Application Note 191-7

HEWLETT PACKARD

HP 5370B Universal Time Interval Counter

High-Speed Timing Acquisition
and Statistical Jitter Analysis

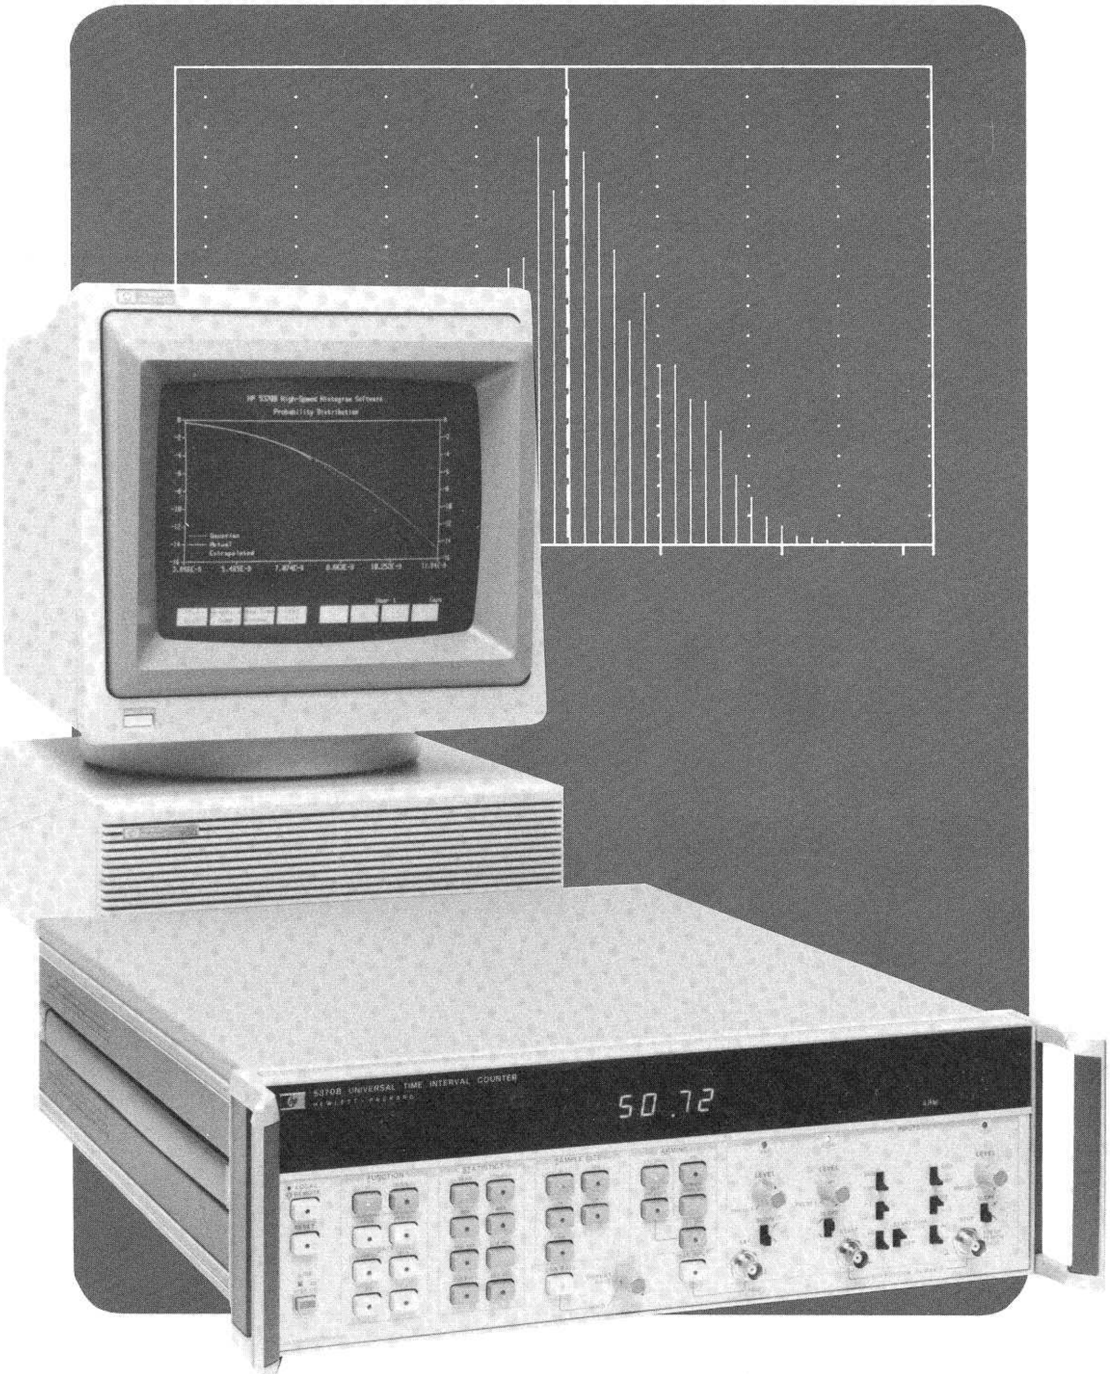

TABLE OF CONTENTS

High-Speed Timing Acquisition and Statistical Jitter Analysis Using the HP 5370B Universal Time Interval Counter

Introduction

The Time Interval Counter is a powerful tool for analyzing timing phenomena in digital or Pulse Code Modulated signals. These signals are susceptible to various forms of timing jitter, and can be found in a wide variety of systems that involve magnetic discs, fiber optics, high-speed modems, digital audio, video, sonar, or satellite communications. The histogram is an effective presentation of timing data, because it can compress a large number of measurements into a small amount of memory, without substantial loss of pertinent information. Also, many statistical analyses performed on grouped data, such as a histogram, run many times faster than those performed on ungrouped data. This Application Note describes how to use the HP 5370B High-Speed Histogram Software to capture and analyze various Time Interval (TI) measurements at a throughput of up to 4200 measurements per second with the HP 5370B Universal Time Interval Counter and HP 9000 Series 200/300 controller.

What Is a Timing Histogram?

A Timing Histogram is a graphical and statistical representation of TI data, in which the data is sorted according to TI values (see Figure 1). The histogram is made up of adjacent bins, each representing a unique TI subrange. A TI measurement is placed into the histogram by determining which TI subrange it falls into, and then incrementing the corresponding bin count by one.

When making jitter measurements, the spread of TIs is usually limited such that every TI encountered falls into one of a few defined bins. Taking more data does not increase the number of bins in the histogram, only the number of counts in each bin. Thus, the amount of memory required to store the histogram is relatively small and independent of the amount of data accumulated. Furthermore, the time required to perform statistical calculations depends only on the number of bins, not on the amount of data. It takes no longer to calculate the standard deviation of one million TIs than it does one hundred.

Graphically, the histogram shows at a glance the distribution of TI data; its shape conveys information as to how the TI data is behaving. Statistically, the histogram represents the probability density function of the TI variable being measured, and contains information that can be extracted via statistical analysis.

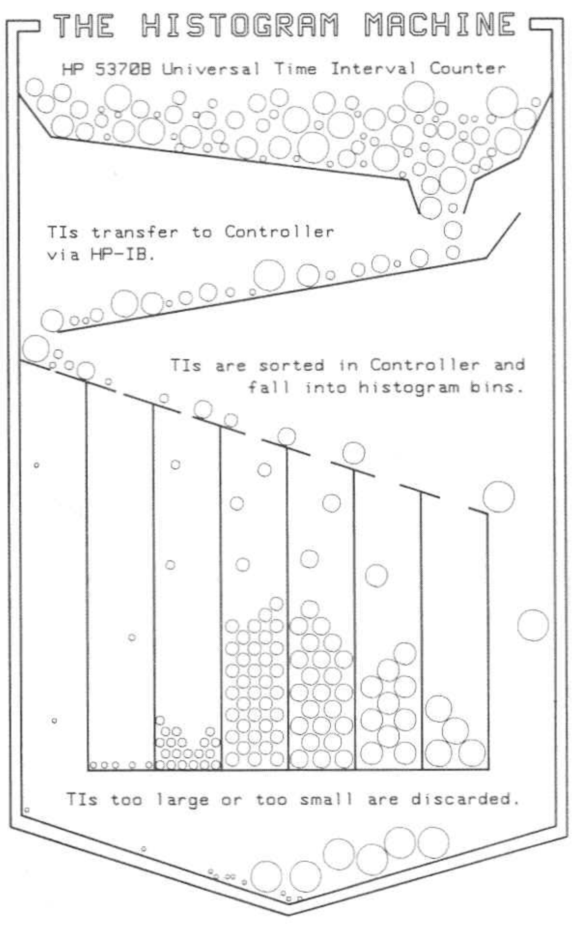

Figure 1. The Histogram sorts TI data by its value.

Histogram Limitations

Because the histogram sorts the TI data, the order in which the TIs are received is lost. Under normal circumstances this is not an issue, unless TI drift over time is of interest. This limitation can be circumvented by storing the data on disc as described in the Main Menu section.

Usually, each histogram bin has the same time width or resolution. When a TI falls into that bin, the bin count is incremented, but the TI is then discarded. When the bin resolution is the same as the counter's resolution, there is exactly one histogram bin for every possible TI that the counter will send, and no information (other than sequential) is lost.

However, quite often the bin resolution is larger (wider) than that of the counter; thus, several close but different TIs might fall into one histogram bin. Once this happens, it is impossible to distinguish these TIs in the bin. It is therefore important to properly configure the histogram so that the bin resolution is large enough to suit your measurement needs, but not so large that pertinent information is lost.

Gaussian Distribution

The Gaussian or Normal distribution is one of the most common distributions which appears when statistically analyzing data, timing jitter in particular. One reason for this is that jitter can be caused by variations of many independent parameters. A random variable composed of other independent variables tends to have a Gaussian distribution. In the real world there are no truly Gaussian variables, because the domain of such is ±infinity; however, the Gaussian distribution provides a sound model for making predictions.

The histogram is a quick way to determine at a glance if a set of data is distributed normally. Data that is skewed will show up as an asymmetric or distorted bell curve, indicating that the jitter is not Gaussian (see Figure 2). Very often this is the first step in characterizing data. Whether or not the data is distributed normally will indicate what other measurements and/or analyses should be made.

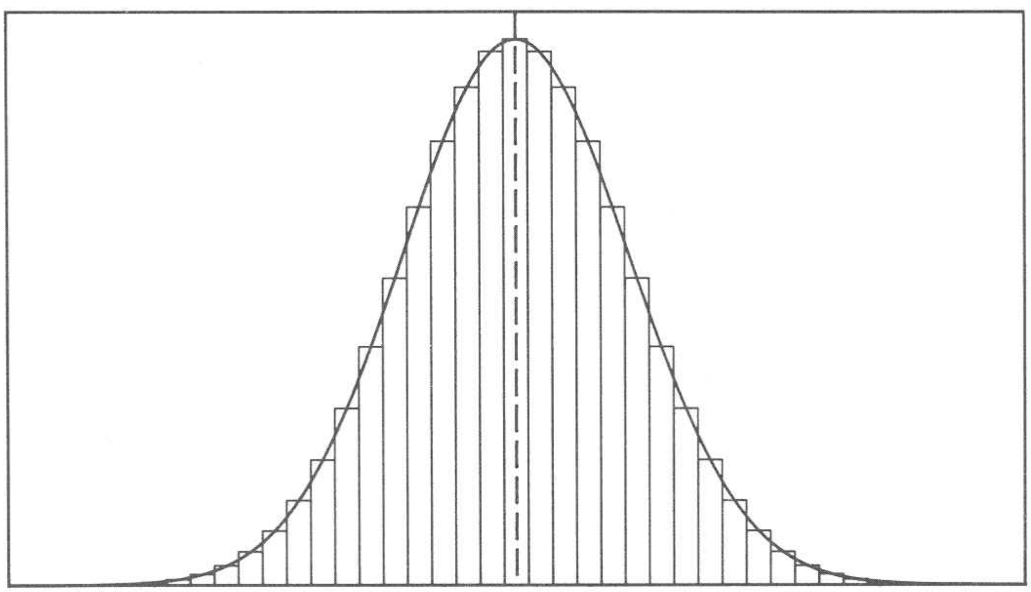

Figure 2a. Gaussian distribution.

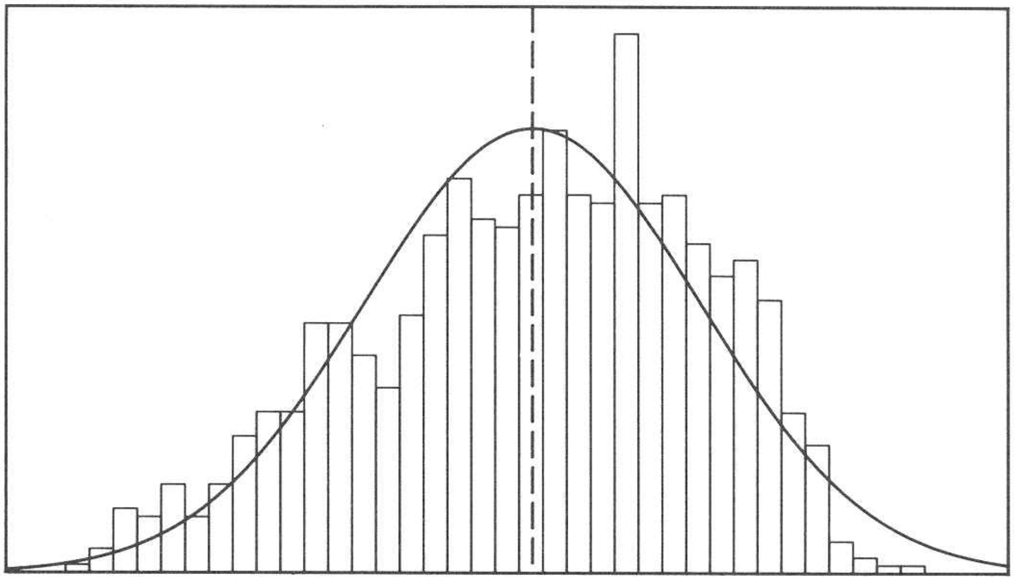

Figure 2b. A skewed distribution.

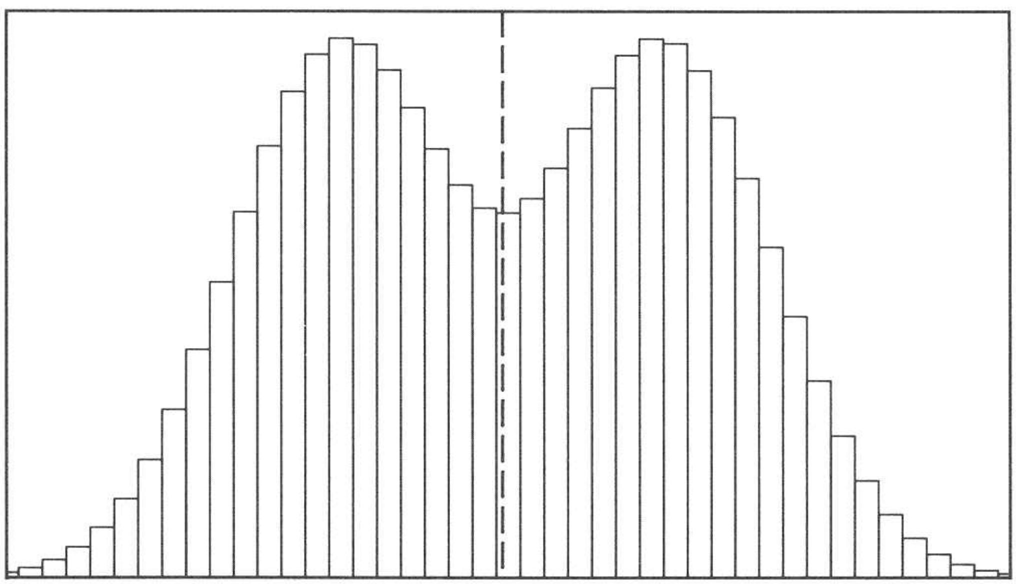

Figure 2c. A bimodal distribution indicates that the jitter has two stable states.

Chi-Square Test

There are numerical tests to determine how closely a particular distribution fits an ideal distribution. The Chi-square test provides a value indicating the "goodness of fit" of histogrammed (grouped) data to a theoretical distribution.

$$\text{Chi-square} = \sum_{i=1}^{N} \frac{(f_i - F_i)^2}{F_i}$$

F_i is the expected count of bin i, assuming an ideal Gaussian distribution with the same mean and standard deviation as that of the data sample, f_i is the actual count of bin i, and N is the total number of bins.

Proper interpretation of Chi-square requires a Chi-square look-up table and the number of degrees of freedom (usually N-3). See the "Statistics Menu" section for constraints on Chi-square implementation.

Using the Probability Distribution For Margin Analysis

There are many ways to analyze a histogram, depending on what information must be extracted. The probability distribution is an analysis of the histogram that can predict the probability of a randomly selected TI falling into, or to the left of, a given TI bin (less than or equal to a given TI). The distribution, denoted as p(x), is calculated by integrating the histogram from -infinity to the given TI bin. The distribution is then normalized such that the integration from -infinity to +infinity produces a probability of 1.

Integrating from the given TI bin to +infinity will predict when a random TI will be greater than or equal to a given TI bin. This produces a 1-p(x) probability distribution, which can be analyzed the same way as the standard p(x) probability distribution.

The probability distribution can be used as a margin analysis tool to predict when timing jitter will cause trouble in a system. In a PCM type of system, data is clocked at a given rate with a specified amount of timing margin to allow for slight timing errors due to jitter on the data line, clock line, or both. If jitter causes the data to stabilize too early or too late with respect to the clock, then a "read error" will occur. By characterizing the jitter and examining the probability distribution, it is possible to verify the timing margin, determine the probability of a read error occurring, and optimize the system performance.

For example, timing jitter on a digital data line relative to a clock is shown in Figure 3a. Jittering TIs are measured from a falling data edge to a falling clock edge. Data must be setup for a finite amount of time before it is strobed by the falling clock line. If the clock occurs before this setup time is satisfied, or the clock occurs after the data is released, then a read error will occur.

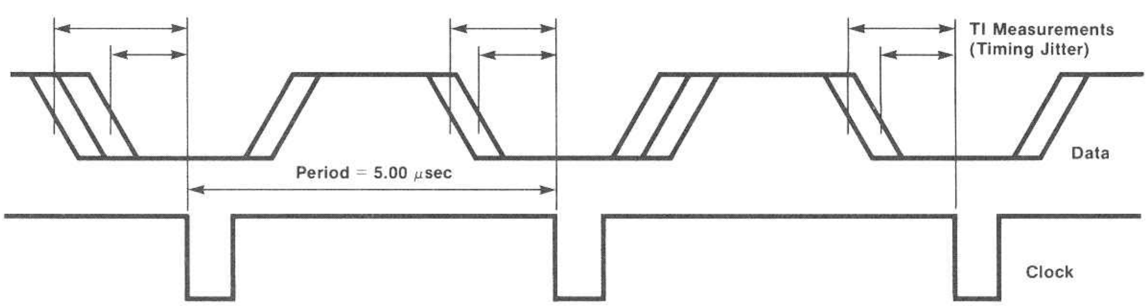

Figure 3a. Example of timing jitter on a data line relative to a clock line.

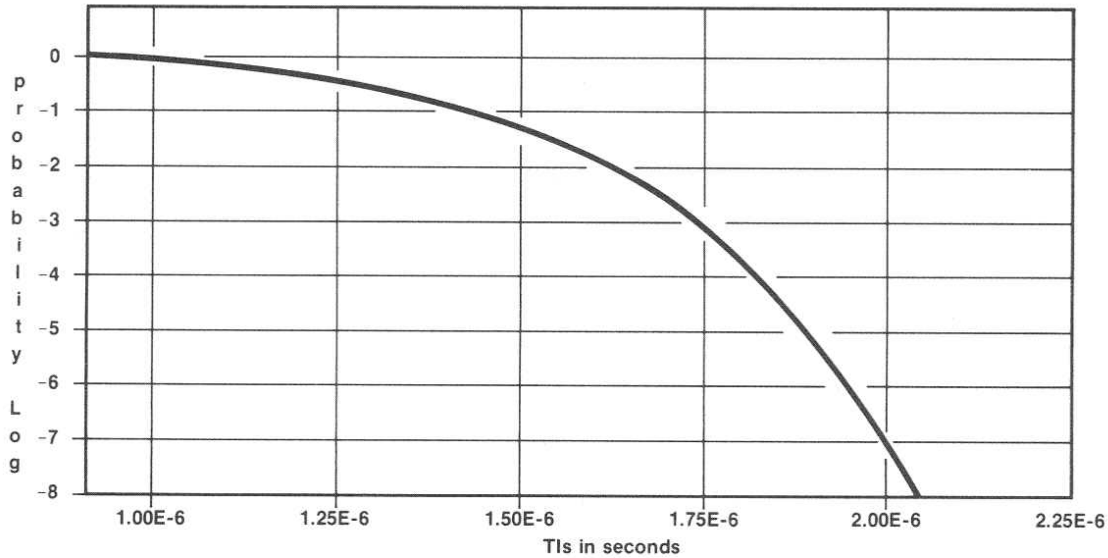

Figure 3b. A 1-p(x) probability distribution of jitter on data/clock lines in Figure 3a.

The probability distribution of the jitter can be used to predict the read error rate of the data/clock system. The plot in Figure 3b indicates that a data-to-clock TI measuring 1.75 μsecs or greater has a probability of occurrence of 10E-3, or about 1 time out of 1000 measurements. A TI measuring 2.00 μsecs or greater has a probability of occurrence of 10E-7, or about 1 time out of ten million measurements.

Timing margin is the amount of timing headroom for a given error rate probability. In Figure 3a, the clock period is 5.00 μsecs; therefore, a 10E-3 error rate leaves a margin of 5.00-1.75 or 3.25 μsecs, and a 10E-7 error rate leaves a margin of 5.00-2.00 or 3.00 μsecs.

Now suppose that the system designer has specified that 2.00 μsecs is plenty of margin for an error rate of 1 part out of 10E+7 data reads. Then there is an extra 1 μsec safety margin in the system. The system could be optimized by reducing the clock period 1 μsec for a 20% increase in performance. This would be the case if the jitter were a function primarily of the hardware, such as gate delays, power supply noise, etc., and not of the clock speed. In any case, the jitter should be recharacterized at the new clock rate to make sure that the 2.00 μsec margin is maintained.

The advantage of calculating the probability distribution is that it greatly reduces the number of measurements required to do margin analysis. Curve fitting to the distribution can change weeks of data acquisition into minutes, because it can statistically, rather than empirically, predict the error rate.

It is important to use extrapolated probability curves with care. If data is not "well behaved," especially out in fringe areas, then extrapolation may be meaningless. In these cases, it is best to take as much data as possible in order to get a reliable probability distribution.

Histogram Software

The HP 5370B High-Speed Histogram Software is a complete data acquisition and analysis package that has been optimized for speed, flexibility, and ease of use.

Speed is obtained by placing most of the statistics and other number-crunching routines in Pascal CSUBs (compiled subprograms) which are callable in BASIC. Also, plotting/display routines are optimized to show only pertinent histogram bins, which can speed up the plotting process by a factor of 15.

Flexibility is demonstrated by the ability to configure and reconfigure any histogram, as well as the ability to configure all setup parameters of the HP 5370B directly from the software. The histogram can be visually updated during the acquisition process to show intermediate results, and can be terminated or aborted at any time. Statistics can be performed on the entire histogram or on selected portions. I/O drivers allow storage and retrieval of configurations, histograms, and unsorted data from disc, and also support graphics printers and plotters.

The software is easy to use, because it is completely menu-driven and user-interactive. Histogram windowing is performed visually via the Knob, Mouse, or Arrow Keys. Two viewports exist to see both the entire histogram and a windowed portion, simultaneously. Statistics are presented in color for emphasis.

Measurement and Analysis Specifications

The software can sort any TI measurement from 0 to ± 330 μsecs, or up to ± 650 μsecs with overranging. (This range limitation is only for Fast Binary output mode on the HP 5370B, which is the only mode in which the software operates. Normal TI measurements can range from 0 to ± 10 seconds). Measurements include pulse width, period, rise time, and fall time.

The data transfer rate between the HP 5370B and the controller is up to 8800 measurements per second, and throughput (continuous transfer and sorting into a histogram) is up to 4200 measurements per second.

The resolution of the HP 5370B is specified at 100 picoseconds rms, typically 30 ps, but the least significant digit is about 20 ps. The software will allow any timing resolution greater than 0.

The histogram array can be as large as 32766 bins, and can contain an essentially unlimited number of measurements (approximately 9E+15 counts per histogram bin).

Analyses include standard deviation, mean, peak-to-peak jitter, Chi-square test for Normality, probability distribution (error rate curve), location and value of maximum bin, and total TI count. All of these functions can be performed on the entire sample of data or on selected portions (windows). This allows the selective analysis of different sets of TIs captured during one measurement session, such as several different pulse widths in a pulse train, or analysis of bimodal jitter.

System Controller

Any HP 9000 Series 200/300 controller with at least 1 megabyte of RAM, graphics capability, and an ALPHA screen width of 80 characters or more can be used. (Essentially all HP 9000 Series 200/300 controllers except the HP 9826. The CSUBs can still be used.)

The best performance will be obtained on systems using a Floating Point Math Card (HP 98635A) or co-processor (Model 320 or better), and a DMA Controller Card (HP 98620B). These systems will see data acquisition rates of up to 8800 measurements per second and throughputs of up to 4200 sorted measurements per second.

The CSUBs will run only on the particular version of BASIC for which they have been compiled. Make sure that the appropriate version is loaded in the controller. (Only BASIC 4.0 or higher versions can be supported.) The following BASIC binaries should be loaded: GRAPH, GRAPHX, IO, TRANS, MAT, PDEV, KBD, ERR, DISC and/or CS80, HPIB, CRTA and/or CRTB, and EDIT (BASIC 5.0 or higher).

Miscellaneous Equipment

Most HP graphics printers, HPGL plotters, and disc drives can be used with the software. The HP 5363B Time Interval Probes and HP 59992A-J06 Time Interval Calibrator can be used, but no drivers currently exist to program them directly from the histogram software. Modifying the software to do so is possible but not documented here.

Setting Up a Time Interval Measurement

Before running the histogram software, it may be wise to first set up the measurement manually in order to get a better idea of how to configure the system. Input impedance, attenuation, coupling, and COM/SEP switches are not programmable, so they must be set manually. Some things to watch out for:

- The attenuation switches are **nominally** divide-by-ten. If a precise attenuation is required, then an external attenuator should be used.
- When the COM/SEP switch is set to COM, set both impedance switches to either 1 megohm or 50 ohms for proper termination. If set to 50 ohms, the signal source will see 50 ohms, but the HP 5370B internal trigger will see 25 ohms; thus, the trigger levels will be halved and should be programmed as such. Also, for consistent triggering, the trigger levels should be set at least 150 millivolts from peak signal amplitudes.

- Making period measurements (rising-to-rising or falling-to-falling edge on one channel) in ±TI arming mode is tricky, because the same signal edge will trigger both Start and Stop channels, resulting in a measurement of the HP 5370B's internal system delay, rather than the signal's period. This type of measurement should be avoided.

  Period measurements greater than 10 ns can be made using +TI ONLY arming. A hint for making more accurate period measurements is to place the COM/SEP switch to SEP, and then specify PERIOD for the Stop channel trigger level in the Configuration Menu. By doing this, the signal will be sent only to the Start channel input amplifier. The Stop channel input amplifier will be disconnected, thus eliminating a potential source of error. Note that the software will perform this period measurement only on the Start channel, not on the Stop channel.

  Another hint for making more accurate period measurements is to measure the delay between Start and Stop trigger circuits by using the ±TI arming mode as described above, and then subtracting that value from all other period measurements made in +TI ONLY arming mode. This is exactly what the SET REF button does to calibrate out system errors, but this register is not read during Fast Binary mode.
- Rise and fall times can be measured, but again the trigger levels must be no closer than 150 mV from either peak for an accurate trigger.
- The minimum pulse width is 5 ns for either single or dual channel measurements.
- The software does not read the current HP 5370B front panel setup. Once the HP 5370B is properly configured manually, make a note of all of the programmable setups, and then duplicate them in the Configuration Menu while running the software. (The setup can be stored on disc and recalled later for similar measurements.)
- If the slew rate of the signal is low, then TIs measured by the software may be slightly different than those on the HP 5370B front panel. This is due to slight differences in the remote trigger level DACs versus the local trigger level pots. Signals with high slew rates will not have this problem.
- Sampling is not continuous; there is approximately 114 μsecs of re-arm time between TI measurements. If an edge occurs during this time, it will not be captured. Also, certain types of pulse trains with dead times of 100 μsecs or greater may be difficult to fully characterize, because the HP 5370B will arm and capture the first pulse after the dead time. Pulses immediately following will be missed because of the counter's re-arm time. In some cases External Holdoff can be used to solve this problem.

Histogram Definition

A histogram is a one-dimensional array in which each element or bin corresponds to a specific and unique TI range. All bins have the same width, or time resolution, and adjacent bins correspond to adjacent TI ranges. There are practically no constraints on the bin resolution. It can be 20 picoseconds or 20 microseconds; both have valid applications, depending on the information to be conveyed by the histogram. There are three parameters that the software uses to explicitly describe any histogram array: Minimum TI, Bin Resolution, and Array Length (see Figure 4).

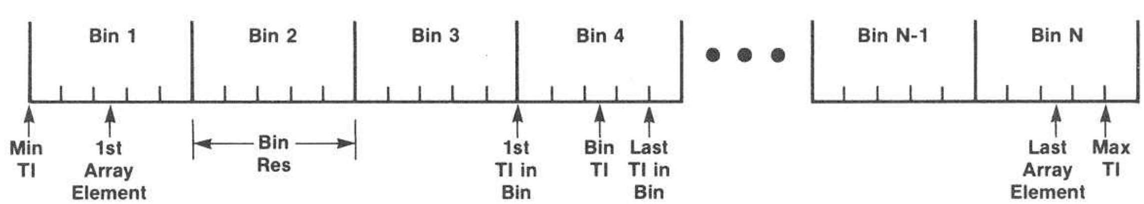

Figure 4. The Blueprints of an N-element histogram.

The left TI boundary of a histogram bin belongs to that bin. The right TI boundary belongs to the next bin. The minimum TI of a bin equals its left boundary; the maximum TI of a bin is less than, but not equal to, its right boundary. Actually, there will be no two different TIs closer than ~20 ps, but because the Minimum TI and Bin Resolution can be expressed as any REAL number, the bin boundaries can be configured to occur essentially anywhere.

The TI associated with each bin is the TI that occurs exactly in the center of that bin. The formula for calculating this average bin TI is:

Bin_ti=(Bin_number-0.5)*Bin_resolution+Min_ti

Conversely, the formula for calculating which bin a particular TI falls into is:

Bin_number=(TI-Min_ti)/Bin_resolution+0.5

Where Bin_number is rounded to the nearest integer.

The ability to select any desired bin resolution and minimum TI allows a great deal of flexibility in configuring the histogram. Wide bin resolutions produce histograms that take up less memory and less processing time, but of course, these histograms lack detail. Fine bin resolutions produce the detail, but also take more memory and analysis time.

Once data has been acquired AND sorted into a histogram, it cannot be re-sorted into a different configuration. Therefore, it is best to know what you want to learn from the histogram before making an important measurement. (It is possible to re-sort raw data, if it is stored on disc during the data acquisition phase. For more information, see the section on "Mass Storage Operations" in the Main Menu.) Some obvious questions to ask regarding a measurement setup:

- What TI resolution do I want to see?
- How many bins are adequate for my distribution?
- What are the minimum and maximum TIs I am interested in?
- How many measurements do I need to make in order to have statistically valid data?
- How much time will I allow for the measurement?
- How often do I want an update on the histogram's progress?
- How much memory does my controller have?
- Will I want to re-sort the data later?
- What is the accuracy of the data?

Obvious questions, though, often do not have obvious answers, and it is likely that a particular configuration will require some fine tuning before it produces desirable results.

Here is a configuration hint that will produce the most accurate histogram representation of TI data: Specify the bin resolution to be a multiple of 19.53125E-12 seconds (5 ns/256). This is the exact least significant digit resolution of the HP 5370B. Even though the actual resolution of the HP 5370B is 100 picoseconds, the math works out better if this detailed resolution is specified.

If a rounded bin resolution were used instead, then there might be an uneven distribution of TIs in the histogram. For example, suppose a histogram were to be configured to measure TIs from 0 up to (but not including) 5 ns in 19.53125 ps bins. Then there would be a total of 256 bins, exactly 1 bin for every possible TI, which would ensure a uniform distribution (see Figure 5a).

Now suppose instead that the bin resolution were to be specified as 20 ps. There would be only 250 bins in which 256 different possible TIs must fall. In this case, 6 bins, spaced approximately 43 bins apart, would contain twice as many TIs as the rest, and would show up as spikes in the otherwise uniform histogram (see Figure 5b).

Another suggestion, not as critical as bin resolution, is to specify the minimum TI to be a multiple of 19.53125 ps, minus 9.765625 ps (0.5*19.53125E-12). This would ensure that the exact center of the bin is the average of all TIs in the bin. (This is true only if bin resolution is specified as above.)

In most cases, these minute details won't matter. However, in situations where the histogram has unexpected spikes or an irregular distribution, these details should be taken into consideration.

Histogram Examples

The following figures show various ways of configuring and sorting data into a histogram. Each figure is a histogram of the same block of TI data. The differences between figures are due solely to different bin resolutions and plotting modes.

Figure 5a/b shows what can happen to a histogram when it is sorted with slightly different bin resolutions. Figure 5a shows a histogram with the 19.53125 ps bin resolution described earlier. Figure 5b shows a histogram of the same data, but sorted into 20 ps bins. Note that some of the bins in Figure 5b have combined, thus doubling their size. The bins show up as spikes, and dramatically alter the appearance of the histogram.

Also note, however, that the mean and standard deviation calculations are about the same for both histograms. These values are based on the entire data distribution, which, numerically, is not significantly different between the two configurations.

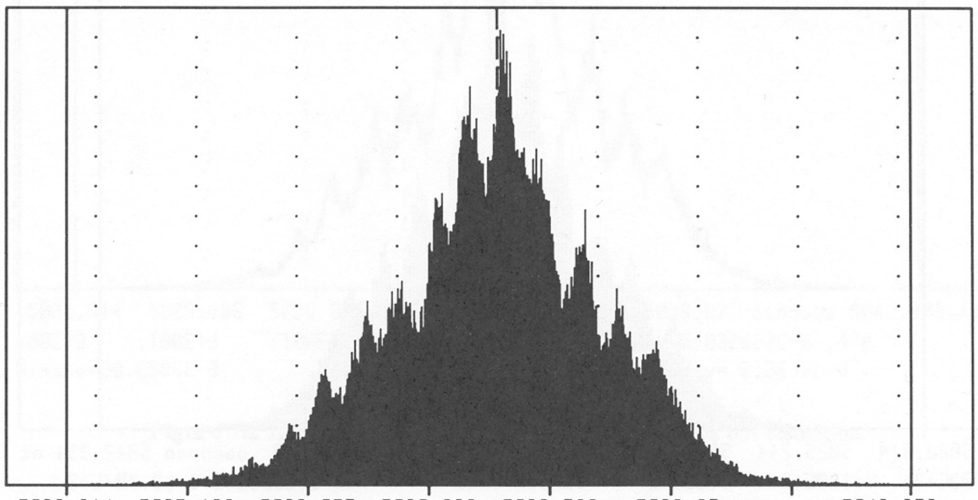

Figure 5a. Histogram in LINE mode with 19.53125 ps bin resolution.

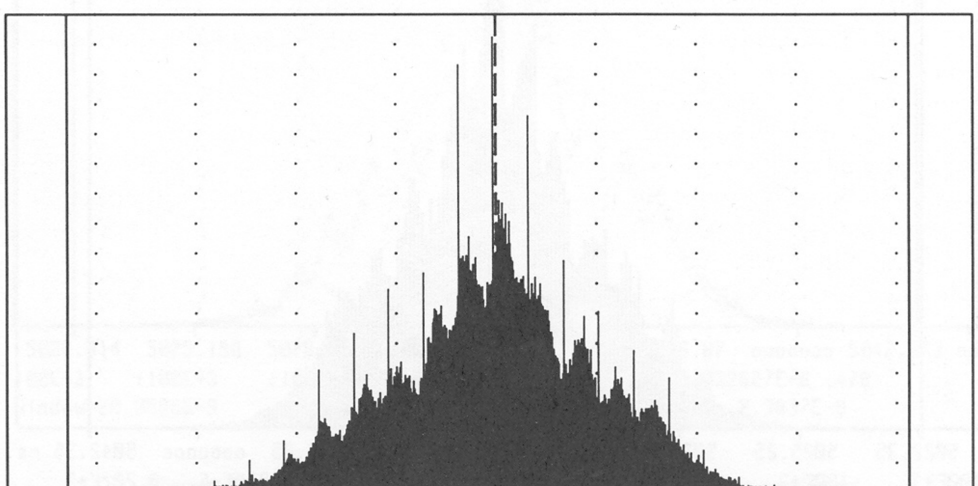

Figure 5b. Histogram in LINE mode with 20 ps bin resolution.

The data from Figure 5a/b is now re-sorted into two new histograms shown in Figure 6a/b. Figure 6a is a histogram with 97.65625 ps bin resolution (5*19.53125 ps). Figure 6b has 100 ps resolution. Unlike the histograms in Figure 5a/b, these two have very similar graphical distributions. Each histogram bin in Figure 6a/b represents the sum of five adjacent bins of the corresponding histogram in Figure 5a/b. Some of the bins in Figure 6b have a spread of six instead of five TIs, which will produce a spike that is only 20% larger than average, rather than 100% larger, which is the case in Figure 5b. Also, as before, the mean and standard deviation calculations are all very similar.

In this case most of the detail in Figure 5a is still preserved in Figure 6a, so the latter is probably a better configuration because it takes up only one fifth as much memory. Again, it all depends on the information required from the histogram.

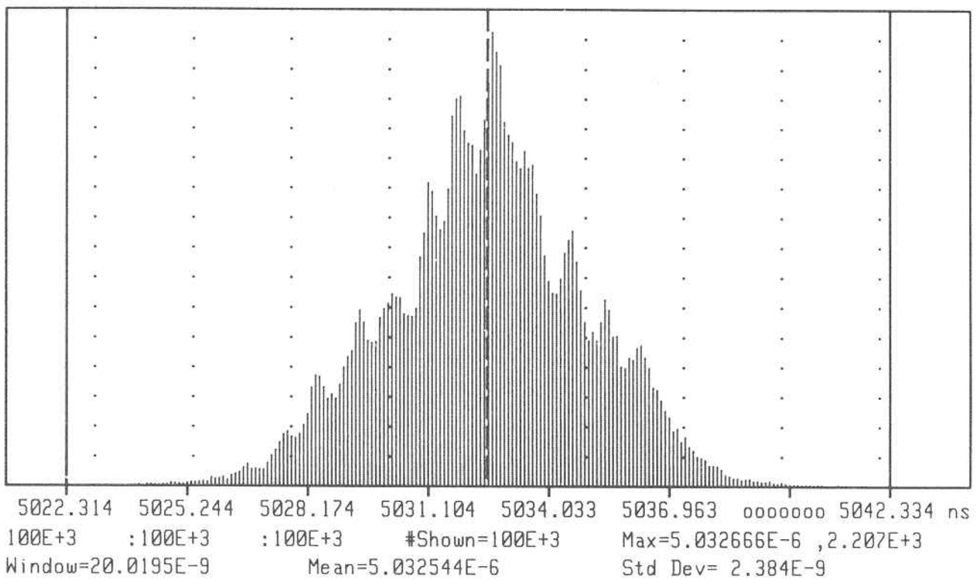

Figure 6a. Histogram in LINE mode with 97.65625 ps bin resolution.

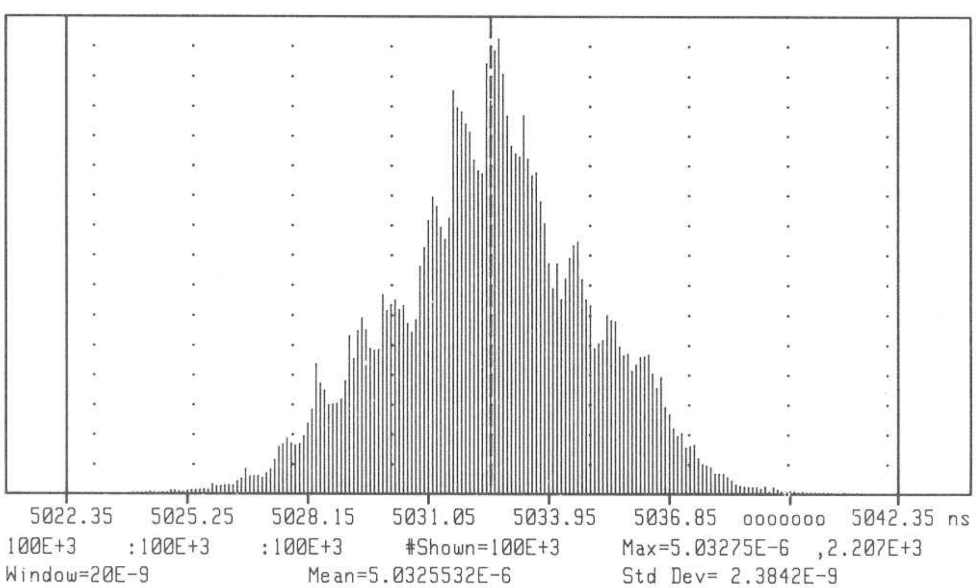

Figure 6b. Histogram in LINE mode with 100 ps bin resolution.

Sometimes visual information can be extracted simply by displaying the histogram differently. Figures 7a and 7b show histograms with the same bin resolution as that of Figure 5a, but in different plotting modes. Figure 7a is in SHAPE mode, which emphasizes the contours of the histogram (and also plots the fastest). Figure 7b is in POINT mode, which helps bring out small bins that are adjacent to large bins, and are thus, frequently overwritten in LINE mode.

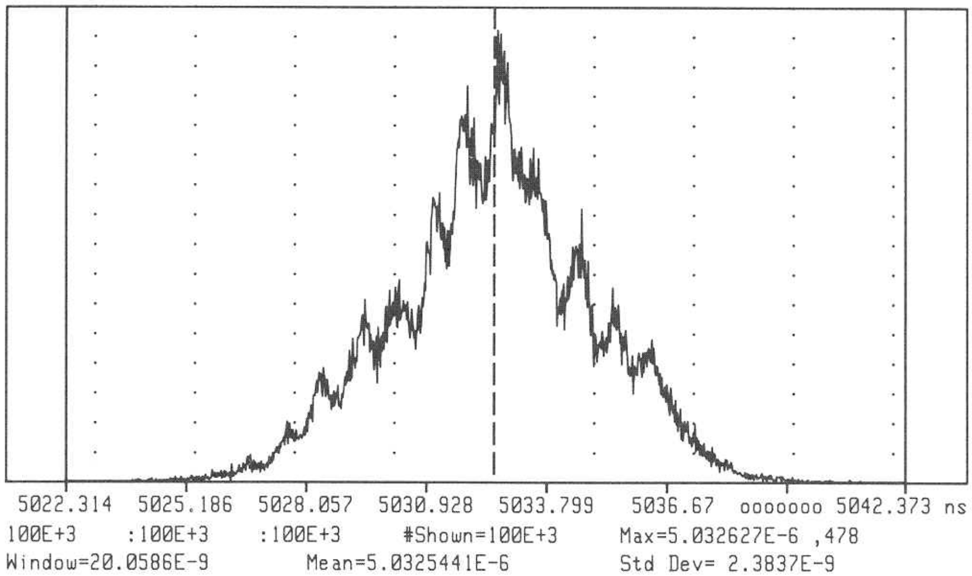

Figure 7a. Histogram in SHAPE mode with 19.53125 ps bin resolution.

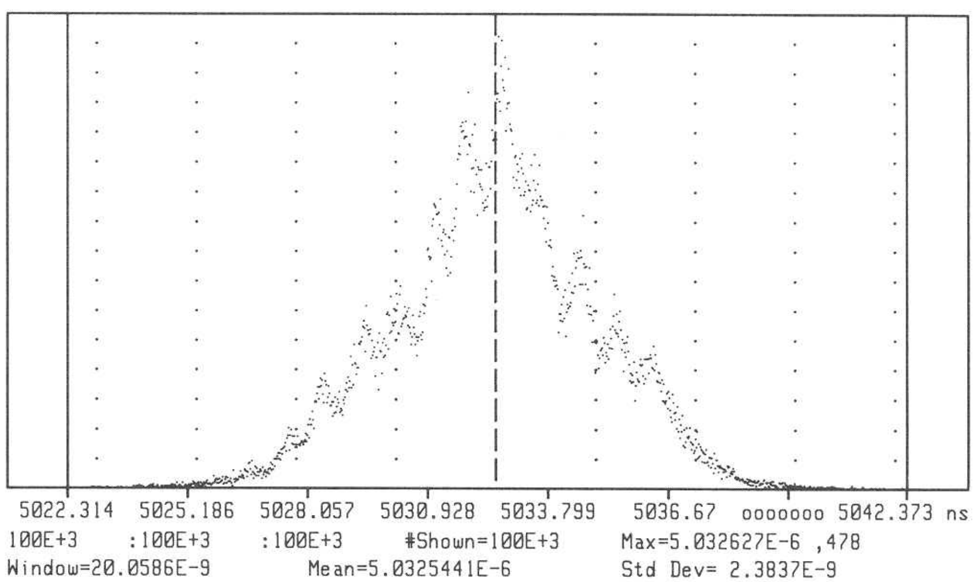

Figure 7b. Histogram in POINT mode with 19.53125 ps bin resolution.

Figure 8a shows the histogram sorted into 400 ps bins. Much of the visual detail is lost as compared to the 100 ps histogram in Figure 6b, but the statistics calculations are still fairly accurate. RECTANGLE plotting mode, shown in Figure 8b, is a good method of checking for empty bins. It is not always obvious in LINE mode whether or not there are empty bins in the centroid of the histogram, because often the line width is less than the bin width on the display (Figure 8a).

Figure 8c shows a histogram in SHAPE mode, which can still show some of the "bumps," even at the coarse 400 ps resolution.

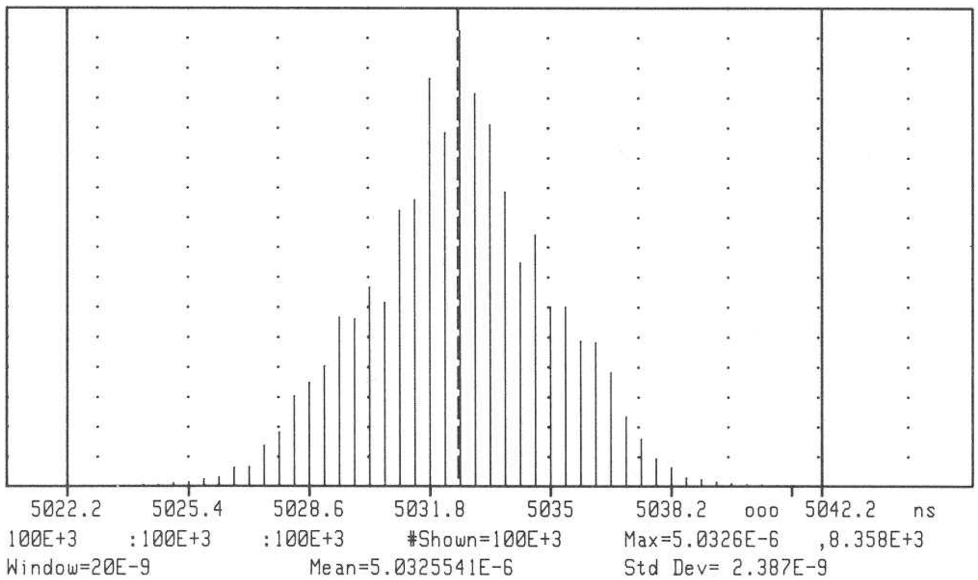

Figure 8a. Histogram in LINE mode with 400 ps bin resolution.

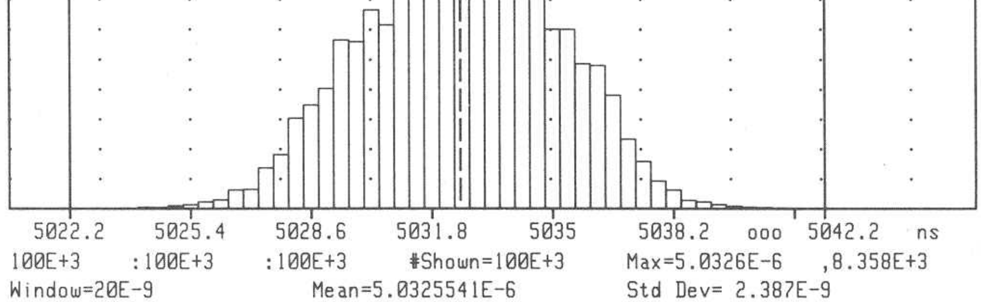

Figure 8b. Histogram in RECTANGLE mode with 400 ps bin resolution.

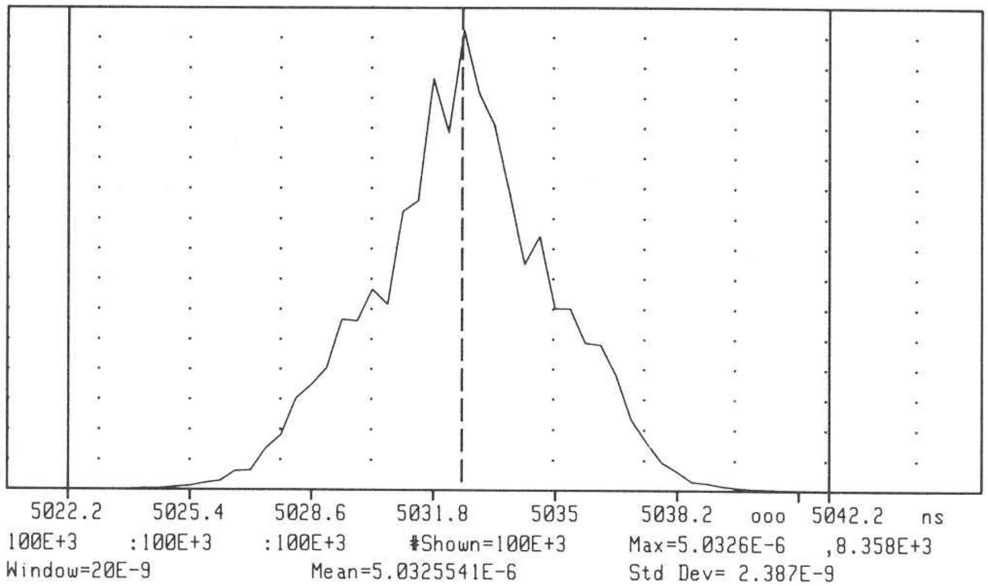

Figure 8c. Histogram in SHAPE mode with 400 ps bin resolution.

Running the Software

To run the software, first load the BASIC binaries shown in "System Controller" section into the controller. Then load and run the histogram software. The HP-IB address of the HP 5370B is expected to be at 703. This address is located in the subprogram Setup_5370, and can be changed there, if necessary.

A Quick Demonstration

The default setup of the software is ideal for capturing a 10 MHz, 50% duty cycle signal, such as the timebase out of the HP 5370B. To set up the system, connect one end of a standard BNC cable to the 10 MHz output at the back of the HP 5370B. Connect the other end to the Start channel input at the front. Next, set the input conditioning switches of both the Start and Stop channels to 50 ohms, divide-by-1 attenuation, DC coupling, and START COM. Finally, connect an HP-IB cable from the controller to the HP 5370B.

Now run the software and press the softkey labeled "Start Meas" (softkey number K3 or f3). The display will start accumulating data, and will update when each block of 1000 measurements has been sorted into the histogram. After 20 updates, all twenty thousand measurements will be shown, and should look similar to Figure 9.

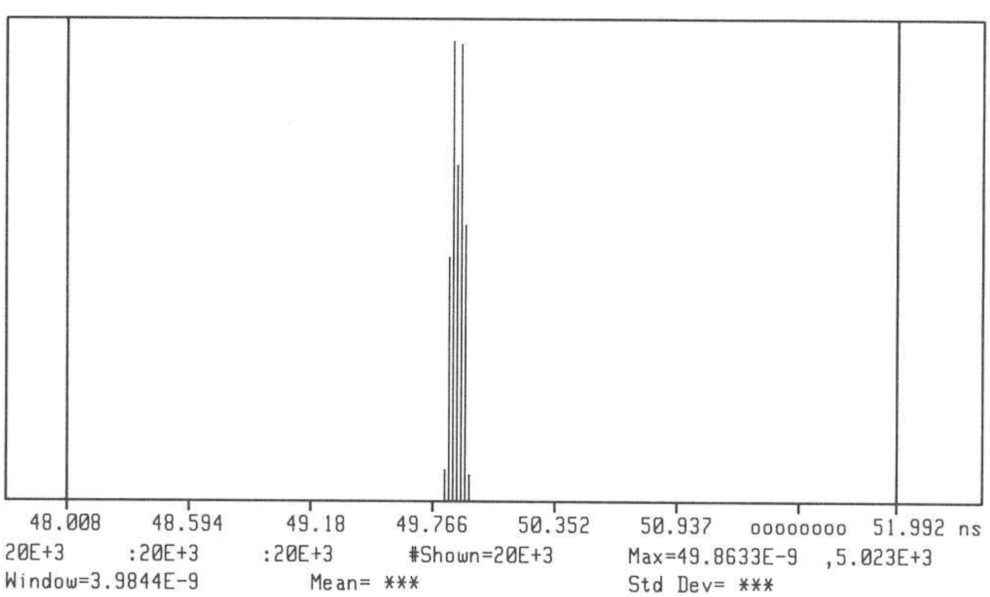

Figure 9. Histogram of the HP 5370B's 10 MHz timebase output.

Note: If the histogram system configuration has been altered, the default setup can be called up by pausing the program, and then executing a "SCRATCH C" command from the keyboard. This will clear any previous setup, and will cause the default setup to be used when the program is rerun.

Often, a histogram may appear to have uneven bin spacing and/or empty bins throughout. This is usually due to the finite graphics resolution of the display, and does not necessarily indicate that some bins are not being filled properly. When in doubt, window in on the section in question to get a better look.

If it turns out that the bins are filling unevenly, check to see that the bin resolution is a multiple of 19.53125 ps, as described earlier in "The Histogram" section.

Histogram Display

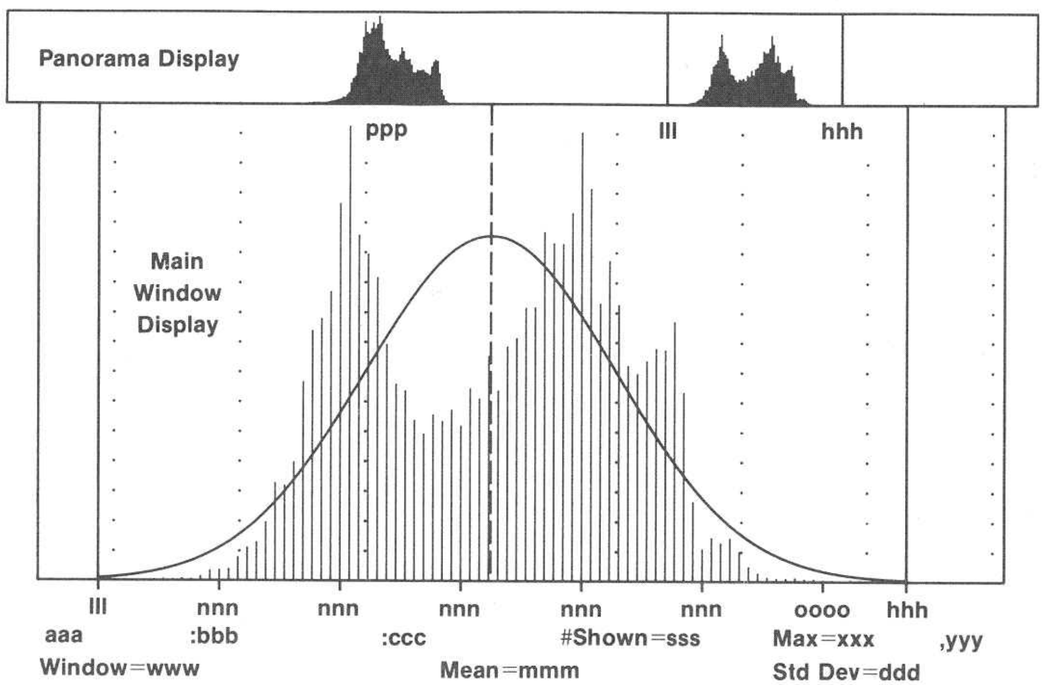

aaa - The total number of measurements in the entire histogram.
bbb - The total number of measurements made (always should be ≥ aaa).
ccc - The total measurement goal when acquisition is to be terminated.
sss - Total number of measurements shown in the window.
xxx - TI bin in seconds that contains the maximum count in the window.
yyy - The count of bin xxx.
www - The width of the window in seconds (shows peak-to-peak jitter).
mmm - Mean of the windowed histogram in seconds.
ddd - Standard deviation in seconds.
(If the window or data is changed, the values for standard deviation and mean are nullified until recalculated. When this happens, three asterisks (***) are put in place of the values.)
ppp - Data in histogram outside of Low/High Window Indices can still be seen in Panorama display.

Figure 10. Detail of the histogram display.

The histogram display contains configuration, status, acquisition, and statistical information about the TI data. Figure 10 shows a breakdown of the display components.

The X-axis on the display is labeled in nanoseconds (**nnn**). The first and last labels correspond to the Low (**lll**) and High (**hhh**) Window Indices. Usually the tick mark to the immediate left of the High Index is labeled with "oooo" which means the window was not evenly divisible by the number of tick marks shown. The distance between this mark and all of the marks to the left of it is constant, so its value can be easily deduced.

The dashed vertical line is the mean of the histogram in the window. The dotted vertical lines show standard deviations from the mean. These are also in different colors for emphasis.

During data acquisition, **aaa** and **bbb** should increase by the block size after each update. If **aaa** is less than **bbb**, or does not change at all, then some or all of the TIs are falling outside the Min and Max specified. However, if <Save Raw> is ON, then none of the data will be lost; all of it can be recalled from the disc and re-sorted into a different histogram later.

Menu/Softkey Descriptions

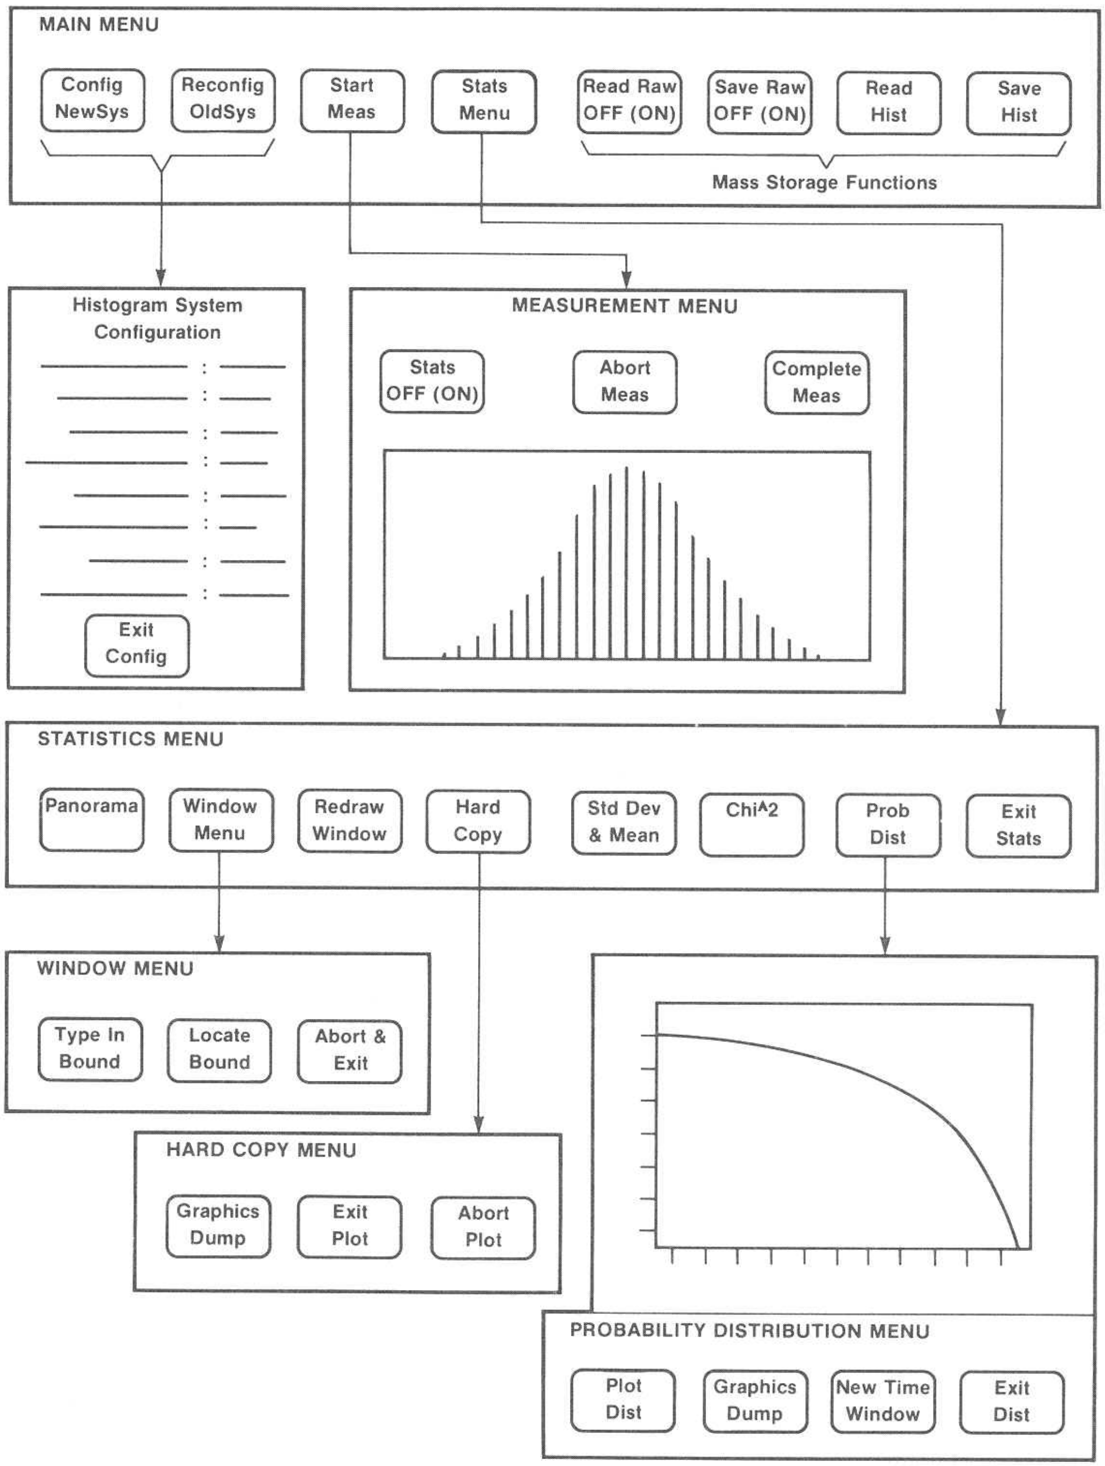

Figure 11. Histogram Software Menu Flowchart

The Main Window Display is always vertically normalized to the Maximum bin count in the window (yyy). The Panorama Display is vertically normalized to the Maximum bin count in the entire histogram.

Figure 11 shows a general flowchart of the menus in the Histogram Software. Following is a detailed description of menus and their specific functions:

Main Menu:

The program starts here. From this menu the system can be configured, a measurement initiated, Statistics Menu accessed, and all mass storage operations executed.

Softkey Description:

Config NewSys — Erases all previous histogram data and enters the Configuration Menu. The current configuration is not erased.

Reconfig OldSys — Enters the Configuration Menu directly, without destroying the current histogram data or configuration. It allows changes to be made in certain viewing and setup parameters for the next measurement which can be added to the current measurement. However, altering the Min/Max TI or Bin Resolution parameters will corrupt the current data.

Start Meas — Enters the Measurement Menu and initiates the data acquisition process. Data will be read from either the HP 5370B or a data file, whichever is specified. The histogram will be displayed and updated after each block of data is read, and when the entire measurement is done.

If nothing appears to be happening, or the screen updates with no data showing, then the system may not be properly configured. When this happens, place the HP 5370B in local mode, and make sure that it makes the desired TI measurement. Check to see that the Start and Stop trigger levels and slopes are the same as those in the Configuration Menu, and that the Min and Max TI bracket the desired measurement. Also make sure that the Bin Resolution is small enough to show adequate detail.

Stats Menu — Enters the Statistics Menu where the histogram window can be modified and analyzed.

Mass Storage Operations

The following keys allow storage and retrieval of various forms of histogram data from data files. If no mass storage unit specifier is included with the data file name, then the default MSUS will be used.

Read Raw OFF (ON) — Specifies the HP 5370B (OFF) or a data file (ON) as the source for data acquisition. If the source is a file, then data is read from the file and sorted into the current histogram configuration just as if it were sent from the HP 5370B. This allows a block of important data to be sorted several different ways without losing any of it. Keep in mind that when using a file as the data source, only histogram configuration parameters will affect the outcome, not HP 5370B parameters. For example, changing the Stop Channel slope will obviously have no effect on data being read from a file, but certainly would affect data being read from the HP 5370B. However, changing the Bin Resolution would affect the resulting histogram, no matter what the source.

When a file is specified as the source, it will remain as such until the softkey is pressed again or the program is rerun. If the END OF FILE is reached before the entire measurement is finished, then the current block is aborted, the histogram is updated, the file pointer is reset, and the program returns to the Main Menu. If <Start Meas> is pressed at this point, then the program starts reading the file again from the beginning.

If the entire measurement finishes before the END OF FILE is reached, then the file remains open, and the next measurement will pick up where the last one left off.

Save Raw OFF (ON) Specifies that the raw data read from the 5370B is to be sent directly to a data file before it is sorted. The same data is then sorted and displayed as usual. This function is active when ON. Whenever the measurement is terminated, the file is closed, and the softkey is turned OFF. This function creates the raw data files that can be read by turning <Read Raw> ON.

NOTE: Data cannot be saved while the acquisition mode is CONTINUOUS (see the Configuration Menu), because there is no defined file size to set up on the disc.

Read Hist Reads a completed histogram and setup file including any statistical data. The histogram and all of the pertinent setup information associated with it is recalled exactly as it was stored on disc, and can be manipulated like any other histogram. All previous data and setup information is lost. The action is immediate.

Save Hist Stores a completed histogram and setup in a data file to be recalled later by the softkey above. Generally, this file takes up less disc space than a raw data file; however, this data cannot be re-sorted. Use this key to archive histograms rather than plot or print them, and/or store histogram configurations and HP 5370B setups.

NOTE: Only one file can be open at any time, since only one file I/O path exists. Thus, selecting any Mass Storage Operation will immediately cancel any other operation that has an open file, and also turn OFF its softkey.

Configuration Menu:

This menu sets up the HP 5370B and histogram parameters. The default setup is shown in Figure 12. With the exception of the <Exit Config> softkey, this menu is not softkey driven. Instead, use the Arrow Keys, Knob, or Mouse to index menu items, and then select them using <ENTER>, <EXECUTE>, <SELECT>, or <RETURN>. Some items require a keyboard entry to change them. Others will rotate forward through a list as they are being selected. Hold the <CTRL> key down to rotate backwards through the list.

WARNING: Changing the configuration and then exiting this menu by any means other than the <Exit Config> softkey may leave the system in an invalid configuration, which may cause errors or erroneous measurements later.

*** * * HISTOGRAM SYSTEM CONFIGURATION * * ***

Minimum Time Interval (seconds)	:	47.99804687E-9
Maximum Time Interval (seconds)	:	52E-9
Histogram Bin Resolution (seconds)	:	19.53125E-12
Low Index of Histogram Window (seconds)	:	48.0078125E-9
High Index of Histogram Window (seconds)	:	51.99218749E-9
Measurement Block Length	:	1E+3
Total # of Blocks (0 for continuous)	:	20
Start Channel Trigger Level (volts)	:	0
Stop Trigger Level (volts, or P for PERIOD)	:	0
External/Start/Stop Channel Slope (+,-)	:	+/+/-
Arming (+TI,±TI)/(START,STOP,INT,EXT)	:	±TI/STOP
Histogram Growth/Auto Window Fit (OFF,ON)	:	ON/OFF
Gauss (OFF,ON)/Plot (LINE,SHAPE,RECT,PNT)	:	OFF/LINE
Histogram Plot Title	<*>	HP 5370B High-Speed

Exit Config

Figure 12. Histogram System Configuration Menu showing the default setup.

Menu item descriptions:

Min/Max Time Intervals: The smallest and largest TIs in seconds to be sorted in the histogram. Any TIs outside these limits will not be sorted. However, the number of counts outside the limits is retained. (See parameter descriptions of Info array in the Appendix.)

As mentioned earlier, there is no exact maximum TI. This number will indicate what the last histogram bin will be. Any TIs that fall into this bin will be included, even if they are greater than the Max TI specified. However, they will never be greater than Max TI plus Bin Resolution. The program will adjust the Max TI entry to be the right-hand boundary of the last bin minus 1 picosecond. The program never alters Min TI.

There are no constraints on the Min/Max TIs except Min TI $<$ Max TI. However, the HP 5370B can only measure up to ± 330 microseconds in its Fast Binary Dump mode. Specifying a TI outside of this range will generate a warning, because the counter will be overranged. An overranged measurement will show up between ~ 330 and ~ 650 μsecs, even if it is larger than 650 μsecs. (The counter starts over every 327.68 μsecs.)

Bin Resolution: The width of one histogram bin in seconds. Any TI within this bin width will increment the bin count by one. The TI for every bin is assigned to the center of the bin. Each element in the histogram array corresponds to one bin.

Low/High Index: The indices of the windowed portion of the histogram. This is the portion that is displayed, plotted, and statistically analyzed. Data can be placed outside of the window during acquisition without being lost, as long as it is still between the Min and Max TI. See the Auto Window Fit note for constraints.

Block Length: The number of measurements to be made between display updates. After each block of data has been acquired, the histogram will be updated to show the status of the entire measurement up to that point. If the block length is greater than the Raw_data buffer size, then data acquisition will actually occur in two or more sub-blocks, each the same length, until the total number of measurements equals the block length. For this reason, the program may alter the block length by a few measurements in order for it to contain an integral number of sub-blocks. Most of this will be transparent to the user.

Note that when a sub-block is transferred, it is sorted immediately in order to clear the buffer, and the HP 5370B is not acquiring data at that time. Usually this does not matter, and the extra time involved to sort several sub-blocks versus 1 whole block is very slight. However, the requirement for sub-blocks means that the data in the whole block was not captured all at once. See Appendix for more information.

Total # of Blocks: The number of display updates before terminating the measurement. (0 goes forever.)

The total number of measurements made equals the block length times the total number of blocks. The fastest acquisition will occur when the total number of blocks is 1, and the total number of measurements equals the block length; however, no intermediate results will be shown.

Start/Stop Trigger Level: HP 5370B trigger levels for the Start and Stop channels (-2.00 to +2.00 volts in 0.01 volt steps). Specifying a "P" for the Stop channel will cause a period measurement to be made on the Start channel at the Start channel's trigger level and slope. This measurement will only go through the Start channel input amplifier; thus, it will have slightly better accuracy than if making a period TI measurement between the Start and Stop channels. (See the section on "Setting Up a Time Interval Measurement" for more information on increasing period accuracy.)

When in PERIOD mode, the Stop trigger slope is changed to match the Start slope, if they are different. Even though the Stop slope is never used in this case, this should remind the user that PERIOD mode is in effect.

If PERIOD mode is used with ±TI arming, the resulting measurement will most likely be the internal system delays of the HP 5370B, not the period of the signal. This can be used to calibrate out the system delays, but may cause some confusion if the user is not aware of what is going on.

Ext/Start/Stop Slopes: The rising (+) or falling (–) edge of the External, Start, and Stop channels that will cause an ARM and/or TRIGGER condition to occur. When the Stop channel trigger level is in PERIOD mode, then the Stop slope will always match the Start slope.

Arming: (+TI,±TI) (START,STOP, INT,EXT) Specifies what channel will arm the HP 5370B to make a TI measurement. There are seven possible arming modes: +TI ONLY/(INTERNAL, EXTERNAL, EXTERNAL HOLD OFF) and ±TI/(START, STOP, INTERNAL, EXTERNAL). Consult the HP 5370B manual for explanations of these arming modes.

Growth/Auto Window Fit: GROWTH ON causes the program to include newly acquired data with the current histogram data. The current data could be a previous block or a completely different measurement. In this mode vast amounts of data can be accumulated and analyzed all at once. The practical limit for the size of any histogram bin is 9E+15, because the resolution of the 8-byte floating point REAL format will be reached. Incrementing this number by one will not change it.

GROWTH OFF causes any previous data to be discarded before each measurement block.

Auto Window Fit ON causes the Low and High Window Indices to track the first and last non-zero bins during data acquisition. After each update, the left and right window boundaries will reside on the first and last non-zero bins in the entire histogram. The result is that the window may start small and then gradually get larger as the distribution of TIs gets wider. However, when any statistics are implemented in the Statistics Menu, this function reduces the window to the first and last non-zero bins in the **current** window, and never increases it. This makes it easy to find peak-to-peak timing jitter on a windowed histogram section. (The Window Width parameter in the histogram display will indicate peak-to-peak jitter.)

When Auto Window Fit is OFF, the Low and High Window Indices remain unchanged during acquisition and statistics calculations.

NOTE: When Auto Window Fit is OFF, the Low/High Index entries must be consistent in order to exit the Configuration Menu, e.g., Low Index ≥ Min TI, High Index ≤ Max TI, and Low Index ≤ High Index. With Auto Window Fit ON, it does not matter what they are set to; the program will take care of them.

Gauss/Plot (ON,OFF), (LINE, SHAPE, RECTANGLE, POINT) When standard deviation and mean are available, GAUSS ON places a theoretical Normal (Gaussian) curve on top of the windowed histogram data. When GAUSS is OFF, or standard deviation is 0, then no curve is plotted.

PLOT displays the histogram as vertical LINEs, POINTs, RECTANGLEs, or SHAPE (outline). SHAPE plots the fastest, because it connects only the maximum non-zero histogram bins. However, it is the least accurate display mode, does not show holes in the histogram, and may change standard deviation colors at locations other than actual standard deviation boundaries.

Plot Title: Any string up to 128 characters to title the histogram. (80 characters for medium resolution CRTs.)

Exit Config Returns to the Main Menu if the configuration is valid. If not, a message is displayed indicating what should be fixed.

Measurement Menu:

Once in this menu, the measurement has already begun. However, from here you can turn statistical calculations ON or OFF during data acquisition, and prematurely terminate the measurement.

In this menu the softkeys are read in the background of the Controller operations, which may delay their response. If the Controller beeps immediately when a softkey is pressed, then the program is currently updating the histogram display which briefly disables the keyboard. If this happens, press the softkey again.

Softkey Description:

Stats OFF (ON) Turns ON or OFF the statistics mode while acquiring data. When statistics are ON, the standard deviation and mean are calculated on the next histogram update, which may slow down the measurement cycle. For histograms with a small number of bins, the delay will be negligible.

Abort Meas Forces immediate termination of data acquisition, discards the current measurement block if incomplete, updates the histogram, and returns to the Main Menu. An abort message will appear, but depending on the size and condition of the histogram, the process may take several seconds to complete. This is the friendliest way to terminate the measurement if something goes wrong, e.g., the HP-IB cable is disconnected, or the signal source goes away.

Complete Meas Terminates the measurement cycle after the next full data block is read and processed, and returns to the Main Menu. No data acquired up to that point is lost.

Statistics Menu:

This menu allows selected portions of the histogram to be windowed and statistics performed on just those portions. Also, a hard copy of the histogram can be obtained from a graphics printer or HPGL plotter.

Softkey Description:

Panorama Creates a viewport above the main display that shows the entire histogram from Min TI to Max TI simultaneously with the windowed portion. This makes it easy to see the location of the window with respect to the entire histogram. The Panorama exists until the Statistics Menu is exited.

Window Menu Enters the Window Menu that allows easy selection of new window boundaries.

Redraw Window Redraws the current histogram to the new window boundaries. Statistics will be displayed only if they were displayed previously and the window did not change.

Hard Copy Allows the current histogram display to be sent to a graphics printer or HPGL plotter. Also, the histogram display width and bin-to-pixel density can be altered here. Once altered, they remain in effect until changed again. The hard copy will exactly resemble the display. After selecting this key, enter either 0 to exit without plotting, 1 to do a graphics dump to a printer, or the HP-IB address of the plotter to do a plot (e.g., 705).

To dump to a printer, enter 1. At this point the display width and pixel density can be modified. Display width is the horizontal length of the histogram in pixels on the CRT. The graphics dump is a dump of these pixels. Some graphics printers do not have as many horizontal pixels as some CRTs; therefore, the display width will have to be reduced. Otherwise, part of the right side of the histogram will be cut off. Reducing the display width will cause the Statistics Information Line to wrap around, possibly breaking in the middle of a number. To avoid this, the recommended display width for medium resolution CRTs is 512 pixels, and for high resolution CRTs is 640 pixels. The ThinkJet printer will work for both of these settings, but the HP 9876 printer will work only with the 512 setting.

The pixel density factor is a means of speeding up the plotting process by not plotting every histogram bin when there are more bins than pixels on the CRT. The visual results are identical, but the speed can be increased by a factor of 15. A pixel density of 1 is the default setting which draws only 1 histogram bin per pixel. A factor of 2 will draw 1 histogram bin per 2 pixels (twice as fast, but holes will show), and a factor of 0.5 will draw 2 bins per pixel. A factor of 0 will force all bins to be drawn. Probably most cases will leave the factor at 1. However, many plotters have better resolution than the CRT, in which case, decreasing the factor may be desired. If specifying a factor less than 1 and greater than 0, the reciprocal of the number should be an INTEGER for best results, e.g., 0.5, 0.25, 0.2, etc.

Remember that once the display width and pixel density are changed, they will remain that way until changed back.

If a plot is desired instead of a graphics dump, then set the pen limits to the desired size, shape and location of the histogram on the plotting medium. The pen limits, P1 and P2, will describe diagonals of a rectangle. P1 must be the lower-left corner and P2 the upper-right, otherwise an error will occur. The resulting plot will fill this rectangle, whatever the shape.

Up to 7 different pens will be used, depending on what colors are on the CRT as indicated in Table 1:

Plotter Pen	CRT Color
1	White
2	Red
3	Yellow
4	Green
5	Cyan
6	Blue
7	Magenta

Table 1. Correspondence between plotter pen number and CRT color for plotting histograms.

If fewer pens are provided, the plotter will automatically substitute. Once the pens and pen limits have been set, enter the plotter's HP-IB address, and the program will begin plotting.

The plot can be aborted by pressing the <Abort Plot> softkey. A graphics dump to a printer cannot be aborted.

The program will turn off the softkeys and print the message, "Key 1 = Graphics Dump, Key 2 = Exit Plot." Pressing softkey 1 will dump graphics to the default DUMP DEVICE. Softkey 2 will return to the Statistics Menu.

Std Dev & Mean Performs standard deviation and mean calculations on the windowed portion of the histogram. The mean is the arithmetic mean of all TIs in the window, and probably will not lie in the center of a bin. It is displayed graphically as a vertical dashed line (- - -).

The standard deviation is calculated relative to the mean, and is shown graphically as vertical dotted lines (. . .). Also, each standard deviation from the mean in either direction is shown in a different color to add emphasis.

Chi^2

Performs a Chi-square calculation on the windowed portion of the histogram (between the low and high indices). If standard deviation and mean have not been calculated, then they are calculated first, and then Chi-square is calculated.

$$\text{Chi-square} = \sum_{i=Low}^{High} \frac{(f_i - F_i)^2}{F_i}$$

Where F_i is the expected histogram count at bin i, assuming a theoretical Gaussian distribution with the given mean and standard deviation, and f_i is the actual count of bin i. Low and High are the histogram window boundaries.

The degrees of freedom for Chi-square are also displayed. This number is simply the total number of bins between the Low and High Window Indices minus 3. If there are fewer than 3 bins, or the standard deviation is zero, then Chi-square is aborted.

Proper implementation of Chi-square is very important in order to get meaningful results. The best results will be obtained when there are no empty bins, and the number of non-zero bins is at least $10*\log_{10}(N)$, where N is the total number of measurements. In order to meet these criteria, it may be necessary to play with the histogram configuration, and sort the data several different ways. The software does not check to see if these criteria are met.

An accurate interpretation of the Chi-square result requires the use of a Chi-square look up table. Generally, the smaller Chi-square is, the better the data resembles the theoretical distribution.

Prob Dist

Enters the Probability Distribution menu.

Exit Stats

Removes the Panorama, if it exists, and returns to the Main Menu.

Window Menu:

In this menu the Low and High window indices can be changed visually in order to select certain portions of the histogram for analysis and display. Also, the value of each bin in the histogram can be quickly determined.

The Mouse, Knob, and Arrow keys can be used to move the delimit bars around, <ENTER>, <EXECUTE>, <SELECT>, or <RETURN> can be used to select the new boundaries. It does not matter which bound is entered first. After two bounds are entered, the lower one is assigned to the Low Window Index, the higher one is assigned to the High Window Index, and the program immediately returns to the Statistics Menu.

The coordinates for the delimit bar position are shown, indicating the TI in seconds for the bin it currently resides on, as well as, the number of counts in that bin. By moving the delimit bar around, you can easily see the value of each bin. The bar can also be moved off screen to show the values of any bins not in the window.

Often it is helpful to have a Panorama up at the same time, in order to quickly see where the window is with respect to the entire histogram.

The Knob, Mouse, and Left and Right Arrow keys move the delimit bars 1 pixel width with respect to the Panorama, whether or not the Panorama is displayed. Thus, if the window is very small compared to the Panorama, the delimit bars may jump several bins at a time in the window. When this is the case, use the Up and Down Arrow keys to move the bars one bin at a time. Also, <SHIFT-LEFT> or <SHIFT-RIGHT> Arrow Keys will place the bars at the far left or right of the entire histogram, respectively.

Softkey Description:

Type In Bound Allows the entry of a boundary by typing a TI in seconds, rather than using a graphics input device. The boundary entered must be between the Min and Max TI.

Locate Bound Places the delimit bar in the middle of the display, in case it was moved outside of the current window.

Abort & Exit Restores the boundaries to what they were when the Window Menu was entered, and returns to the Statistics Menu.

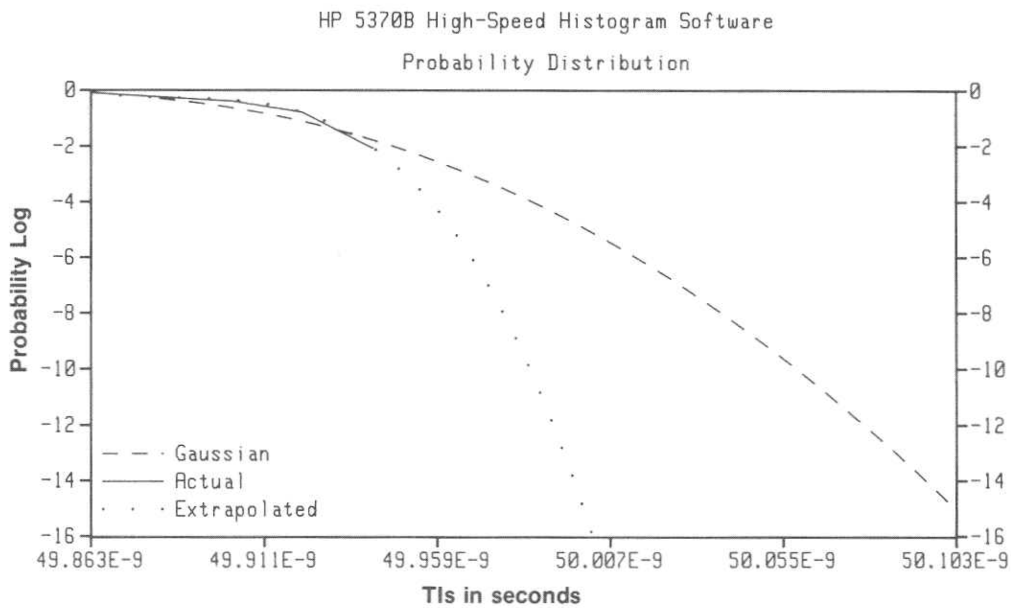

Figure 13. A 1-p(x) probability distribution of the HP 5370B's 10 MHz timebase in Figure 9.

Probability Distribution:

This menu plots the 1-p(x) probability distribution of the histogram data, along with that of a theoretical Gaussian distribution. The data is also extrapolated via a least-squares polynomial curve fit to predict the probability distribution beyond the given histogram window. The probability distribution is the integral of the histogram from the mean to the right side of the time window, and shows the probability of any measurement falling into, or to the right of, a particular histogram bin. The horizontal axis represents TI bins in seconds; the vertical axis represents the log (base 10) of the probability of a TI being greater than, or equal to, a given TI bin. Figure 13 shows the probability distribution of the histogram of the HP 5370B's 10 MHz timebase in Figure 9.

Upon entering this menu, the user must enter three values to configure the probability distribution plot: the logarithm (base 10) of the desired probability, the order (maximum exponent) of the polynomial curve fit, and the percentage of the amount of data to be used for the curve fitting routine.

The probability log must be less than 0 and greater than or equal to -50. This value is approximately the minimum Y-axis value of the probability distribution plot (max is 0), and determines how far the program is to extrapolate the data. The default is -15.

The order of the polynominal curve fit determines the shape of the extrapolated curve. The default order is three, which is a good approximation for extrapolating near-Gaussian distributions.

A word of caution when specifying higher-order curve-fitting: the algorithm used is excellent for fitting a polynomial function between data points. However, beyond the given data, there is nothing to control the direction of the curve, unless some assumptions are made about the distribution. Specifying an order higher than three may result in a better fit where the data exists, but also may produce unexpected results in the extrapolated area.

The percentage of curve fit determines how much of the histogram data is to be included in the curve-fitting algorithm. This percentage pertains only to the bins in the histogram window, and is relative to the left side of the window. For example, 75% would indicate that only the first 3/4 of the histogram window should be used for constructing a curve fit. The reason for limiting the data is so that "fringe bins" at the far right of the histogram do not throw the curve off. Specifying less than 100% will leave these bins out of the calculation, if desired. The default is 95%.

Softkey Description:

[Plot Dist] Sends the probability distribution to a plotter in much the same way as the HARD COPY softkey does in the Statistics Menu.

[Graphics Dump] Sends the probability distribution to a graphics printer.

[New Time Window] Returns to the beginning of this procedure in order to enter a new probability log, order, and/or curve fit percent on the same histogram data.

[Exit Dist] Returns to the Statistics Menu.

Appendix — Data Structures and CSUBs

This section is for those who want to use the CSUBs in other programs, such as a production test system. It describes the variables and array constructs to be used, and how they are read by CSUBs.

The HP 5370B raw TI data is stored in a multi-dimensional INTEGER BUFFER array called Raw_data. The total size of this array determines the total number of TIs that can be acquired in one shot. If the block size is larger than this array, then the block is broken up into smaller sub-blocks that will fit into it. The extra time required to process sub-blocks as opposed to one block is slight, but measurable. If acquisition must occur as fast as possible, or all at once, then this array should be declared large enough to accomodate one block. Because it is accessed only as a BUFFER, it does not matter how many dimensions it has. The software declares it as a 2-dimensional array.

Each element in the Raw_data array is two bytes long; however each raw TI measurement from the HP 5370B is 5 bytes long. The second dimension of Raw_data is declared from 1 to 2500, which corresponds to 1000 TI measurements. The first dimension is declared from 1 to 10, for a maximum of 10,000 TIs in one block. To change the size of the buffer, it is recommended that, for purely aesthetic reasons, just the first dimension be changed. This keeps the buffer a multiple of 1000 TIs long.

For example, to increase the buffer size to 50,000 TIs, change the first dimension to 50, and leave the second dimension at 2500. If memory is scarce, however, the minimum recommended buffer size is 1000 TIs, so the first dimension can be set to 1. The program will still run, but with slightly reduced performance.

The histogram is stored in a 1-dimensional REAL array called Hist_data, where each element of the array corresponds to one histogram bin. It is dimensioned from 1 to 32766, which is the maximum allowable number of bins. (Larger sizes would require a 2-dimensional array declaration, which the CSUBs can handle, but would require substantial changes in the BASIC software.) The declaration can be reduced to save memory space, if the user does not require that many bins.

All of the parametric values of the histogram are stored in a 16-element REAL array called Info. The Info and Hist_data arrays tell the entire histogram story. Table 2 describes what each element of Info does, which CSUB it is used in, and whether it is used as an Input or Output parameter. For the CSUB entry, H = Histogram, M = Ministats, S = Stats, E = End_bins, and C = Chi_squared.

There are two sets of CSUBs included in the software, only one of which is needed, depending on the system. Each set is comprised of six CSUBs: Histogram, Ministats, Stats, End_bins, Chi_squared, and Fill_matrix. Each CSUB also has a suffix in its name of "10" or "20". The suffix "10" means that the CSUB is compiled for systems using the 68000/10 processor, and "20" is compiled for 68020 processors with co-processors. The "10" CSUBs will run with the 68020, but not as fast as the "20" CSUBs. The "20" CSUBs will not run with the 68000/10 processors.

The program automatically selects which set of CSUBs to use, and ignores the other set. It is okay to delete the unused set to save memory.

Histogram CSUB

Histogram(INTEGER Raw_data(*),REAL Hist_data(*),Info(*),
OPTIONAL Ti_data(*),Ti_start)

This CSUB converts the HP 5370B raw TI data into real numbers and sorts them into a histogram at the same time. Usually, the raw data is discarded once it is sorted; however, it can be retained by two different means: by saving raw data on disc during the acquisition cycle, or by passing an optional array called Ti_data to the CSUB. Ti_data is an n-dimensional REAL array that will contain the converted raw TI data in the order in which it was captured. Each entry will be the actual TI in seconds. If the Ti_data array is not large enough to contain an entire block of data, then Error 17 is returned before any sorting is done.

By passing Ti_data to the Histogram CSUB, it is possible to preserve the sequential information as real numbers before it is sorted. However, each measurement takes up 8 bytes of memory, so a sufficient amount of RAM must be present to retain both the histogram and the TI data.

Info Element	CSUB Used	I/O Parm	Description
1	H	I	Beginning of Raw_data buffer pointer. This corresponds to the buffer empty pointer (in BASIC, STATUS @Buffer,5;Empty). WARNING: The Histogram CSUB does not wrap around the array the way BASIC buffer pointers do. To be safe, always reset the Raw_data buffer before filling it, so that the the first byte of the measurement is at Empty pointer location 1. Then this parameter will always be 1.
2	H	I	Length of Raw_data buffer in bytes (in BASIC, STATUS @Buffer, 4;Bytes). Usually this number will be 5 times the block/sub-block length.
3	HS	I	Minimum TI in seconds. Anything less than this will not be sorted in the histogram.
4	HS	I	Histogram Bin Resolution in seconds.
5*	H	O	Number of TIs encountered that were less than the Minimum TI specified in Info (3).
6*	H	O	The number of TIs encountered that were greater than the Maximum TI specified. The Maximum TI is not explicitly specified, but is determined by by the size of Hist_data. REDIMensioning Hist_data will set the Max TI according to the following formula: Max_ti < SIZE (Hist_data,1)*Bin_res+ Min_ti. A TI greater than or equal to this value is considered outside the Max TI boundary.
7*	H	O	Total number of TIs sorted, excluding those outside the Min/Max TI boundaries.
8	MSEC	IO	As an Input parameter, this is the Low Window Index of the Hist_data section to perform statistical analysis on. As an Output parameter, this is the first non-zero bin within the Low/High Indices. (Used as an Output parameter in End_bins only.)
9	MSEC	IO	Same as above, except as an Input, High Index, and as an Output, last non-zero bin.
10	SC	O	Mean calculation of windowed histogram.
11	SC	O	Standard deviation of windowed histogram.
12	MS	O	Index of the bin with the most counts between Low and High Window Indices.
13	MSC	O	Total number of TIs (counts) between indices.
14**	MS	I	Width of a single graphics pixel in bins.
15**	MS	I	Graphics pixel starting reference in bins.
16	C	O	Chi-square.

*These parameters as Inputs should be set to 0 before the CSUB is called, if and only if Histogram Growth is OFF.
**These parameters must be calculated in BASIC based on the current VIEWPORT, WINDOW, and CRT characteristics. They are needed only if an Index array is passed to the CSUB in order to do bin-to-pixel mapping for accelerated graphics. Otherwise, they are not required.

Table 2. Description of Info array subscripts.

Ti_start indicates which array element in Ti_data is to be the first element where the TI block is to be placed. After conversion, Ti_start will point to the last element in the block plus 1. It acts just like a buffer fill pointer, except that it points to array elements, not bytes. Because it updates automatically, several TI blocks can be placed in Ti_data without adjusting Ti_start, as long as Ti_data is large enough. If Ti_start plus the block length overflows Ti_data, then Error 17 is returned before any sorting is done. If Ti_start is not specified, then Ti_data will always be filled starting with its first element.

Ti_data can be declared from 1 to 6 dimensions. When it is multi-dimensional, the CSUB fills it by incrementing the right-most subscript the fastest, and the left-most subscript the slowest. For example, if it were declared as Ti_data(7:9,2:5), it would fill up in the order shown in Figure 14:

(7,2) 1	(7,3) 2	(7,4) 3	(7,5) 4
(8,2) 5	(8,3) 6	(8,4) 7	(8,5) 8
(9,2) 9	(9,3) 10	(9,4) 11	(9,5) 12

Figure 14. A two-dimensional array will fill up the rows first from left to right, and then from top to bottom. Rows are indexed by the left subscript; columns are indexed by the right subscript.

Ti_start is referenced to the BASE of the right-most subscript and always increments by the block length after each call to the CSUB. For example, using the same declaration above and a block length of 5 TIs, Ti_start should be set to 2 to point to Ti_data(7,2), the first element in Ti_data. (2 is the base of the right-most subscript.) After the first block of 5 TIs is sorted, Ti_start will be set to 7, which points to Ti_data(8,3). After the second block, Ti_start will be set to 12 which points to TI_data(9,4). If a third block is attempted, Error 17 will result because there is not enough room for 5 more TIs in Ti_data at the Ti_start position.

Stats CSUB

Stats(Hist_data(*),Info(*),OPTIONAL Index(*))

This CSUB performs all of the statistical analyses and plotting calculations within the histogram window. Analyses include mean, standard deviation, the bin with the largest count (its index and count), and the total number of counts in the window. The window is set by the Low and High Indices in Info(8) and Info(9), respectively. If these indices are inconsistent, i.e., switched or outside the current histogram boundaries, then Error 17 is returned.

For accelerated graphics, an optional Index array can be passed which will contain the indices of the maximum bins for each pixel. The first element in the Index array contains the number of bins that should be plotted, the rest contain the pertinent Hist_data bin indices. Plotting just these indices will produce the same result on the CRT as plotting all of the indices. The difference is that the histogram can contain 32k bins, but the Index array may contain only 400 indices. Plotting 400 lines runs much faster than plotting 32766 lines.

The Index array can be multi-dimensional; however, there is probably no need for more than one dimension. The lower bound is recommended to be 0, since it will contain the total number of indices to be plotted, and not an index of Hist_data. The next element will be the first Hist_data bin to plot, if Index(0) is greater than 0.

The size of Index must be declared large enough to contain (Display_width/Pixel_density)+2 indices. Otherwise, Error 17 will be returned. See the <Hard Copy> softkey description in the Statistics Menu for more information.

Ministats CSUB

Stats(Hist_data(*),Info(*),OPTIONAL Index(*))

This CSUB does everything that the Stats CSUB does except the mean and standard deviation calculations. All other calculations are identical.

End_bins CSUB

End_bins(Hist_data(*),Info(*))

This CSUB increases the Low Window Index, Info(8), and decreases the High Window Index Info(9), until both indices reside on a non-zero Hist_data bin. The result can be interpreted as a timing jitter measurement in which the window is reduced to show just the peak-to-peak jitter. If any index already resides on a non-zero bin, then it is not changed. If there is only one non-zero bin in the window, then both indices are placed on that bin. If there are no bins in the window, then both bins are placed at the lower base declaration minus one. This is the only CSUB that alters the values of Info(8) and Info(9).

Calling this CSUB before calling either Stats or Ministats may increase the speed of these CSUBs, because any leading or trailing zero bins in the window will automatically be skipped. Make sure to restore the Info(8) and Info(9) parameters if the window change was not meant to be permanent.

Chi_squared CSUB

Chi_squared(Hist_data(*),Info(*))

This CSUB performs the Chi-square calculation on the histogram between Low and High Window Indices, Info(8) and Info(9). It does not perform an End_bins calculation, and requires that the standard deviation be calculated first. Otherwise, an error will occur.

Fill_matrix CSUB

This CSUB is part of the Curve_fit subprogram. List Curve_fit for a complete explanation of how to use it.

Using the CSUBs With Multi-Dimensional Arrays

All of the CSUBs have been designed to accept any size array passed to them. The array must be of the proper type, e.g., REAL or INTEGER, but can be any dimension from 1 to 6. BASIC allows an array to have a maximum of 6 dimensions, the maximum length of each dimension being 32767 indices. In order to index more than 32767 elements, the array must have more than 1 dimension. Actually, two dimensions are all that is needed to index more elements than there is memory in the controller; however, up to six can be used.

The histogram CSUB section describes how the Ti_data array can be more than one dimension and how Ti_start can index it. Hist_data and Index arrays work the same way. If Hist_data were declared with more than one dimension, then it could contain more than 32766 bins. The 1-dimensional limitation was set to make the BASIC program a little simpler. Min_ti is always in the first bin of the histogram, which is always the first element of Hist_data, which is always the element whose indices are the lower base of each dimension. The next bin is the element whose right-most index is increased by one. The last bin is the element whose indices are all equal to the upper bound of each of their dimensions.

For a two-dimensional array declared as A(X,Y), here is a general method of indexing the array:

```
    X_base=BASE(A,1)
    X_size=SIZE(A,1)
    Y_base=BASE(A,2)
    Y_size=SIZE(A,2)
Total_entries=X_size*Y_size
```

The Nth element in A(X,Y) will be in the range,

Y_base ≤ N ≤ Y_base+Total_entries-1 .

To determine which X and Y values index the Nth element,

X=((N-Y_base) DIV Y_size)+X_base   Y=N-Y_size*(X-X_base) .

Conversely, to determine which element is indexed by X and Y,

N=Y_size*(X-X_base)+Y .

Note that if X=X_base, (or all indices other than the right-most index equal their lower base declarations,) then N=Y. That is, the Nth element is the same as the Nth index, which is always true for one-dimensional arrays. Info(8), Info(9), and Info(12) are all Hist_data indices that will follow this rule.

Histogram Software Menu Flowchart

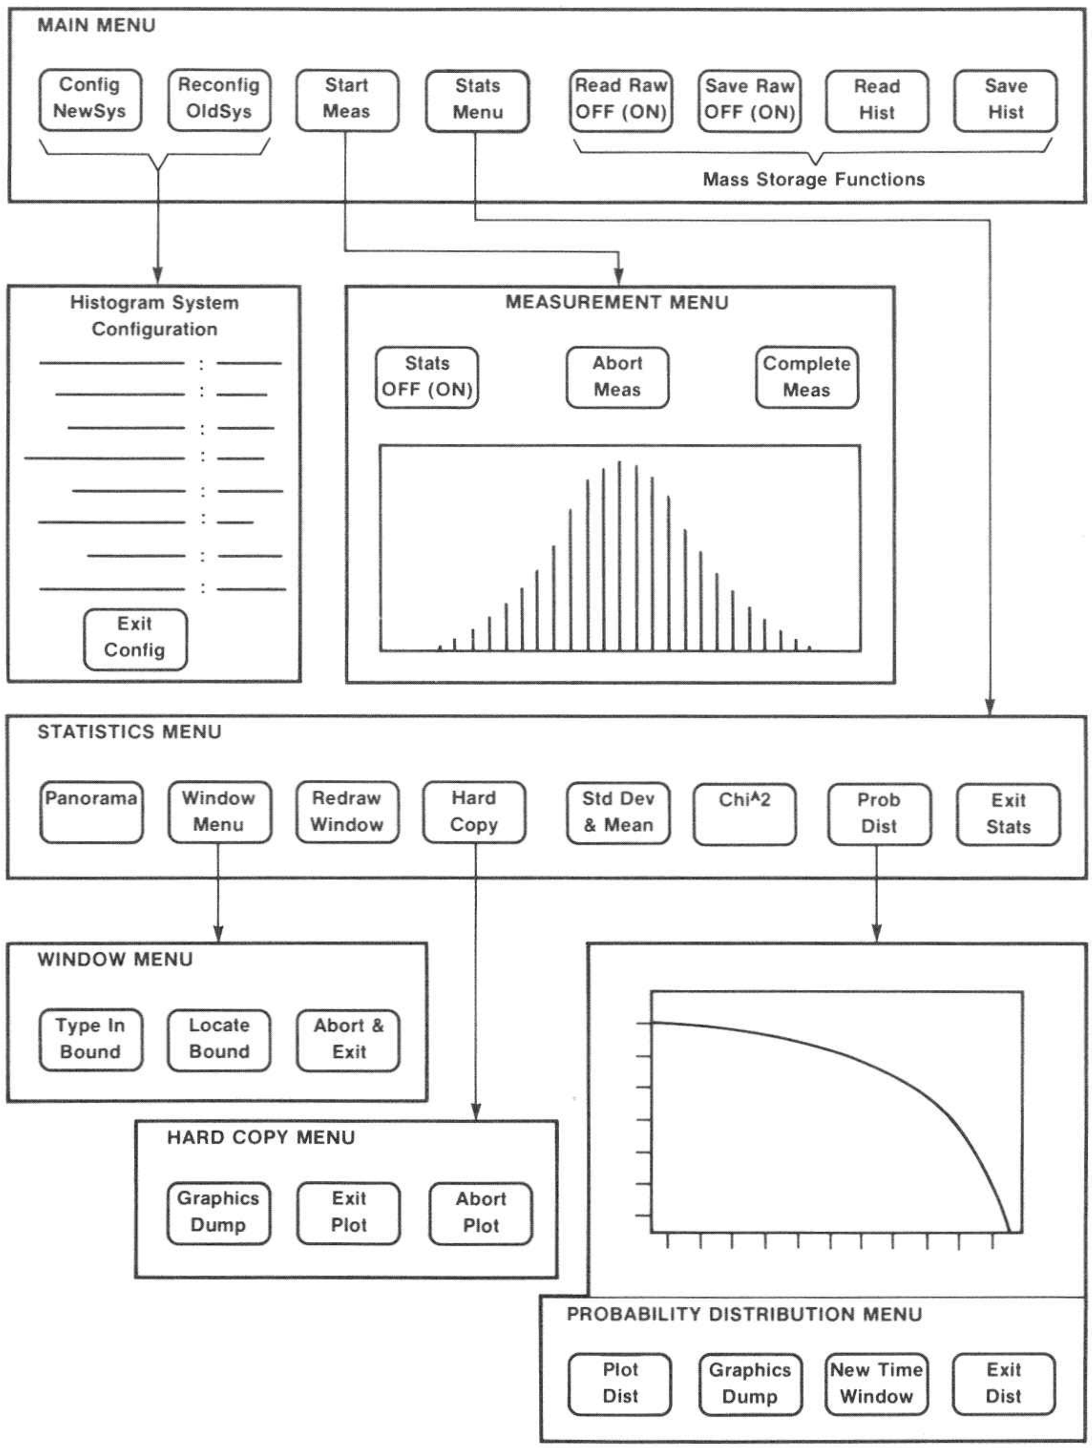

For more information, call your local HP sales office listed in your telephone directory or an HP regional office listed below for the location of your nearest sales office. Ask for the Electronic Instrument Department.

United States:
Hewlett-Packard Company
4 Choke Cherry Road
Rockville, MD 20850
(301) 258-2000

Hewlett-Packard Company
5201 Tollview Dr.
Rolling Meadows, IL 60008
(312) 255-9800

Hewlett-Packard Company
5161 Lankershim Blvd.
No. Hollywood, CA 91601
(818) 505-5600

Hewlett-Packard Company
2000 South Park Place
Atlanta, GA 30339
(404) 955-1500

Canada:
Hewlett-Packard Ltd.
6877 Goreway Drive
Mississauga, Ontario L4V1M8
(416) 678-9430

Japan:
Yokogawa-Hewlett-Packard Ltd.
29-21, Takaido-Higashi 3-chome
Suginami-ku, Tokyo 168
(03) 331-6111

Latin America:
Hewlett-Packard Company
3495 Deer Creek Rd.,
Palo Alto, CA 94304 USA
(415) 857-1501

Australia/New Zealand:
Hewlett-Packard Australia Ltd.
31-41 Joseph Street,
Blackburn, Victoria 3130
Melbourne, Australia
(03) 895-2895

Far East:
Hewlett-Packard Asia Ltd.
47/F China Resources Building
26 Harbour Road,
Hong Kong
(5) 833-0833

Germany:
Hewlett-Packard GmbH
Hewlett-Packard-Strasse
6380 Bad Homburg
West Germany
(49) 6172/400-0

France:
Hewlett-Packard France
Parc d'activité du Bois Briard
2, avenue du Lac
91040 Evry Cedex, France
(33) 1/60778383

United Kingdom
Hewlett-Packard Ltd.
Miller House—The Ring
Bracknell
Berkshire RG12 1XN, England
(44) 344/424898

Italy:
Hewlett-Packard Italiana S.A.
Via G. di Vittorio, 9
20063 Cernusco S/N (MI)
Milan, Italy
(39) 2/923691

Northern Europe:
Hewlett-Packard S.A.,
P.O. Box 999,
1180 AZ Amstelveen,
The Netherlands
(31) 20/437771

Southeast Europe/Africa/ Middle East:
Hewlett-Packard S.A.
1217 Meyrin 1, Geneva
Switzerland
(41) 22/989651

Or Write To:

United States:
Hewlett-Packard Company
P.O. Box 10301,
Palo Alto, CA 94303-0890

Europe/Middle East/Africa:
Hewlett-Packard Company
Central Mailing Department,
P.O. Box 529,
1180 AM Amstelveen,
The Netherlands

For all other areas:
Hewlett-Packard Company
Intercontinental Headquarters
3495 Deer Creek Rd.,
Palo Alto, CA 94304

HEWLETT PACKARD

Data Subject to Change
Printed in U.S.A.
02-5952-7908
June 1987