APPLICATION NOTE 162-1

TIME INTERVAL AVERAGING

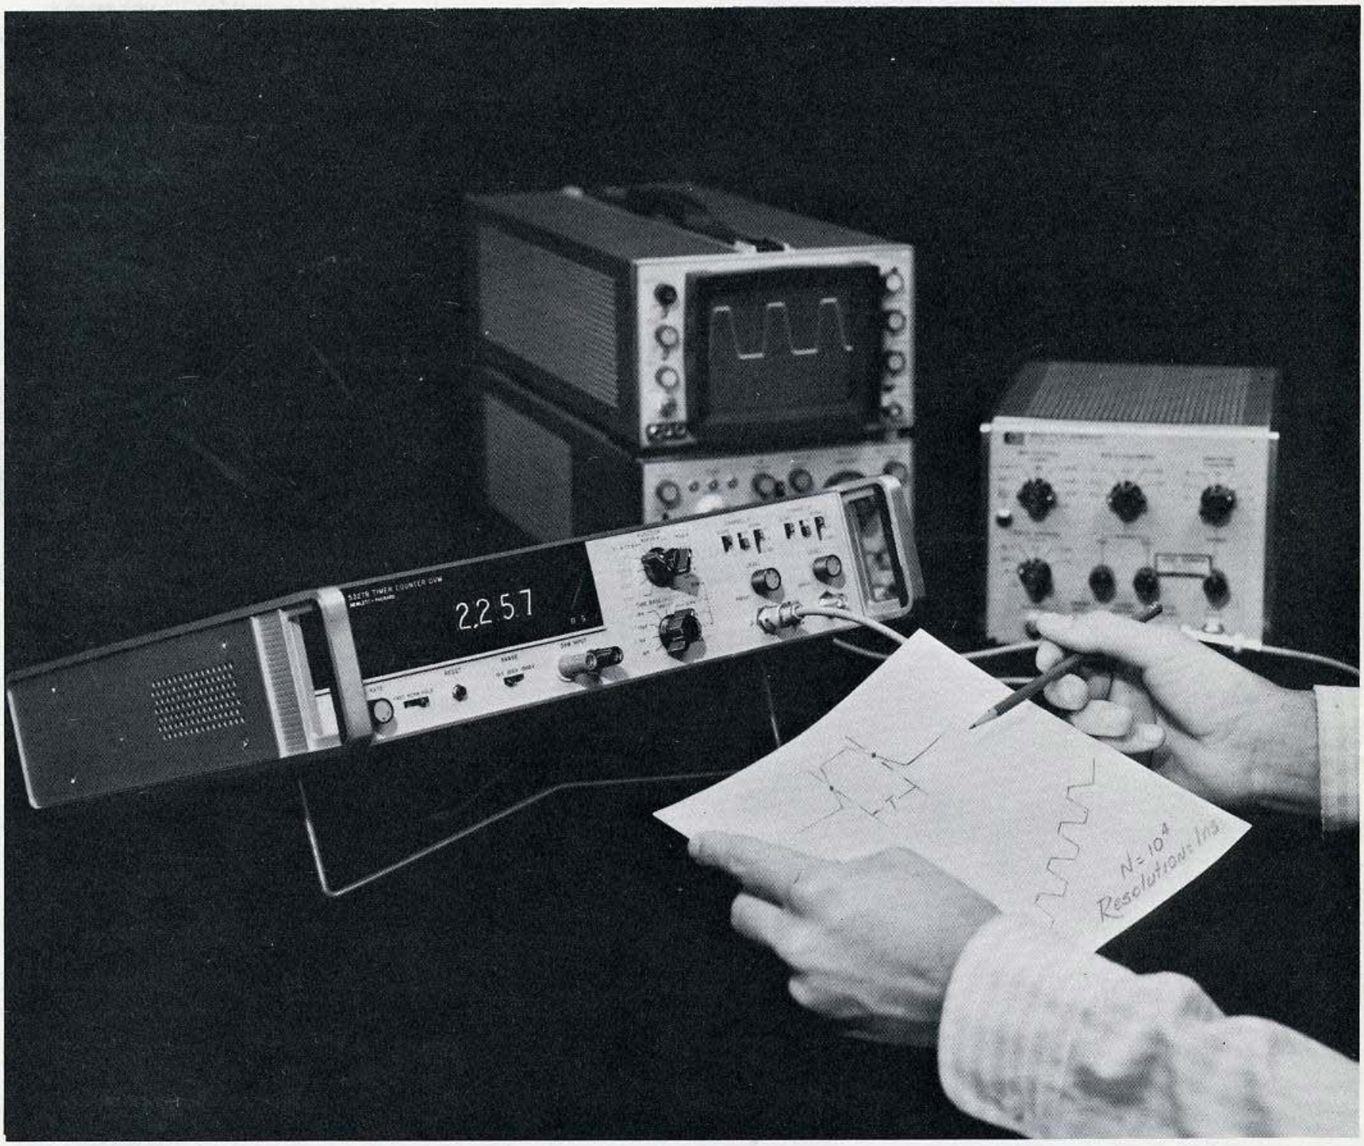

HEWLETT hp PACKARD

FOREWORD

Time interval averaging provides a powerful but unusual measurement capability. Averaging has a statistical basis, unlike other universal counter measurements. Nevertheless, time interval averaging can be used as easily as all other counter functions. Typical applications include logic timing measurements, pulse generator calibration and IC testing where very short time intervals must be measured.

This application note gives a summary of the major factors in making and evaluating time interval averaging measurements. Naturally many of these factors apply equally to conventional single shot time interval measurements. In addition, this application note provides the basic theoretical foundation of time interval averaging, showing how averaging obtains such significant improvement in accuracy and resolution. A section on the evaluation of T.I.A. measurements shows the origin of the basic accuracy specification, as well as the cases where accuracy can be even better than the general specification. For those who wish to study the statistics of averaging, the appendices contain derivations of the major results.

TABLE OF CONTENTS

LIST OF FIGURES

SECTION 1: A SYNOPSIS OF TIME INTERVAL AVERAGING

This section gives an overview of time interval averaging for the person who wants a working knowledge in summary form. The synopsis not only tells when time interval averaging can be useful, but also identifies the factors that affect the accuracy of a time interval measurement. For those who want to study T.I.A. in more detail, the remaining sections develop the basis of averaging and show when the accuracy and resolution of an averaging measurement can be even better than the general specification.

What is Time Interval Averaging?

Time interval averaging, found in the Hewlett-Packard 5326/5327 series of universal counters, provides a powerful yet economical method of greatly increasing the accuracy and resolution of time interval measurements on repetitive signals. With the HP 5326/5327 counters, time interval averaging (T.I.A.) makes possible the measurement of intervals as short as 150 ps with resolution to 100 ps — a thousandfold improvement over conventional one shot time interval measurements with little or no cost increase over counters that make only the conventional measurements.

The basis of time interval averaging is the *statistical reduction of the ±1 count error* inherent in digital measurements. As more and more intervals are averaged, the measurement will tend toward the true value of the unknown time interval. This application note describes some of the major considerations of time interval averaging measurements.

Figure 1 shows the elements of a typical time interval measurement set up. Table 1 gives some examples. Note that the driving source can also be the device under test; the sensors can simply be wire leads.

Table 1-1. Elements of a Measurement Set Up

	Example 1	Example 2
Driving Source	Oscillator	Pulse Generator
Device Under Test	Amplifier	Pulse Generator
Sensors	Wire Leads	Coax Cable
Value to be Measured	Propagation Delay	Pulse Width

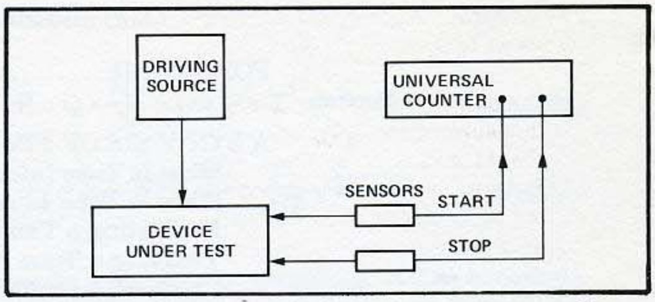

Figure 1. Elements of a Time Interval Measurement Setup.

When is Time Interval Averaging Useful?

Time interval averaging is useful when

- ±1 count error from a single time interval measurement significantly degrades the accuracy or resolution of a time interval measurement; and
- the time interval is repetitive.

Example:

The width of a repetitive pulse is approximately 1 μs. The ±1 count error in a pulse width measurement using conventional one-shot techniques is 100 ns (the period of the counter's clock). The ±1 count error for the one-shot measurement represents 10% of the time interval; however, averaging 10^4 time intervals can produce 1 ns resolution.

EVALUATING A TIME INTERVAL AVERAGING MEASUREMENT

Time interval averaging is a statistical process. After a T.I.A. measurement is made, we want to determine not only the best estimate for the true value of the measured time interval, but also a level of confidence in the measurement. The best estimate of the measured time interval is, of course, the value of the counter reading.

In general for a T.I.A. measurement with N intervals averaged, the following expression gives the accuracy at a very high level of confidence.

Accuracy of T.I.A. Measurements:

$$\pm \frac{1}{\sqrt{N}} (\pm 1 \text{ count error} + \text{internal trigger error})$$

$$+ \text{ time base error}$$

$$+ \text{ systematic error}$$

The accuracy expression considers four basic types of errors as follows:

1. *±1 count:* This is the intrinsic ambiguity of digital measurements. For time interval measurements the ±1 count error equals the period of the clock. In general, averaging reduces this error by a factor of $1/\sqrt{N}$. Often the improvement in resolution and accuracy can be even greater.

2. *Internal Trigger Error:* Internal noise in a counter's amplifier-trigger circuits can cause the counter to randomly start or stop a time interval measurement slightly early or late. Internal trigger error is generally much less than ±1 count error and so can be virtually ignored for most measurement situations.

3. *Time Base Error:* The accuracy of the counter's reference time base, or clock, can limit the accuracy of a measurement for long time intervals. In averaging situations, time base error is rarely a significant factor.

4. *Systematic Error:* Systematic error includes differences in the propagation times of the start and stop sensors, differential delays in the start and stop channel amplifiers of the counter, and errors in trigger level settings of the start and stop channels of the counter. For a given measurement set up, with constant waveforms, the systematic error is fixed. Therefore, systematic error affects accuracy but has no effect on resolution. Most systematic errors can be virtually eliminated by calibrating the measurement set up.

Verifying a Time Interval Averaging Measurement

To test the validity of a T.I.A. measurement, check the following points.

1. Does the counter have synchronized gating? Check the counter's data sheet to find out. Synchronized gating (found on the HP 5326/5327 family) is necessary for unbiased measurements.

2. Is the repetition rate of the time interval asynchronous with the counter clock? If the repetition rate is synchronous with the clock, the average may not converge to the true value of the time interval.

   If the repetition rate of the driving source can be varied, then make the T.I.A. measurement at two slightly different frequencies. Do the results agree within the computed limits for time interval averaging? This test can exclude the possibility of a synchronous relationship between the repetition rate of the driving source and the counter's clock.

3. Is the measurement reasonable? The counter reading should agree with other evidence such as scope displays or calibrated dial settings. Gross differences indicate a problem with systematic errors. Averaging reduces only the ±1 count and internal trigger errors.

CHECKING FOR SYNCHRONOUS REPETITION RATES

Time interval averaging produces valid and useful results in the vast majority of cases. However, if the repetition rate of the time interval is synchronous with the clock, then averaging may not improve resolution as much as expected. For the repetition rate to be synchronous, the driving source must not only be set at a synchronous frequency, but must also be exceptionally stable in period during the measurement.

Symptoms of a Synchronous Repetition Rate

Either of the following may indicate that the repetition rate is synchronous with the clock.

1. The counter reading seems to "hang up" on some value, particularly a reading that is an integral multiple of the clock period.
2. Averaging more intervals does not increase the resolution of the measurement.

Avoiding a Synchronous Repetition Rate

Any of the following methods effectively break a synchronous relationship between the repetition rate and the clock.

1. Change the repetition rate if possible. This method offers the easiest solution. Figure 2 shows the offset required from the synchronous frequencies as a function of the number of intervals averaged.

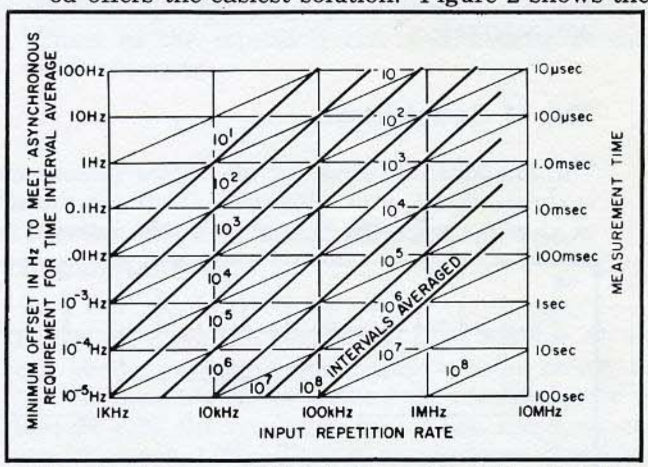

Figure 2. Frequency Offset to Ensure Repetition Rate is Asynchronous with the Counter's Clock.

2. Check the stability of the repetition rate. Average enough intervals that the relative instability of the repetition rate and the clock will effectively break the synchronous relationship.

3. Introduce phase jitter on the repetition rate. Sufficient phase jitter essentially breaks a phase locked relationship between the repetition rate and the counter clock.

SECTION 2: BASIC TIME INTERVAL MEASUREMENT

A brief review of the basic digital time interval measurement technique will help to identify the factors involved in determining the accuracy of a time interval measurement. After describing the sources of error, we can concentrate our attention more clearly on the benefits of time interval averaging for increasing the accuracy and resolution of a time interval measurement.

Figure 3 shows the basic circuit arrangement of a typical universal counter when making a time interval measurement. The time base generates a pulse train with a very accurate and stable period. To measure

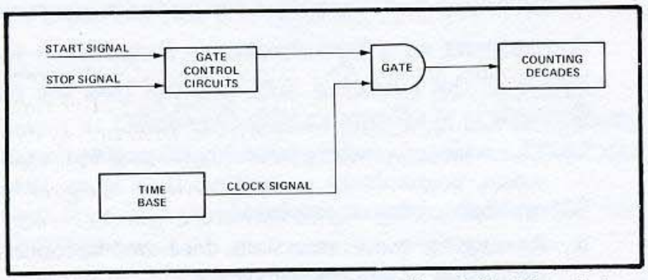

Figure 3. Block Diagram of Counter Circuit for Time Interval Measurement.

the unknown time interval, the gate must open when the start signal is received and close when the stop signal is received. While the gate is open, the time base (clock) signal passes through the gate and is counted. The total count provides the measure of the time interval.

The ±1 Count Error

The fundamental error for all digital measurements is the ±1 count error. Figure 4 illustrates the source of ±1 count error. In Figure 4a, a time interval equal

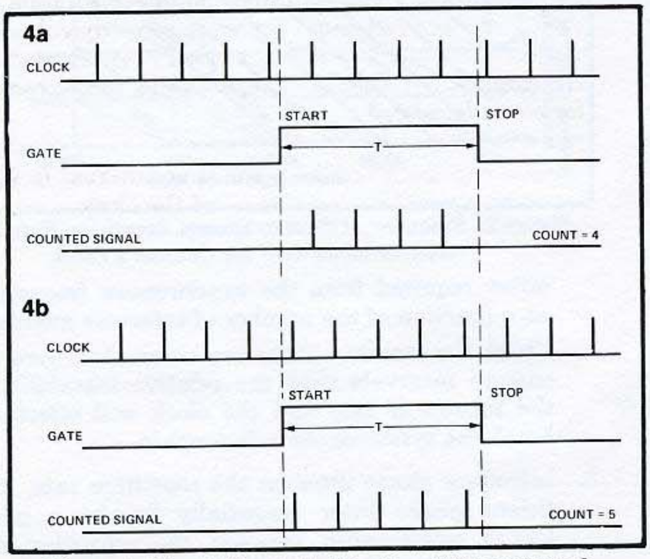

Figure 4. The ±1 Count Ambiguity of Time Interval Measurements.

to 4.5 clock periods opens the gate to allow 4 counts. In Figure 4b, the same time interval allows 5 counts through the gate. The only difference in the two situations is the phase of the clock relative to the start of the time interval. In this example, the ±1 count error is a significant percent of the measurement.

In discussing time interval measurements, this application note expresses both the measured time interval, T, and the resultant counter reading, R, in units of the counter's clock period. Thus the clock period equals one count. A time interval equal to 4.5 clock periods (4.5 counts) will produce a counter reading of 4 counts or 5 counts.

Example: The HP 5326B Counter uses a clock with a 100 ns period. Thus a time interval of 450 ns is

$$T = 4.5 \text{ counts}$$

The counter reading for a single time interval measurement will be

$$R = 4 \text{ counts}$$

or

$$R = 5 \text{ counts}$$

depending on the phase relationship between the clock and the start of the time interval.

As indicated in Figure 5, T can be broken into an integral and fractional part.

$$T = Q + F \qquad Q = 0, 1, 2, \ldots \qquad 0 \leq F < 1 \tag{1}$$

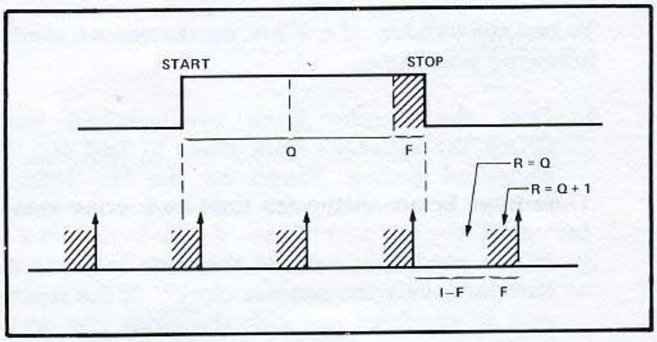

Figure 5. Measurement Zones.

After making a time interval measurement, the counter will display either

$$R = Q$$

or

$$R = Q + 1$$

depending on just when the START signal occurs. In either case, however, the counter will register at least Q counts on every measurement.

Fortunately, the expected[1] counter reading will be equal to the measured time interval. Looking at Figure 5, we can see that

$$\text{Prob}(R = Q) = 1 - F \tag{2a}$$

$$\text{Prob}(R = Q + 1) = F \tag{2b}$$

If $\overline{R}$ is the expected counter reading then

$$\overline{R} = Q + F = T \tag{3}$$

This result is essential for unbiased time interval averaging.

Other Sources of Measurement Error

In general, three additional types of errors can affect the accuracy of digital time interval measurements.

Internal Trigger Error: Internal noise in a counter's amplifier and trigger circuits can cause the start and stop channels to trigger at random slightly above or below the actual trigger level setting. Figure 6 shows the effect of trigger error in a pulse width measurement. For most measurements, this random trigger error is much smaller than the ±1 count error. However, if internal trigger error should be significant, time interval averaging reduces the error just as it reduces ±1 count error.

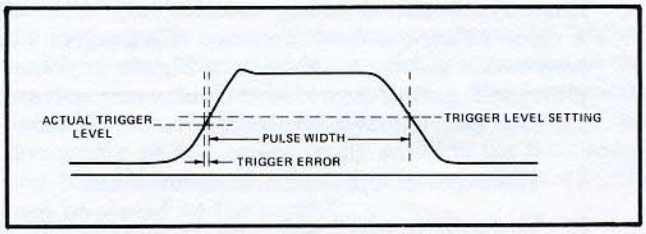

Figure 6. Trigger Error in Time Interval Measurement.

Time Base Error: Although time base error may be a factor in the measurement of relatively long time intervals, it is at most a minor factor for time interval averaging situations.

The error results from the difference between the actual time base oscillator frequency and its nominal frequency. For most high quality universal counters the time base is accurate at least to parts in 10^6, with far more accurate time bases available.

A simple example shows why time base error is usually far less than the ±1 count error when measuring short time intervals. Suppose we use a counter with a 10 MHz clock to measure a repetitive time interval of approximately 1 μs. The ±1 count error of 0.1 μs represents ±1 part in 10. Time interval averaging can reduce this error to ±1 part in 10^4, a thousandfold improvement. However, the time base error of parts in 10^6 is still negligible.

Systematic Error: Any slight mismatch between the start channel and the stop channel amplifier risetimes and propagation delays results in internal systematic errors. Mismatched probe leads or cable lengths introduce external systematic errors. Trigger levels set at the wrong point introduce another source of systematic error. Figure 6 shows how improperly set trigger levels will cause measurement error.

Thus precise measurement set up is essential for accurate results. The HP 5326B and 5327B provide a built-in DVM so that trigger levels can be measured and set with digital accuracy by reading the levels directly on the display. The built-in DVM of the 5326B and 5327B provides a far more convenient and accurate method of setting trigger levels than other available techniques.

Fortunately, for a given measurement set up with constant waveforms, the systematic error is fixed and will be repeated in every measurement. Thus systematic errors can reduce the accuracy of a measurement but will have no effect on the resolution. In fact, if the measurement set up is properly calibrated, systematic errors can be virtually eliminated (see HP Application Note 129, Logic Timing Measurements, for more details).

[1] Assuming the START of the time interval is equally likely to occur at all phases of the clock.

SECTION 3: COUNTER GATING

Time interval averaging requires special techniques for control of the counter's main gate. Traditional direct control methods of gating can introduce significant bias into T.I.A. measurements. However, synchronized gating, developed by Hewlett-Packard for the 5326/5327 series of counters, produces reliable and accurate measurements.

Direct Gating

Direct gating can cause an unacceptable bias in time interval measurement by truncating clock pulses. Figure 7 shows what can happen using direct gating.

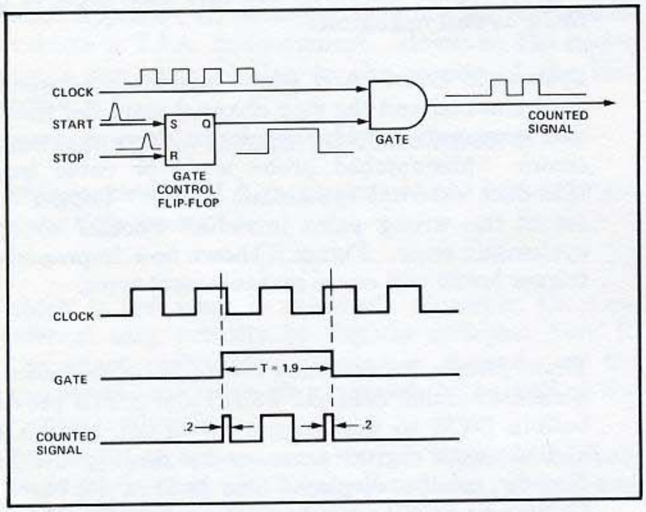

Figure 7. Direct Gating.

The clock signal is actually a pulse train. When the gate opens it may truncate some fraction of a clock pulse. When closing, the gate may again truncate a clock pulse. The counter does not know which of the truncated pulses should be counted. In Figure 7, if the minimum countable pulse width of the counter is less than 0.2 then the counter will display

$$R = 3$$

which produces an error of *greater than 1 count.* Such errors can produce a significant bias in the expected counter reading.

In summary, direct gating has the following disadvantages for time interval averaging measurements.

1. Truncation of clock pulses can produce more than 1 count error.
2. Time interval measurements will be biased.
3. The time interval can be too short. The counter will never count intervals shorter than the minimum countable pulse width.

Appendix I gives more detail about the expected bias produced by direct gating.

Synchronized Gating

Synchronized gating solves the problem of bias in time interval measurements. In fact, synchronized gating is a key feature of the exceptional time interval measurement capability of the HP 5326/5327 family.

Figure 8 shows a representative synchronized gating circuit and resultant gate timing. In practice several variations of the circuit may be used.[2] The gate is "synchronized" to the clock. The start and stop signals properly arm the gate to either open or close; an *edge of a clock pulse* actually switches the gate control flip-flop. Thus only integral clock pulses can pass the gate. No clock pulses are truncated.

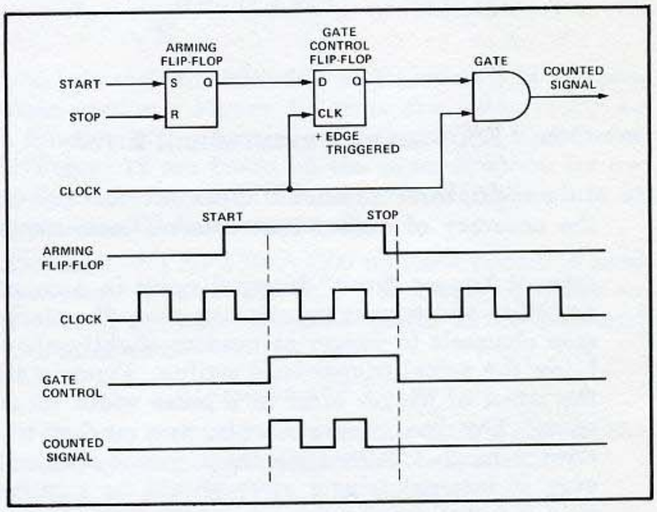

Figure 8. Synchronized Gating.

Since synchronized gating operates only on an edge of a clock pulse, the clock becomes effectively a train of zero width pulses, as shown in Figure 9. Thus synchronized gating provides the following advantages for time interval measurements:

1. The expected measurement is unbiased.
2. The synchronized gate can be designed to measure time intervals even shorter than the minimum countable pulse width of the counter.

These advantages make synchronized gating essential for time interval averaging.

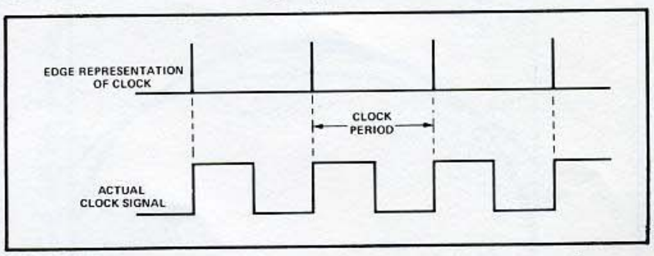

Figure 9. Edge Representation of Clock Signal.

[2]See the HP Journal, April 1970, p. 6-7.

SECTION 4: TIME INTERVAL AVERAGING

To average N time interval measurements, a counter accumulates the counts for the individual measurements until N intervals have passed. Since N is usually selectable in decade steps (N = 1, 10, 100, ...) the displayed reading is the total count with a positioned decimal point.

Example:

Using an HP 5326B Counter, we are measuring a pulse width of 225 ns. The 5326B has a clock period of 100 ns. Therefore each time interval measurement will yield either 2 or 3 counts, that is, a reading of 200 ns or 300 ns. But if we average 100 measurements, we may obtain the following results:

Counts During Interval	Number of Intervals	Total Counts
2	75	150
3	25	75
Total:	100	225

The total count is 225 for 100 intervals, an average of 2.25 counts per measurement. Thus the counter will display 225 ns.

The Expected Value of an Averaging Measurement

In the example above, averaging produces the "right" answer. But averaging is a statistical process, so the answer is not entirely guaranteed. Two questions come to mind. First, is the expected value of an averaging measurement really equal to the time interval being measured? Second, what confidence limits can be placed on the result?

Appendix II develops the statistical basis of time interval averaging when the repetition rate is asynchronous with the clock. The major results are:

1. The expected counter reading equals the time interval being measured. That is

$$\overline{R} = T \tag{4}$$

   This result is obviously critical for valid time interval averaging.

2. The standard deviation of the counter reading is proportional to $1/\sqrt{N}$.

Recall that the time interval, T, can be expressed as an integral part, Q, plus a fractional part, F

$$T = Q + F \qquad Q = 0, 1, 2, \ldots$$

$$0 \leq F < 1$$

After averaging N time intervals, the counter reading will have a binomial probability distribution between Q and Q + 1 (the two possible results of the individual measurements). The actual probability distribution depends on the number of intervals averaged and the fractional part of the time interval being measured. The standard deviation of the averaging measurement will be

$$\sigma_R = \frac{1}{\sqrt{N}}\sqrt{F(1-F)} \tag{5}$$

Thus the effective resolution of the measurement depends on both $1/\sqrt{N}$ and F. Figure 10 shows the standard deviation of the measurement as a function of F. Note that in the worst case

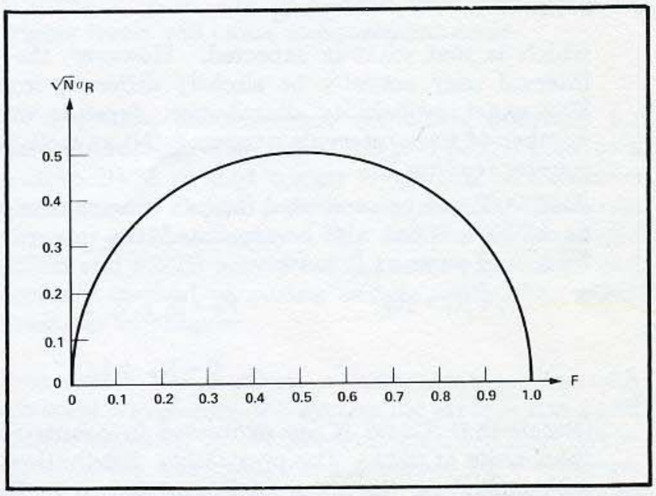

Figure 10. Standard Deviation of the Counter Reading as a Function of the Fractional Part of the Time Interval.

$$\sigma_R\ (F = 1/2) = \frac{1}{2\sqrt{N}}$$

This result is well within the $1/\sqrt{N}$ specification for T.I.A. measurements. For other values of F, the confidence limits on the counter reading will be even higher.

SECTION 5: EVALUATING A T.I.A. MEASUREMENT

The previous section discussed the expected counter reading for a time interval averaging measurement assuming that the time interval is known. In practice, the problem is just the reverse. Given a counter reading what do we know about the true value of the measured time interval?

In the synopsis of time interval averaging in Section 1, the general accuracy specification shows the ±1 count error is reduced by a factor of $1/\sqrt{N}$. Such a general specification requires using limits that are valid in all circumstances where the repetition rate is not synchronous with the clock. Yet in many cases, resolution can be far better than the $1/\sqrt{N}$ factor may indicate. Appendix III develops the statistics required to evaluate a T.I.A. measurement. However, the major results are very similar to the results of the last section.

Estimating the Measured Time Interval

After making a T.I.A. measurement, the best estimate of the time interval is

$$\hat{T} = R \qquad (6)$$

which is just what is expected. However, the time interval may actually be slightly different from R. The exact probability distribution depends on the number of time intervals averaged (N) as well as the counter reading R.

Just as T can be separated into an integral and fractional part, R can also be separated into integral and fractional parts as follows:

$$R = R_I + R_F \qquad R_I = 0, 1, 2, \ldots$$

$$0 \leq R_F < 1 \qquad (7)$$

(Recall that T and R are expressed in counts rather than units of time.) The probability distribution of T will depend on the value of R_F as shown in Figure 11. This figure is drawn for the average of 10 time intervals. Naturally, if more intervals are averaged, the probability curves will be much more sharply peaked. Those who recognize the curves as belonging to a beta probability function may wish to review Appendix III.

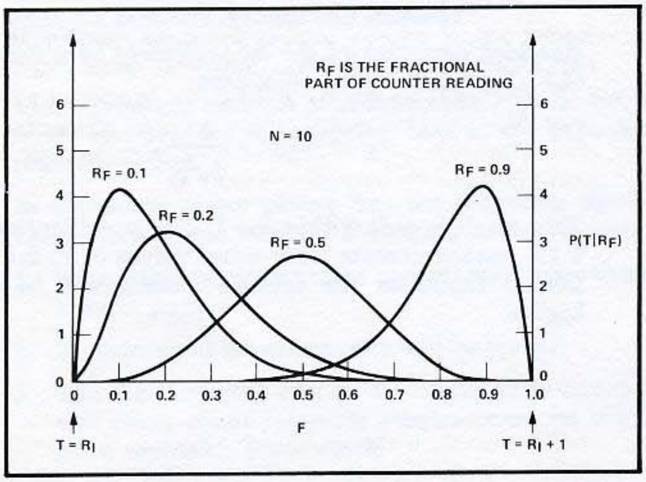

Figure 11. Probability of F for Several Values of R_F.

Setting Confidence Limits on a T.I.A. Measurement

The standard deviation of the probability distribution for T is an effective estimate of the resolution of a T.I.A. measurement. Not too surprisingly, the value of σ_T takes a familiar form. When 100 or more intervals are averaged, an excellent approximation of σ_T is

$$\sigma_T \simeq \frac{1}{\sqrt{N}}\sqrt{R_F\,(1 - R_F)} \qquad R_F \neq 0 \qquad (8)$$

which is virtually identical to Equation 5 in the previous section. Figure 12 plots the value of σ_T as a function of R_F for several values of N. The curves in Figure 12 are based on the exact equation for σ_T as found in Appendix III. An example shows how to evaluate a T.I.A. measurement.

Example: An HP 5327A (100 ns clock period) is used to average 10^4 time intervals. The counter displays an answer of 225 ns. The best estimate of the measured time interval is

$$T_{measured} = 225 \text{ ns}$$

To find the standard deviation for T, first convert the counter reading from nanoseconds into counts

$$225 \text{ ns} = 2.25 \text{ counts}$$

Take the fractional part of the counter reading

$$R_F = 0.25 \text{ counts}$$

Now use the curves in Figure 12 or the Equation 8 to find σ_T as an estimate of measurement resolution.

$$\sigma_T = \frac{1}{\sqrt{10^4}}\sqrt{.25\,(1 - .25)}$$

$$\sigma_T = .0043 \text{ counts}$$

$$\sigma_T = 0.43 \text{ ns}$$

Thus averaging provides a striking improvement over the ±100 ns resolution of single-shot measurements.

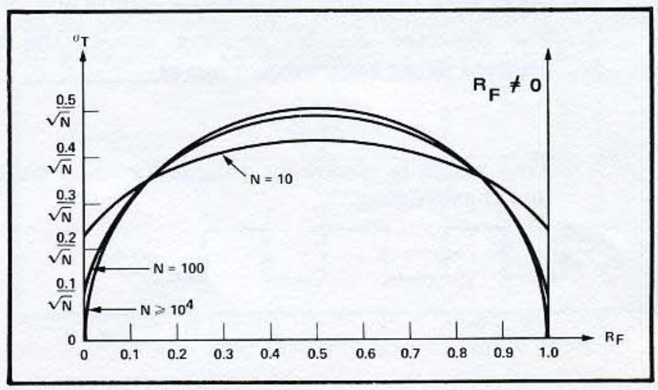

Figure 12. σ_T as a Function of the Fractional Part of the Counter Reading.

SECTION 6: THE CASE OF SYNCHRONOUS REPETITION RATES

So far all the results for averaging are based on a repetition rate that is asynchronous with the counter clock. However, in some cases a synchronous repetition rate can actually improve the resolution of a T.I.A. measurement; in other cases the synchronous repetition rate limits the resolution of the measurement.

The synopsis of time interval averaging lists some rules of thumb for recognizing a synchronous repetition rate and proposes solutions for any difficulties that may be encountered. This section examines more closely just what a "synchronous" repetition rate is and what effect it has on an averaging measurement.

A repetition rate is synchronous with the clock if the start of each time interval always occurs at one particular phase of the clock, or at some limited number of points during the phase of the clock. However, T.I.A. is based on the start of N time intervals tending to occur uniformly throughout the clock period. Figure 13 compares an asynchronous repetition rate with a synchronous repetition rate. There are an infinite number of repetition rates that produce perfectly valid averaging; there are only a finite number of cases where synchronous repetition rates can cause difficulties.

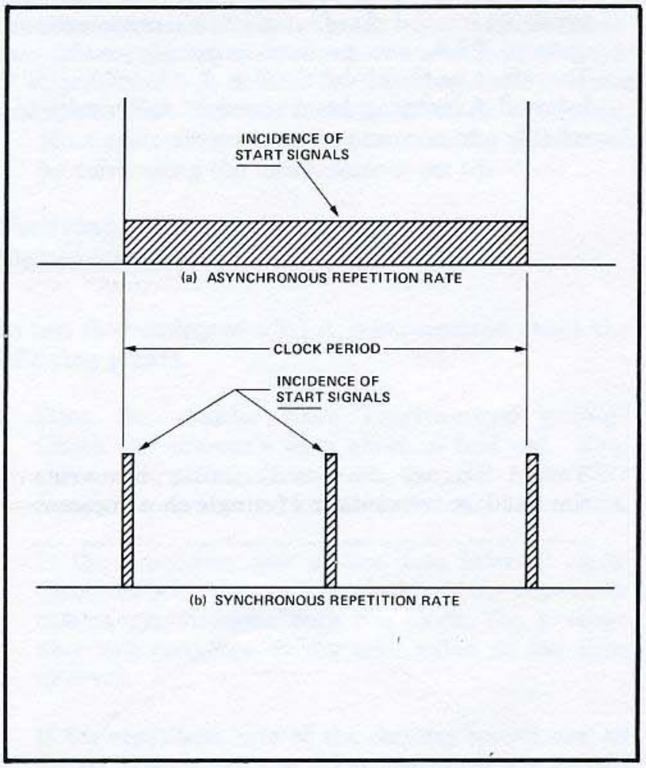

Figure 13. Comparison of Synchronous and Asynchronous Repetition Rates.

Example:

A counter has a 100 ns clock. A repetitive sequence of time intervals starts every 250 ns. If the first start signal occurs simultaneously with a clock pulse, then the second start signal will occur exactly between two clock pulses. Each successive start will occur either at a clock pulse or midway between clock pulses, but at no other time. Thus the repetition rate is synchronous with the clock.

The Class Number "M"

In the example above, the repetition rate has class number M = 2. That is, the repetitive start signals occur only at two points in the clock period. In general, if the repetitive start signals occur only at M points during the clock period, then the repetition rate has class M.

The theoretical significance of M is a key issue. *No matter how many time intervals are averaged, the resolution of the time interval measurement can never by increased by more than a factor of 1/M.* In theory, if M = 1 then averaging provides no benefit at all.

In practice, averaging can still improve a measurement considerably. Neither the counter clock nor the repetition rate are ever perfectly stable. Unless the driving source and the counter clock are actually phase locked together, short term fluctuations and phase jitter from both can break the synchronous relationship. If enough periods are averaged, instabilities in the repetition rate will usually produce valid averaging.

The Benefit of a Synchronous Repetition Rate

Better yet, if we can control the class number M, we can obtain averaging that improves more nearly as $1/N$ than as $1/\sqrt{N}$. By setting the repetition rate so that M = N, the resolution of a T.I.A. measurement can approach $1/N$. Thus a synchronous repetition rate can actually improve averaging.

Finding the Class Number

When using an HP 5326/5327 Universal Counter for T.I.A. measurements, a simple procedure can determine the class number M.

1. Switch the function selector from time interval average to period average[3] to measure the repetition rate.
2. Divide the period of the repetition rate by the clock period. Look only at the fractional part, P_F, of the result.
3. If P_F can be expressed as

$$P_F = \frac{L}{M} \qquad L, M = 1, 2, 3, \ldots$$

$$L < M \qquad (9)$$

where L and M are integers and are co-prime, then the repetition rate has class M.

Example:

Step 1 — Using period averaging with the HP 5326B, the repetition rate of a time interval measures 166.7 ns.

Step 2 — The period of the counter clock is 100 ns. Therefore

$$\frac{166.7}{100} = 1.667 \qquad P_F = .667$$

Therefore the repetition rate has a class of M = 3.

[3]Period averaging is not a statistical process. Thus period averaging provides a certain improvement in resolution proportional to 1/N.

Step 3 —

$$P_F = .667 = \frac{2}{3} = \frac{L}{M}$$

A few classes are shown in Table 6-1 below;

Table 6-1. Some Synchronous Classes

P_F	M
0.0	1
0.5	2
0.333, 0.667	3
0.25, 0.75	4

Summary

Only in rare cases does a synchronous repetition rate become a problem in T.I.A. measurements. Unless the repetition rate is exceptionally stable (e.g., a synthesizer) and fixed at a critical synchronous frequency, T.I.A. can produce extremely useful results. The steps outlined in Section 1, "Verifying a Time Interval Averaging Measurement" will easily lead to reliable and accurate measurements.

APPENDIX I

MEASUREMENT BIAS CAUSED BY DIRECT GATING

Section 3 of this application note discusses how direct gating introduces a bias into time interval measurements. This bias depends on both the duty cycle, d, of the clock and the minimum countable pulse width, m, for the counter. For any counter,

$$m < d$$

or the counter would not count at all. If we define the minimum countable fraction of the clock pulse to be "r", then

$$r = \frac{m}{d} \qquad 0 \leq r < 1$$

Direct gating will introduce a bias in the time interval measurements which is just

$$\text{Bias (in counts)} = d(1-2r)$$

as shown in Figure A-1. Only with the parameter r = 1/2 can direct gating produce unbiased measurements. However the value of "r" is virtually impossible to control in constructing a counter.

Synchronized gating effectively reduces the clock pulse width to zero by actually counting only on a clock *edge*. Thus synchronized gating yields unbiased measurements regardless of the value of "r".

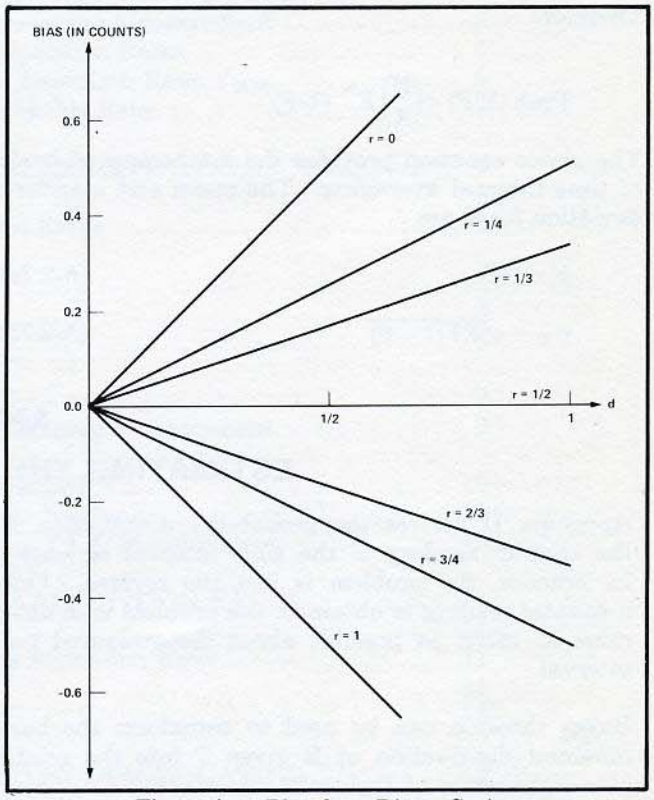

Figure A-1. Bias from Direct Gating.

APPENDIX II

THE PROBABILITY DISTRIBUTION FOR TIME INTERVAL AVERAGING MEASUREMENTS

The expected value and the standard deviation of a T.I.A. measurement depend on the probability distribution of the counter reading about the true value of the measured time interval, T.

In Section 2 we found that for a time interval with

$$T = Q + F \qquad Q = 0, 1, 2, \ldots \qquad 0 \leq F < 1$$

the counter will measure either

$$R = Q \quad \text{or} \quad R = Q + 1$$

For simplicity, in each measurement i of the N measurements to be averaged

$$R_i = Q + k_i \qquad k_i = 0, 1 \qquad i = 1, 2, \ldots, N \qquad \text{(A-2.1)}$$

Now the average of N individual measurements is just

$$R = \frac{1}{N}\sum_{i=1}^{N} R_i$$

$$R = Q + \frac{1}{N}\sum_{i=1}^{N} k_i \qquad k_i = 0, 1 \qquad \text{(A-2.2)}$$

If we let

$$K = \sum_{i=1}^{N} k_i \qquad K = 0, 1, 2, \ldots, N \qquad \text{(A-2.3)}$$

Then

$$R = Q + K/N \qquad K = 0, 1, 2, \ldots, N \qquad \text{(A-2.4)}$$

Thus the counter reading depends on K, or the number of times that the count for a single measurement was (Q+1) rather than just Q.

The variable K turns out to have a simple binomial distribution* because it is the sum of N variables k_i, each with

$$\text{Prob}(k_i = 0) = 1 - F \quad \text{(A-2.5a)}$$

$$\text{Prob}(k_i = 1) = F \quad \text{(A-2.5b)}$$

Therefore

$$\text{Prob}(K|F) = \binom{N}{K} F^K (1-F)^{N-K} \quad \text{(A-2.6)}$$

The above equation provides the mathematical basis of time interval averaging. The mean and standard deviation for K are

$$\overline{K} = NF \quad \text{(A-2.7a)}$$

$$\sigma_K = \sqrt{NF(1-F)} \quad \text{(A-2.7b)}$$

Thus we find that

$$\overline{R} = Q + \frac{NF}{N} = Q + F = T \quad \text{(A-2.8a)}$$

$$\sigma_R = \frac{1}{N}\sqrt{NF(1-F)} = \frac{1}{\sqrt{N}}\sqrt{F(1-F)} \quad \text{(A-2.8b)}$$

The results obtained in Equations A-2.8a,b contain the major aspects of time interval averaging as discussed in Section 4.

*This assumes that the start of each time interval occurs with equal probability over all phases of the clock. For the vast majority of cases this assumption is perfectly valid.

APPENDIX III

ESTIMATING THE MEASURED TIME INTERVAL

Appendix II derives the probability distribution for the counter reading if the time interval is known. In practice, the problem is just the reverse. Once a counter reading is obtained, the problem is to determine as much as possible about the measured time interval.

Bayes theorem can be used to transform the basic binomial distribution of R given T into the related beta distribution of T given R. Recall that

$$T = Q + F \qquad Q = 0, 1, 2, \ldots$$

$$0 \leq F < 1$$

and that from Appendix II the counter reading after averaging N intervals will be

$$R = Q + K/N \quad \text{(A-3.1)}$$

with K binomially distributed with

$$\text{Prob}(K|F) = \binom{N}{K} F^K (1-F)^{N-K} \quad \text{(A-3.2)}$$

Assuming no prior knowledge about T or F, Bayes theorem yields

$$p(F|K) = (N+1)\binom{N}{K} F^K (1-F)^{N-K} \qquad K = 1, 2, \ldots, N-1$$

$$0 \leq F < 1 \quad \text{(A-3.3)}$$

which is a beta distribution for F. To use this distribution it is necessary to find K from the counter reading. In Section 5, R was separated into integral and fractional parts as follows

$$R = R_I + R_F \qquad R_I = 0, 1, 2, \ldots$$

$$0 \leq R_F < 1 \quad \text{(A-3.4)}$$

Since R_F corresponds to the term K/N in Equation A-3.1,

$$K = N \cdot R_F \quad \text{(A-3.5)}$$

K will always be an integer.

Figure 11 in Section 5 plots a few examples of Equation A-3.3. Two key finds come from the distribution for F.

1. The best estimate of the measured time interval is simply the counter reading

$$\hat{T} = R \quad \text{(A-3.6)}$$

2. The standard deviation of T is**

$$\sigma_T = \sqrt{\frac{(K+1)(N-K+1)}{(N+2)^2(N+3)}} \qquad K = 1, 2, \ldots, N-1 \quad \text{(A-3.7)}$$

However, for large values of N, a very good approximation of σ_T is

$$\sigma_T \simeq \frac{1}{\sqrt{N}}\sqrt{R_F(1-R_F)} \qquad 0 < R_F < 1 \quad \text{(A-3.8)}$$

Equation A-3.8 clearly shows the $1/\sqrt{N}$ factor used for evaluating T.I.A. measurements. In the worst case

**The case of K = 0 will be mentioned shortly.

$$\sigma_T\ (R_F = 1/2) = \frac{1}{2\sqrt{N}}$$

As plotted in Figure 12, for other values of R_F, the standard deviation is less.

A Special Circumstance

Sometimes we may obtain a counter reading that is an integral multiple of the clock period.† That is

$$R_F = 0$$

Now R_F can become zero in either of two equally likely ways:

Case I

$$R = R_I + \frac{K}{N} \qquad \text{with } K = 0;\ Q = R_I$$

Case II

$$R = (R_I - 1) + K/N \qquad \text{with } K = N;\ Q = R_I - 1$$

Analysis shows that when $R_F = 0$, the statistics are as follows:

1. The best estimate for the measured time interval is still simply the counter reading

   $\hat{T} = R$

2. The standard deviation of T is

$$\sigma_T\ (R_F = 0) = \sqrt{\frac{2}{(N+2)\,(N+3)}} \qquad \text{(A-3.9)}$$

which is well within the factor of $1/\sqrt{N}$. For large N the standard deviation is approximately $\sqrt{2}/N$ and is much smaller than $1/\sqrt{N}$.

†The case of $R_F = 0$ can certainly occur as the result of a perfectly valid T.I.A. measurement. However, it can also be a symptom of a synchronous repetition rate. Check for a synchronous repetition rate by following the steps listed in Section 1.

HEWLETT hp PACKARD

Printed in USA