User Manual

Tektronix

DAS 9200
Performance Analysis

070-8499-01

This document supports 92PA software version 5.20 and above.

Please check for change information at the rear of this manual.

First Printing: May 1993

Instrument Serial Numbers

Each instrument manufactured by Tektronix has a serial number on a panel insert or tag, or stamped on the chassis. The first letter in the serial number designates the country of manufacture. The last five digits of the serial number are assigned sequentially and are unique to each instrument. Those manufactured in the United States have six unique digits. The country of manufacture is identified as follows:

B010000	Tektronix, Inc., Beaverton, Oregon, USA
E200000	Tektronix United Kingdom, Ltd., London
J300000	Sony/Tektronix, Japan
H700000	Tektronix Holland, NV, Heerenveen, The Netherlands

Instruments manufactured for Tektronix by external vendors outside the United States are assigned a two digit alpha code to identify the country of manufacture (e.g., JP for Japan, HK for Hong Kong, IL for Israel, etc.).

Tektronix, Inc., P.O. Box 500, Beaverton, OR 97077

Printed in U.S.A.

WARRANTY

Tektronix warrants that this software product will conform to the specifications in the documentation provided with the product, when used properly in the specified operating environment, for a period of three (3) months. The warranty period begins on the date of shipment, except that if the program is installed by Tektronix, the warranty period begins on the date of installation or one month after the date of shipment, whichever is earlier. If this software product does not conform as warranted, Tektronix will provide remedial services as described in the documentation provided with the product. Tektronix does not warrant that the functions contained in this software product will meet Customer's requirements or that operation of the programs will be uninterrupted or error-free or that all errors will be corrected.

In order to obtain service under this warranty, Customer must notify Tektronix of the defect before the expiration of the warranty period and make suitable arrangements for such service in accordance with the instructions received from Tektronix. If Tektronix is unable, within a reasonable time after receipt of such notice, to provide remedial services, Customer may terminate the license for this software product and return this software product and any associated materials to Tektronix for credit or refund.

This warranty shall not apply to any software product that has been modified or altered by Customer. Tektronix shall not be obligated to furnish service under this warranty with respect to any software product a) that is used in an operating environment other than that specified or in a manner inconsistent with the User Manual and documentation or b) when the software product has been integrated with other software if the result of such integration increases the time or difficulty of analyzing or servicing the software product or the problems ascribed to the software product.

THIS WARRANTY IS GIVEN BY TEKTRONIX WITH RESPECT TO THE LISTED PRODUCT IN LIEU OF ANY OTHER WARRANTIES, EXPRESSED OR IMPLIED. TEKTRONIX AND ITS VENDORS DISCLAIM ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. TEKTRONIX' RESPONSIBILITY TO PROVIDE REMEDIAL SERVICE WHEN SPECIFIED, REPLACE DEFECTIVE MEDIA, OR REFUND CUSTOMER'S PAYMENT IS THE SOLE AND EXCLUSIVE REMEDY PROVIDED TO THE CUSTOMER FOR BREACH OF THIS WARRANTY. TEKTRONIX AND ITS VENDORS WILL NOT BE LIABLE FOR ANY INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES IRRESPECTIVE OF WHETHER TEKTRONIX OR THE VENDOR HAS ADVANCE NOTICE OF THE POSSIBILITY OF SUCH DAMAGES.

Contents

List of Figures

List of Tables

Preface

The Digital Analysis System (DAS) 9200 documentation package provides the information necessary to install, operate, maintain, and service the DAS 9200. DAS 9200 documentation consists of:

- A **system user manual** that includes a beginning user orientation, a discussion of DAS 9200 system-level operation, and reference information such as installation procedures, specifications, error messages, and a glossary of terms
- A **programmatic command language user manual** that describes the set of programmatic commands available for remotely controlling the DAS 9200
- A series of **module user manuals** that cover each of the DAS 9200 acquisition, pattern generation, and optional I/O modules
- An **on-line documentation** package that includes a location-dependent system of technical notes
- A series of **application software user manuals** that describe the various application software packages
- A series of microprocessor-specific **microprocessor support instructions** (designed to accompany the *92A60/90 Module User Manual* and *92A96 Module User Manual*) that describe the various microprocessor support packages
- Service documentation that help a qualified technician isolate DAS 9200 problems to the individual module level and determine corrective action (including on-site removal and replacement of modules)

How to Use this Manual

This *DAS 9200 Performance Analysis User Manual* is designed to help you learn and use the performance analysis application software package for the DAS 9200. If you are unfamiliar with the operator's controls and display conventions, read the orientation in Section 2 of the *DAS 9200/SE System User Manual.*

This manual consists of the following sections:

- **Section 1: Introduction.** Provides basic information about the performance analysis application software package and compatibility information with the DAS 9200.
- **Section 2: Getting Started.** Provides basic information for installing the performance analysis software on a DAS 9200 and briefly describes some of the DAS 9200 Setup menus used with the package.
- **Section 3: Using Named Ranges.** Provides information on using the Named Ranges application. The Named Ranges menus, fields, and function keys are described in detail.
- **Section 4: Single Event.** Provides information on using the Single Event application. The Single Event menus, fields, and function keys are described in detail.
- **Section 5: Application Information.** Provides guidelines in adjusting the histogram display, transferring the data files to a host computer, and provides simple troubleshooting suggestions.
- **Appendix A: PA Data Files.** Provides descriptive information on the contents of the performance analysis data files for use with host computer applications.

Introduction

The DAS 9200 performance analysis (92PA) software allows you to use a DAS 9200/SE mainframe with the 92A96, 92A60, or 92A90 Data Acquisition Modules to perform a variety of software analysis tasks. You can use performance analysis to determine the level of activity within various distribution ranges that you define. You can analyze how memory is being used, determine the relative execution times of various subroutines or program modules, or determine where to focus your attention when optimizing code.

The DAS 9200 performance analysis software has two different applications, the Named Ranges application and the Single Event application. The Named Ranges application lets you monitor the activity of the address group and display the results in a histogram. The Single Event application lets you time events with a specific timer or counter and display the results in a histogram. Both applications are discussed in detail in this manual.

Basic Information

To use this product, you need to have the following:

- This manual
- A DAS 9200/SE mainframe with a 92A96, 92A60, or 92A90 Data Acquisition Module
- Other DAS 9200 mainframe and module user documentation
- Knowledge of your specific DAS 9200/SE mainframe configuration and its operation
- Knowledge of your system under test (SUT)

DAS 9200 System Software Compatibility

The DAS 9200 92PA software version 5.20 is compatible with DAS 9200/SE mainframes and System Software Release 3, Version 1.00 and up; it is not compatible with previous versions of software or mainframes. The software is compatible with DAS 9200/SE mainframes running 92XTerm System Software Release 3, Version 1.10.

About this Manual

The following conventions are used in this manual:

- The terms performance analysis and PA are used interchangeably throughout this document
- The term DAS 9200/SE refers to the DAS 9200/SE (DAS 9221) mainframe; it does not refer to earlier versions of the mainframe (DAS 9220 or DAS 9219)
- The terms system under test and SUT are used interchangeably when referring to the user's system that is being tested
- References to the 92A96 Data Acquisition Module include all versions of that module
- References to the 92A60 and 92A90 Data Acquisition Modules include all versions of those modules
- All menu illustrations in this document show menus displayed on a 9202XT terminal or workstation. If your DAS 9200 system has a 9201T terminal, your menus will not show the on-screen buttons displayed on the bottom of the menus; however, the main contents of the menus will be similar.

Other Necessary Documentation

Before using this product, you should be familiar with the operation of a DAS 9200/SE mainframe and the data acquisition module you are using. For general instructions on using the DAS 9200/SE mainframe, refer to the *DAS 9200/SE System User Manual* and to the data acquisition module's user manual. For information on specific microprocessor disassemblers, refer to the manual that came with your disassembler product.

DAS 9200 Configuration

To use the DAS 9200 Performance Analysis software, your DAS 9200/SE mainframe must be equipped with at least one 92A96, 92A90, or 92A60 Data Acquisition Module. Follow the configuration guidelines as described in Appendix C of the *DAS 9200/SE System User Manual*.

Getting Started

This section briefly describes how to install the Performance Analysis (PA) software on the DAS 9200 and how to start the software for simple applications.

Installing the PA Software

The 92PA software comes on a 5.25-inch floppy disk. The floppy disk contains all the software necessary to install and run performance analysis on the DAS 9200.

You must install the PA software on the DAS 9200 hard disk prior to using it; the DAS 9200 cannot execute the software from the floppy disk.

To install the PA software onto the hard disk, follow these steps:

1. Power on the DAS 9200 and display the Menu Selection overlay.
2. Select the Disk Services menu.
3. Select F6: MOVE TO UTILITY.
4. When the Disk Services Menu appears, select Install Applications in the Operation field.
5. Select F8: EXECUTE OPERATION and follow the on-screen prompts.

If there is inadequate disk free space available on the hard disk, you must use the Remove Application or Delete File function of the Disk Services menu to free up enough disk space to install the software. The approximate space required to install the software is listed on the label of the floppy disk.

After the DAS 9200 successfully copies the software onto the hard disk, the following message displays on your screen: **Application Installation complete with no errors.** Remove the floppy disk and store it in a safe place in case you need to reinstall the software.

Connecting to the System Under Test

Refer to the *DAS 9200/SE System User Manual* for specific information on installing the 92A96, 92A60, or 92A90 Data Acquisition Modules in the DAS 9200. Refer to the individual module user manuals for details on connecting to the system under test. If you are using one of the microprocessor disassemblers, refer to the microprocessor instruction manual for information on connecting the DAS 9200 to the system under test via a probe adapter.

Setting Up the DAS 9200

Prior to using the PA software for the first time, you must select either a 92A96, 92A90, or 92A60 Module and define the setup menus. The Performance Analysis software will not work with other DAS 9200 modules.

You can use the PA software with one of the DAS 9200 microprocessor disassemblers, or you can use it for general purpose performance analysis.

Depending on whether you use the Named Ranges application or the Single Event application, the initial module setups may differ. Such differences, where they occur, are explained in the following paragraphs.

Defining the Setup Menus Using DAS 9200 Microprocessor Disassemblers

After specifying the acquisition module, you should define the individual setup menus. If you use one of the DAS 9200 microprocessor disassembler support packages, some of the setup menus are automatically defined when you select the microprocessor support in the Config menu.

Config Menu. Use the Config menu to specify the software support (microprocessor disassemblers installed on the DAS 9200). You can also specify the length or depth of your acquisition module's memory. This affects the number of samples processed with each acquisition in the Named Ranges application and the update time interval between acquisitions.

The configuration memory depth and the trigger position can affect the number of samples stored. After you change the memory depth, you should check the trigger position to ensure the desired number of samples will be captured per acquisition.

Channel Menu. In most cases, when you use one of the microprocessor disassemblers, the Channel menu will be defined for you; there should be no need to change the channel setups further.

Clock Menu. When used with the microprocessor disassemblers, the Clock menu is automatically set up. However, you may want to specify some of the clocking choices to further control the data sampling. For example, you may want to include or exclude DMA Cycles. Refer to your microprocessor disassembler user manual for more information on the individual clocking choices available with your microprocessor disassembler software.

Trigger Menu. Use the Trigger menu to look for specific occurrences of data (events). You can look for a specific event or use signals with other acquisition or pattern generation modules.

When using the Named Ranges application, you can also use storage qualification to acquire only code fetches or other specific information into memory while ignoring other unwanted data. You should verify that the trigger program and trigger position selections are set such that the desired amount of acquisition memory will be filled before the module stops. For most applications (using the default Trigger menu), the module will trigger immediately and stop when the selected amount of post-trigger data has been acquired. The best trigger position for most applications is "T————", to fill the module's memory with post-trigger data.

If you use the Named Ranges application with a 92A96 with memory depths greater than 128 K (such as the 92A96SD Module), you should observe the limitations of the trigger positions. The trigger location is always within the last 128 K locations of the acquisition memory. If you want to store more samples after the trigger than the Defined field selection allows, you can use a counter in the trigger program to delay the actual trigger by the desired number of sample cycles. You can also use the "Trigger on Anything with Extended Postfill" trigger library selection to store more samples after an event is recognized.

When using the Single Event application, you must set up one or more counters or timers in the Trigger menu. You can set up the Trigger menu to look for a specific event. When the module detects the event, it can start a counter or timer to measure an event or the time between events. When the module stops and triggers, the values of the counter or timers are used by the PA software. You can also use counters or timers to trigger immediately after a measurement or when a specific event has occurred a specific number of times. However, the module must trigger before the data can be used by the PA software.

You should verify that the trigger program and trigger position selections are set such that the desired amount of acquisition memory will be filled before the module stops. For most applications (using the default Trigger menu), the module will trigger immediately and stop. The best trigger position for Single Event applications is "————T." The Single Event application only cares about the final counter or timer values, and not about the amount of data stored in memory.

The easiest way to set up the Trigger menus to look and act on trigger events is to use the acquisition module's trigger libraries. Refer to your module user manual for information on using trigger libraries.

Defining the Setup Menus Using General Purpose DAS 9200 Setups

How you define the Setup menus for general purpose use of the DAS 9200 depends on how you want to use the PA software with your specific application. This section provides only a brief overview in setting up the setup menus; refer to your acquisition module user manual for specific information on each of the setup menus.

Config Menu. Use the Config menu to specify the length or depth of your acquisition module's memory. This menu also describes how input and output signals are used with the acquisition module. This affects the number of samples processed with each acquisition and the update time interval between acquisitions.

The configuration memory depth and the trigger position can affect the number of samples stored. After you change the memory depth, you should check the trigger position to ensure the desired number of samples will be captured per acquisition.

Channel Menu. Use the Channel menu to organize your acquisition module's individual channels into meaningful groups for display. The default channel setup depends on the acquisition module you use. You can also assign a name for your channel group (in most cases this will be your module's address channel group) and that name will appear in the Named Ranges menu.

When using the Named Ranges application, channel order is important. Channels in the address group must be in normal hardware order. The channel menu arranges the channels to appear as shown in Table 2–1 (default order at power-on).

Table 2–1: 92A96 and 92A60/90 Channel Order

Section (n= slot number)	Channels
n_A3	7 6 5 4 3 2 1 0
n_A2	7 6 5 4 3 2 1 0
n_A1	7 6 5 4 3 2 1 0
n_A0	7 6 5 4 3 2 1 0

If you use fewer than 32 channels in a 92A96 Module, first remove the channels from the channel 7 (most significant bit) end of Section n_A3; the remaining channels must be contiguous. If you use the PA software with a multicard 92A96 Module, you can only use the first 32 channels of the address group of the first module (that is, you cannot use Sections A3 through A0 in the second or subsequent 92A96 modules).

If you use fewer than 32 channels in a 92A60 or 92A90 Module, first remove the channels from the channel 0 end (least significant bit) of Section n_A0; the remaining channels must be contiguous.

When using the Single Event application, channel order can be ignored. This is because you are interested in measuring events rather than specific data patterns.

Clock Menu. The clock menu specifies when the DAS 9200 samples data. Use the Clock menu to specify the clocking requirements for your application.

Trigger Menu Use the Trigger menu to look for specific occurrences of data. You can look for a specific event or use signals with other acquisition or pattern generation modules.

When using the Named Ranges application, you can determine or qualify the amount of data to sample. You should verify that the trigger program and trigger position selections are set such that the desired amount of acquisition memory will be filled before the module stops. For most applications (using the default Trigger menu), the module will trigger immediately and stop when the selected amount of post-trigger data has been acquired. The best trigger position for most applications is "T————", to fill the module's memory with post-trigger data.

If you use the Named Ranges application with a 92A96 with memory depths greater than 128 K (such as the 92A96SD Module), you should observe the limitations of the trigger positions. The trigger location is always within the last 128 K locations of the acquisition memory. If you want to store more samples after the trigger than the Defined field selection allows, you can use a counter in the trigger program to delay the actual trigger by the desired number of sample cycles. You can also use the "Trigger on Anything with Extended Postfill" trigger library selection to store more samples after an event is recognized.

You should verify that the trigger program and trigger position selections are set such that the desired amount of acquisition memory will be filled before the module stops. For most applications (using the default Trigger menu), the module will trigger immediately and stop. The best trigger position for Single Event applications is "————T." The Single Event application only cares about the final counter or timer values, and not about the amount of data stored in memory.

When using the Single Event application, you must set up one or more counters or timers in the Trigger menu. You can set up the Trigger menu to look for a specific event. When the module detects the event, it can start a counter or timer to measure an event or the time between events. When the module stops and triggers, the values of the counter or timers are used by the PA software. You can also use counters or timers to trigger immediately after a measurement or when a specific event has occurred a specific number of times. However, the module must trigger before the data can be used by the PA software.

The easiest way to set up the Trigger menus to look and act on trigger events is to use the acquisition module's trigger libraries. Refer to your module user manual for information on using trigger libraries.

Using the PA Software

To use the PA software, follow these basic steps:

1. Enter the PA application menu (PA-Named Rng or PA-Single Ev). If necessary, go to the Define Setup overlay and define the PA setups.
2. Start the PA. The DAS 9200 will begin displaying the data in the form of histograms (bar graphs) and tables. You can scroll through the displayed data as it is analyzed by the DAS 9200.
3. Analyze the data. You can either analyze the data on the DAS 9200 or you can save the data to an ASCII file and transfer the file to a host computer for further analysis. Appendix A of this manual provides information on the format and contents of the data files.

The previous steps are meant to give you a brief overview of using performance analysis. More detailed information is provided later in this manual.

Using Symbol Tables With the PA Software

The Named Ranges application is most often used with address ranges to get a clear picture of how much time is spent in different sections of code. In reality, the address ranges or sections of code are specific areas in memory. These areas of memory are commonly referred to by symbolic names (for example, strlen, io_memory) instead of specific addresses. The PA software lets you use symbolic names as well as numeric address ranges. The symbolic names are defined in a symbol table stored on the DAS 9200's hard disk.

You can also use symbol tables with the Single Event application. However, symbol tables are more apt to be used with the Named Ranges application where you are more interested in specific ranges than measuring events.

You can use symbol tables with the PA software to clearly define the ranges you are interested in for your specific application. You can create symbol tables to use with the PA software using one or more of the following methods:

- Use the LA-LINK tool (available from Tektronix) to create a symbol table from the output of various compilers and assemblers and download the symbol table to the DAS 9200.
- Create the symbol table on a host computer with a text editor and download it to the DAS 9200 using one of the DAS 9200's file-transfer tools. If you create a symbol table in this manner, make sure that you place it in the Symbol Table directory on the DAS 9200 hard disk.
- Use the DAS 9200's Symbol Table Editor (refer to the *DAS 9200/SE System User Manual* for more information).

Because the PA software deals with ranges, instead of specific data patterns, you can only use range symbol tables with performance analysis; pattern symbol tables are not supported. Refer to the *DAS 9200/SE System User Manual* for more information on range and pattern symbol tables.

- Save the PA range definitions as a symbol table using the F6: SAVE AS SYM TBL function key.

Using Named Ranges

Performance Analysis is a useful tool for designing, fine tuning, and debugging software. Running software for the first time with performance analysis can help you determine if you need to cut features, optimize code, or in some cases, redesign the hardware. Using the Named Ranges application of performance analysis with a microprocessor disassembler product can help you determine what the software is doing and whether it executes the code in the right amount of time.

For best results, performance analysis is best used for measuring repetitive tasks. As you run your software in a repetitive mode, your sampling errors will decrease over time as you acquire more data samples. The longer you run the software, the more accurate your results will be; this is evident when the displayed histogram becomes stable.

A key advantage of using the Named Ranges application is the large number of ranges available. With the DAS 9200 you can define up to 5000 ranges.

The Named Ranges application provides you with the flexibility of changing the size for the acquisition buffer (in the Config menu of the acquisition module). You can acquire smaller buffers of information more often, or larger buffers less often; you will need to analyze your application to determine which method is best.

The Named Ranges application works with address ranges to get a clear picture of how much time is spent in a certain section of code. You can also determine which sections of code (such as subroutines) are accessed or not accessed. As long as you acquire enough data for relevant results, the PA software will provide you with the necessary information to help you analyze your code.

You can use the Trigger menu of your acquisition module to qualify the acquired data to store in memory. For example, you may only want to capture all code fetches into memory within a certain address range. When working with microprocessors, you can set up the Trigger menu to look for data in specific address spaces or instruction modes (for example, Supervisor or User).

The Named Ranges application acquires and processes data in the following manner:

1. The DAS 9200 acquires and tests data as defined in the acquisition module's trigger setups.
2. Valid data fills the acquisition memory and the acquisition stops.
3. The data is processed. Each time data execution occurs within the ranges you defined, the PA software advances a counter for that range.

4. The Named Ranges menu displays the data in a table and histogram.

5. The DAS 9200 fills the acquisition memory again and repeats the entire process.

The results of the performance analysis help you identify sections of code that are not performing according to specifications.

When using the Named Ranges, you usually begin optimizing your code by studying the activity for your entire system-under-test. If you observe unexpected activity in a range (for example, too much or too little time within a certain range), then you can study that range in more detail.

For example, you can quickly create a series of ranges that span an address region. The Named Ranges application then produces a histogram showing the time spent in the ranges. This information can help determine if a particular subroutine executes too frequently and slows down the overall performance of your system.

Defining Setup Menus

Before starting the application, be sure that you have defined your acquisition module's setups in the Setup menus. Refer to *Setting Up the DAS 9200* in Section 2 for a brief overview on setting up the menus.

Setting Up Named Ranges

After defining your setups, from the Menu Selection overlay, move the cursor to the Applications column and select PA-Named Rng. Select function key F7: MOVE TO APPLICTN to enter the Named Ranges application.

Figure 3-1 shows the default Named Ranges menu; details of certain fields are described below the menu. At this point you can either start the PA software or you can select a symbol table or specify other parameters in the Setup Definition overlay.

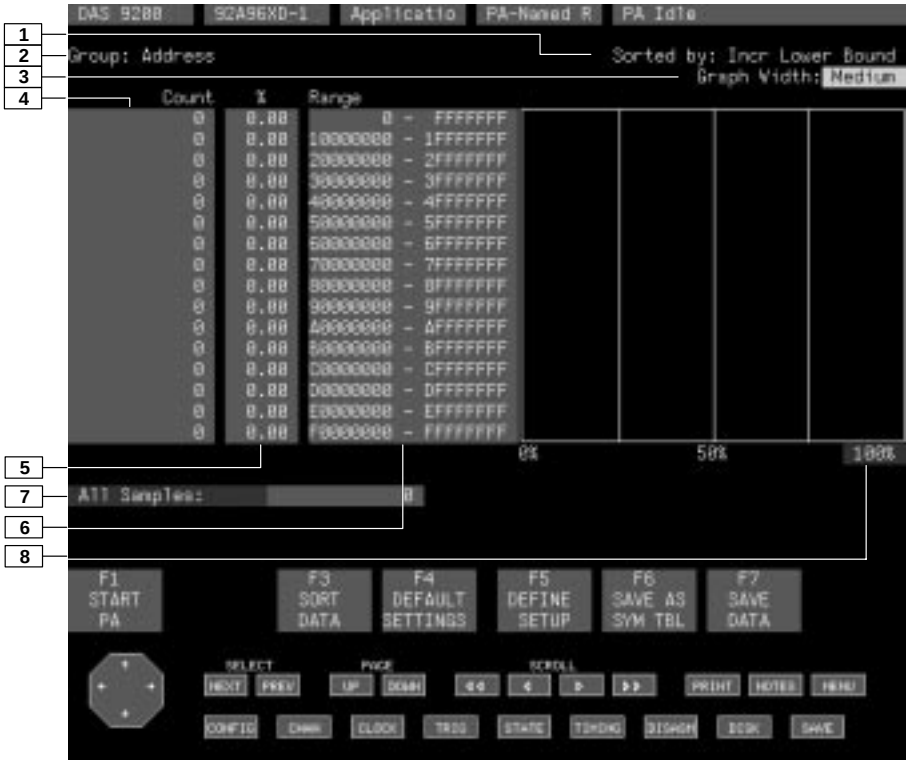

Figure 3-1: Default Named Ranges menu

1 **Sorted by.** This label shows how the displayed ranges are sorted. Use the F3: SORT DATA function key to specify how to sort the data in the display.

2 **Group.** This label shows the name of the address group as specified in the Channel menu setup. If there is no valid address group, the DAS 9200 will display an error message describing the recommended action.

3 **Graph Width.** Specifies the width of the histogram: Min (minimum), Small, Medium (default), Large, or Max (maximum).

4 **Count.** Shows the number of samples within each range. Upon first entering the menu, the count values will be zero.

5 **Range Percent.** Lists the percentage of samples within each range. The percentage is calculated by dividing the count by the number of samples (all samples, or selected samples).

6 **Range Names.** Lists the name for each range. The name of each range can be a symbol from the current symbol table, or it can be numeric low and high boundaries.

7 **Samples Selection.** Specifies the reference for the data collected. The "All Samples" selection is a total of all the samples from the acquisition. The "Selected Samples" selection displays the total number of samples contained in the defined ranges. Changing this field changes the percentage values of the ranges and the length of the histogram bars; the count values for each range remain unchanged.

8 **Histogram Magnification.** Specifies the upper value of the histogram. Changing this value changes the magnification of data in the histogram. If you change this value, the label in the center of the histogram changes proportionally.

Function Keys

F1: START PA. Starts the module (or cluster). Any previous range counts will be cleared. When the PA software is running, the function key legend changes to STOP PA.

F2: CONTINUE PA. Resumes the PA acquisition cycles. This function key does not appear in the default menu; it appears when you stop the PA. Any existing range counts are preserved as if the PA software had never stopped.

F3: SORT DATA. Displays a pop-up window allowing you to choose how you want to sort the displayed data. Valid selections are: Sort by Increasing Lower Bound, Sort by Decreasing Lower Bound, Sort by Decreasing Count, and Sort by Increasing Alphanumeric Range Names (only valid for ranges defined in a symbol table). Select your choice and press either the Return key or the Open/Close key. The changes are immediately applied to the menu and the Sorted By field. See Section 5 for more information on sorting data.

F4: DEFAULT SETTINGS. Changes all fields to their power-on default settings. You will be prompted to confirm your choice.

F5: DEFINE SETUP. Calls the Setup Definition overlay.

F6: SAVE AS SYM TBL. Saves the current range settings as a symbol table in the Symbol Table directory on the DAS 9200's hard disk. You will be prompted for a file name; the file name can consist of up to 11 alphanumeric characters. If the file already exists, you will be prompted for confirmation to overwrite the file. You can access the file with the Symbol Table editor, the Disk Services menus, or any of the DAS 9200's file-transfer tools. You can edit the new symbol table with the DAS 9200's Symbol Table editor and then select the symbol table in the Setup Definition overlay.

F7: SAVE DATA. Saves the acquired PA data in an ASCII file in the PA_NR_Data directory on the DAS 9200's hard disk. You will be prompted for a file name; the file name can consist of up to 11 alphanumeric characters. If the file already exists, you will be prompted for confirmation to overwrite the file. The new file will be saved as ASCII and you can access the file with the Disk

Services menus or any of the DAS 9200's file-transfer tools. Appendix A of this manual contains information on the format and contents of the file.

Although you can start the performance analysis upon entering the menu, in most cases you will want to go to the Setup Definition overlay to further define your setup. Select F5: DEFINE SETUP to enter the overlay.

Setup Definition Overlay

The Setup Definition overlay lets you define the range boundaries (bounds). You can specify a symbol table that has the boundaries predefined, or you can use Linear Generation to create equally sized ranges.

Linear Generation Setups

The PA software uses the Linear Generation mode to generate equally sized ranges based on the total number of ranges and on the maximum and minimum boundaries. The default setup for the Named Ranges application is Linear Generation.

Figure 3-2 shows the default Setup Definition overlay when you use Linear Generation.

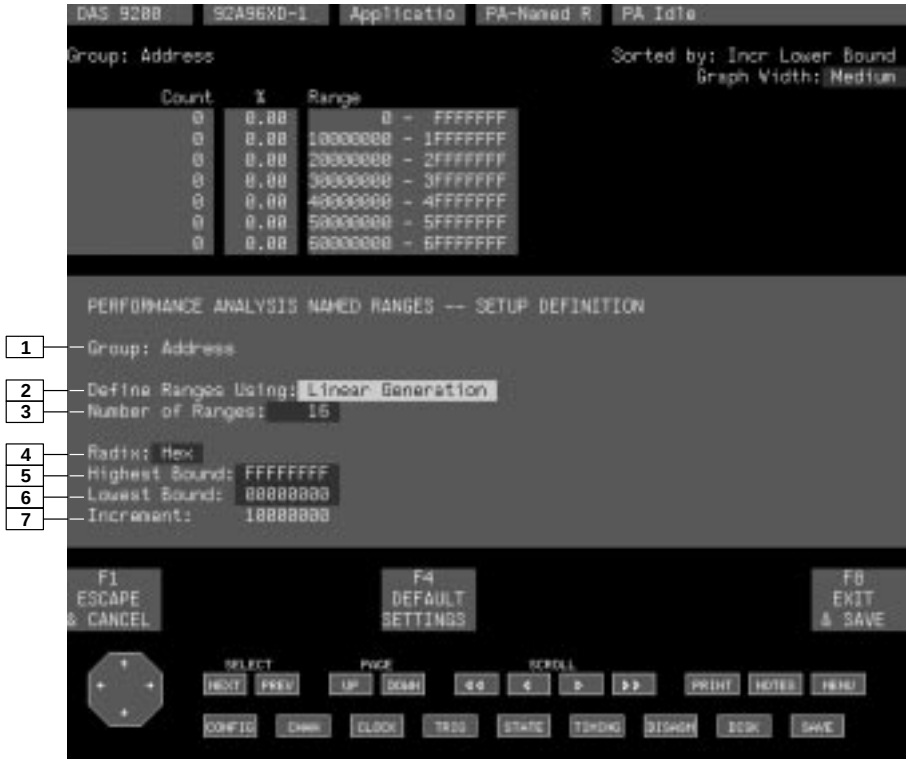

Figure 3-2: Default Setup Definition overlay

1 **Group.** This label shows the name of the address group as specified in the Channel menu setup.

2 **Define Ranges Using.** Specifies whether the ranges will be defined by Linear Generation or based on values in a symbol table. The Symbol Table selection is only available if a legal range symbol table exists.

3 **Number of Ranges.** Specifies the number of ranges to use. The default value is 16 for Linear Generation; you can specify any number of ranges from 1 up to 5000.

4 **Radix.** Specifies the radix of the range boundaries in the overlay as well as in the Named Ranges menu. Valid selections are Hex (hexadecimal), Dec (decimal), and Oct (octal).

5 **Highest Bound.** Specifies the highest (upper) boundary of the ranges. For Linear Generation, enter the value of the highest boundary you want to use (this value must be higher than the Lowest Bound value).

6 **Lowest Bound.** Specifies the lowest (smallest) boundary for the ranges. For Linear Generation, enter the value of the lowest boundary you want to use (this value must be lower than the Highest Bound value).

7 **Increment.** Reports the size of the range in Linear Generation. The PA software will try to keep the individual range sizes the same (linear) over the entire set of defined ranges. If the span defined by the Highest Bound and Lowest Bound fields is not evenly divisible by the number of ranges, then the first "n" ranges (n = remainder) will be one larger than the increment.

Function Keys

F1: ESCAPE & CANCEL. Leaves the overlay without changing the range definitions. The main menu is unaffected when you leave the overlay by selecting this key.

F4: DEFAULT SETTINGS. Changes all fields to their power-on default settings. The changes will only be applied when you select the F8: EXIT & SAVE function key.

F8: EXIT & SAVE. Leaves the overlay and applies all changes made in the overlay. If there were any changes, the range counts, range percentages, and samples selection (All or Selected) will be set to zero.

Symbol Table Setups

When you use a symbol table with performance analysis, the Setup Definition overlay changes from the Linear Generation mode to Symbol Table mode as shown in Figure 3-3.

Specifying a symbol table in the overlay adds other information (fields) to the Named Ranges menu; more on this later.

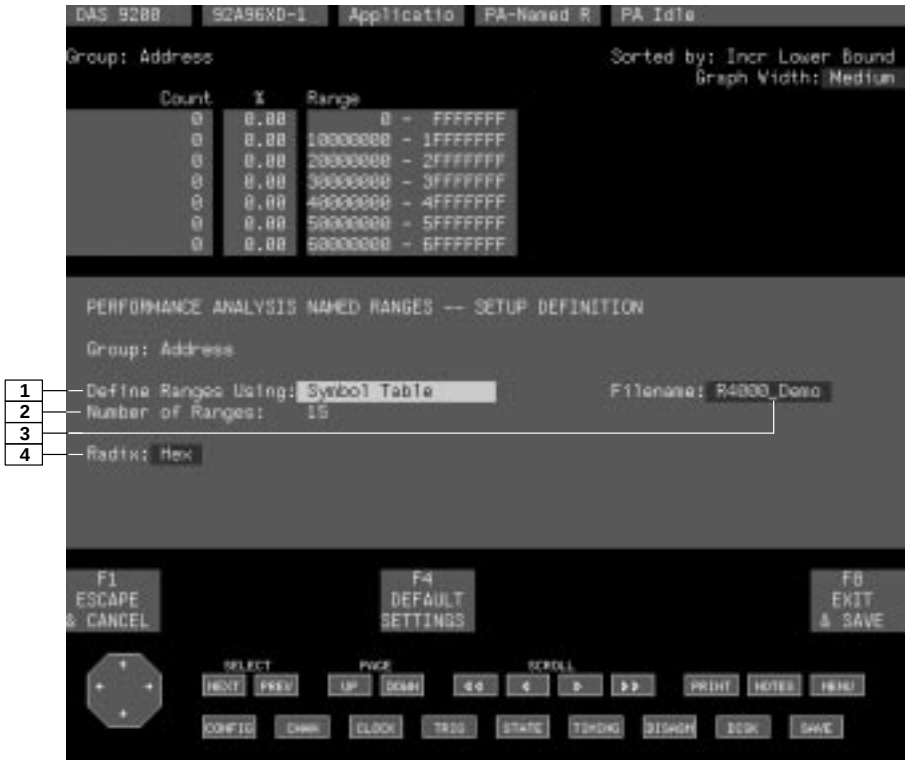

Figure 3-3: Setup Definition overlay using symbol tables

1 **Define Ranges Using: Symbol Table.** Specifies that the ranges will be defined by values in a symbol table. This selection is only available if a legal symbol table exists. The Filename field appears to the right allowing you to specify the name of the symbol table.

2 **Number of Ranges.** The number of ranges depends on the number of symbol table ranges (at least one) in the symbol table.

3 **Filename.** Specifies the name of the range symbol file (pattern symbol tables are not supported with the PA software). The file must reside in the Symbol_Table directory on the DAS 9200's hard disk. Use the Open/Close key to open the field and select other valid (legal) file names.

4 **Radix.** Specifies the radix of the range boundaries. The radix setting overrides the radix of the symbol table file values. Valid selections are Hex (hexadecimal), Dec (decimal), and Oct (octal). For example, if you specify octal, and you view the symbol table values in the Bound mode in the Named Ranges menu (rather than by symbol names), the bounds values are displayed in octal (see Figure 3-5).

After defining your setups, close the overlay by selecting F8: EXIT & SAVE. The setups are immediately applied to the Named Ranges menu.

Be aware that any changes you make to the overlay will clear or flush any data in the Named Ranges menu. In other words, the PA values will be reset.

Acquiring Data

Before starting to acquire data, you may want to consider how you want to make the best use of the information on the screen. You can fine-tune the information in the display as follows:

- Sort the displayed data by selecting F3: SORT DATA.
- Adjust the width of the graph by changing the value of the Graph Width field. You can either select a value in the field or use the Select Next or Prev keys (with the cursor in the Graph Width field) to cycle through the different selections.
- Change the scale of the graph by changing the value of the Magnification field.
- Scroll to ranges that are off the screen using the Joydisk or the Scroll keys; you can use the Scroll and Shift keys together to scroll a half page (window) of data. Arrows at the top or bottom of the left side of the Count column indicate that there are more ranges off the screen.
- Change from viewing all data samples to selected samples and vice versa.
- Change from viewing the symbol table ranges from symbolic names to bound values and vice versa.

You can make changes to the displayed information at any time before, during, or after making PA acquisition cycles. Section 5 of this manual provides some examples of adjusting the menu.

To start the performance analysis, select F1: START PA. The DAS 9200 will begin acquiring data and display the results on the screen. The status field in the upper-right corner of the screen indicates that the PA software is running.

Figure 3-4 shows an example of the Named Ranges menu while the PA software is acquiring cycles.

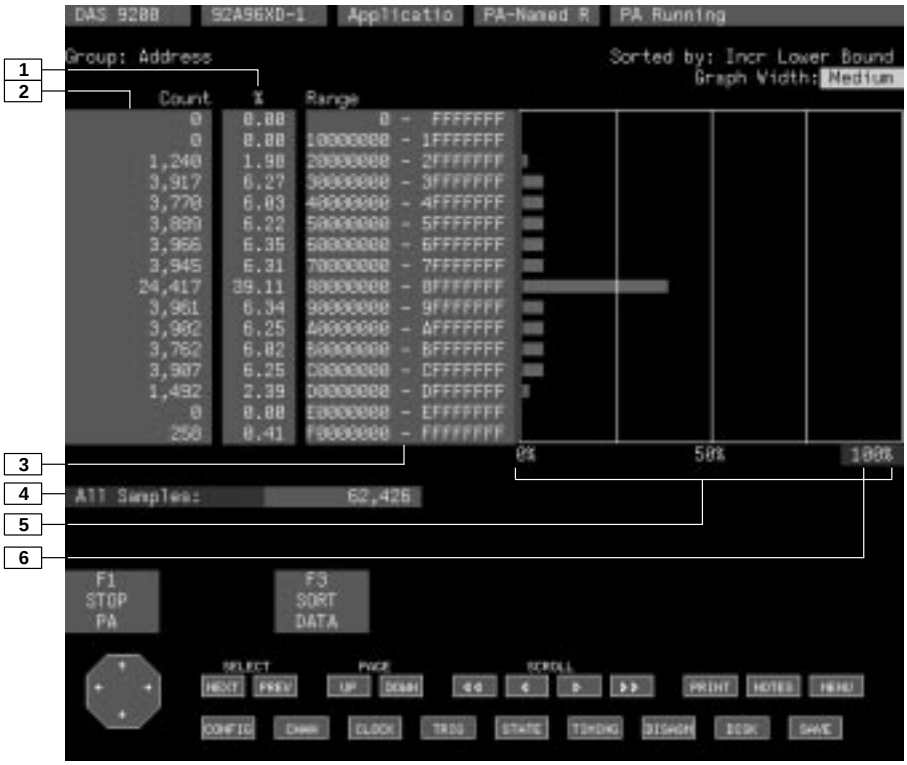

Figure 3-4: Performance analysis data (Linear Generation)

1 **Count.** Shows the number of samples within each range.

2 **Range Percent.** Shows the percentage of samples within each range (relative to the Samples Selection field).

3 **Range Names.** Shows the names or boundaries of each range.

4 **Samples Selection.** Specifies the reference for the data collected. The All Samples selection is a total of all the samples from the acquisition. The Selected Samples selection displays the total number of samples contained in the defined ranges. Changing this field changes the percentage values of the ranges and the length of the histogram bars; the count values for each range remain unchanged.

5 **Histogram Data.** Displays the acquired data in a histogram; one bar for each range. To view ranges with smaller percentage values, you can magnify the display by changing the value of the Histogram Magnification field.

6 **Histogram Magnification.** Specifies the upper value of the histogram. Changing this value changes the magnification of data in the histogram. If you change this value, the label in the center of the histogram changes proportionally.

Ranges Defined As Symbol Tables

If you acquire data with the ranges defined from a symbol table, other fields are displayed in the menu. Figure 3-5 shows these fields. In this example the Symbol field and the Show field are only displayed when the range definitions were loaded from a range symbol table file in the Setup Definition overlay.

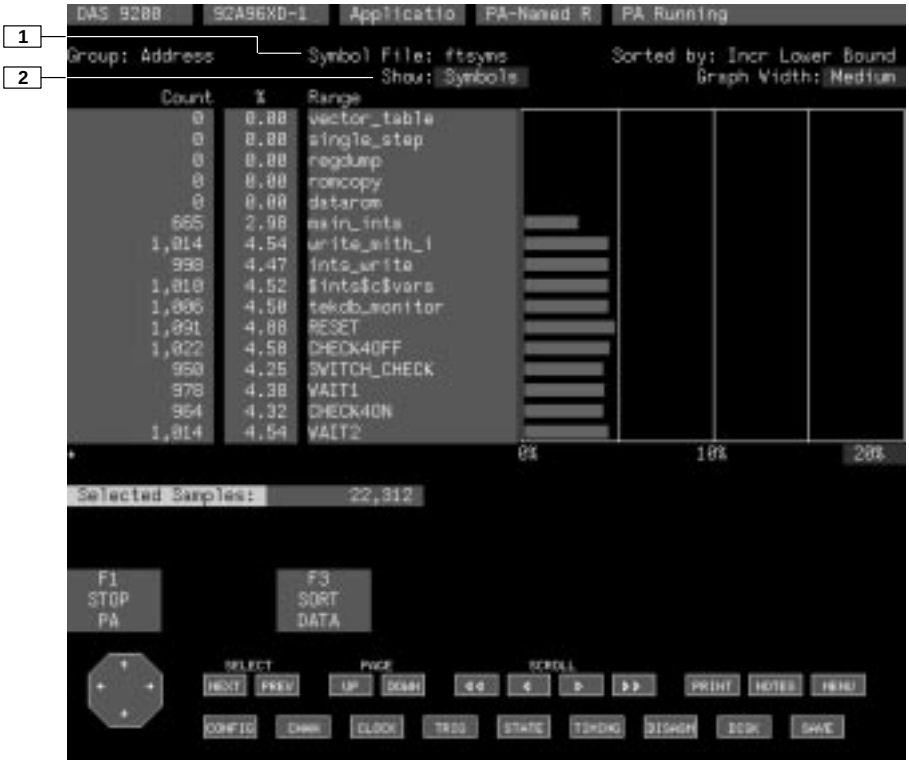

Figure 3-5: Performance analysis data (Symbol Table)

1 **Symbol File.** Specifies the name of the current symbol table.

2 **Show.** Specifies whether the ranges will be displayed as either Symbols or Bounds. If you view the ranges as bounds, the bounds will be displayed in the radix value specified in the Setup Definition overlay when the ranges were loaded from the symbol table.

After you have acquired data, you can analyze the data in the menu. You can also save the data in an ASCII file and transfer it to a host computer for further analysis. Refer to Section 5 for more application information. Refer to Appendix A for information on the details of the contents of the PA data file.

Using Single Event

The Single Event Application of the PA software allows you to display the range of execution times or range of items counted during a given event. An event can be defined as one of several conditions that can be detected by the acquisition module. Examples of events can be a specific word recognizer value, a specific counter value, or signals recieved from another module or device connected to the logic analyzer.

You can set up the Trigger menu to program a counter or timer to measure a specific event. For example, you may want to time interrupt routines or exception handlers. This lets you wait for the desired event to occur, measure the event with a counter or timer, and then graphically display the execution times.

You may want to use the Single Event application to validate the performance of an interrupt routine or a data handling function. You may have a system specification that requires the routines to execute within a certain amount of time under all conditions. You can use a range symbol table to define the acceptable, nominal, and critical ranges and then graphically display the execution times.

You can also use the Single Event application for error analysis. You can monitor the exception handler and see how much time your system spends handling error routines.

In each application, you would program the Trigger menu to detect and time the event. You would then run your software under all test conditions. The resultant display will show the average, minimum, and maximum amount of time that the routine required to execute. From this information you can determine if you are within the specifications and by how much.

When programming the Trigger menu, you can build a trigger program from the default menu settings, or you can use and build on some of the module's trigger library selections. For example, to count the occurrences between two events, you can use the selection "Count occurrences of C between A and B." Alternately, to measure the amount of time between two events, use the selection, "Measure time spent between A and B." The important thing to remember is that in order for the Single Event Application to work, the acquisition module must trigger before the counter and timer values can be used.

Defining Setup Menus

Before starting the application, be sure that you have defined your acquisition module's setups in the Setup menus. Refer to *Setting Up the DAS 9200* in Section 2 for a brief overview on setting up the menus.

Setting Up Single Event

After defining your setups, from the Menu Selection overlay, move the cursor to the Applications column and select PA-Single Ev. Select function key F7: MOVE TO APPLICTN to enter the Single Event application.

Figure 4-1 shows the default Single Event menu; details of certain fields are described below the menu. At this point you can either start the PA software or you can specify other parameters in the Setup Definition overlay.

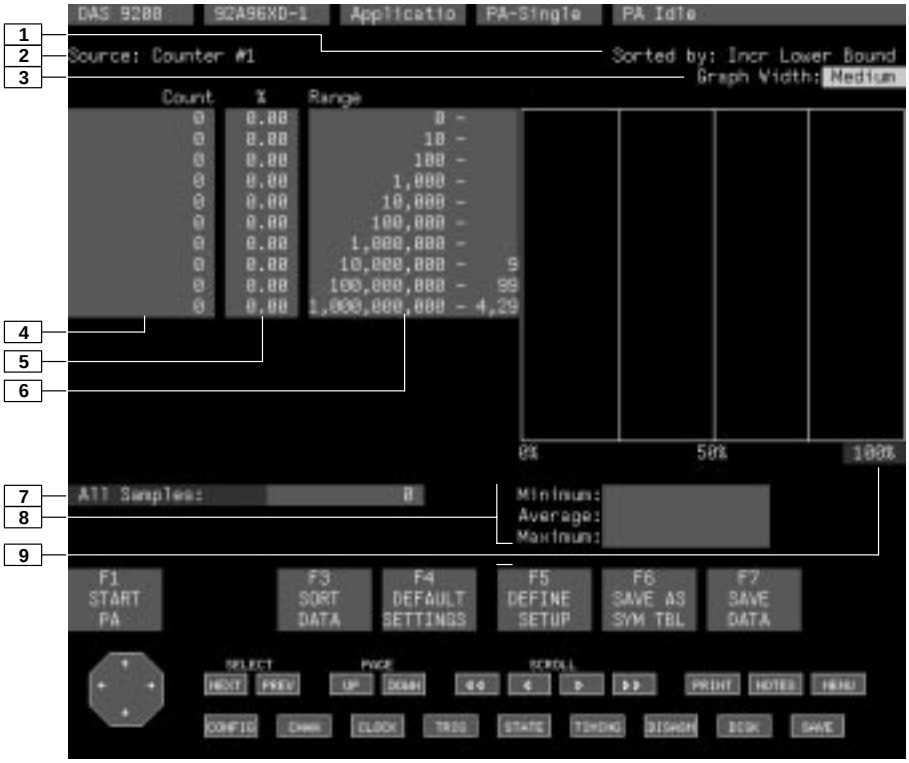

Figure 4-1: Default Single Event menu

1 **Sorted by.** This label shows how the displayed ranges are sorted. Use the F3: SORT DATA function key to specify how to sort the data in the display.

2 **Measurement Source.** This label shows the source of the data. The source of the data is the final value of a counter or timer. If the source is unused in the Trigger menu, the DAS 9200 will display an error message describing the recommended action. Use the Setup Definition overlay to define the measurement source.

3 **Graph Width.** Specifies the width of the histogram: Min (minimum), Small, Medium (default), Large, or Max (maximum).

4 **Count.** Shows the number of samples within each range. Upon first entering the menu, the count values will be zero.

5 **Range Percent.** Lists the percentage of samples within each range. The percentage is calculated by dividing the count by the number of samples (all samples, or selected samples).

6 **Range Names.** Lists the name for each range. The name of each range can be a symbol from the current symbol table, or it can be numeric low and high boundaries. If the measurement source is a timer, an abbreviation for the time unit is added to the end of the range name.

Timer measurements consist of real numbers. 92A96 timers have two digits to the right of the decimal while 92A60/90 timers have one digit to the right of the decimal.

7 **Samples Selection.** Specifies the reference for the data collected. The All Samples selection is a total of all the samples from the acquisition. The Selected Samples selection displays the total number of samples contained in the defined ranges. Changing this field changes the percentage values of the ranges, the length of the histogram bars, and the values of the Minimum, Average, and Maximum fields. The count values for each range remain unchanged.

8 **Minimum, Average, Maximum Readouts.** Displays the minimum, average, and maximum counter or timer values over the set of acquisition cycles. When the Samples Selection field is set to All, all counter or timer values are displayed. When the Samples Selection field is set to Selected, the counter or timer values within the defined ranges will be displayed. Timer values are rounded to the nearest units.

9 **Histogram Magnification.** Specifies the upper value of the histogram. Changing this value changes the magnification of data in the histogram. If you change this value, the label in the center of the histogram changes proportionally.

Function Keys

F1: START PA. Starts the module (or cluster). Any previous range counts will be cleared. When the PA software is running, the function key legend changes to STOP PA. If the counter or timer specified as the measurement source is not used in the Trigger menu, an error message will be displayed and the PA will not start.

F2: CONTINUE PA. Resumes the PA acquisition cycles. This function key does not appear in the default menu; it appears when you stop the PA. Any existing range counts are preserved as if the PA software had never stopped.

F3: SORT DATA. Displays a pop-up window allowing you to choose how you want to sort the displayed data. Valid selections are: Sort by Increasing Lower Bound, Sort by Decreasing Lower Bound, Sort by Decreasing Count, and Sort by Increasing Alphanumeric Range Names (only valid for ranges defined in a symbol table). Select your choice and press either the Return key or the Open/Close key. The changes are immediately applied to the menu and the Sorted By field. See Section 5 for more information on sorting data.

F4: DEFAULT SETTINGS. Changes all fields to their power-on default settings (Log10 range generation is used over the entire domain of Counter #1). You will be prompted to confirm your choice.

F5: DEFINE SETUP. Calls the Setup Definition overlay.

F6: SAVE AS SYM TBL. Saves the current range settings as a symbol table in the Symbol Table directory on the DAS 9200's hard disk. You will be prompted for a file name; the file name can consist of up to 11 alphanumeric characters. If the file already exists, you will be prompted for confirmation to overwrite the file. You can access the file with the Symbol Table editor, the Disk Services menus, or any of the DAS 9200's file-transfer tools. You can edit the new symbol table with the DAS 9200's Symbol Table editor and then select the symbol table in the Setup Definition overlay.

F7: SAVE DATA. Saves the acquired PA data in an ASCII file in the PA_SE_Data directory on the DAS 9200's hard disk. You will be prompted for a file name; the file name can consist of up to 11 alphanumeric characters. If the file already exists, you will be prompted for confirmation to overwrite the file. The new file will be saved as ASCII and you can access the file with the Disk Services menus or any of the DAS 9200's file-transfer tools. Appendix A of this manual contains information on the format and contents of the file.

Although you can start the performance analysis upon entering the menu, in most cases you will want to go to the Setup Definition overlay to further define your setup. Select F5: DEFINE SETUP to enter the overlay.

Setup Definition Overlay

The Setup Definition overlay lets you define the range boundaries (bounds). You can specify a symbol table that has the boundaries predefined, you can use Linear Generation to create equally sized ranges, or you can use Log10 Generation to create logarithmically sized ranges. Log10 Generation is the default setup when you first enter the overlay.

Log10 Generation Setups

The PA software uses the Log10 Generation mode to generate logarithmically sized ranges based on the total number of ranges and on the maximum and minimum boundaries. The default setup for the Named Ranges application is Log10 Generation.

Figure 4-2 shows the Setup Definition overlay when you use Log10 Generation. The default measurement source is Counter #1. Figure 4-2 shows the measurement souce as a timer (with the additional Units field).

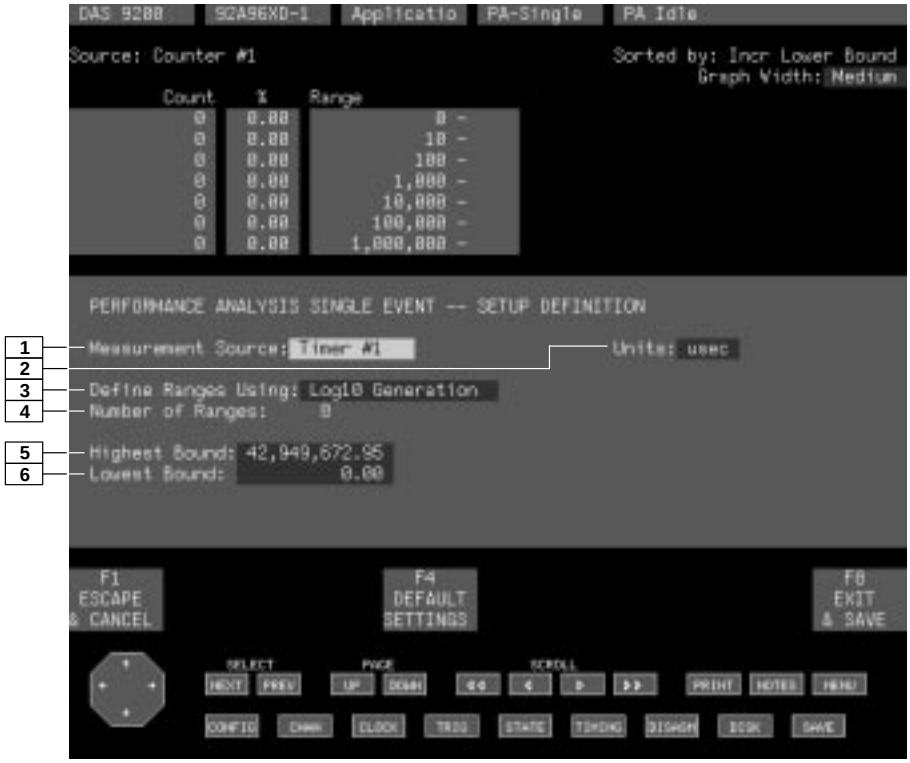

Figure 4-2: Setup Definition overlay (Log10 Generation)

1 **Measurement Source**. Specifies the name of the counter or timer to use as the measurement source. A counter is the default measurement source. If you specify a timer, a Units field will be displayed on the right side of the overlay.

2 **Units.** Specifies the time units for the ranges and the Minimum, Average, and Maximum readout fields. If you select msec or sec, the least significant digits do not appear. When generating ranges, these digits are zeros for a low bound and nines for a high bound. This field is only displayed when a timer is specified as the measurement source. The default unit is µsec.

3 **Define Ranges Using.** Specifies whether the ranges will be defined by Log10 Generation, based on values in a symbol table, or by Linear Generation. The Symbol Table selection is only available if a legal range symbol table exists. Log10 Generation is the default selection.

4 **Number of Ranges.** Specifies the number of ranges used. When using Log10 Generation or Symbol Table, this field is a label only. For Log10 and Linear Generation, the number of ranges depend on the Highest Bound and Lowest Bound values; the closer the Lowest Bound value to the Highest Bound value, the smaller the number of ranges.

5 **Highest Bound.** Specifies the highest (upper) boundary of the ranges. For Log10 Generation or Linear Generation, enter the value of the highest boundary you want to use (this value must be higher than the Lowest Bound value).

6 **Lowest Bound.** Specifies the lowest (smallest) boundary for the ranges. For Log10 Generation or Linear Generation, enter the value of the lowest boundary you want to use (this value must be lower than the Highest Bound value).

Function Keys

F1: ESCAPE & CANCEL. Leaves the overlay without changing the range definitions. The main menu is unaffected when you leave the overlay by selecting this key.

F4: DEFAULT SETTINGS. Changes all fields to their power-on default settings (Log10 range generation is used over the entire domain of Counter #1). The changes will only be applied when you select the F8: EXIT & SAVE function key.

F8: EXIT & SAVE. Leaves the overlay and applies all changes made in the overlay. If there were any changes, the range counts, range percentages, and samples selection (All or Selected) will be set to zero.

Symbol Table Setups

When you use a symbol table with performance analysis, the Setup Definition overlay changes from the Log10 Generation mode to Symbol Table mode as shown in Figure 4-3. Most of the fields in the overlay function the same way as with Log10 Generation. The differences are described after Figure 4-3.

Specifying a symbol table in the overlay changes some of the information (fields) to the Single Event menu; more on this later.

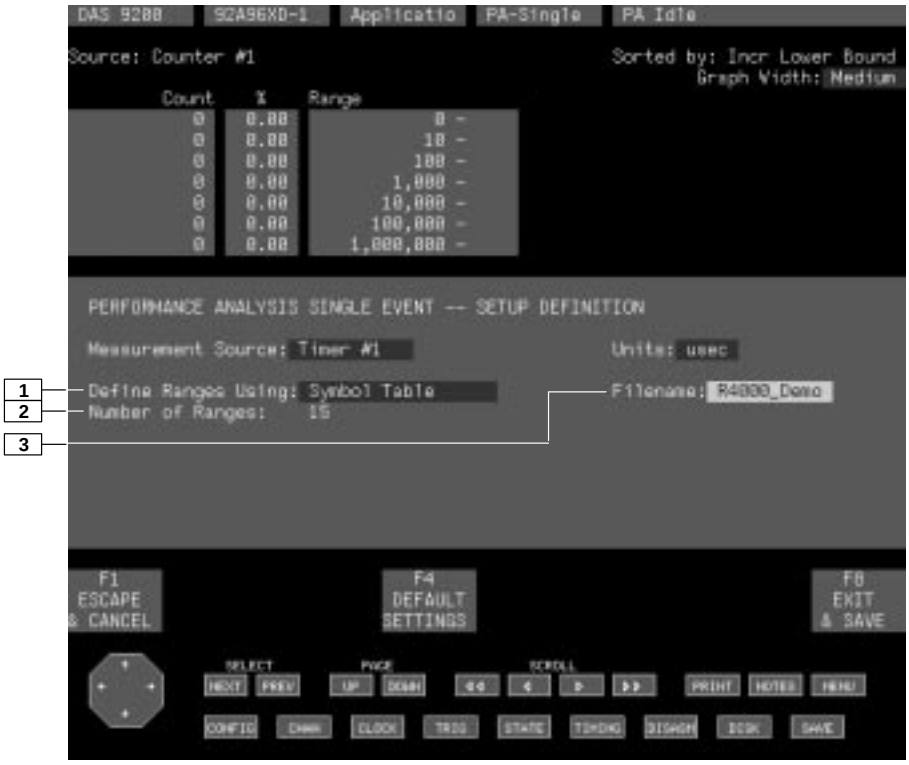

Figure 4-3: Setup Definition overlay using symbol tables

1 **Define Ranges Using: Symbol Table.** Specifies that the ranges will be defined by values in a symbol table. This selection is only available if a legal symbol table exists. The Filename field appears to the right allowing you to specify the name of a symbol table.

2 **Number of Ranges.** The number of ranges depends on the number of symbol table ranges (at least one) in the symbol table.

3 **Filename.** Specifies the name of the range symbol file (pattern symbol tables are not supported with the PA software). The file must reside in the Symbol_Table directory on the DAS 9200's hard disk. To display a list of valid (legal) file names, open the field and select a file name.

When using symbol tables with a 92A96 Module's timer, the symbol table bounds are scaled to 0.01 μs regardless of the units selection. When used with a a 92A60 or 92A90 Module, the bounds are scaled to 0.1 μs. Table 4–1 summarizes the scaling used with each of the modules.

Table 4–1: Symbol Tables and Timer Scaling

Symbol Table Bound	92A60/90 Module Timer	92A96 Module Timer
0	0.0 µs	0.00 µs
1	0.1 µs	0.01 µs
9	0.9 µs	0.09 µs
10	1.0 µs	0.10 µs

Linear Generation Setups

Besides displaying Single Event data using Log10 Generation or Symbol Tables, you can also display data using Linear Generation. Figure 4-4 shows the Setup Definition overlay when you select Linear Generation.

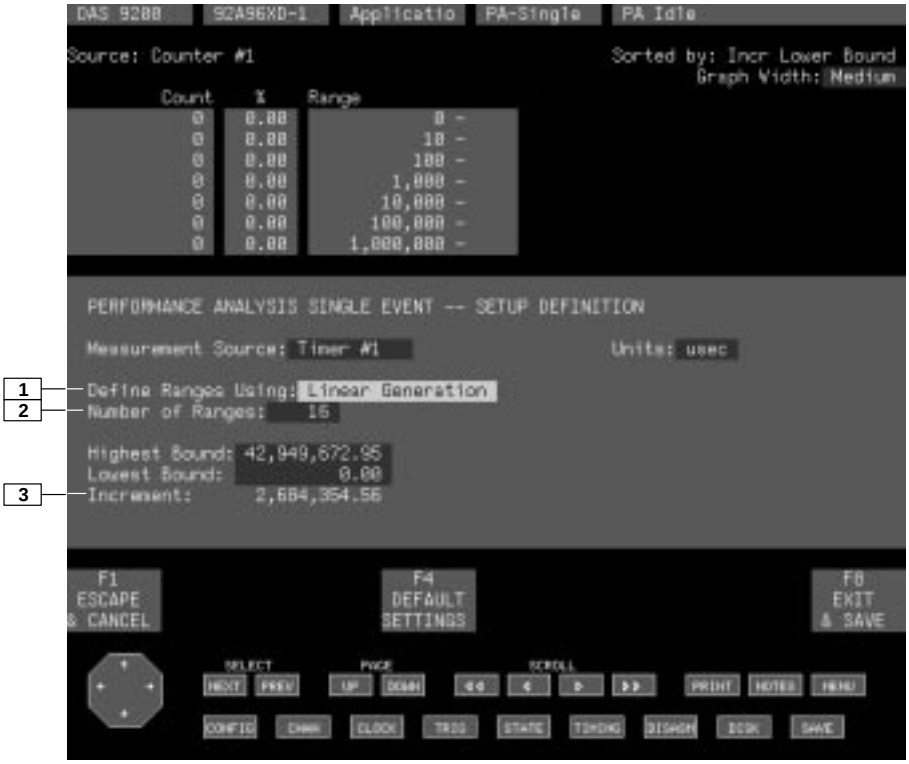

Figure 4-4: Setup Definition overlay using Linear Generation

1 **Define Ranges Using: Linear Generation.** Specifies that the ranges will be defined by values in Linear Generation.

2 **Number of Ranges.** Specifies the number of ranges used. The default value is 16 for Linear Generation. You can specify any number of ranges from 1 up to 5000.

3 **Increment.** Reports the size of the range in Linear Generation. The PA software will try to keep the individual range sizes the same (linear) over the entire set of defined ranges. If the span defined by the Highest Bound and Lowest Bound fields is not evenly divisible by the number of ranges, then the first "n" ranges (n = remainder) will be one larger than the increment.

After defining your setups, close the overlay by selecting F8: EXIT & SAVE. The setups are immediately applied to the Named Ranges menu.

Be aware that any changes you make to the overlay will clear or flush any data in the Named Ranges menu. In other words, the PA values will be reset.

Acquiring Data

Before starting to acquire data, you may want to consider how you want to make the best use of the information on the screen. You can fine-tune the information in the display as follows:

- Sort the displayed data by selecting F3: SORT DATA.
- Adjust the width of the graph by changing the value of the Graph Width field. You can either select a value for the Graph Width field, or use the Select Next or Prev keys (with the cursor in the Graph Width field) to cycle through the selections.
- Change the scale of the graph by changing the value of the Magnification field.
- Scroll to ranges that are off the screen using the Scroll keys, scroll buttons, or the Joydisk; use the Scroll and Shift keys together to scroll a half page (window) of data. Arrows at the top or bottom of the left side of the Count column indicate that there are more ranges off the screen.
- Change from viewing all data samples to selected samples and vice versa.
- Change from viewing the symbol table ranges from symbolic names to bound values and vice versa.

You can make changes to the displayed information at any time before, during, or after making PA acquisition cycles. Section 5 of this manual provides some examples of adjusting the menu.

To start the performance analysis, select F1: START PA. The DAS 9200 will begin acquiring data and display the results on the screen. The status field in the upper-right corner of the screen indicates that the PA software is running.

Figure 4-5 shows an example of the Single Event menu while the PA software is acquiring cycles.

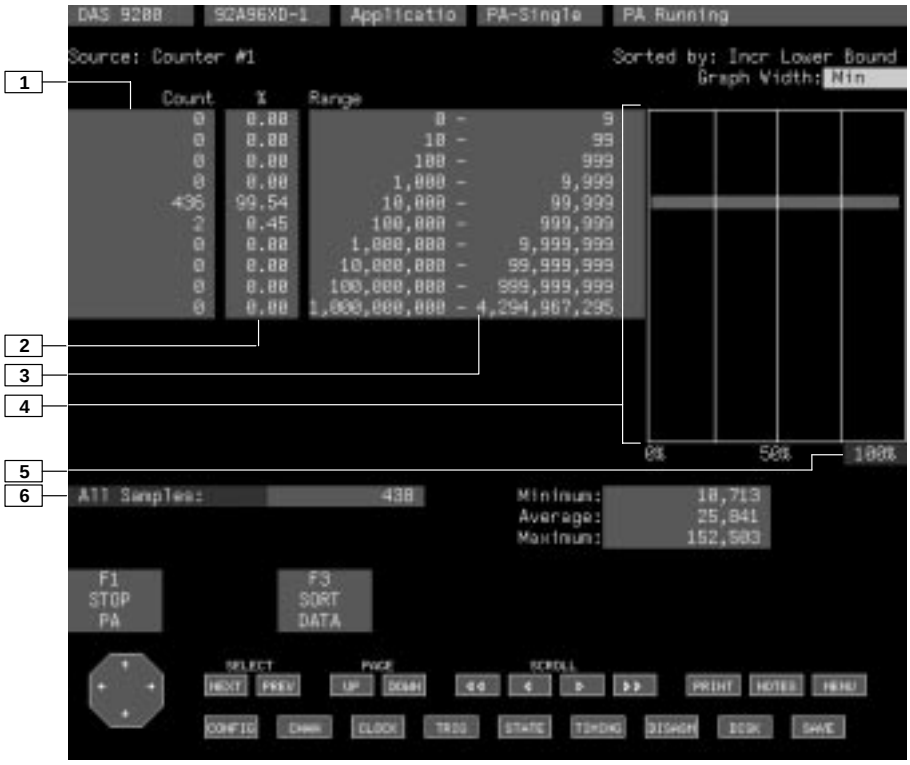

Figure 4-5: Performance analysis data (Log10 Generation)

1 **Count.** Shows the number of samples within each range.

2 **Range Percent.** Shows the percentage of samples within each range (relative to the Samples Selection field).

3 **Range Names.** Shows the names or boundaries of each range.

4 **Histogram Data.** Graphically shows the acquired data in a histogram; one bar for each range. To view ranges with smaller percentage values, you can magnify the display by changing the value of the Histogram Magnification field.

5 **Histogram Magnification.** Specifies the upper value of the histogram. Changing this value changes the magnification of data in the histogram. If you change this value, the label in the center of the histogram changes proportionally.

6 **Samples Selection.** Specifies the reference for the data collected. The All Samples selection is a total of all the samples from the acquisition. The Selected Samples selection displays the total number of samples contained in the defined ranges. Changing this field changes the percentage values of the

ranges and the length of the histogram bars; the count values for each range remain unchanged.

Ranges Defined As Symbol Tables

If you acquire data with the ranges defined from a symbol table, a few other fields are available in the Named Ranges menu. Figure 4-6 shows these fields. In this example the Symbol field and the Show field are only displayed when the range definitions were loaded from a range symbol table file in the Setup Definition overlay.

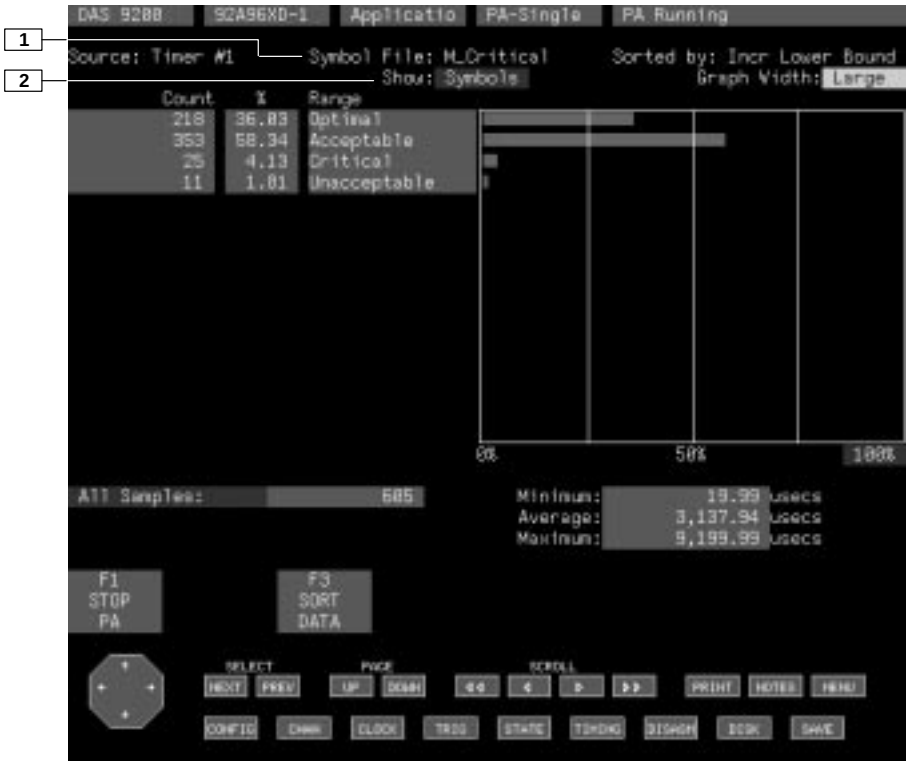

Figure 4-6: Performance analysis data (Symbol Table)

1 **Symbol File.** Specifies the name of the current symbol table.

2 **Show.** Specifies whether the ranges will be displayed as either Symbols or Bounds. If you view the ranges as bounds, the bounds will be displayed in decimal.

After you have acquired data, you can analyze the data in the menu. You can also save the data in an ASCII file and transfer it to a host computer for further analysis. Refer to Section 5 for more application information. Refer to Appendix A for information on the details of the contents of the PA data file.

Application Information

This section provides application information common to the Named Ranges and Single Event applications. This section explains how you can optimize the menus for displaying PA data, save the PA data to a file and transfer it to a host computer for more detailed analysis, and how to resolve some of the problems that you may encounter while using the applications.

Optimizing the Menus

The PA software provides you with the flexibility of adjusting the appearance of the Named Ranges and the Single Event menus to best suit your needs. You can make these adjustments while you acquire data as well as before or after you acquire data.

The PA software allows you to optimize the menus in the following ways:

- Sort the displayed data
- Change the width of the histogram
- Magnify the displayed data
- Display selected data samples

Sorting Displayed Data

The PA software provides different ways of sorting displayed data. To sort the displayed data, select F3: SORT DATA, select one of the sort methods, and close the sort pop-up window.

The PA software will sort the ranges as soon as you close the sort selection window. If the PA software is running and you sort by the decreasing count values, the ranges will be sorted again when you explicitly stop the PA software.

Sort by Incr Lower Bound. Ranges are sorted based on increasing low bound range values. This is the default sort method. When you define the ranges from a symbol table, you can view the data as bounds by changing the Show field to Bounds.

Sort by Decr Lower Bound. Ranges are sorted based on decreasing low bound range values. When you define the ranges from a symbol table, you can view the data as range bounds by changing the Show field to Bounds.

Sort by Decreasing Count. Ranges are sorted based on decreasing count values. Using this sort method, the range where the code spends most of its time will appear at the top of the histogram.

When the PA software is running, the ranges will be sorted whenever you reselect the Sort by Decreasing Count method using the F3 function key. When the PA software is stopped (idle), the ranges are always sorted by the last selected sort method.

Depending on how many ranges are defined, unused ranges will be displayed either off the screen or near the bottom of the display. Use the scroll keys to access these ranges.

Figure 5-1 shows an example of sorting ranges by decreasing range counts.

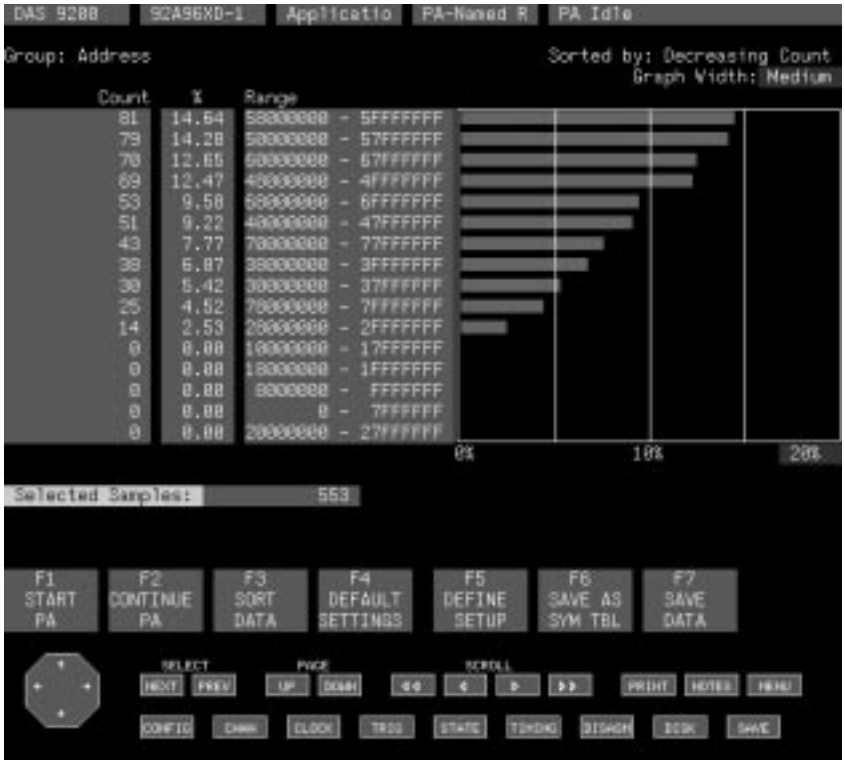

Figure 5-1: PA data sorted by decreasing range counts

Sort by Increasing Alpha. This selection is only available when you define the ranges from a symbol table file. Ranges are sorted by their symbolic names in alphanumeric order.

Changing the Histogram Width

The PA software allows you to change the width of a histogram while you view the data. The width of a histogram is controlled by the Graph Width field. The wider your histogram, the narrower the range field.

Magnifying the Displayed Data

In addition to varying the width of a histogram, you can also change its magnification. Use the Magnification field to specify the upper value of the histogram display. You can display 100% of the data or as little as 1%.

By changing the Magnification field, you can graphically view the difference between two ranges.

Displaying Selected Samples

The PA software lets you scale your range counts against either the sum of all the range counts (Selected Samples) or against the total number of samples acquired, regardless of how they fell into the ranges. Displaying Selected Samples vs. All Samples affects the percentage lists and the length of the histogram bars; range counts are unaffected.

Figure 5-2 shows an example of Named Ranges data showing only the selected samples. Figure 5-3 shows similar information except that all samples are included; note the difference in the range percentages and the lengths of the histogram bars. In order to view the data on the screen, the magnification field in Figure 5-3 is set to 20%.

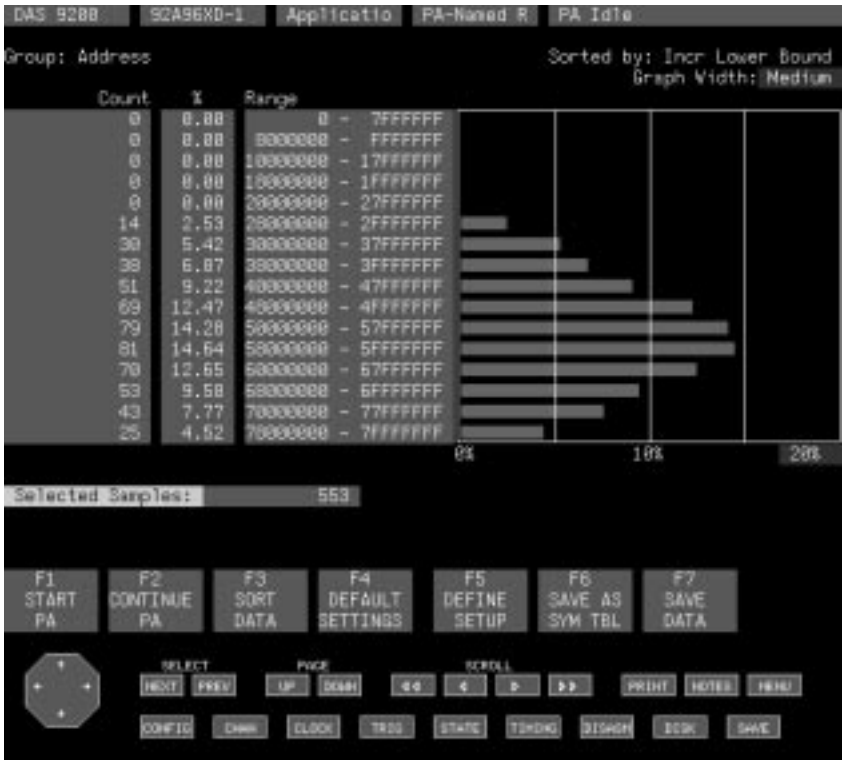

Figure 5-2: Named Ranges data with Selected Samples

For example, your acquisition module may acquire more samples that fall outside of the defined ranges. When you display the ranges with All Samples, the range percentages appear very small (see Figure 5-3). When you display Selected Samples, you only focus on those samples within the defined ranges. Therefore, the percentages and the lengths of the histogram bars will change as shown in Figure 5-2.

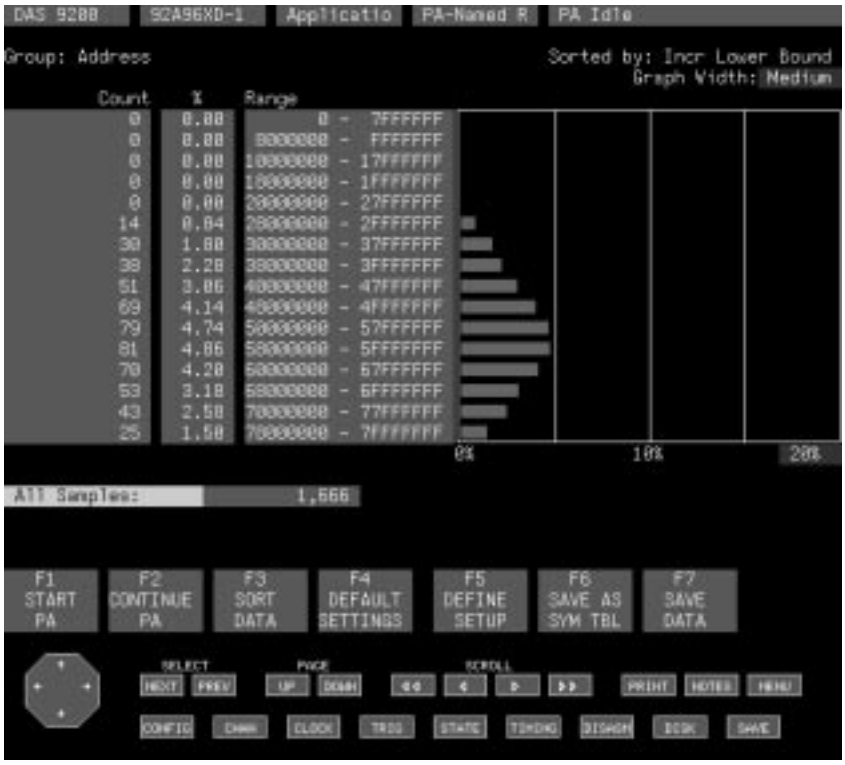

Figure 5-3: Named Ranges data with All Samples

Saving Performance Analysis Information

The PA software provides different ways of saving PA information for either future use or for further analysis. You can save the range definitions as a symbol table file for immediate use with the DAS 9200's Symbol Table Editor or with other DAS 9200 applications. You can also save the PA data in an ASCII file and transfer the file to a host computer to do additional analysis.

Saving the Range Definitions as a Symbol Table

There may be times when you want to save the current ranges as a new or separate symbol table. There may also be times when you want to edit a copy of the current symbol table to fine-tune it for your application.

The PA software provides a means to save the ranges in a symbol table with the F6: SAVE AS SYM TBL function key. Press the function key and enter a file name at the prompt. You can then use the Symbol Table Editor to edit the new symbol table. You can also transfer the new symbol table to a host computer for editing with a text editor, and then return the file to the DAS 9200.

To use the new symbol table with the PA software, you must select it in the Setup Definition overlay.

Be aware that if you change the current, selected, symbol table (using the F6: SAVE AS SYM TBL function key and the DAS 9200's Symbol Table Editor, or one of the DAS 9200's file transfer tools), the PA software will rebuild the ranges from the updated symbol table and clear all range counts.

Saving Data for Analysis on a Host Computer

Sometimes you may want to do further analysis of the data on a host computer or print the PA data. The PA software allows you to save the resultant data in a file and transfer it to the host computer.

After acquiring data, stop the PA software and press F7: SAVE DATA. You will be prompted for a file name to save the data under.

The data is stored as an ASCII file. The PA-specific information such as range counts and percentages are written as comments within the file. All defined ranges are saved to the file (including ranges that are scrolled off the display).

The format of the file allows you to use it as an ASCII range symbol table file. After editing the file on the host computer, you can transfer the file back to the DAS 9200, store it in the Symbol_Table directory, and use it as one of the symbol tables.

The PA data is stored in a file in the PA_NR_Data or the PA_SE_Data directory. You can use any of the DAS 9200's file transfer capabilities to transfer the file to the host computer (ftp, 92LANP, PCL, Kermit, DASdisk, etc).

Transferring PA Data to a Host Computer

The following paragraphs provide some examples on using the DAS 9200's file transfer tools. In some cases you can only use these file transfer tools if they are available on your DAS 9200/SE mainframe. For example, you cannot use ftp to transfer files if your DAS 9200/SE mainframe does not contain a LAN (local area network) module.

These examples assume that you have properly set up the hardware and software with each associated file transfer tool. If necessary, refer to the appropriate documentation for information on the hardware and software setups. In each of the examples, the name of the PA data file is "my_data."

FTP Example

Transferring a file from a DAS 9200/SE mainframe to a host computer is relatively easy (assuming you have set up the ftp information in the LAN overlay to the Communications menu). To use ftp, perform the following steps:

1. Log on to your host computer.
2. Access the DAS 9200/SE with ftp (using the name and password, if any, of your DAS 9200 as defined in the LAN overlay.)

3. Use the CWD or CD command to change to the PA_NR_Data or PA_SE_Data directory.

4. Get the file using the GET command:

For more information on using ftp, refer to your LAN module's user manual or to your host computer's documentation.

92LANP Example

The easiest way to transfer the PA data file with the 92LANP (Ethernet) application is to use the pclconnect command with the FILE? query. To use the 92LANP application, perform the following steps:

1. Use the pclconnect command to establish a connection with the DAS 9200/SE mainframe.

2. Use the PCL FILE? query and redirect the output to a file on the host computer. For example:

```
FILE? "PA_NR_Data/my_data" > das1_pa_data
```

Refer to the *92LANP User Manual* for more information on using the 92LANP application. Refer to the *DAS 9200 Programmatic Command Language User Manual* for more information on the FILE query.

PCL Example

The easiest way to transfer the PA data file using GPIB is with the DAS 9200 Programmatic Command Language (PCL). The GPIB syntax used with your host computer depends on your GPIB software. This example only lists the PCL syntax and assumes that you want to transfer the PA_NR_Data file.

1. From the host computer, issue the following PCL command to the DAS 9200:

```
FILE? "PA_NR_Data/my_data"
```

2. Use your GPIB software to copy the contents of the my_data file to another file on the host computer.

Refer to the *DAS 9200 Programmatic Command Language User Manual* for more information on the FILE query.

Kermit Example

The Kermit application is best used to transfer a file over an RS-232C connection. There are different variations of Kermit available; the exact steps that you want to use depend on your version of Kermit.

Use the following general steps when using Kermit to transfer a file from the DAS 9200/SE mainframe:

1. Set up Kermit on the host computer to accept an ASCII file.

2. Set up the DAS 9200 to send a file using Kermit by entering the following information in the DAS 9200's Communications menu:

```
Operation: Send to Host
File Type: PA_NR_Data or PA_SE_Data
File Name: my_data
```

3. Start Kermit from the host computer.

4. Send the file from the DAS 9200 by selecting the F8: SEND FILE function key.

Refer to the *DAS 9200/SE System User Manual* for more information on using Kermit with the DAS 9200.

DASdisk Example

The DASdisk utility lets you copy the contents of a DAS 9200 floppy disk to a PC. To use DASdisk, perform the following steps:

1. Go to the Copy File operation of the Disk Services menu on the DAS 9200 and use the setups in Table 5–1. Copy the file to the floppy disk with the F8: EXECUTE OPERATION function key.

Table 5–1: DASdisk Copy Operation Setups

Field Label	Field Selection
Source Disk	Hard Disk
Source File Type	PA_NR_Data or PA_SE_Data
Source File Name	my_data
Destination Disk	Floppy Disk

2. Start the DASdisk utility on the PC, and install the floppy disk.

3. When using DASdisk for the first time with the 92PA files, use option 11 to create a new file type. Enter the name PA_NR_Data or PA_SE_Data and answer yes to the next two prompts.

4. Use the Read operation (option 3) to read the new file type and copy the contents of the my_data file to the PC.

Refer to the *DAS 9200/SE System User Manual* for more information on using the DASdisk utility.

When Things Go Wrong

This section describes how you can resolve some of the problems that you may encounter while using the PA software.

The PA software is a subset of the powerful tools available with the DAS 9200. When using the PA software with multiple modules in a DAS 9200, some problems may occur that may not be evident while running the PA application. It may be necessary to temporarily leave the PA application and acquire the data using one of the acquisition module's display menus. The DAS 9200 may display status messages in other menus that may not be displayed in any of the PA menus.

The following steps provide a few ideas on solving problems with the SUT or DAS 9200 setups that can occur when using the PA software. You may also want to refer to other DAS 9200 documentation for more help. For example, if you are experiencing disassembly problems, you may want to refer to your disassembler's instruction manual for clues to solving disassembly problems. You may also want to refer to Appendix D of the *DAS 9200/SE System User Manual* for help with DAS 9200 system specific error messages.

1. Temporarily leave the PA application and go to the System Monitor menu. Start the DAS 9200 and observe the System Monitor menu.

The System Monitor menu provides an overview of the entire operation of the DAS 9200. If you use more than one module in a cluster, you can verify that the modules are running (not stopped). You can also view messages displayed to the System Monitor menu that may not be displayed in the PA menus.

If you use clustered modules, the PA software forces the other modules in the cluster to stop when the main PA module has filled with data. The System Monitor menu will show you the status of any clustered modules. For more information on the status of the individual module, you should look at that module's Monitor menu.

2. Verify that your acquisition module's Setup menus are set up properly. If you are using Named Ranges, you may want to ask some of the following questions.

- Is the correct clocking being used?
- Is too much or too little data being qualified out?
- Is the acquisition module triggering on the correct data?
- Is the buffer space being used efficiently?
- Are the correct channels in the address group set up?

- Is the memory depth and the trigger position set up properly for the acquisition module to acquire the desired number of data samples?

If you are using Single Event, you may want to ask some of the following questions.

- Is the correct counter or timer being used?
- Is the counter or timer being incremented in the Trigger menu?
- Is the timer being correctly stopped in the Trigger menu?

If the DAS 9200 does not trigger, you can view the status of your acquisition module with the module's Monitor Menu. If necessary, refer to your module's user manual for information on the Monitor menu.

3. If you are using the PA software with one of the microprocessor disassemblers, you should verify that caching, prefetching, and bursting do not cause a false address that may not be executed to come across the address bus. Refer to your microprocessor disassembler user manual for more information.

4. Verify that your DAS 9200 is properly connected to your system under test. Check for loose or faulty probe connections.

Appendix A: PA Data Files

The PA software allows you to save PA data in one of two files, the PA_NR_Data file or the PA_SE_Data file. The files are nearly identical. The PA_NR_Data file is created from the Named Ranges application, and the PA_SE_Data file is created from the Single Event application. The files are created from the respective applications with the F7: SAVE DATA function key.

The main purpose of the files is to provide a means to transfer the PA data to a host computer for printing or for further analysis. The files contain information on the PA setups as well as the PA data. The file is saved in ASCII format in the PA_NR_Data or PA_SE_Data directory on the DAS 9200's hard disk.

You can edit the ranges in the range symbol table section and then download the entire file to the DAS 9200 to be used as a range symbol table (assuming that you place it in the Symbol_Table directory on the DAS 9200).

Figure A-1 shows an example of a PA Named Ranges data file. Figure A-2 shows an example of a PA Single Event data file. The actual data in the files is 110 characters wide; the data in Figures A-1 and A-2 is condensed to display on a page.

The PA data files are divided into two main sections, the PA Header section and the Range Symbol Table section. Both of these sections are described in more detail in the following paragraphs.

```
#~ FILETYPE:               PA_NR_SAVE_DATA
#~ VERSION:                1.0
#~ DATE/TIME:              Thu Apr  8 10:31:44 1993
#~
#~ MODULE_NAME:            92A96XD-1
#~ GROUP_NAME:             Address
#~ GROUP_WIDTH:            32
#~ GROUP_CHANNELS:         A3:76543210   A2:76543210   A1:76543210   A0:76543210
#~ CHANNEL_POLARITY:       A3:++++++++   A2:++++++++   A1:++++++++   A0:++++++++
#~

RANGE HEX HEX 0

#~ DEF_USING:              Linear
#~ SYMTAB_NAME:            <none>
#~ LOW_BOUND:              0
#~ HIGH_BOUND:             FFFFFFFF
#~ INCREMENT:              10000000
#~ TOTAL_RANGES:           16
#~
#~ SORT_METHOD:            Incr_low_bound
#~ COUNT_RADIX:            Dec
#~ SELECTED_SAMPLES:       13566
#~ ALL_SAMPLES:            13566
#~
#~ PA_SYMTAB_SECTION_FOLLOWS:
#~
#~            Range       Low_Bound   High_Bound           Count     %_Sel        %_All
#~            -----       ----------   ----------          -----     -----        -----
"       0 -  FFFFFFF"             0    FFFFFFF  #~           247      1.82         1.82
"10000000 - 1FFFFFFF"      10000000   1FFFFFFF  #~           115      0.84         0.84
"20000000 - 2FFFFFFF"      20000000   2FFFFFFF  #~           212      1.56         1.56
"30000000 - 3FFFFFFF"      30000000   3FFFFFFF  #~           345      2.54         2.54
"40000000 - 4FFFFFFF"      40000000   4FFFFFFF  #~           436      3.21         3.21
"50000000 - 5FFFFFFF"      50000000   5FFFFFFF  #~           529      3.89         3.89
"60000000 - 6FFFFFFF"      60000000   6FFFFFFF  #~           686      5.05         5.05
"70000000 - 7FFFFFFF"      70000000   7FFFFFFF  #~           799      5.88         5.88
"80000000 - 8FFFFFFF"      80000000   8FFFFFFF  #~           901      6.64         6.64
"90000000 - 9FFFFFFF"      90000000   9FFFFFFF  #~           993      7.31         7.31
"A0000000 - AFFFFFFF"      A0000000   AFFFFFFF  #~          1084      7.99         7.99
"B0000000 - BFFFFFFF"      B0000000   BFFFFFFF  #~          1185      8.73         8.73
"C0000000 - CFFFFFFF"      C0000000   CFFFFFFF  #~          1289      9.50         9.50
"D0000000 - DFFFFFFF"      D0000000   DFFFFFFF  #~          1449     10.68        10.68
"E0000000 - EFFFFFFF"      E0000000   EFFFFFFF  #~          1591     11.72        11.72
"F0000000 - FFFFFFFF"      F0000000   FFFFFFFF  #~          1705     12.56        12.56
```

Figure A-1: PA Named Ranges ASCII data file

```
#~ FILETYPE:                PA_SE_SAVE_DATA
#~ VERSION:                 2.0
#~ DATE/TIME:               Wed Mar 31 13:40:07 1993
#~
#~ MODULE_NAME:             92A90-1
#~ MEASUREMENT_SOURCE:      Timer #1
#~

RANGE DEC DEC 0

#~ UNITS:                   Usec
#~ DEF_USING:               Log
#~ SYMTAB_NAME:             <none>
#~ LOW_BOUND:               0.0
#~ HIGH_BOUND:              1677721.5
#~ INCREMENT:               0.0
#~ TOTAL_RANGES:            7
#~
#~ SORT_METHOD:             Incr_low_bound
#~ COUNT_RADIX:             Dec
#~ SELECTED_SAMPLES:        100
#~ SELECTED_MAXIMUM:        9998.9
#~ SELECTED_AVERAGE:        9998.9
#~ SELECTED_MINIMUM:        9998.9
#~ ALL_SAMPLES:             100
#~ ALL_MAXIMUM:             9998.9
#~ ALL_AVERAGE:             9998.9
#~ ALL_MINIMUM:             9998.9
#~
#~ PA_SYMTAB_SECTION_FOLLOWS:
#~
#~  Range                         Low_Bound High_Bound               Count        %_Sel         %_All
#~  -----                         --------- ----------               -----        -----         -----
"        0.0 -         9.9 us"          0         99  #~                 0         0.00          0.00
"       10.0 -        99.9 us"        100        999  #~                 0         0.00          0.00
"      100.0 -       999.9 us"       1000       9999  #~                 0         0.00          0.00
"    1,000.0 -     9,999.9 us"      10000      99999  #~               100       100.00        100.00
"   10,000.0 -    99,999.9 us"     100000     999999  #~                 0         0.00          0.00
"  100,000.0 -   999,999.9 us"    1000000    9999999  #~                 0         0.00          0.00
"1,000,000.0 - 1,677,721.5 us"   10000000   16777215  #~                 0         0.00          0.00
```

Figure A-2: PA Single Event ASCII data file

PA Header Section

The PA header section of the data file represents all the PA data values excluding the PA range definitions and data. Each line of information consists of a comment with a key word (label) and a value.

Comments generated by the PA software begin with a pound sign and a tilde sign (#~). By using comments in the PA header in addition to the range symbol table data, all information is available within a single file.

Host computer applications can also treat the entire file as a range symbol table file while ignoring the PA header section with the comments.

Table A–1 describes the contents of the PA header section of a Named Ranges file. Table A–2 describes the contents of the PA Header section of a Single Event file. Each line in the PA Header section begins with a comment.

The last line in the PA header section separates the Range Symbol Table section from the PA header section.

Table A–1: PA Named Ranges Data File Description

Label	Description
FILETYPE:	Specifies the file type. This value will always be PA_NR_SAVE_DATA for Named Ranges data files
VERSION:	Version number of the file
DATE/TIME:	Date and time the file was created
MODULE_NAME:	Name of the data acquisition module the file was created under
GROUP_NAME:	Name of the address group
GROUP_WIDTH:	Width of the address group
GROUP_CHANNELS:	Section number and the address channels used in the group
CHANNEL_POLARITY:	Specifies the polarity of each channel in the group. A plus sign indicates positive polarity, and a minus sign indicates negative polarity. A period indicates an unused channel
RANGE:	Indentifies the files as a range symbol table (for the occasions when you want to use the file as a symbol table). The radixes match the radix used with the PA software; the offset is always 0
DEF_USING:	Specifies whether the PA software was set up using symbols or Linear Generation
SYMTAB_NAME:	Name of the symbol table file. No name will be listed if the file was saved under Linear Generation
LOW_BOUND:	Defines the low bound when used with Linear Generation. The low bound value will be zero when symbol tables are used

Table A–1: PA Named Ranges Data File Description (Cont.)

Label	Description
HIGH_BOUND:	Defines the high bound when used with Linear Generation. The high bound value will be zero when symbol tables are used
INCREMENT:	Defines the size of a range when used with Linear Generation. This value will be zero when symbol tables are used
TOTAL_RANGES:	Total number of ranges
SORT_METHOD:	Current sorting method
COUNT_RADIX:	Radix of the count; this value is always decimal
SELECTED_SAMPLES:	Number of samples acquired within the defined ranges
ALL_SAMPLES:	Total number of samples acquired
PA_SYMTAB_SECTION_ FOLLOWS	This key word indicates that the information after the comments in the file contains the actual range symbol table data

Table A–2: PA Single Event Data File Description

Label	Description
FILETYPE:	Specifies the file type. This value will always be PA_SE_SAVE_DATA for Single Event data files
VERSION:	Version number of the file
DATE/TIME:	Date and time the file was created
MODULE_NAME:	Name of the data acquisition module the file was created under
MEASUREMENT_ SOURCE:	Specifies what counter or timer used as the measurement source
RANGE:	Indentifies the files as a range symbol table (for the occasions when you want to use the file as a symbol table). The radixes be decimal; the offset is always 0
UNITS:	Units used for timers (Usec, Msec, or Sec) when used as the measurement source. Counters have no units
DEF_USING:	Specifies whether the PA software was set up using Log10 Generation, symbol tables, or Linear Generation
SYMTAB_NAME:	Name of the symbol table file. No name will be listed if the file was saved under Log10 or Linear Generation
LOW_BOUND:	Defines the low bound when used with Log10 or Linear Generation. The low bound value will be 0 with symbol tables
HIGH_BOUND:	Defines the high bound when used with Log10 or Linear Generation. The high bound value will be 0 with symbol tables
INCREMENT:	Defines the size of a range when used with Linear Generation. This value will be 0 with Log10 Generation or symbol tables

Table A–2: PA Single Event Data File Description (Cont.)

Label	Description
TOTAL_RANGES:	Total number of ranges
SORT_METHOD:	Current sorting method
COUNT_RADIX:	The radix of the count; this value is always decimal
SELECTED_SAMPLES:	The number of samples acquired within the defined ranges
SELECTED_MAXIMUM:	Maximum sample value acquired within the defined ranges. If you use a timer and the time units are Sec or Msec, the value is rounded
SELECTED_AVERAGE:	Average sample value acquired within the defined ranges. If you use a timer and the time units are Sec or Msec, the value is rounded
SELECTED_MINIMUM:	Minimum sample value acquired within the defined ranges. If you use a timer and the time units are Sec or Msec, the value is rounded
ALL_SAMPLES:	Total number of samples acquired
ALL_MAXIMUM:	Maximum sample value acquired. If you use a timer and the time units are Sec or Msec, the value is rounded
ALL_AVERAGE:	Average sample value acquired. If you use a timer and the time units are Sec or Msec, the value is rounded
ALL_MINIMUM:	Minimum sample value acquired. If you use a timer and the time units are Sec or Msec, the value is rounded
PA_SYMTAB_SECTION_ FOLLOWS	This key word indicates that the information after the comments in the file contains the actual range symbol table data

Range Symbol Table Section

The Symbol Table section of the PA data file contains all the information of a legal symbol table. The first two lines in this section contain the following column heads (see Figures A-1 and A-2):

- Range
- Low Bound
- High Bound
- Count
- %_Sel
- %_All

Each subsequent line in the file contains specific information about each range. If the symbol name contains leading or trailing blanks, imbedded double quotes, or a leading pound sign (the comment character), the entire symbol name will be surrounded by double quotes.

Because the PA software encloses range names containing quotes, blanks, and leading pound character signs with double quotes, it is possible to produce extremely long character strings. The net effect of adding the quotes causes the data in the symbol table section not to line up properly. It is therefore a good idea to avoid using spaces, quotes, and pound characters in symbol names.

Notice that the values under the Count, %_Sel, and %_All columns are preceded by the PA comments. This is because these data values are not used with ranges in a symbol file.

Index

Numbers

A

B

C

D

E

F

G

H

I

K

L

M

N

O

P

Q

R

S

T

W