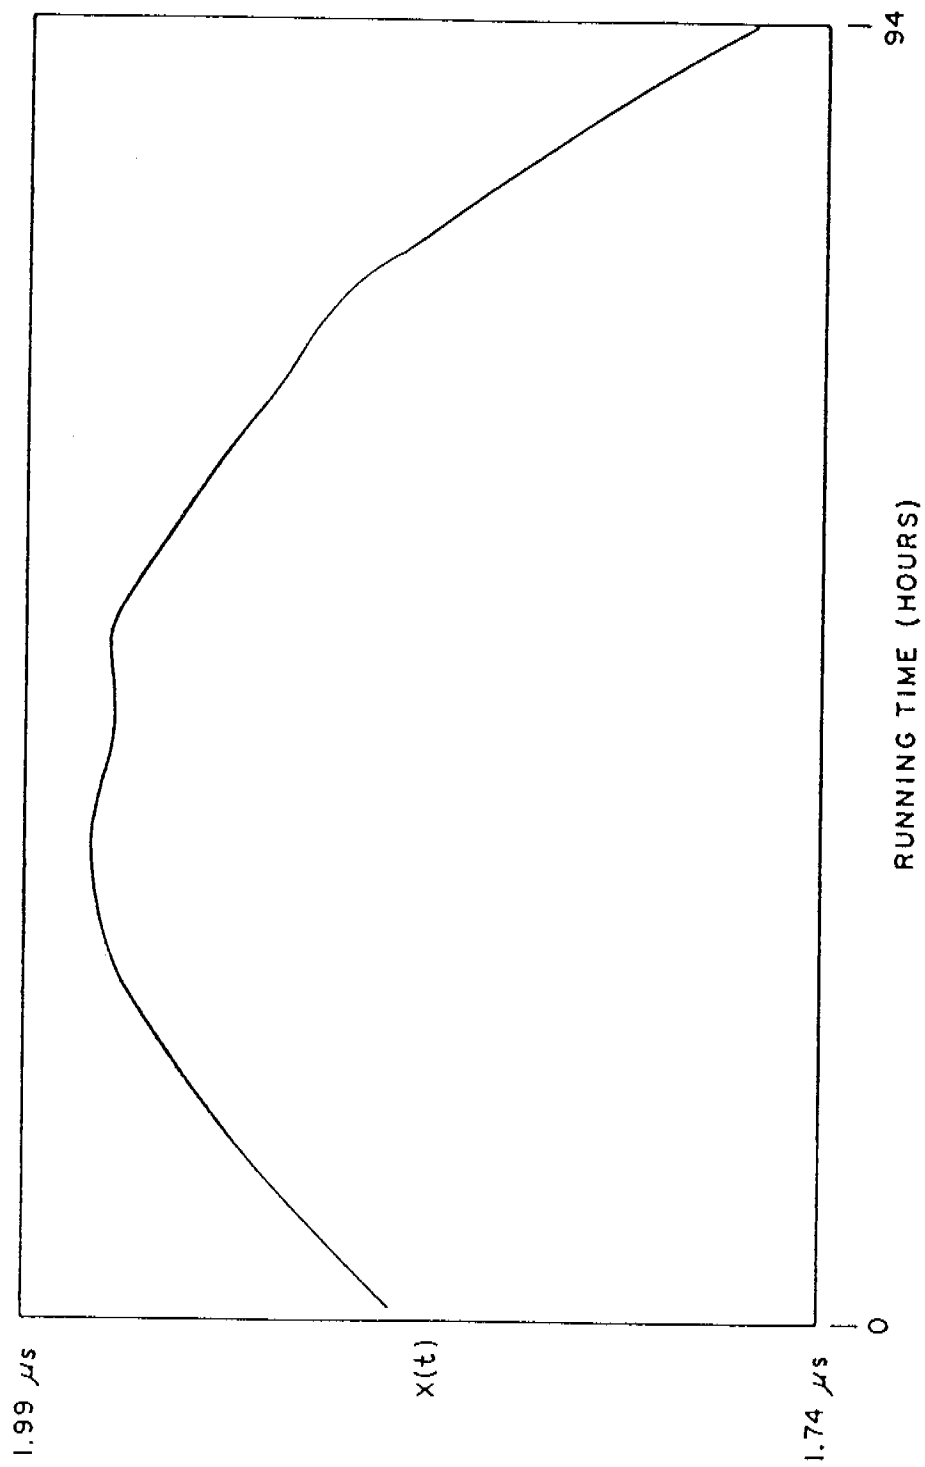$\hat{c}$ = -3.7534E-16 (sec/sec^2)   t-ratio = -143.51

553

Fig. 1 - Phase difference

$$\hat{Dr} = 2\hat{c} = -7.507E-16 \quad (\text{about } -6.5E-11 \text{ per day})$$
$$\text{Std. Error} = 0.05231E-16$$

(Note: $\hat{a}$ is the value estimated for the a-parameter, etc.)

The Analysis of Variance, Table 2, above suggests an impressive fit of the data to a quadratic function, with 99.71% of the variations in the data "explained" by the regression. The estimated drift coefficient, $\hat{Dr}$, is $-7.507E-16$ (sec/sec^2) or about $-6.5E-11$ per day --- 143 times the indicated standard error of the estimate. However, Figure 2, a plot of the residuals, reveals significant auto-correlations even visually and without sensitive tests. (The autocorrelations can be recognized by the essentially smooth variations in the plot. See Fig. 5 as an example of a more nearly white data set.) It is true that the regression reduced the peak-to-peak deviations from about 18 microseconds to less than one microsecond. It is also true that the drift rate is an unbiased estimate of the actual drift rate, but the model assumptions are NOT consistent with the auto-correlation visible in Fig. 2. This means that the confidence intervals for the parameters are not reliable. In fact, the analysis to follow will show just how extremely optimistic these intervals really are.

At this point we can consider at least two other simple analysis schemes which might provide more realistic estimates of the drift rate and its variance. Each of the two analysis schemes has its own implicit model; they are:

(2) Regress the beat frequency on a straight line.
      (Model: Linear frequency drift and white FM.)

(3) Remove a simple average from the second difference of the phase.
      (Model: Linear frequency drift and random walk FM.)

Continuing with scheme 2, above, the (average) frequency, $\overline{Y}(t)$, is the first difference of the phase data divided by the time interval between successive data points. The regression model is:

$$\overline{Y}(t) = b + Dr \cdot t + \varepsilon(t) \tag{3}$$

where $\varepsilon(t) = [\phi(t + \tau_0) - \phi(t)]/\tau_0$. Following standard regression procedures as before, the results are summarized in another Analysis of Variance Table, Table 3.

TABLE 3. ANALYSIS OF VARIANCE LINEAR FIT TO FREQUENCY
(Units: $\text{sec}^2/\text{sec}^2$)

SOURCE	SUM OF SQUARES	d.f.	MEAN SQUARE
Regression	1.879E-18	2	
Residuals	2.084E-20	91	2.29E-22
Total	1.899E-18	93	2.042E-20

(Note: Taking the first differences of the original data set reduces the number of data points from 94 to 93.)
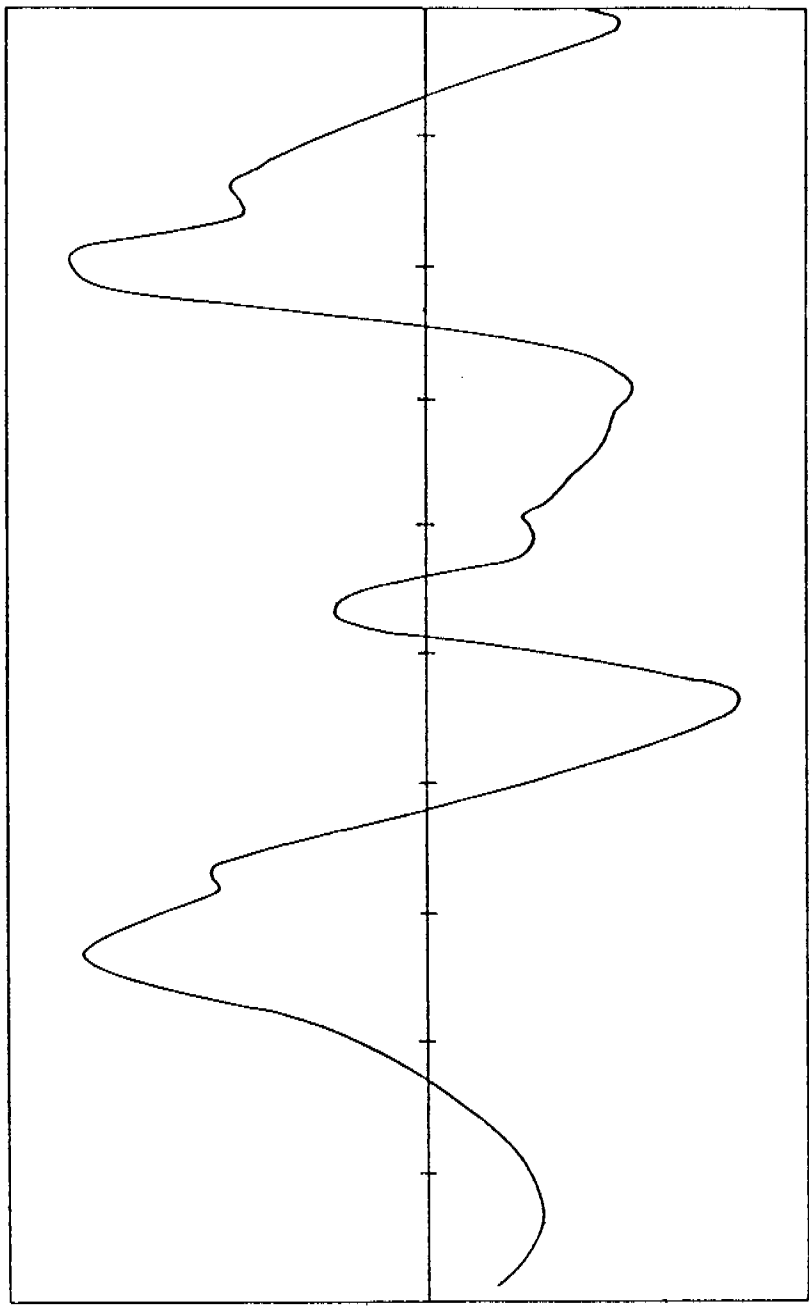
Fig. 2 — Residuals from quadratic fit to phase

RUNNING TIME (HOURS)

5.12 ×10⁻⁷ SEC

$x(t)$

O

−4.61 ×10⁻⁷ SEC

0

94

Coefficient of simple determination   0.9605

Parameters:

$$\hat{b} = 1.049\text{E-}10 \text{ (sec/sec)} \qquad \text{t-ratio} = 3.32$$

$$\hat{Dr} = -7.635\text{E-}16 \text{ (sec/sec}^2) \qquad \text{t-ratio} = -4.70$$
$$\text{Std. Error} = 0.1624\text{E-}16 \text{ (sec/sec}^2)$$

While the drift rate estimates for the two regressions are comparable in value (-7.507E-16 and -7.635E-16), the standard errors of the drift estimates have gone from 0.052E-16 to 0.162E-16 (a factor of 3). The linear regression's coefficient of simple determination is 96.05% compared to 99.17% for the quadratic fit. Figure 3 shows the residuals from the linear fit and they appear more nearly white. A cumulative periodogram[4] is a more objective test of whiteness, however. The periodogram, Fig. 4, does not find the residuals acceptable at all.

IV. DRIFT AND RANDOM WALK FM

In the absence of noise, the second difference of the phase would be a constant, $Dr \cdot \tau_0^2$. If one assumes that the second difference of the noise part is white, then one has the classic problem of estimating a constant (the drift term), in the presence of white noise (the second difference of the phase noise). Of course, the optimum estimate of the drift term is just the simple mean of the second difference divided by τ_0^2. The results are summarized below, Table 4:

TABLE 4. SIMPLE MEAN OF SECOND DIFFERENCE PHASE

Simple mean $\hat{Dr}$ = -6.709E-16   t-ratio = -2.45

Degrees of Freedom   = 91

Standard Deviation $\hat{s}$ = 26.2405E-16

Standard Deviation of the Mean = 2.7358E-16

Figure 5 shows the second difference of the phase after the mean was subtracted. Visually, the data appear reasonably white, and the periodogram, Fig. 6, cannot reject the null hypothesis of whiteness. Now the standard error of the drift term is 2.735E-16, 52 times larger than that computed for the quadratic fit! Indeed, the estimated drift term is only 2.45 times its standard error.

V. SUMMARY OF TESTS

The analyses reported above were all performed on a single data set. In order to verify any conclusions, all three analyses used above were performed a total of four times on four different data sets from the same pair of oscillators. Table 5 summarizes the results:
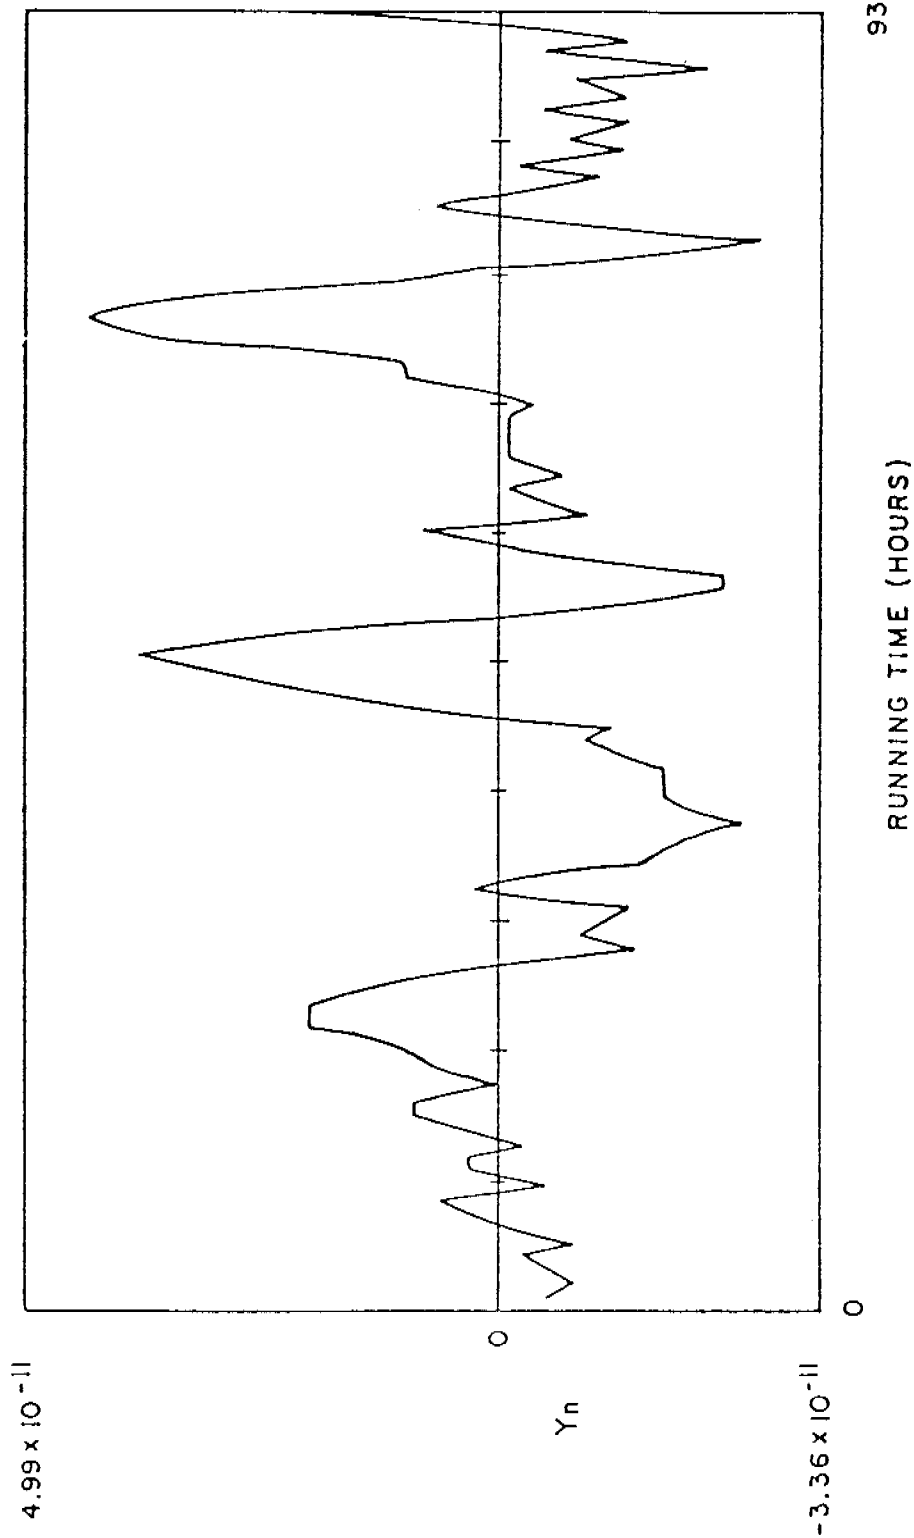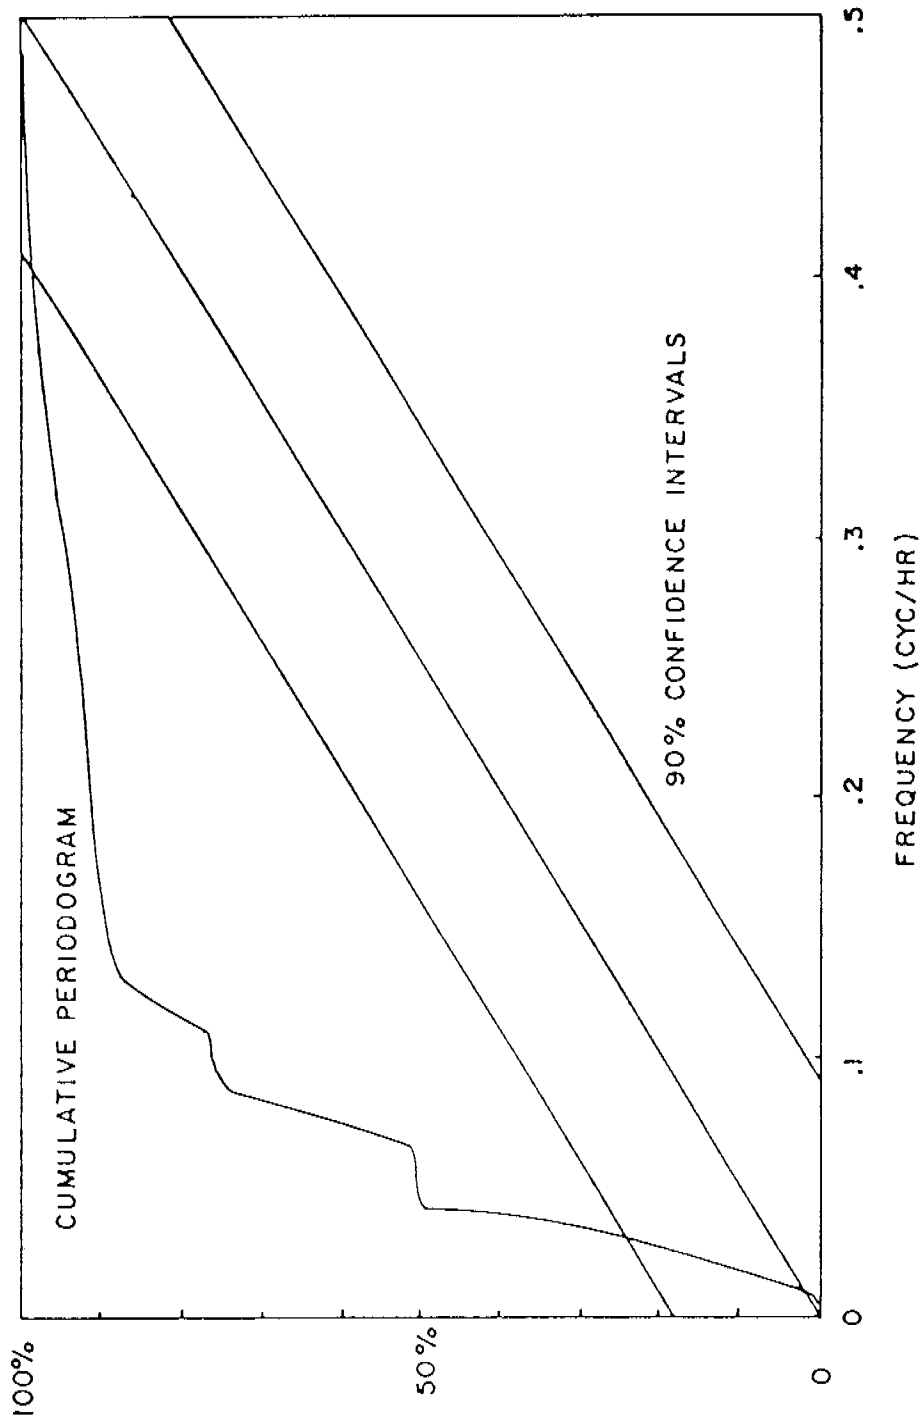
Fig. 3 - Frequency residual
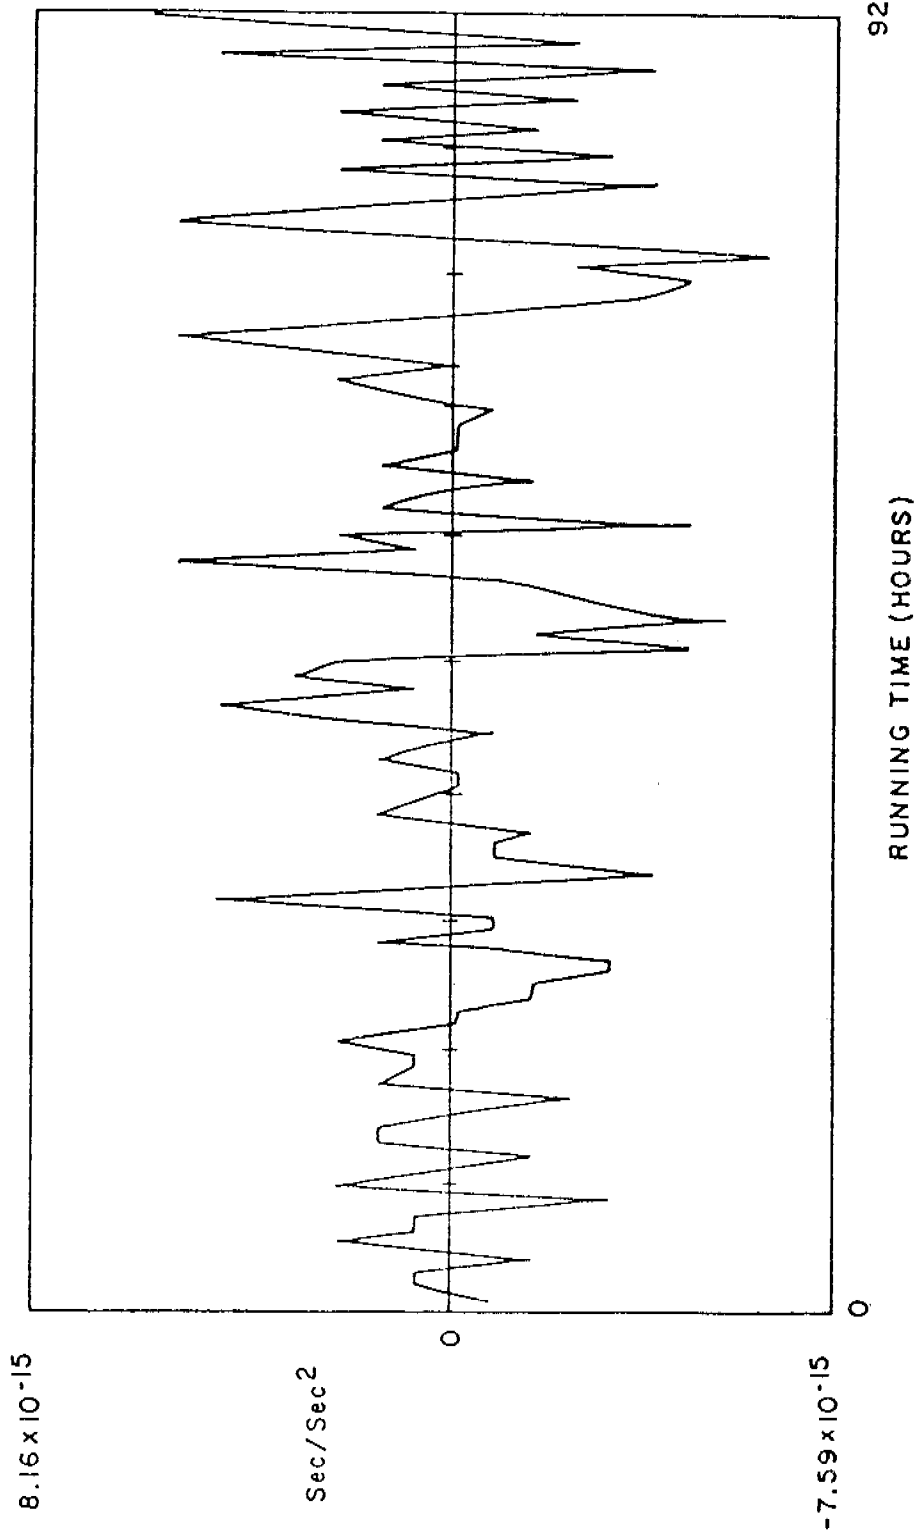(1st diff-lin)

558

Fig. 4 - First diff - linear

Fig. 5 — Second diff of phase

CUMULATIVE PERIODOGRAM

90% CONFIDENCE INTERVALS

FREQUENCY (CYC/HR)

100%

50%
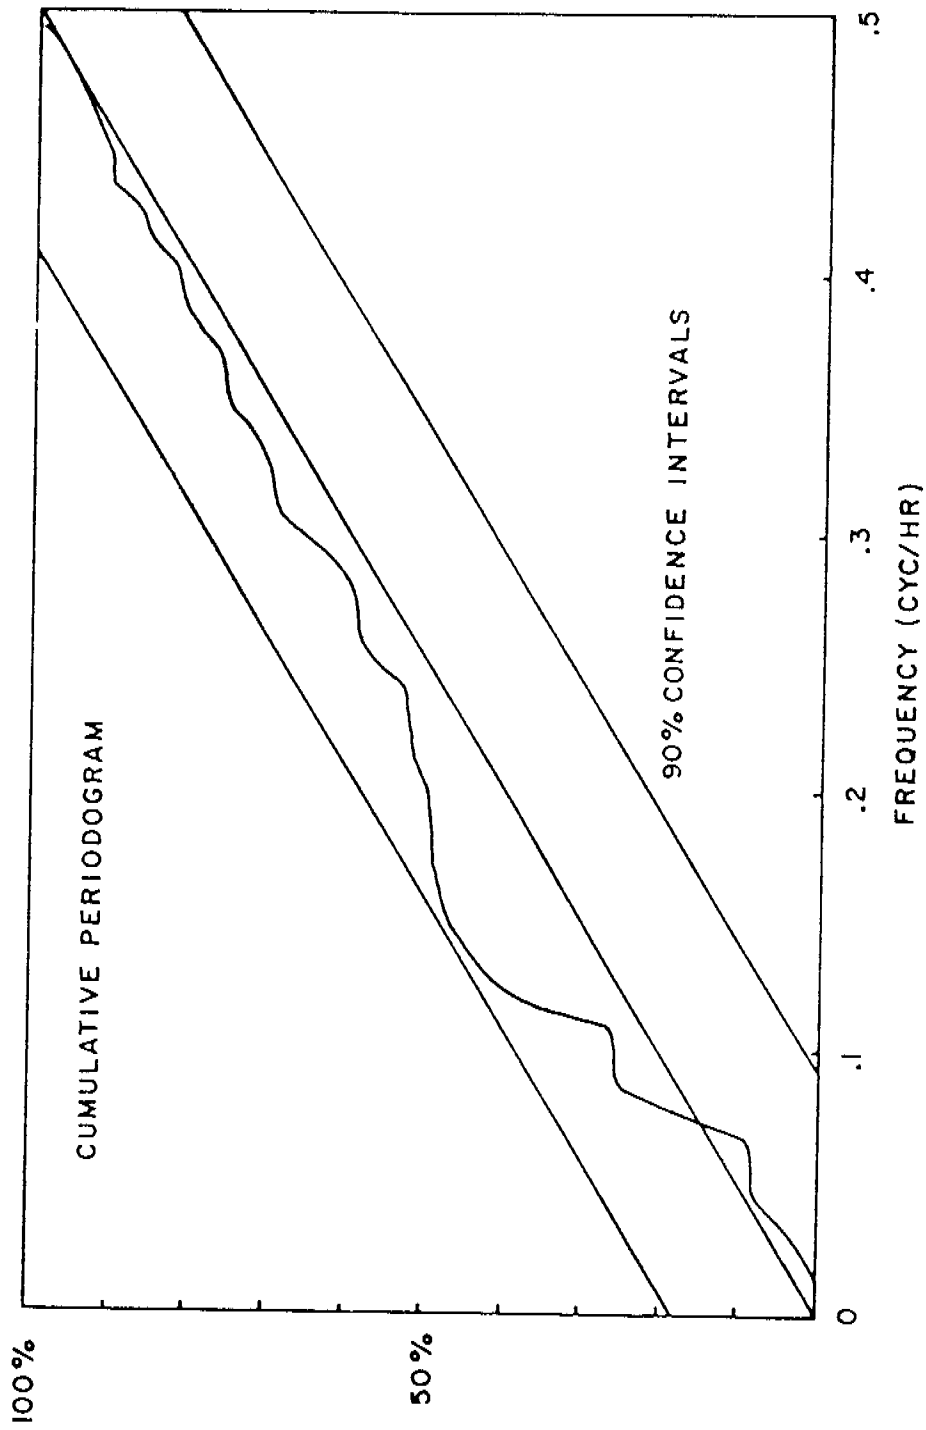
0   .1   .2   .3   .4   .5

Fig. 6 – Second diff – avg.

TABLE 5. SUMMARY OF DRIFT ESTIMATES
(Units of $1.E-16$ sec/sec^2)

ESTIMATION PROCEDURE & MODEL	DRIFT ESTIMATE	COMPUTED STANDARD ERROR	PASS WHITENESS TESTS?
Quad Fit (White PM and Drift)	-7.507	.0523	No
	-8.746	.0493	No
	-6.479	.0645	No
	-6.468	.0880	No
1st Difference Linear Fit (White FM and Drift)	-7.635	.162	No
	-8.558	.206	No
	-6.443	.192	No
	-6.253	.295	No
	-6.710	2.736	Yes
Second Difference Less Mean (Random Walk FM and Drift)	-7.462	9.335	No
	-6.870	3.424	Yes
	-6.412	3.543	Yes

One can calculate the sample means and variances of the drift estimates for each of the three procedures listed in Table 5, and compare these "external" estimates with those values listed in the table under "Computed Standard Error," the "internal" estimates. Of course the sample size is small and we do not expect high precision in the results, but some conclusions can be drawn. The comparisons are shown in Table 6.

It is clear that the quadratic fit to the data displays a very optimistic internal estimate for the standard deviation of the drift rate. Other conclusions are not so clear cut, but still some things can be said. Considering Table 5, the "2nd Diff – Mean" residuals passed the whiteness test three times out of four. The external estimate of the drift standard deviation lies between the internal estimates based on the first and second differences. Since the oscillators under test were crystal oscillators, one expects flicker FM to be present at some level. One also expects the flicker FM behavior to lie between white FM and random walk FM. This may be the explanation of the observed standard deviations, noted in Table 6.

TABLE 6.   STANDARD DEVIATIONS

PROCEDURE (Model)	EXTERNAL ESTIMATE (Std. Dev. of Drift Estimates from Col. #2, Table 5)	INTERNAL ESTIMATE (RMS Computed Std. Dev. Col. #3, Table 5)
Quad Fit (White PM)	1.08	0.065
1st Diff - Lin (White FM)	1.08	0.203
2nd Diff - Mean (Rand Wlk FM)	0.44	5.45

VI. DISCUSSION

In all three of the analysis procedures used above, more parameters than just the frequency drift rate were estimated. Indeed, this is generally the case. The estimated parameters included the drift rate, the variance of the random (white) noise component, and other parameters appropriate to the specific model (e.g., the initial frequency offset for the first two models). If these other parameters could be known precisely by some other means, then methods exist to exploit this knowledge and get even better estimates of the drift rate. The real problems, however, seem to require the estimate of several parameters in addition to the drift rate, and it is not appropriate to just ignore unknown model parameters.

To this point, we have considered only three, rather ideal oscillator models, and seldom does one encounter such simplicity. Typical models for commercial cesium beam frequency standards include white FM, random walks FM, and frequency drift. Unfortunately, none of the three estimation routines discussed above are appropriate to such a model. This problem has been solved in some of the recent work of Jones and Tryon[5],[6]. Their estimation routines are based on Kalman Filters and maximum likelihood estimators and these methods are appropriate for the more complex models. For details, the reader is referred to the works of Jones and Tryon.

Still left untreated are the models which, in addition to drift and other noises, incorporate flicker noises, either in PM or FM or both. In principle, the methods of Jones and Tryon could be applied to Kalman Filters which incorporate empirical flicker models[7]. To the author's knowledge, however, no such analyses have been reported.

VII. CONCLUSIONS

There are two primary conclusions to be drawn:

   (1) The estimation of the linear frequency drift rates of oscillators
        and the inclusion of realistic confidence intervals for these

estimates are critically dependent on the adequacy of the model used and, hence the adequacy of the analysis procedures.

(2) The estimation of the drift rate must be carried along with the estimation of any and all other model parameters which are not known precisely from other considerations (e.g., initial frequency and time offsets, phase noise types, etc.)

More and more, scientists and engineers require clocks which can be relied on to maintain accuracy relative to some master clock. Not only is it important to know that on the average the clock runs well, but it is essential to have some measure of time imprecision as the clock ages. For example, the uncertainties might be expressed as, say,, 90% certain that the clock will be within 5 microseconds of the master two weeks after synchronization. Such measures are what statisticians call "interval estimates" (in contrast to point estimates) and their estimations require interval estimates of the clock's model parameters. Clearly, the parameter estimation routines must be reliable and based on sound measurement practices. Some inappropriate estimation routines can be applied to clocks and oscillators and give dangerously optimistic forecasts of performance.

APPENDIX A

REGRESSION ANALYSIS
(Equally Spaced Data)

We begin with the continuous model equation:

$$X(t) = a + b \cdot t + c \cdot t^2 + \phi(t) \tag{A1}$$

We assume that the data is in the form of discrete readings of the dependent variable $X(t)$ at the regular intervals given by:

$$t = n\tau_0 \tag{A2}$$

Equation (A1) can then be written in the obvious form:

$$X_n = a + \tau_0 \, b \cdot n + \tau_0^2 \, c \cdot n^2 + \phi(n\tau_0) \tag{A3}$$

for $n = 1, 2, 3 \ldots, N$.

Next, we define the matrices:

$$\underline{N} = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 2 & 4 \\ 1 & 3 & 9 \\ " & " & " \\ " & " & " \\ 1 & N & N^2 \end{bmatrix} \qquad \underline{X} = \begin{bmatrix} X_1 \\ X_2 \\ X_3 \\ " \\ " \\ X_N \end{bmatrix}$$

$$\underline{T} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \tau_0 & 0 \\ 0 & 0 & \tau_0^2 \end{bmatrix}$$

$$(\underline{NT})'\underline{X} = \begin{bmatrix} \sum_1^N X_n \\ \tau_0 \sum_1^N X_n \, n \\ \tau_0^2 \sum_1^N X_n \, n^2 \end{bmatrix} = \underline{T} \cdot \begin{bmatrix} \sum_1^N X_n \\ \sum_1^N X_n \, n \\ \sum_1^N X_n \, n^2 \end{bmatrix} = \underline{T} \cdot \begin{bmatrix} S_x \\ S_{nx} \\ S_{nnx} \end{bmatrix}$$

where four quantities must be calculated from the data:

$$S_x = \sum_{n=1}^{N} X_n \qquad\qquad S_{nnx} = \sum_{n=1}^{N} X_n \, n^2$$

$$S_{nx} = \sum_{n=1}^{N} X_n \, n \qquad\qquad S_{xx} = \sum_{n=1}^{N} X_n^2$$

Define

$$\underline{B} = \begin{bmatrix} a \\ b \\ c \end{bmatrix}$$

With these definitions, Eq. A3 can be rewritten in the matrix form:

$$\underline{X} = \underline{N}\,\underline{T}\,\underline{B} + \underline{\varepsilon} \tag{A4}$$

and the coefficients, B, which minimize the squared errors are given by:

$$\underline{\hat{B}} = \begin{vmatrix} \hat{a} \\ \hat{b} \\ \hat{c} \end{vmatrix} = \underline{T} \cdot (\underline{N'N})^{-1} \cdot (\underline{N'X}) \tag{A5}$$

The advantage of evenly spaced data for these regressions is that, with a bit of algebra, the matrix, $(\underline{N'N})^{-1}$, can be written down in closed form:

$$(N'N)^{-1} = \begin{bmatrix} A & B & C \\ B & D & E \\ C & E & F \end{bmatrix} \cdot 1/G \tag{A6}$$

where

$A = 3\,[3\,(N + 1) + 2]$

$B = -18\,(2N + 1)$

$C = 30$

$D = 12\,(2N + 1)\,(8N + 11) / [(N + 1)\,(N + 2)]$

$E = -180 / (N + 2)$

$F = 180 / [(N + 1)\,(N + 2)]$

and

$$G = N (N - 1) (N - 2)$$

Also, the inverse of T is just:

$$\underline{T}^{-1} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1/\tau_0 & 0 \\ 0 & 0 & 1/\tau_0^2 \end{bmatrix}$$

The complete solution for the regression parameters can be summarized as follows:

There are four quantities which must calculate from the data:

$$S_x = \sum_{n=1}^{N} X_n \qquad\qquad S_{nnx} = \sum_{n=1}^{N} X_n n^2$$

$$S_{nx} = \sum_{n=1}^{N} X_n n \qquad\qquad S_{xx} = \sum_{n=1}^{N} X_n^2$$

for $n = 1, 2, 3, \ldots, N$. Based on these four quantities, the regression parameters are calculated from the seven following equations:

$$\hat{a} = (A\, S_x + \tau_0\, B\, S_{nx} + \tau_0^2\, C\, S_{nnx}) / G$$

$$\hat{b} = (B\, S_x + \tau_0\, D\, S_{nx} + \tau_0^2\, E\, S_{nnx}) / (G\, \tau_0)$$

$$\hat{c} = (C\, S_x + \tau_0\, E\, S_{nx} + \tau_0^2\, F\, S_{nnx}) / (G\, \tau_0^2)$$

$$\hat{\sigma}^2 = (S_{xx} - \hat{a}\, S_x - \tau_0\, \hat{b}\, S_{nx} - \tau_0^2\, \hat{c}\, S_{nnx}) / (N - 3)$$

$$\hat{\sigma}_a^2 = \hat{\sigma}^2\, A / G$$

$$\hat{\sigma}_b^2 = \hat{\sigma}^2\, D / (G\, \tau_0^2)$$

$$\hat{\sigma}_c^2 = \hat{\sigma}^2\, F / (G\, \tau_0^4)$$

where the coefficients A, B, C, etc., are given by:

$$A = 3 [3N (N + 1) + 2]$$

$$B = -18 (2N + 1)$$

$$C = 30$$

$$D = 12 (2N + 1) (8N + 11) / [(N + 1) (N + 2)]$$

$$E = -180 / (N + 2)$$

$$F = 180 / [(N + 1) (N + 2)]$$

and

$$G = N (N - 1) (N - 2).$$

In matrix form, the error variance for forecast values is:

$$\text{Var} (\hat{X}_k) = \hat{\sigma}^2 [1 + \underline{N}_k{}' (\underline{N}'\underline{N})^{-1} \underline{N}_k]$$

where $N_k' = [1 \ n_0 \ n_0]$ and $\tau_0 \ n_0$ is the date for the forecast point, $\hat{X}_k$. That is, $n_0 = N + K$ and K is the number of lags past the last data point at lag N.

REGRESSIONS ON LINEAR AND CUBIC FUNCTIONS

The matrixes $(\underline{N'N})^{-1}$ for the linear fit and cubic fit, which correspond to Eq. A6 in Appendix A are as follows:

For the linear fit:

$$(N'N)^{-1} = \begin{bmatrix} A & B \\ B & C \end{bmatrix} 1/D$$

where

$A = 2(2N + 1)$

$B = -6$

$C = 12/(N + 1)$

and

$D = N(N - 1)$

For the cubic fit:

$$(N'N)^{-1} = \begin{bmatrix} A & B & C & D \\ B & E & F & G \\ C & F & H & I \\ D & G & I & J \end{bmatrix} 1/K$$

where

$A = 8(2N + 1)(N + N + 3)$

$B = -20(6N^2 + 6N + 5)$

$C = 120(2N + 1)$

$D = -140$

$E = 200(6N^4 + 27 N^3 + 42 N^2 + 30 N + 11)/L$

$F = -300(N + 1)(3N + 5)(3N + 2)/L$

$G = 280(6N^2 + 15N + 11)/L$

$$H = 360 \ (2N + 1) \ (9N + 13) \ / \ L$$

$$I = -4200 \ (N + 1) \ / \ L$$

$$J = 2800 \ / \ L$$

and

$$K = N \ (N - 1) \ (N - 2) \ (N - 3)$$

$$L = (N + 1) \ (N + 2) \ (N + 3).$$

The restrictions on these equations are that the data is evenly spaced beginning with $n = 1$ to $n = N$, and no missing values. For error estimates (and their distributions) to be valid, the residuals must be random, uncorrelated, (i.e., white).

REFERENCES

1. D.W. Allan, "The Measurement of Frequency and Frequency Stability of Precision Oscillators," Proc. 6th PTTI Planning Meeting, p109 (1975).

2. J.A. Barnes, et al., "Characterization of Frequency Stability," Proc IEEE Trans. on Inst. and Meas., IM-20, p105 (1971).

3. N.R. Draper and H. Smith, "Applied Regression Analysis," John Wiley and Sons, N.Y., (1966).

4. G.E.P. Box and G.M. Jenkins, "Time Series Analysis," Holden-Day, San Francisco, p294, (1970).

5. P.V. Tryon and R.H. Jones, "Estimation of Parameters in Models for Cesium Beam Atomic Clocks," NBS Journal of Research, Vol. 88, No. 1, Jan-Feb. 1983.

6. R.H. Jones and P.V. Tryon, "Estimating Time from Atomic Clocks," NBS Journal of Research, Vol. 88, No. 1, Jan-Feb. 1983.

7. J.A. Barnes and S. Jarvis, Jr., "Efficient Numerical and Analog Modeling of Flicker Noise Processes," NBS Tech. Note 604, (1971).
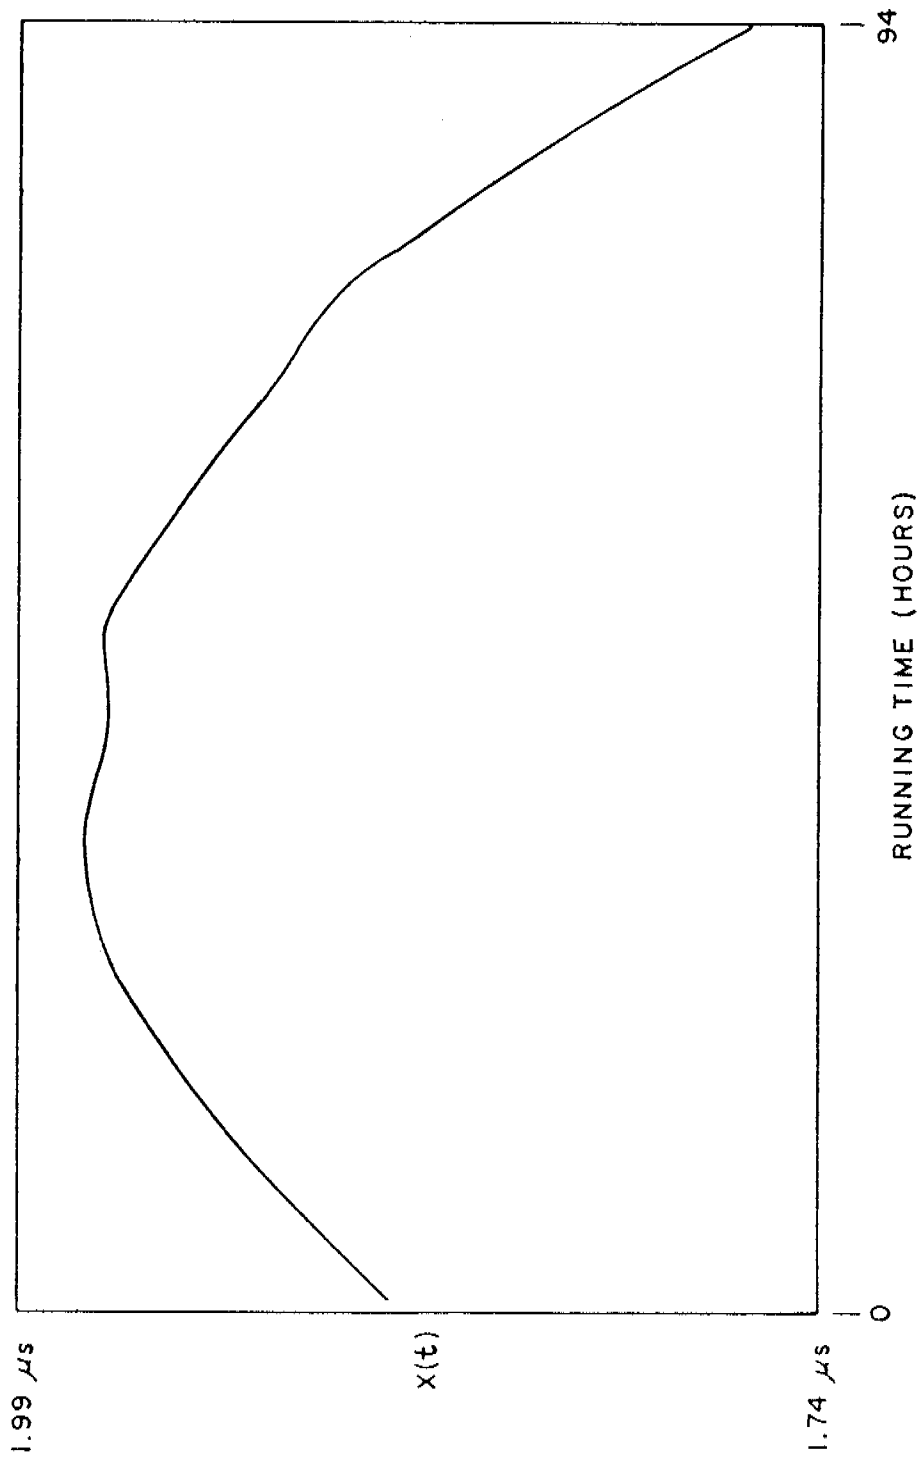
RUNNING TIME (HOURS)

Fig. 1 — Phase difference

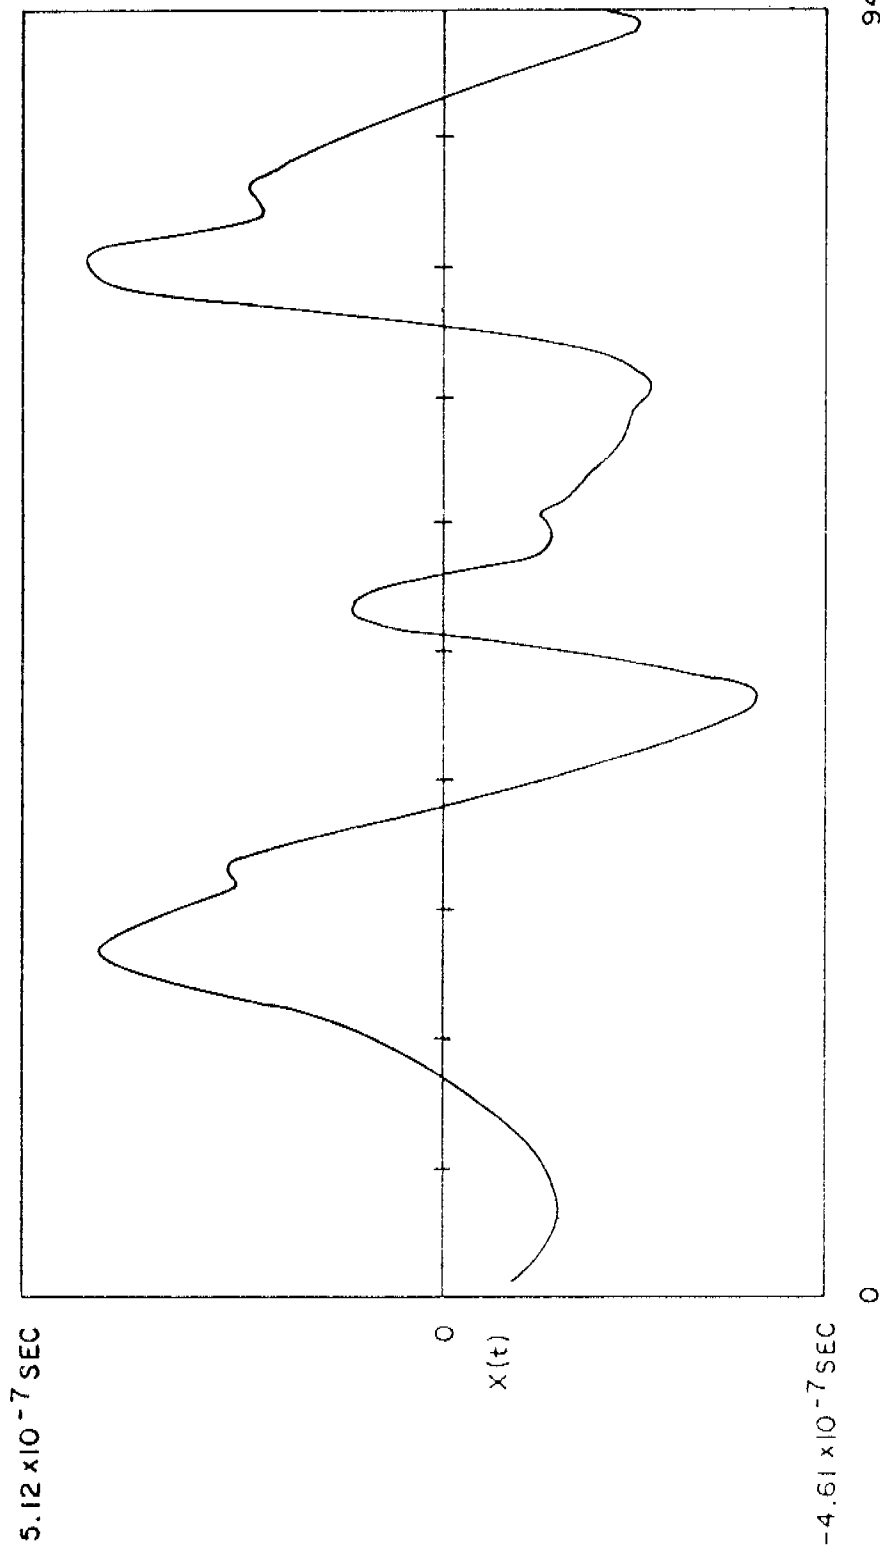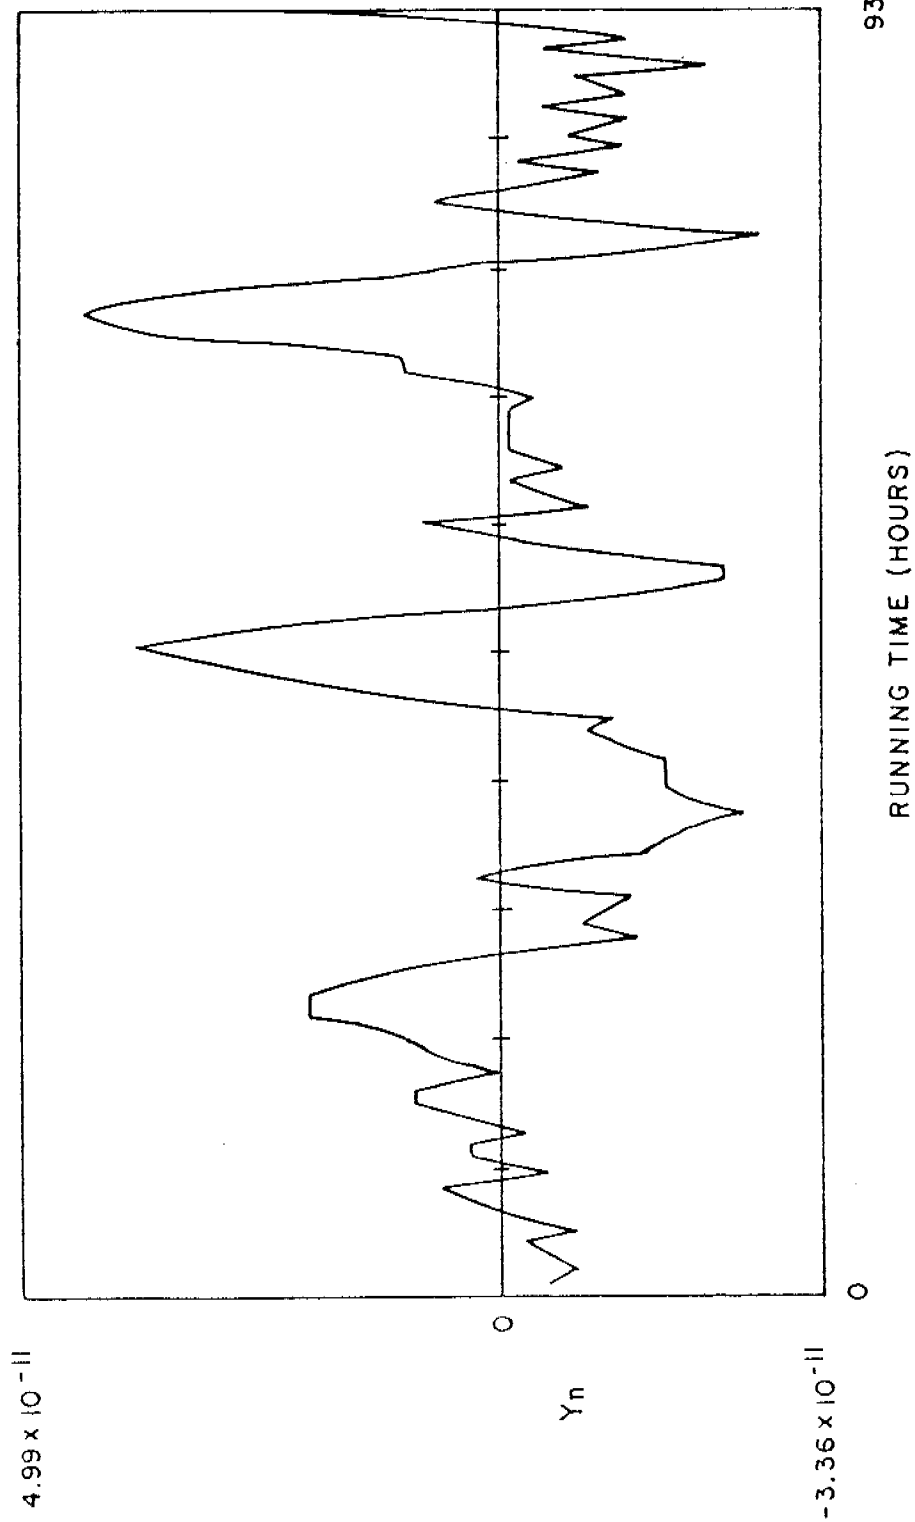1.99 μs

x(t)

1.74 μs

0

94

RUNNING TIME (HOURS)

Fig. 2 – Residuals from quadratic fit to phase
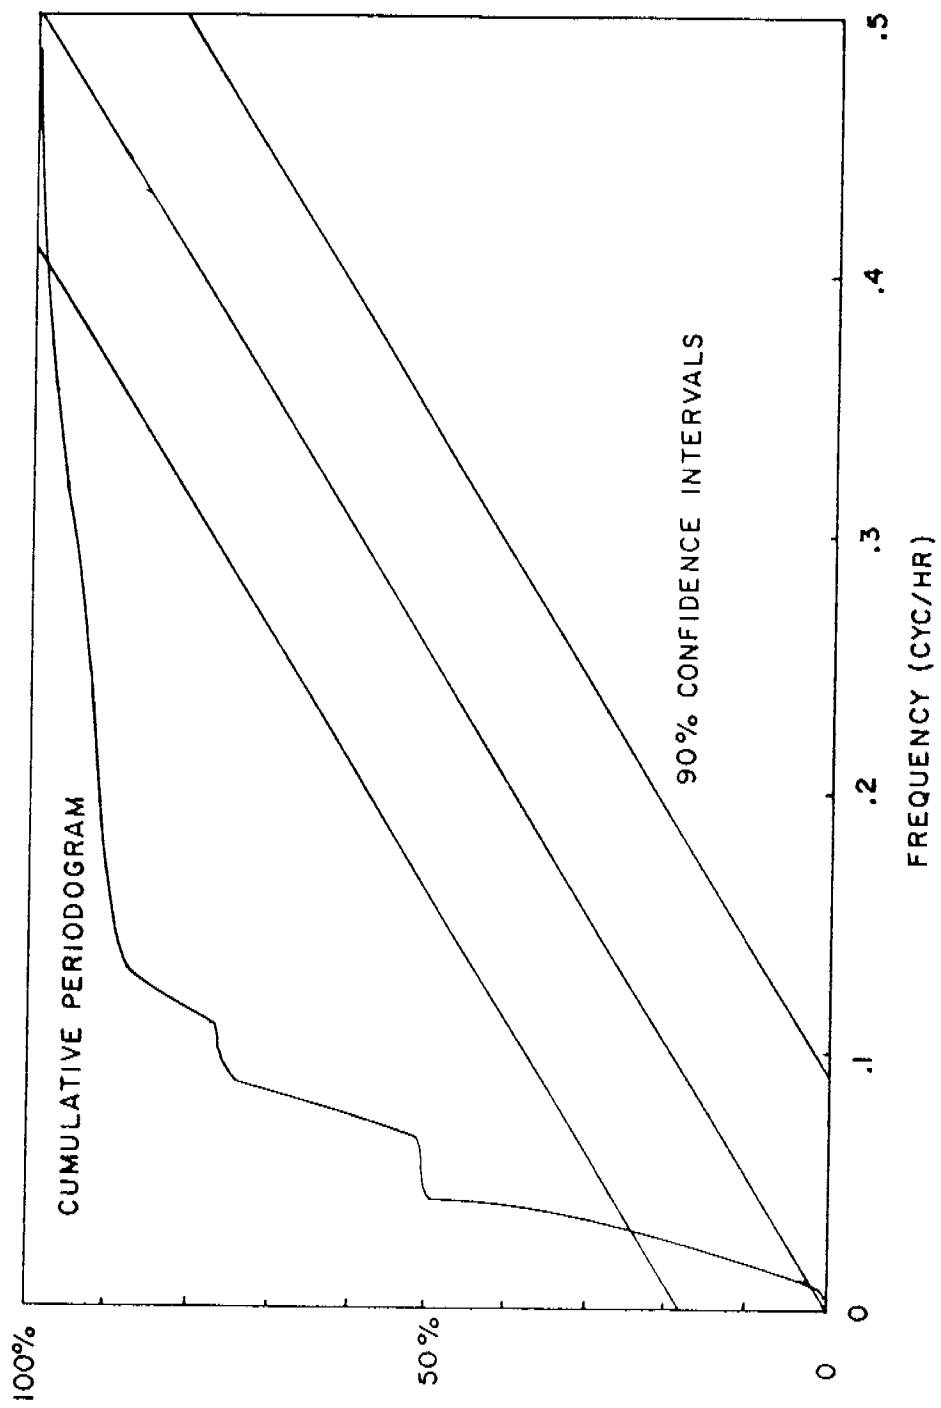
RUNNING TIME (HOURS)

Fig. 3 – Frequency residual
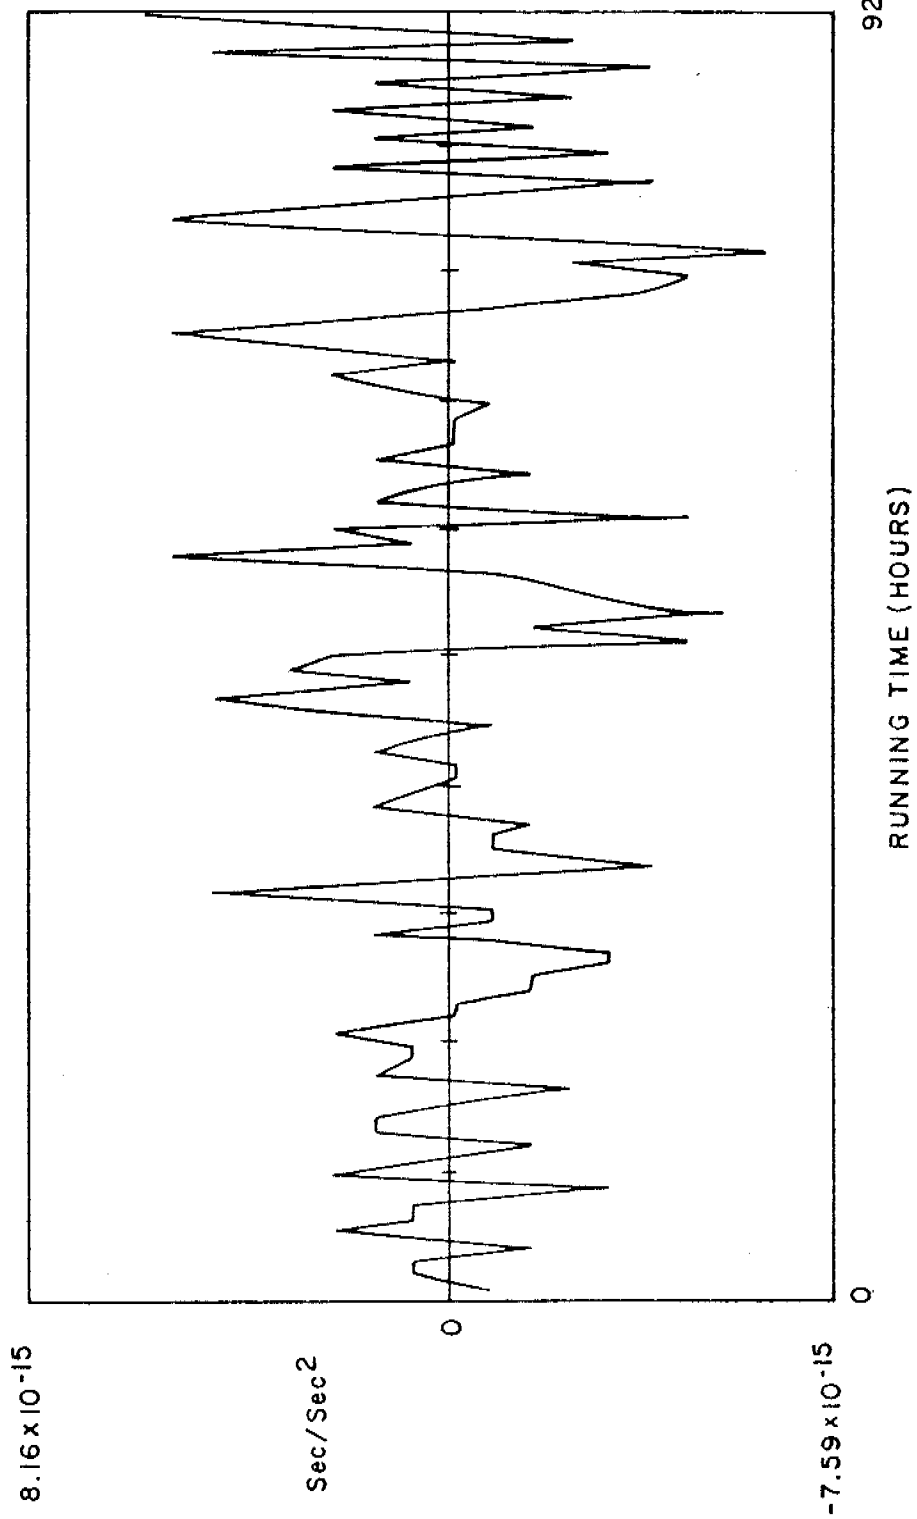(1st diff-lin)

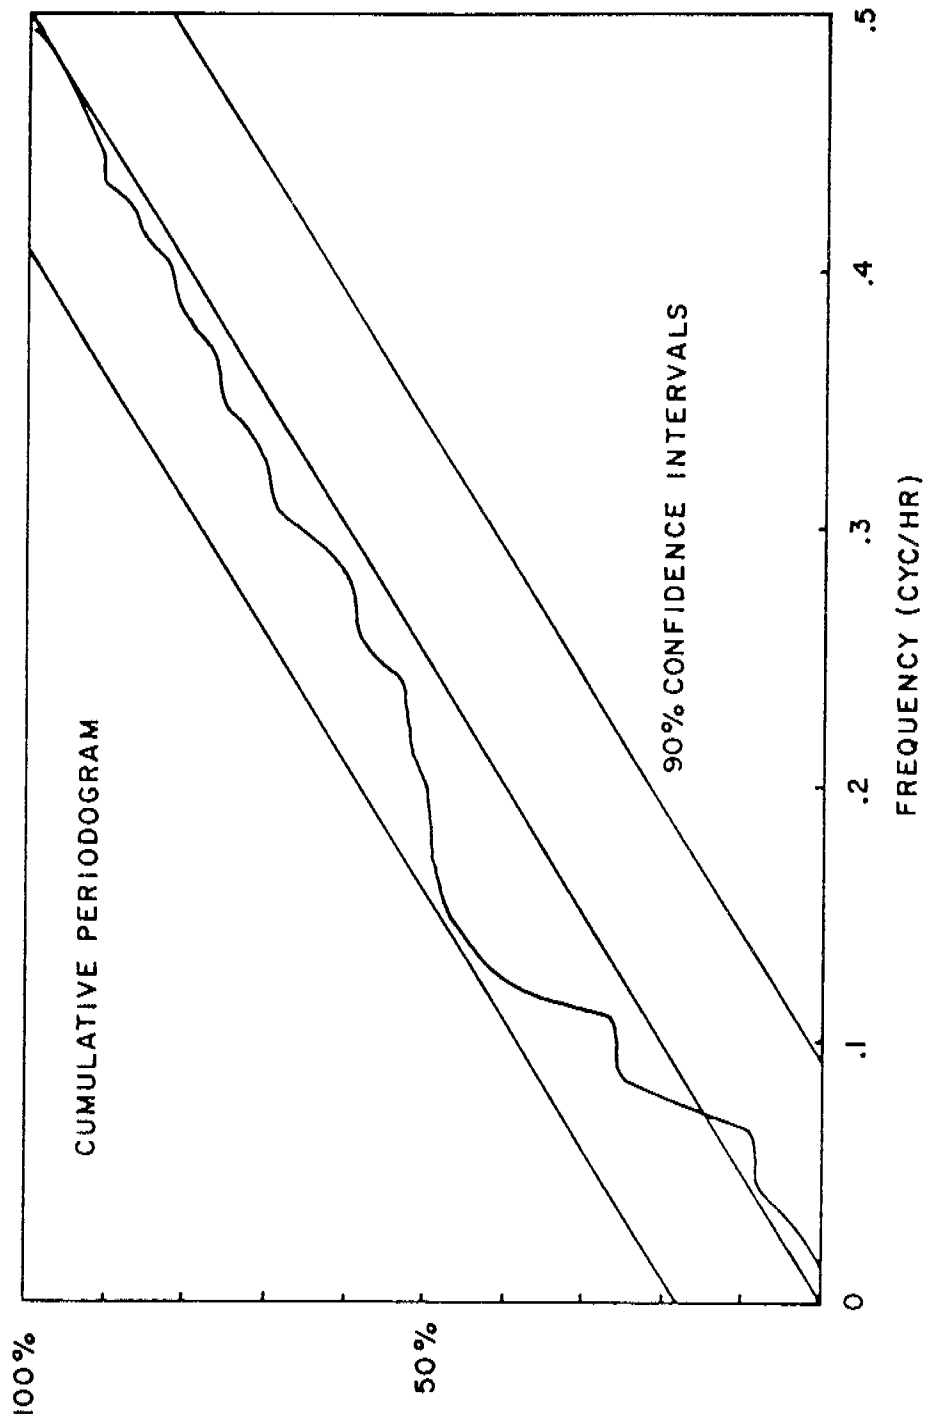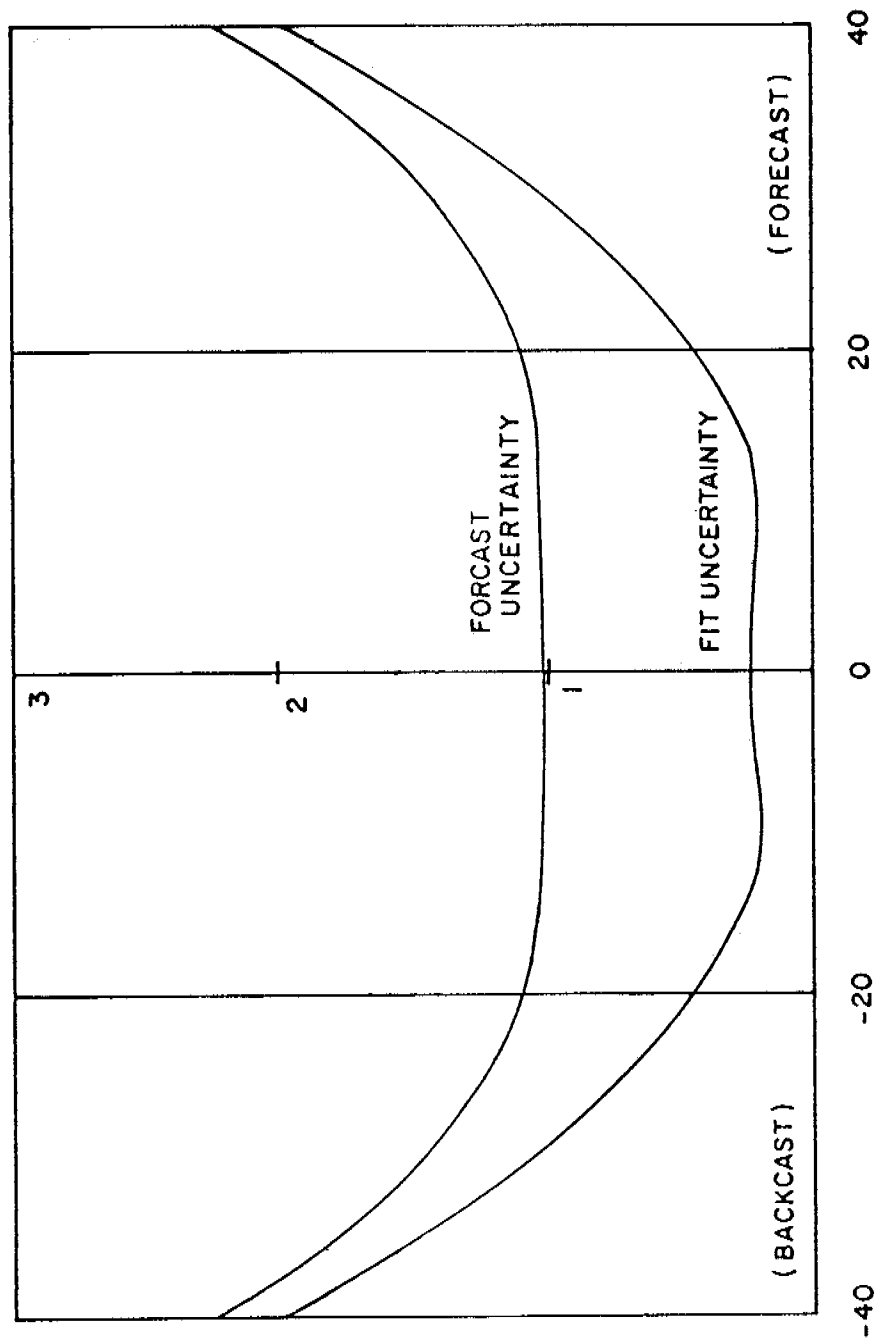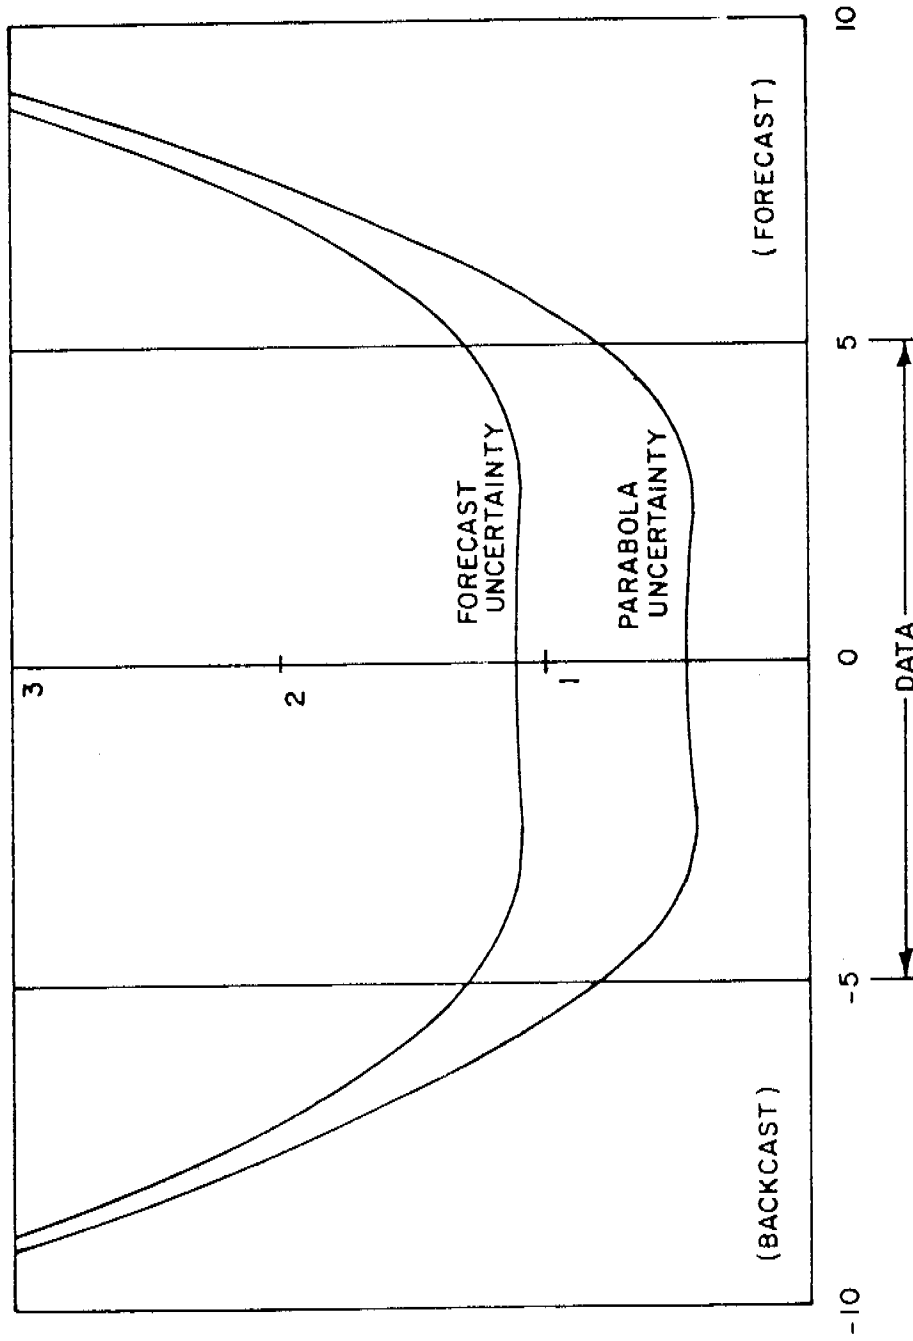Fig. 4 – First diff – linear

Fig. 5 — Second diff of phase

Fig. 6 - Second diff - avg.

FREQUENCY DRIFT AND WHITE PM

STANDARD DEVIATIONS

FORCAST UNCERTAINTY

FIT UNCERTAINTY

(BACKCAST)

(FORECAST)

FREQUENCY DRIFT AND WHITE PM

STANDARD DEVIATIONS

QUESTIONS AND ANSWERS

MR. ALLAN:

We have found the second difference drift estimator to be useful for the random walk FM process. One has to be careful when you apply it, because if you use overlapping second differences, for example, if you have a cesium beam, and you have white noise FM out to several days and then random walk FM beyond that point, and you are making hourly or 2 hour measurements, if you use the mean of the second differences, you can show that all the middle terms cancel, and in fact you are looking at the frequency at the beginning of the run and the frequency at the end of the run to compute the frequency drift and that's very poor.

DR. BARNES:

My comment would be: There again the problem is in the model, and not in the arithmatic. The models applied here were 3 very simple models, very simple, simpler than you will run into in life. It was pure random walk plus a drift or pure frequency noise plus a drift, or pure random walk of frequency noise plus a drift and it did not approach at all any of the noise complex models where you would have both white frequency and random walk frequency and a drift. That's got to be handled separately, I'm not even totally sure how to perform it in all cases at this point.

MR. McCASKILL:

I would like to know if you would comment on the value of the sample time and the reason why, of course, is that the mean second difference does depend on the sample time? For instance, if you wanted to estimate the aging rate or change in linear change of frequency, what value of sample time would you use in order to make that correction. So, really, the question is, how does the sample time enter into your calculations?

DR. BARNES:

At least I will try to answer in part. I don't know in all cases, I'm sure, but if you have a complex noise process where you have at short term different noise behavior than in long term, it may benefit you to take a longer sampling time and effectively not look at the short term, and then one of the simple models might apply. I honestly haven't looked in great detail at how to choose the sample time. It is an interesting question.

MR. McCASKILL:

Well, let me go further, because we had the benefit of being out at NBS and talking with Dr. Allan earlier and he suggested that we use the mean second difference, and the only problem is if we want to calculate or correct for our aging rate at a tau of five days or ten days, and you calculate the mean second difference, you come up with exactly the Allan Variance. What appears is that in order to come up with the number for the aging rate, you have to calculate the mean second difference using a sample time whenever you take your differences of longer than, let's say, ten days sample time. So we use, of course, the regression model, but we use on the order of two or three weeks in order to calculate an aging rate correction for, say, something like the rubidium in the NAVSTAR 3 clock. It looks like in order to come up with a valid value for the Allan Variance of five or ten day sample time you have to calculate the aging rate at a longer, maybe two or three times longer sample time.

DR. BARNES:

Dave Allan, do you think you can answer that?

MR. ALLAN:

Not to go into details, but if you assume that in the longer term you have random walk frequency modulation as the predominant noise process in a clock, which seems to be true for rubidium, cesium and hydrogen, you can do a very simple thing. You can take the full data length and take the time at the beginning, the time in the middle and the time in the end and construct a second difference, and that's your drift.

There is still the issue of the confidence interval on that. If you really want to verify your confidence interval, you have to have enough data to do a regression. You need enough data to test to be sure the model is good.

DR. WINKLER:

That argument is fourteen years old, because we have been critized here. I still believe that for a practical case where you depend on measurements which are contaminated and maybe even contaminated by arbitrarily large errors if they are digital. These, theoretical advantages that you have outlined may not be as important as the benefits which you get when you make a least square regression. In this case you can immediately identify the wrong data. Otherwise, if you put the data into an algorithm you may not know how much your data is contaminated. So for practical applications, the first model even though theoretically it is poor, it still gives reasonable estimates of the drift and you have residuals which let you identify wrong phase values immediately.

DR. BARNES:

I think that's true and you may have different reasons to do regression analysis, if your purpose is to measure a drift and understand the confidence intervals, then I think what has been presented is reasonable, if you have as your purpose to look--to see if there are indications of funny behavior in a curve that has such strong curvature or drift that you can't get it on graph paper without doing that, I think it is a very reasonable thing to do. I think looking at the data is one of the healthiest things any analyst can do.