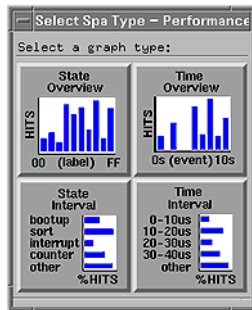


© 1996-2002 Agilent Technologies. All rights reserved.

Toolsets: System Performance Analyzer (Agilent Technologies B4600B)

Using the System Performance Analyzer



The System Performance Analyzer (SPA) helps you analyze, debug, and optimize code based on acquired data. Use the SPA tools to display occurrences of state and timing events in vertical and horizontal distribution formats, and to view detailed information about the acquired events and execution times. All four tools support data acquired in both time and state acquisition modes.

System Performance Analyzer Tools

- “Using State Overview” on page 13
- “Using State Interval” on page 62
- “Using Time Overview” on page 82
- “Using Time Interval” on page 96

See Also

- “Opening and Saving SPA Configurations” on page 10
- “Print the SPA Window” on page 11
- “Print SPA Data to a File” on page 12
- Main System Help (see the *Agilent Technologies 16700A/B-Series Logic Analysis System* help volume)

Contents

Using the System Performance Analyzer

1 Using the System Performance Analyzer

Opening and Saving SPA Configurations 10

Print the SPA Window 11

Print SPA Data to a File 12

Using State Overview 13

Uses of State Overview 14

Defining State Overview Ranges 15

Selecting a Label 15

Defining State Overview Buckets 16

If Bucket Number Exceeds Maximum 16

Bucket Information 17

Why Defining Buckets is Useful 17

Highlighting a Bucket 17

Qualifying State Overview Data 18

Qualify the Event 19

When to Qualify Data 19

Setting Thresholds 21

Threshold Hits and Color 21

Selecting Another Base 22

Running and Stopping a Measurement 23

Analyzing State Overview 24

Y Axis Is Rescaled Constantly 25

Contents

State Overview Information Areas	26
State Overview Display and Bucket Information	27
Working with Graphs	29
Selecting Print Options	30
Closing the SPA Menu	31
Accumulating Data	31
Viewing Acquisition Time	32
Showing Comments	32
Setting a SPA Tool Type	33
Clearing and Replotting a Graph	33
Setting State Overview X-Axis Range	34
Setting Time Overview X-Axis Range	35
Setting State and Time Intervals X-Axis Ranges	36
Changing Attributes	36
Sorting the Data	37
Sorting by Number of Hits	37
Sorting by Name	37
Sorting by Range	38
Getting Help	38
Rescaling Graphs	39
Dragging and Rescaling a Graph	39
Data is Rescaled	40
Options for Rescaling a Graph	41
Expanding a Bucket	42
Placing Markers	43
Resizing a Dialog	43
Use Information Areas in Graphs	45
Message Area	45
Analyze SPA Graphs	46

Contents

Data Selection Criteria	47
Held Values when Using Multiple Datasets	47
State Search Criteria	48
Time Overview Single Event Search Criteria	49
Time Interval and Dual-Event Time Overview Search Criteria	50
Bus Mode Search Criteria	52
Identifying Routines that Take the Most CPU Cycles	56
Storing Values	57
Probing Signals	58
Identifying Exceptions	59
Using State Interval	62
Including Other Data Values	63
Uses of State Interval	64
Defining State Interval Ranges	65
Qualifying State Interval Data	67
Analyzing State Interval	68
Display Percentage Field	70
State Interval Information Areas	71
State Interval Display and Bucket Information	72
Using the Object File Symbol Selector	74
Filtering the Symbol List	75
List of Symbols	75
Selecting Symbols	75
Adding Symbols to the Defined Ranges List	76
To Delete Symbols from the Defined Ranges List	76
Examples of Regular Expressions	76
Selecting an Object File	77

Contents

Zooming in on Potential Software Bottlenecks	78
Characterizing Microprocessor CPU Cycles	80
Using Time Overview	82
Reading a Marker Data Value	83
Uses of Time Overview	83
Defining Time Overview Ranges	85
Qualifying Time Overview Data	86
Qualifying Bus Transitions	87
Defining Time Overview Buckets	89
Time Overview Bucket Information	89
Analyzing Time Overview	91
Y Axis Is Rescaled Constantly	92
Time Overview Information Areas	93
Time Overview Display and Bucket Information	94
Using Time Interval	96
Uses of Time Interval	96
Defining Time Interval Ranges	98
Using Auto Range	99
Including Other Time Ranges	100
Qualifying Time Interval Data	102
Analyzing Time Interval	104
Specifying Time Event Ranges	105

Contents

Time Interval Information Areas	106
Time Interval Display and Bucket Information	107
Characterizing Timing of Interrupt Routines	109

Glossary

Index

Contents

Using the System Performance
Analyzer

Opening and Saving SPA Configurations

SPA tool settings can be saved to a configuration file, along with the tools connected to it, and loaded from a previously saved configuration file.

- Loading Configuration Files (see the *Agilent Technologies 16700A/B-Series Logic Analysis System* help volume)
- Saving Configuration Files (see the *Agilent Technologies 16700A/B-Series Logic Analysis System* help volume)

NOTE:

The *Load Configuration* window can be accessed via File->Load Configuration.

The *Save Configuration* window can be accessed via File->Save Configuration.

Print the SPA Window

The print windows operation lets you print the SPA tool window. Use this operation if you want a hardcopy bitmap file of the currently displayed SPA window.

NOTE:

Only the currently displayed viewing area of the SPA window is printed. If any data or configuration fields appear offscreen, scroll the desired data or configuration fields into the window's viewing area before printing.

1. Optional - configure the Print options (see the *Agilent Technologies 16700A/B-Series Logic Analysis System* help volume).
Print options include print destination, file format type, filename autoincrement, and color/b&w; pixel mapping.
2. In the SPA tool menu bar, select *File*, then select *Print this window*. The print output will be as configured in the Print options in step 1.

See Also

Set Print Options (see the *Agilent Technologies 16700A/B-Series Logic Analysis System* help volume)

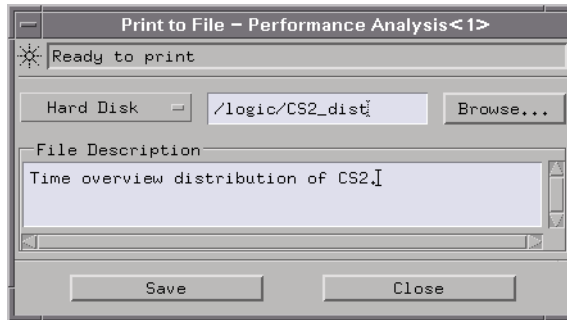
Set Up the Printer (see the *Agilent Technologies 16700A/B-Series Logic Analysis System* help volume)

Print SPA Data to a File


This operation prints the SPA display and bucket information. The resulting file saves information about the currently displayed SPA tool only.

To Print to a File

1. From the SPA tool's menu bar select *File*, then select *Print to file*.
2. Select either *Hard Disk* or *Flexible Disk*.
3. Enter the path of the filename or select *Browse...* button to open the *File Selection* dialog from which to select the path.
4. Optional: enter text in the *File Description* area.
5. Select *Save*. Filenames can only include the following alphanumeric characters:
0-9, a-z, A-Z, (-), (_), (.), (/), (\$), (:), and (+).



Using State Overview



State Overview displays a vertical distribution of qualified states for an analyzer label. State Overview allows you to make high-level state measurements for all occurrences of patterns on a selected label. This tool is most valuable during the discovery phase of performance analysis. Because State Overview measurements show all qualified occurrences, no assumptions are made as to the analysis of the system under test.

By studying the distribution results, you can quickly identify regions of high activity. These regions provide the raw information needed to determine which routines or operations may be responsible for throughput bottlenecks. You can then use this information to begin a productive low-level systems analysis using the State Interval tool.

- Application Examples**
- “Identifying Routines that Take the Most CPU Cycles” on page 56
 - “Identifying Exceptions” on page 59

See Also

- “Defining State Overview Ranges” on page 15
- “Qualifying State Overview Data” on page 18
- “Selecting Another Base” on page 22
- “Running and Stopping a Measurement” on page 23
- “Analyzing State Overview” on page 24
- Uses of State Overview (see page 14)
- “State Overview Information Areas” on page 26
- “Working with Graphs” on page 29
- Data Selection Criteria (see page 47)

“Opening and Saving SPA Configurations” on page 10

“Print the SPA Window” on page 11

“Print SPA Data to a File” on page 12

Uses of State Overview

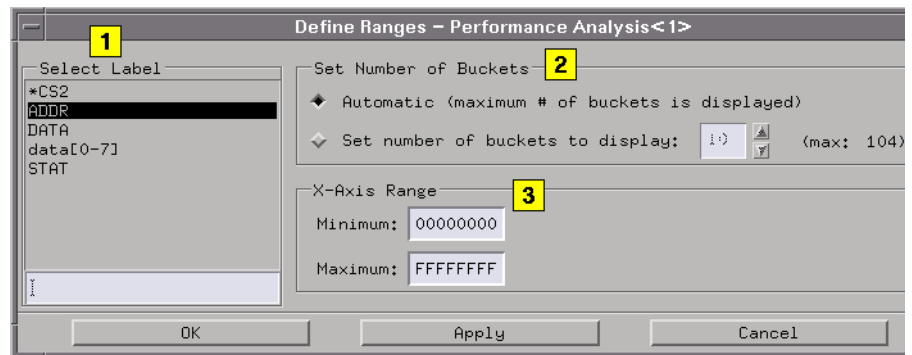
You can use State Overview to:

- Find the address that is initiating a defined event.
- See an overview of system bus activity.
- Display overall system or subsystem usage.
- View the distribution of interrupt types.
- Verify which interrupts are being generated.
- Identify overview distribution of STATus lines.
- Identify routines that take the most CPU cycles.
- Identify exceptions.

Defining State Overview Ranges

Defining ranges will customize the display for a selected label.

1. Select the Label (see page 15) to display.
2. Select the number of buckets (see page 16) to display.
3. Set the X-Axis Range (see page 34).
4. Select *Apply* to apply the changes or *OK* to apply and exit the dialog.



See Also

State Search Criteria (see page 48)

Selecting a Label

Selecting a label in the Define Ranges menu results in the display of qualified data values for the selected label.

1. Select a label in the Select Label area and select *Apply*.
2. Or, enter a label in the Select Label entry area and select *Apply*.

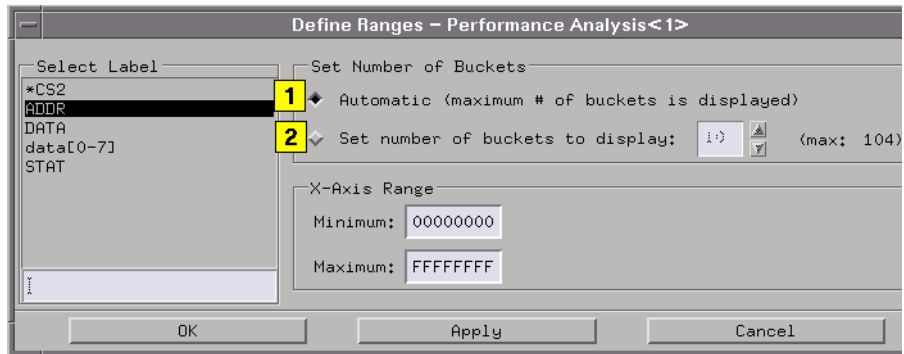
See Also

“Data Selection Criteria” on page 47

Defining State Overview Buckets

You can select a specific number of buckets (see page 17) to hold and display data or have SPA automatically calculate the maximum number of buckets.

1. To display the maximum number of buckets, select *Automatic*.
2. To display another number of buckets (see page 16), select *Set number of buckets to display* and enter a number or select the up and down arrows.
3. Select *Apply* to apply the changes or *OK* to apply and exit the dialog.



See Also

State Overview Bucket Information (see page 17)

State Search Criteria (see page 48)

If Bucket Number Exceeds Maximum

If you enter a number of buckets that is greater than the maximum allowed, your number is overwritten with the maximum bucket count. To see more buckets, rescale the main window.

When resizing the window, if the number of buckets exceeds the maximum number allowed, the number of displayed buckets will be set to the new maximum number of buckets allowed.

Bucket Information

Bucket information is updated for the highlighted bucket in the graph. When you highlight a bucket (see page 17), you will see:

Hits (%d) - displays the number of hits in selected bucket, and percentage of those hits relative to the qualified hits shown on the display.

Hits (%t) - displays number of hits in selected bucket, and percentage of those hits relative to the total qualified hits.

Range - shows low and high data values for selected bucket.

Min/Max - shows minimum and maximum sampled data values in selected bucket.

Why Defining Buckets is Useful

Defining buckets is useful in the interpretation of SPA data results. For example, for an 8-bit label (maximum range 0-F), you may want to specify 16 buckets, so that there is a single data value per bucket. Or, if you are viewing a 100-us time interval, you may want 10 buckets, resulting in each bucket representing 10 us of time.

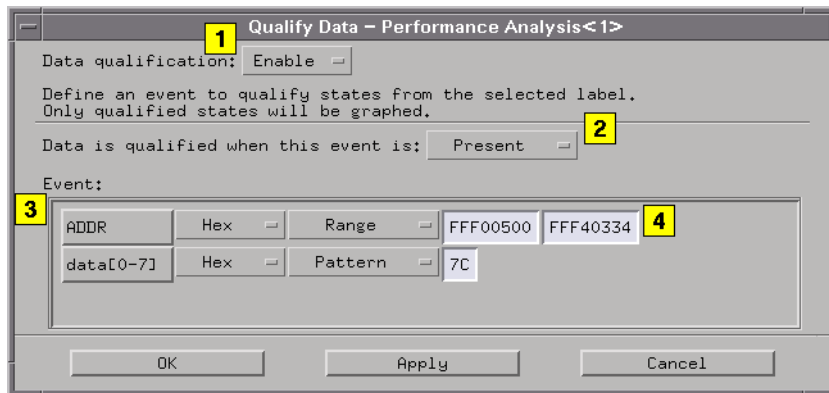
Highlighting a Bucket

When you move the cursor over the buckets in the graph, the highlighted bucket tracks the cursor movement by default. To highlight a specific bucket, select the bucket of interest. This will stick the highlight on the selected bucket.

To return to the default cursor tracking mode, place the cursor over the highlighted bucket and select again. Rescaling and other events that result in the rebucketization of data will return the tool to the default cursor tracking mode.

Qualifying State Overview Data

1. Enable or disable the *Data qualification*. By default, data qualification is disabled.
2. If the data qualification is enabled, select when to qualify the data acquisition (see page 19).
3. Define the event (see page 19) to qualify by selecting the *Label*, *Base*, and *Pattern/Range* buttons.
4. Enter a pattern or range.
5. Insert, delete, or replace as many labels as are needed to qualify the event.
6. Select *Apply* to apply the changes or *OK* to apply and exit the dialog.



Qualifying data will set up the State Search Criteria (see page 48) for the label (see page 15) selected in the Define Ranges (see page 15) menu. For example, if label ADDR is selected in the Define Ranges menu, and a qualifying event is defined in the Define Ranges menu, then only the ADDR data values that met the qualification and event criteria will be displayed (for example, ADDR=0-143 and STATE COUNT=FF). All other ADDR data values will be discarded since they do not meet the qualified event criteria.

NOTE:

Event labels are ANDed to create a single event term.

See Also

Selecting a Label (see page 15)

“Defining State Overview Ranges” on page 15

Qualifying Bus Transitions (see page 87)

Qualify the Event

When the data qualification (see page 19) is disabled, all states or times are qualified, and all of the data is used in the analysis.

When the data qualification is enabled, you can define an event to qualify the states or times. Only values that meet the qualified event criteria are displayed in the graph. The number of hits for a particular label value is determined by the State Search Criteria (see page 48).

When you have multiple labels for data qualification, the labels are ANDed to create a single event condition. The field definitions are:

- *Label Name* - selects Insert Label, Delete Label (when two or more labels are defined, Replace Label, Label Width (Short, Unique, and Full). At least one term is defined at all times.
- *Base Field* - selects Binary, Octal, Hexadecimal, Decimal, ASCII, Symbol, or Twos Complement.
- *Pattern/Range option* - selects Pattern, Range, Not Pattern, and Not Range.
- *Pattern/Range value* - lets you enter text for a search pattern.

You can set these parameters by selecting the field of interest.

See Also

Bus Mode Search Criteria (see page 52)

When to Qualify Data

This field only appears when data qualification is enabled. Data for the selected label is based on when the data is qualified, and the event definition.

- *Entered* - when a transition into the defined event occurs, the data value that existed for the selected label when this transition occurred is used in the distribution hit count.
- *Exited* - when a transition out of the defined event occurs, the data value that existed when this transition occurred is used in the distribution hit count.
- *Not Present* - when the defined event is not present, the data values for the selected label are used in the distribution count.
- *Present* - when the defined event is present, the data values for the selected label are used in the distribution count.
- *Present >* - after the defined event has existed for greater than the specified time, the hit count on the selected label is tracked. If you select this option and enter 50 us in the time field, the first state that occurred on the selected label when the event existed for > 50 us is stored.
- *Present <* - after the defined event has existed for greater than the specified time, the hit count on the selected label is tracked. If you select this option and enter 50 us in the time field, the last state that occurred when the event existed for < 50 us is stored.
- *Present ><* - all states that occur when the event is true within the specified time range are stored in the selected label hit count.

See Also

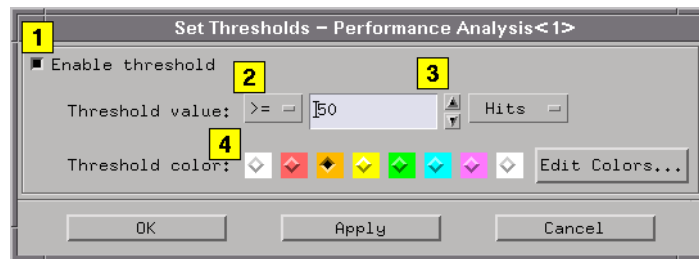
“Data Selection Criteria” on page 47

Setting Thresholds

You can customize the graph by specifying a threshold value (see page 21) . You can specify a number of hits or a percentage of hits for the threshold, and a color.

After seleting *Graph* and *Set thresholds* from the SPA tool menu bar, you are ready to set a threshold value and color. This feature is useful when you want to highlight buckets that have occurrence counts that exceed expected values.

1. Select *Enable threshold*.
2. Select the *Threshold value* option button.
3. Enter number, or select the up and down arrow keys to change the value. Then, select *Hits* (or *%Hits*, depending on the last operation) and select threshold value type.
4. Select a *Threshold Color* button to set a color for data bars that meet the threshold criteria.
5. Select *Apply* to apply the changes or *OK* to apply and exit the dialog.




Threshold Hits and Color

When you select a number of hits for the threshold, any data value bar that meets the threshold criteria will be displayed in the threshold color.

Selecting Another Base

You can set the X-axis base to binary, octal, hexadecimal, and decimal.

1. Select the *Base* button in the Display Information area and select the option you want.




The base selection also updates the selected base for Bucket Information and the X-axis range.

Running and Stopping a Measurement

If you are not changing the run options, simply select the *Run* icon. to run a measurement. To cancel a *Run Repetitive* run, select the *Cancel* field.

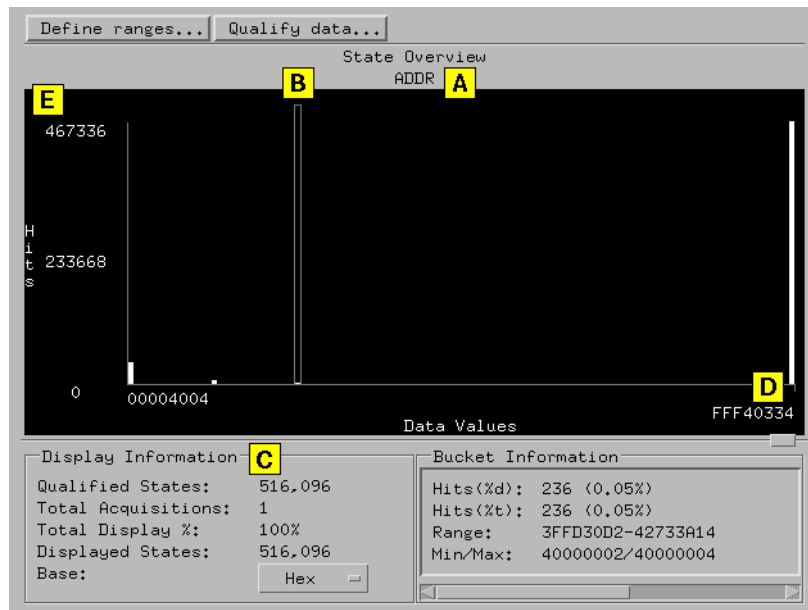
For more information on making measurements using multiple instruments or machines, refer to The Intermodule Window (see the *Agilent Technologies 16700A/B-Series Logic Analysis System* help volume)

Analyzing State Overview



State Overview displays a bar chart of a label's state value versus the relative number of occurrences of each value in the defined range of the label. State Overview is available for any defined label.

How to Interpret the State Overview



- A. The selected label (ADDR) is at the top of the graph.
- B. The highlighted bucket's information is displayed in the Bucket Information (see page 27) area.
- C. The Display Information (see page 27) is updated after every run. The information is based on the qualified data.
- D. The X axis is the defined range for the specified label.

E. The Y axis is the relative number of occurrences in each bucket.

The maximum value of the Y axis is rescaled constantly (see page 25) to reflect the number of occurrences in the bucket that has the most occurrences, and is displayed as the maximum count.

The total number of qualified states sampled for the selected label is presented as the total count, and includes states that may be outside the user-specified X-axis range.

Effects of Defining Ranges

Defining ranges (see page 15) allows you to select the label of interest, the number of buckets, and the X-axis range.

Effects of Qualifying States

Qualified states are the states used in the analysis. You can set up when to qualify data (see page 19) and define a qualifying event term (see page 19) . If you define an event on which to qualify the acquired data, only those qualified states are counted. Those states are then bucketized and displayed.

See Also

“State Overview Information Areas” on page 26

“Working with Graphs” on page 29

“Setting Thresholds” on page 21


Changing Attributes (see page 36)

“Data Selection Criteria” on page 47

Y Axis Is Rescaled Constantly

As successive acquisitions are acquired and qualified events are filtered, the display is constantly rescaled vertically so that the upper limit of the Y axis represents the largest number of occurrences in any bucket.

State Overview Information Areas



State Overview information areas include messages, comments, the data graph, and Display and Bucket Information areas. When you select anywhere in the graph area, a list of rescale choices are presented that allows you to rescale the data on the graph.

Message Area

The message area to the right of the *Run* icon helps you understand interactions allowed as you move the cursor in the entire dialog.

Comment Area

The comment area appears when you select *Options* and *Show Comments* in the menu bar. Use the comment area to document setup information and data results. The size of the comments file is not limited. Comments are saved when you *Print to File*.

Data Graph Area

The data area plots the number of hits for each qualified data value attached to a particular label. Data is plotted based on the defined events, accumulate status, number of buckets, graph attributes, thresholds, rescaling operations, the number of values represented on the X axis.

Display and Bucket Information Areas

The Display and Bucket Information Areas (see page 27) are just below the graph.


See Also

“Rescaling Graphs” on page 39

“Working with Graphs” on page 29

State Overview Display and Bucket Information

State Overview shows display and bucket information for the acquisition. When Accumulate mode is on, data is accumulated in the graph from run to run, and related information is updated.



Display Information

Qualified States - are the number of qualified states in the distribution analysis. If the acquisition is set to collect 8K states, and you define an event on which to qualify the data (see page 19), only those qualified states count. Those states are then bucketized and displayed.

Total Acquisitions - when Accumulate mode is on, this number reflects the number of accumulated data runs. When Accumulate mode is off, the Total Acquisitions field will be 0 if no runs occurred, or 1 if the displayed data represents a single run.

Total Display % - shows the percentage of all data values displayed. If you rescale the graph to show only a portion of the data, this number reflects the new percentage of displayed samples versus the total qualified samples.

Displayed States - displays the total number of states acquired.

Base - lets you select binary, octal, hexadecimal, and decimal. The base is reflected in the Display Information and Bucket Information areas, the X-axis Range, and the X-axis display.

Bucket Information

When you highlight a specific bucket (see page 17), the following information about that bucket displays in the Bucket Information area.

Hits (%d) - displays the number of hits in selected bucket, and percentage of those hits relative to the qualified hits shown on the display.

Hits (%t) - displays the number of hits in selected bucket, and percentage of those hits relative to the total qualified hits.


Range - shows the low and high data values represented in the selected bucket.

Min/Max - shows the minimum and maximum data values that were acquired in the selected bucket.

See Also

“Data Selection Criteria” on page 47

Working with Graphs



Use menu bar *Graph* and *Option* picks to work with the SPA graphs. When an option is not available, it appears in light grey, and you cannot select it.

Graph Picks

- Setting a SPA Tool Type (see page 33)
- Clearing and Replotting a Graph (see page 33)
- “Setting Thresholds” on page 21
- Changing Attributes (see page 36)
- Viewing Acquisition Time (see page 32)
- Setting State Overview X-Axis Range (see page 34)
- Setting Time Overview X-Axis Range (see page 35)
- Setting State and Time Intervals X-Axis Ranges (see page 36)

Option Picks

- Accumulating Data (see page 31)
- Showing Comments (see page 32)
- Qualifying Bus Transitions (see page 87)
- Changing the Trigger Reference (see the *Listing Display Tool* help volume)

See Also

- “Analyze SPA Graphs” on page 46
- “Rescaling Graphs” on page 39
- “Use Information Areas in Graphs” on page 45
- “Opening and Saving SPA Configurations” on page 10

Sorting the Data (see page 37) (State Interval, Time Interval)

Selecting Print Options (see page 30)

“Print the SPA Window” on page 11

“Print SPA Data to a File” on page 12

Closing the SPA Menu (see page 31)

Getting Help (see page 38)

Selecting Print Options

Print options

This prints the standard print dialog.

1. From the SPA tool menu bar, select *File*.
2. Select *Print options*.

Print this window

This prints the current window, using settings in the *Print options* dialog.

1. From the SPA tool menu bar, select *File*.
2. Select *Print this window*.

Print to file

This saves the number of hits for each bucket to a file.

1. From the SPA tool menu bar, select *File*.
2. Select *Print to file*.
 - When saving SPA data, bucket information for all buckets is saved. Current display information is also saved.
 - When printing to a file, the file information area is empty by default. This area has no size limit.

- For the file to be saved, you must select *Save* before you select *Close*.
- To close the SPA dialog, select *File* and *Close*.

See Also

“Print SPA Data to a File” on page 12

“Print the SPA Window” on page 11

Closing the SPA Menu

To close the SPA dialog:

1. From the SPA tool menu bar, select *File*.
2. Select *Close* to close the current SPA tool window.

Accumulating Data

When *Accumulate* is *On*, data is accumulated in the graph from run to run, and related information is updated. By default, *Accumulate* is *Off*, which means that each new run overwrites data from the previous run.

Any operation that results in the repositioning or rebucketization of data will cause accumulated data to be lost. Due to memory and speed requirements, only data from the last acquisition is internally accessible for replotting.

1. From the SPA tool menu bar, select *Options*.
2. Select *Accumulate*, then select *On*.
3. Select the *Run* icon to run a new measurement. As data accumulates, information in the Display and Bucket Information areas are updated with accumulated information.

See Also

“Data Selection Criteria” on page 47

“State Overview Information Areas” on page 26

“State Interval Information Areas” on page 71


“Time Overview Information Areas” on page 93

“Time Interval Information Areas” on page 106

Viewing Acquisition Time

You can view the acquisition time for the data acquired in the last measurement run. The Acquisition Time dialog shows the start and end times, and the duration for the acquisition.

1. From the SPA tool menu bar, select *Graph*.
2. Select *View acquisition time*.



Showing Comments

By default, SPA does not display comments. When you display comments, a text area appears below the message line. You can enter a message of any length into the area. Comments are specific to each tool display.

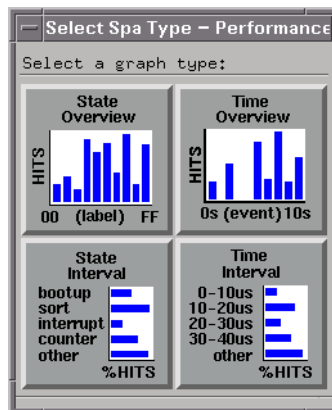
1. From the SPA tool menu bar, select *Options*.
2. To display the comments area, select *Show Comments*.
3. Select *On*. Enter comments into the area if you want.
4. To clear all existing comments for this window, Select *Options, Show Comments, Clear text*.
5. To turn off comments, select *Options, Show Comments, Off*.

When you select *File*, *Print to file*, or *Save Configuration*, comments are saved.

Setting a SPA Tool Type

You can use any one of the four SPA display tool types to view your data: State Overview, State Interval, Time Overview, or Time Interval.

1. From the SPA tool menu bar, select *Graph*.
2. Select *Set SPA Type*. Four buttons appear showing the four SPA tools.
3. Select the SPA tool you want to use to view your data.



Clearing and Replotting a Graph

Clearing a Graph

You can clear the graph display bars and reset the number of hits to 0. This also resets the graph information areas.

1. From the SPA tool menu bar, select *Graph*.
2. Select *Clear Graph*. The bars are erased and information areas reflect what is displayed.

3. To replot the graph, select *Replot Graph*. This redisplay the bars and information areas.

Replotting a Graph

You can redisplay the graph display bars and information areas.

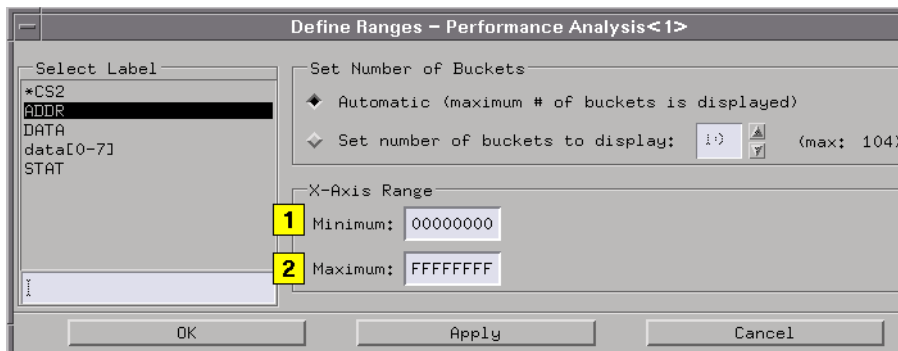
1. From the SPA tool menu bar, select *Graph*.
2. Select *Replot Graph*. This replots the data from the last acquisition. Any data that was accumulated is lost.

Setting State Overview X-Axis Range

You can set the minimum and maximum values to display on the X axis. The base for the X-axis values appears in the State Overview Display Information area. If you try to insert a value that is out-of-range, the minimum or maximum possible value is inserted. Follow the rules (see page 35) when entering values.

After selecting the State Overview *Define ranges* button, you are ready to set the X-axis range.

1. Select the *Minimum* field and enter the minimum value.
2. Select the *Maximum* field and enter the maximum value.
3. Select *Apply* to apply the changes or *OK* to apply and exit the dialog. This will rescale the graph and recalculate the buckets based on the new minimum and maximum values.



The base for the X-axis values is set in the State Overview Display Information area.

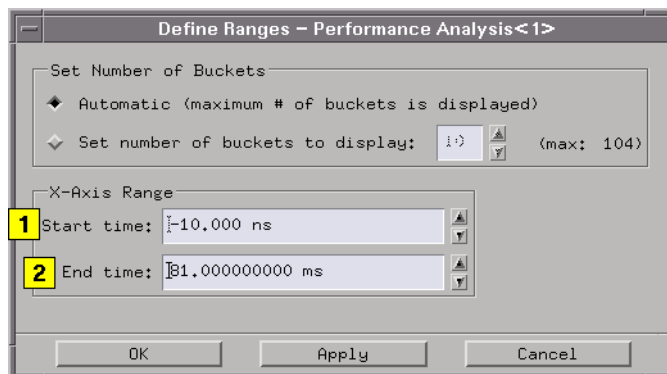
Rules for Entering Values

- Do not enter don't care values in the fields.
- Do not enter a negative minimum value.
- Do not let the minimum X-axis range value exceed the minimum value based on the number of bits in the label (for example, if the label has 8 bits, the value cannot exceed 255).
- Do not let the maximum X-axis range value exceed the maximum data value possible based on the number of bits.

Setting Time Overview X-Axis Range

You can enter values (see page 36) to set the start and end times to display on the SPA Time Overview X axis.

1. Select the *Start time* field and enter the starting time.
2. Select the *End time* field and enter the ending time.
3. Select *Apply* to apply the changes or *OK* to apply and exit the dialog. This will rescale the graph and recalculate the buckets based on the new minimum and maximum values.



Entering X-Axis Range Values

- If you try to insert a value that is out-of-range, the maximum possible value is inserted.
- You can specify a negative start time; a negative value reflects the time before the trigger.
- Do not let the minimum time value exceed the maximum time value. Time values must be between -1 Ms and +1 Ms.

Setting State and Time Intervals X-Axis Ranges

You can set the X-axis range percentage (1-100%) for State Interval and Time Interval.

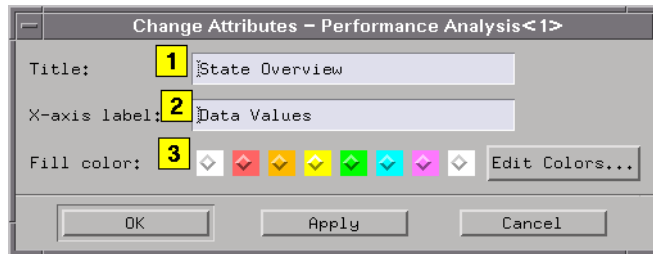
1. Select the *Maximum range percent* field, and enter a number from 1 to 100.
2. Select *Apply* to apply the changes or *OK* to apply and exit the dialog.

Changing Attributes

You can customize the SPA State Overview display by changing the title, X-axis label, and fill color for the data distribution bars.

After selecting *Graph* and *Change Attributes* from the SPA tool menu bar, you are ready to change attributes.

1. To change the title, select the *Title* field and enter a new title.
2. To change the X-axis label, select the *X-axis label* field and enter a new label.
3. To change the fill color, select one of the 8 color buttons.
4. Select *Apply* to apply the changes or *OK* to apply and exit the dialog.



Sorting the Data

Sort choices only apply when you are using the *State Interval* and *Time Interval* tools. You can sort the buckets according to the number of hits, alphabetically by name (in *State Interval* only), or by the defined range.

1. From the SPA tool menu bar, Select *Sort*.
2. To sort by the number of hits (see page 37), select *Hits*.
3. To sort by name (see page 37), select *Alphabetic*.
4. To sort by the range (see page 38), select *Range*.
5. To sort by the current sort criteria, select the *Sort* button.

Sorting by Number of Hits

This sort shows which buckets are receiving the greatest number of hits, in descending order. Each run can change the ordering. As new buckets are defined or old buckets are changed or deleted, the ordering will change depending on the number of hits in the bucket.

Sorting by Name

This sorts the defined buckets according to the bucket name, in

ascending order. This sort is *not* updated automatically on every run, but does update automatically when buckets are added, deleted, or changed. This is only valid for State Interval.

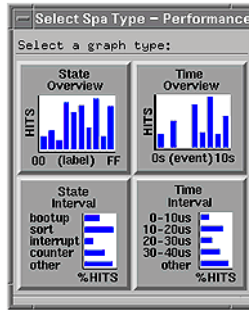
Sorting by Range

This sorts the buckets according to the defined range, in ascending order. Buckets that have a pattern definition, rather than a range definition, are sorted by replacing don't cares with zeros, and placing them into their proper ascending position. This sort is *not* updated automatically on every run, but does update automatically when buckets are added, deleted, or changed.

Getting Help

1. From the SPA tool menu bar, select *Help*.
2. To access help information about the SPA tools, select *On This Window*.
3. To display the main system help page, select *On Main System*.
4. When you have accessed the help, to display additional information about using the help system select *Help, Overview - Using Help*. This help shows you how to use hyperlinks, browse help topics, use the keyword index, print help topics, and re-visit topics.
5. To display the version of online help, select *Help, About Help*.

Rescaling Graphs



You can rescale the SPA tool graphs to view data in greater detail. You can rescale a graph, then undo the scaling, autoscale the graph, set the graph to full scale, and expand buckets. You can also resize dialogs.

When you rescale a graph, you are not just zooming the information onscreen. The data buckets are actually rescaled to new values, and data is rebucketized.

- Dragging and Rescaling a Graph (see page 39)
- Options for Rescaling a Graph (see page 41)
- Expanding a Bucket (see page 42)
- Resizing a Dialog (see page 43)
- Placing Markers (see page 43)

See Also


“Working with Graphs” on page 29

Dragging and Rescaling a Graph

You can expand a section of the displayed graph data for analysis.

1. Place the cursor over the starting data of interest.

2. Drag the cursor to the ending data of interest. This creates a dash-lined rectangle.
3. Release the cursor, and data is rescaled (see page 40) as shown in the following diagram.



NOTE:

When you rescale a graph, you are not just zooming the information onscreen. The data buckets are actually rescaled to new values, and data is rebucketized. Accumulated data is lost during a rescale operation.

Data is Rescaled

Data inside the rectangle is rescaled to reflect only the data values inside the rectangle. The X axis is also rescaled to reflect only these data values. If the number of data values is less than the number of buckets, data will continue to rescale until a minimum of one data value is represented in the graph area. The X axis will reduce accordingly.

Options for Rescaling a Graph

Depending on the selected SPA tool, you can choose one or more of these options.

Auto Scaling

Auto Scaling sets the X-axis values to the minimum and maximum values found after a *Run*. For example, if a label was 16 bits wide, and the data values found after a *Run* were 55-1904, the X-axis minimum and maximum values would be 55 and 1904.

Undo Scaling

Undo Scaling reinstates previous scale settings and replots the graph using the old scale values.

Full Scale

Full Scale sets the X-axis values to the minimum and maximum values found after a *Run*.

- For example, In State Overview, if a label is 16 bits wide, and the data values found after the run are 55-1904, the X-axis minimum and maximum values would be 0 and FFFF.
- In Time Overview, if after an acquisition the first acquired sample occurred at -10 ms, and the last acquired sample occurred at 50 ms, the X-axis minimum and maximum values would be rescaled to -10 ms and 50 ms, respectively.

Using a Sash

When there is a small button (sash) in the SPA window, usually on the right side of the screen, you can select it and drag to resize that portion of the window.


Expand Bucket

Expand Bucket expands the highlighted bucket, rescaling the X axis so that the bucket low range becomes the minimum X-axis value and the bucket high range becomes the maximum X-axis value.

Expanding a Bucket

In State Overview, expanding a highlighted bucket places the bucket's low value as the X-axis minimum, and the bucket's high value as the X-axis maximum. In Time Overview, if a bucket has acquired data, expanding a highlighted bucket places the bucket's minimum value as the X-axis minimum, and the bucket's maximum value as the X-axis maximum.

1. Select the bucket of interest.
2. Select *Expand Bucket*.
3. Select the bucket again to expand it even further as shown in the following diagram.



See Also

Highlighting a Bucket (see page 17)

Placing Markers

Markers are used in Time Overview only.

Goto Marker displays a choice of active markers. For example, if G1 and x are the only active markers, those two picks are displayed, along with the *Trigger* pick.

Place Marker displays a list of active markers. This selection uses the cursor position to select the cursor options and places any of the active markers. When you select an active marker, it will be placed at the cursor position. There is no *Trigger* pick.

Center About displays all possible choices of marker pairs (for example, G1 & G2 or G1 & x). When selected, the overview distribution is rescaled so that the first marker is placed at the 10% position of the display, and the second marker is placed at the 90% position.

NOTE:

When you rescale a graph, you are not just zooming the information onscreen. The data values are actually rescaled to new values. Accumulated data is lost during a rescale operation.

Resizing a Dialog

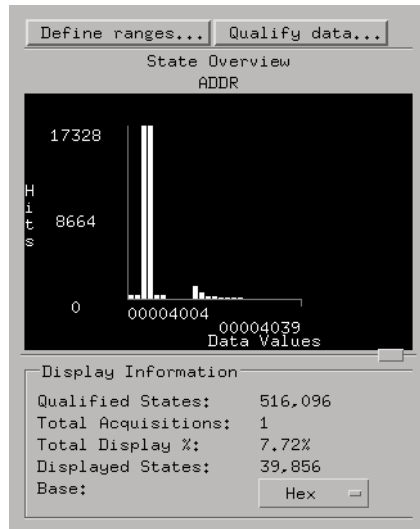
1. Select a dialog border.
2. Drag to the desired size.

NOTE:

Be aware that if you size the dialog too small *you may lose some important information* displayed on screen (for example, this resized State Overview window does not show Bucket Information).

Chapter 1: Using the System Performance Analyzer


Rescaling Graphs



See Also

“Working with Graphs” on page 29

Use Information Areas in Graphs



SPA tool dialogs have information areas to help you customize and analyze acquisition results. They include message areas, comment areas, data graph areas, and display and bucket information areas.

- “State Overview Information Areas” on page 26
- “State Interval Information Areas” on page 71
- “Time Overview Information Areas” on page 93
- “Time Interval Information Areas” on page 106

See Also

“Working with Graphs” on page 29

“Analyze SPA Graphs” on page 46


Message Area (see page 45)

Showing Comments (see page 32)

Message Area

The message area to the right of the *Run* key helps you understand interactions allowed as you move through the entire dialog.

Analyze SPA Graphs



The following topics summarize each of the SPA tools and show you how to interpret the SPA tool graphs. They also describe the effects that defining ranges and qualifying states have on the displayed graph results.

- “Analyzing State Overview” on page 24
- “Analyzing State Interval” on page 68
- “Analyzing Time Overview” on page 91
- “Analyzing Time Interval” on page 104

See Also

“Working with Graphs” on page 29

“Use Information Areas in Graphs” on page 45

Data Selection Criteria

These methods are used for searching data:

- State Search Criteria (see page 48) (used in State Overview and State Interval)
- Time Overview Single-Event Search Criteria (see page 49)
- Time Interval and Time Overview Dual-Event Search Criteria (see page 50)
- Bus Mode Search Criteria (see page 52)

See Also

Held Values when Using Multiple Datasets (see page 47)

Held Values when Using Multiple Datasets

The examples used to describe the various data selection criteria methods show a single *data set* with a 4-bit label called *NIBBLE*. The concepts are easily extended to multiple data sets, taking into account held values. Multiple data sets exist when labels from two or more acquisition machines are sourced into a single analysis or display tool (SPA, for example).

NIBBLE (data set 1)		BYTE (data set 2)		
0	4E	0	FF	4 ns
1	4F	1	00	8 ns
(no state)		2	11	12 ns
2	50	3	22	16 ns
3	7F	4	33	20 ns

The NIBBLE value of 4F is held until 50 is seen.

During a search, if the qualifying event was set to *When Present* and the qualifying term is BYTE with the value of 11, the value plotted for the NIBBLE label would be the last value seen on that label at the time the qualifying state came in (for example, 4F). Even though there is no state sampled at the time of BYTE=11, the held value is still plotted.

State Search Criteria

The following criteria is used in State Overview and State Interval.

Examples

These examples are marked on the right side of the table.

1. When entering 6 (returns first state when criteria met)
2. When exiting 6 (returns first state AFTER criteria met)
3. When present 6 (returns state)
4. When 6 present > 4 ns (returns first state AFTER criteria met)
5. When 6 present > 2 ns (returns first state AFTER criteria met)
6. When 6 present < 5 ns (returns last state BEFORE criteria met)
7. When 6 present > 5 ns & 6 present < 12 ns (returns every state within criteria)
8. When 6 present > 1 ns & 6 present < 12 ns (returns every state within criteria)
9. When 6 present > 5 ns | 6 present < 1 ns (returns every state within criteria)

State	Nibble	Time	Example #																
			1	2	3	4	5	6	7	8	9								
0	B	0 s																	
1	5	4 ns																	
2	F	8 ns																	
3	6	12 ns	x		x				x					x					
4	8	16 ns		x				x											
5	A	20 ns																	
6	6	24 ns	x		x									x	x	x			
7	6	28 ns			x									x	x	x			
8	F	32 ns			x														
9	F	36 ns																	
10	9	40 ns																	
11	C	44 ns																	
12	7	48 ns																	
13	7	52 ns																	
14	D	56 ns																	
15	9	60 ns																	
16	A	64 ns																	
17	0	68 ns																	
18	A	72 ns																	
19	4	76 ns																	
20	1	80 ns																	
21	F	84 ns																	
22	3	88 ns																	

23	6	92 ns	x	x	x	x
24	C	96 ns	x		x	
25	0	100 ns				
26	5	104 ns				
27	0	108 ns				

See Also

Held Values when Using Multiple Datasets (see page 47)

Qualifying Bus Transitions (see page 87)

Bus Mode Search Criteria (see page 52)

Time Overview Single Event Search Criteria

The following search criteria is used in Time Overview when a single event is defined.

Examples

These examples are marked on the right side of the table.

1. When entering 6 (returns first time when criteria met)
2. When exiting 6 (returns first time AFTER criteria met)
3. When present 6 (returns time)
4. When 6 present > 4 ns (returns first time when criteria met) (1)
5. When 6 present > 2 ns (returns first time when criteria met) (1)
6. When 6 present < 5 ns (returns first time when criteria met) (1)
7. When 6 present > 5 ns & 6 present < 12 ns (returns first time when criteria met) (1)
8. When 6 present > 1 ns & 6 present < 12 ns (returns first time when criteria met) (1)
9. When 6 present > 5 ns | 6 present < 1 ns (returns first time when criteria met) (1)

(1) This search returns the time interval start time that met the criteria. If both start and end events are enabled, Time Interval search criteria is used.

State	Nibble	Time	Example #								
			1	2	3	4	5	6	7	8	9
0	B	0 s									
1	5	4 ns									
2	F	8 ns									
3	6	12 ns	x		x		x	x		x	
4	8	16 ns		x							
5	A	20 ns									
6	6	24 ns	x		x	x	x		x	x	
7	6	28 ns			x						
8	F	32 ns		x							
9	F	36 ns									
10	9	40 ns									
11	C	44 ns									
12	7	48 ns									
13	7	52 ns									
14	D	56 ns									
15	9	60 ns									
16	A	64 ns									
17	0	68 ns									
18	A	72 ns									
19	4	76 ns									
20	1	80 ns									
21	F	84 ns									
22	3	88 ns									
23	6	92 ns	x		x		x	x		x	
24	C	96 ns		x							
25	0	100 ns									
26	5	104 ns									
27	0	108 ns									

See Also

- Held Values when Using Multiple Datasets (see page 47)
- Qualifying Bus Transitions (see page 87)
- Bus Mode Search Criteria (see page 52)

Time Interval and Dual-Event Time Overview Search Criteria

The following search criteria is used in Time Interval and in Time Overview (if two events are enabled, where Event 1 is the start event, and Event 2 is the end event). Because there can be many variations (start/end event combinations) of this method, how to determine the time is described using a simple example.

For all searching methods except exiting, the time of the first sample, when the criteria is met, is used. For exiting, the time of the first sample *after* the criteria is met is used. In the following, arrows point to the samples whose time will potentially be used.

Assume that As are used for the search criteria. The first B on the left

side starts at time 0 s, and each sample is taken at a 4-ns increment.

```

Entering          BAAA12CCAEFAAB56A777AAA
                  ^   ^   ^   ^   ^
Exiting          BAAA12CCAEFAAB56A777AAA
                  ^   ^   ^   ^   ^
Present          BAAA12CCAEFAAB56A777AAA
                  ^   ^   ^   ^   ^
Present > 4 ns   BAAA12CCAEFAAB56A777AAA
                  ^   ^   ^   ^   ^
Present < 5 ns   BAAA12CCAEFAAB56A777AAA
                  ^   ^   ^   ^   ^
Present > 5 ns < 9 ns BAAA12CCAEFAAB56A777AAA
                  ^   ^   ^   ^   ^

```

The following example assumes that the start event is qualified as *When Present F*, and the end event is qualified as *When Exiting 6*.

State	Nibble	Time	
0	B	0 ns	
1	5	4 ns	
2	F	8 ns	Start event #1
3	6	12 ns	
4	8	16 ns	End event #1
5	A	20 ns	
6	6	24 ns	
8	F	32 ns	
7	6	28 ns	
8	F	32 ns	Start event #2
9	F	36 ns	
10	9	40 ns	
11	C	44 ns	
12	7	48 ns	
13	7	52 ns	
14	D	56 ns	
15	9	60 ns	
16	A	64 ns	
17	0	68 ns	
18	A	72 ns	
19	4	76 ns	
20	1	80 ns	
21	F	84 ns	
22	3	88 ns	
23	6	92 ns	
24	C	96 ns	End event #2
25	0	100 ns	
26	5	104 ns	
27	0	108 ns	

The time difference End event #1 - Start event #1 = 8 ns, and the time difference End event #2 - Start event #2 = 64 ns. These are the times that will be bucketized.

See Also

Held Values when Using Multiple Datasets (see page 47)

Qualifying Bus Transitions (see page 87)

Bus Mode Search Criteria (see page 52)

Bus Mode Search Criteria

Bus Mode search criteria may be used in State Overview, State Interval, Time Overview, and Time Interval. This mode allows transition events on a bus (a label with ≥ 1 channel) to be qualified in a single event term. Because there can be many variations of data qualification in bus mode, two examples are provided. The first example uses Time Interval to plot Setup Times, defined as the time between Address Strobe (AS#) going true and the first subsequent transition to occur on the Address bus (ADDR). The second example identifies the number of times a particular type of DRAM is accessed using the Time Overview tool.

In the Qualify Data menu, two specifications have a significant impact on how data or events are qualified:

- Data is qualified when the event is *[terms]* and *Pattern/Not Pattern*.
- Data is qualified when the event is *[Entered, Exited, Present, Not Present, Present >, Present <, Present ><]* is used to decide when to qualify the bus.

In general, an event term label using *Pattern* is most often used when you want to qualify on unstable conditions on the bus, for example on two consecutive values are different. *Not Pattern* is most often used when you want to qualify on stable conditions on the bus, for example, on two consecutive values are the same. When bus mode is active, labels in an event term with a pattern of all don't cares will result in a search for a specific transition condition across all lines on the bus.

Search Criteria when *Pattern* is Specified

The following information is used to describe how data or events are qualified when the bus Event: label is set to *Pattern*. Typically, this search criteria looks for instability on the bus lines. For a simple example, let's assume we have a bus made up of 4 lines, with a valid range from 0 to 15. The (^) indicates which samples or events are used in the search. Assume that the pattern has been defined as *X*, and therefore bus transitions will dictate which states to choose in the analysis. A bus is considered unstable when two consecutive bus values are different.

The effects of *when* events are qualified as follows:

Entered - qualifies transitions from a stable bus to an unstable bus.

Exited - qualifies the first sample into a stable bus condition.

Present - qualifies all transitions.

Not Present - qualifies all stable bus samples.

Present > - qualifies if time from the *Entered* transition and an *Exited* condition is > than the user-defined time value.

Present < - qualifies if time from the *Entered* transition and an *Exited* condition is < than the user-defined time value.

Present >< - qualifies if time from the *Entered* transition and an *Exited* condition is inside the user-defined time range.

Let's assume that each sample is taken at 4-ns intervals. Events marked with a (^) are qualified. Present >, Present <, and Present >< searching looks for the time between samples (0 and 1) to determine if the specified time duration is satisfied.

Qualified When:

```

Entered          BAAA12CCADEF AAB56A777ABB
                ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^
Exited          BAAA12CCADEF AAB56A777ABB
                ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^
Present         BAAA12CCADEF AAB56A777ABB
                ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^
Not Present     BAAA12CCADEF AAB56A777ABB
                ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^
Present > 4ns   BAAA12CCADEF AAB56A777ABB
                01 0  10    10    1  0  1
                ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^
Present &< 5ns   BAAA12CCADEF AAB56A777ABB
                01 0  10    10    1  0  1
                ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^
Present >4ns &< 13 ns BAAA12CCADEF AAB56A777ABB
                01 0  10    10    1  0  1
                ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^  ^
    
```

Search Criteria when Not Pattern is Specified

The following information is used to describe how data or events are qualified when the bus Event: label is set to *Not Pattern*. Typically, this search criteria looks for stability on the bus lines. For a simple example, let's assume we have a bus made up of 4 lines, with a valid range from 0 to 15. The (^) indicates which samples or events are used in the search. Assume that the pattern has been defined as X, and therefore bus transitions will dictate which states to choose in the

analysis. A bus is not considered stable until two consecutive bus values are identical.

The effects of *when* events are qualified as follows:

Entered - qualifies the first sample into a stable bus condition.

Exited - qualifies transitions from a stable bus to an unstable bus.

Present - qualifies all stable bus samples.

Not Present - qualifies all bus transitions.

Present > - qualifies if time from the *Entered* transition and an *Exited* condition is > than the user-defined time value.

Present < - qualifies if time from the *Entered* transition and an *Exited* condition is < than the user-defined time value.

Present >< - qualifies if time from the *Entered* transition and an *Exited* condition is inside the user-defined time range.

Let's assume that each sample is taken at 4-ns intervals. Events marked with a (^) are qualified. Present >, Present <, and Present >< searching looks for the time between samples (0 and 1) to determine if the specified time duration is satisfied.

Qualified When:

```

Entered          BAAA12CCADEF AAB56A777ABB
                ^      ^      ^      ^      ^
Exited          BAAA12CCADEF AAB56A777ABB
                ^      ^      ^      ^      ^
Present         BAAA12CCADEF AAB56A777ABB
                ^      ^      ^      ^      ^
Not Present     BAAA12CCADEF AAB56A777ABB
                ^      ^      ^      ^      ^
Present > 4ns   BAAA12CCADEF AAB56A777ABB
                0 1 01 01 0 1 0
                ^
Present &< 5ns   BAAA12CCADEF AAB56A777ABB
                ^ 0 1 01 01 0 1 0
                ^
Present >4ns &< 13 ns BAAA12CCADEF AAB56A777ABB
                ^ 0 1 01 01 0 1 0
                ^
    
```

NOTE:

When more than one event is specified, such as Event 1 and Event 2 in Time Overview, or *Start of event* and *End of Event* in Time Interval, the search for the second event begins where the first event was qualified.

See Also

State Search Criteria (see page 48)

Time Overview Single Event Search Criteria (see page 49)

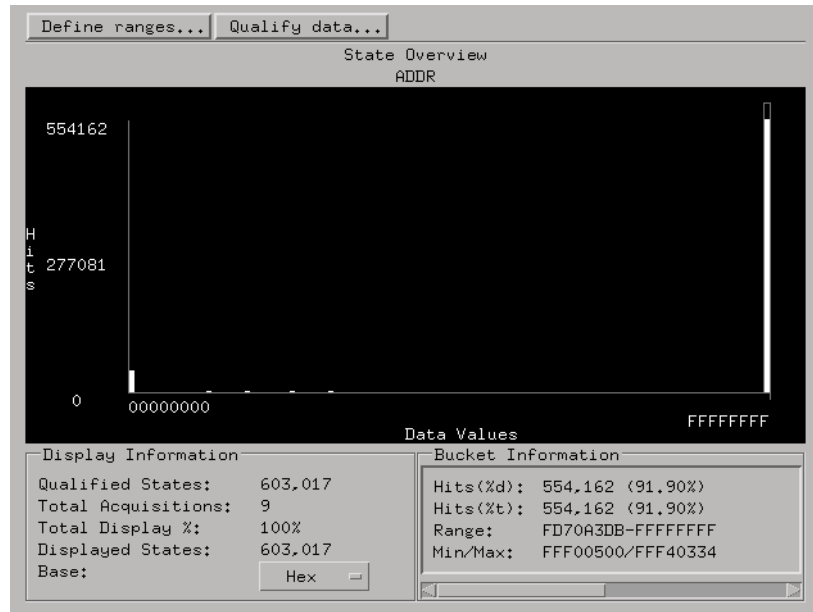
Time Interval and Time Overview Dual-Event Search Criteria (see page 50)

Identifying Routines that Take the Most CPU Cycles

To analyze the performance of a software-based system, it is important to understand the CPU activity across the entire address range. The regions of high activity provide the raw information needed to determine which routines may be responsible for throughput bottlenecks. By knowing the regions of high activity, you can then use State Interval to view the relative utilization of routines in detail. This is an example measurement.

1. Set up the appropriate probing (see page 58) for the ADDRESS and STATUS lines using the Agilent Technologies 1655x acquisition card.
2. Configure a logic analyzer tool and a SPA tool on the workspace.
3. Create the ADDR and STAT labels in the logic analyzer.
4. If you are using an analysis probe, load the appropriate configuration file supplied with the analysis probe.
5. Configure the logic analyzer to store only valid addresses (see page 57) .
6. Activate the SPA State Overview tool.
7. Select the label (see page 15) named ADDR in State Overview.
8. Turn *Accumulate* ON to accumulate data (see page 31).
9. Run Repetitive (see page 23).

The result resembles:



After acquiring much data, we see that the majority of CPU cycles is occurring in two regions. Software designers can use this information to reference routines that occupy these regions. A more detailed analysis is then warranted, where the start and end addresses of these routines become the ranges for more focused State Interval measurements.

Designers can optimize the routines and rerun the State Interval measurements to see the effects of the code optimizations on system throughput.

See Also

Uses of State Overview (see page 14)

Storing Values

A qualified storage of valid values, such as addresses and exceptions, is typically determined by processor status label values.

Probing Signals

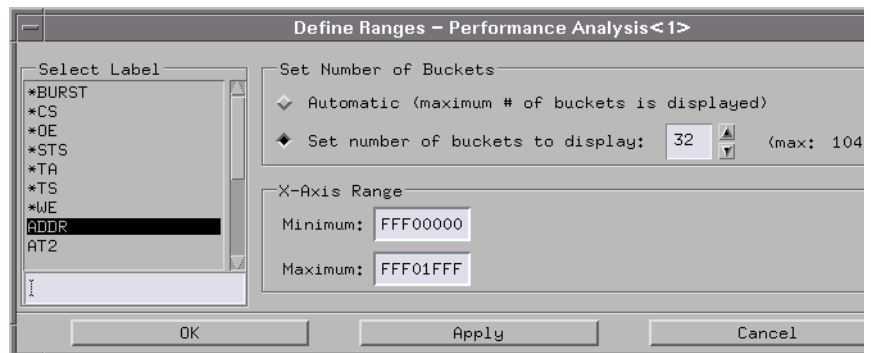
Whenever possible, probing is best done using an analysis probe. Analysis probes provide a secure connection between the targets system microprocessors and the logic analyzer.

Identifying Exceptions

Hardware designers often want a quick overview of system exceptions. An unexpected number of exceptions, or type of exception, may signal a design flaw.

In this example, assume the processor exception vector table is from 0xFFF00000 to 0xFFF01FFF and the start addresses of the exception vector service routines are every 0x100 bytes (that is, 0xFFF00000, 0xFFF00100, 0xFFF00200, etc.).

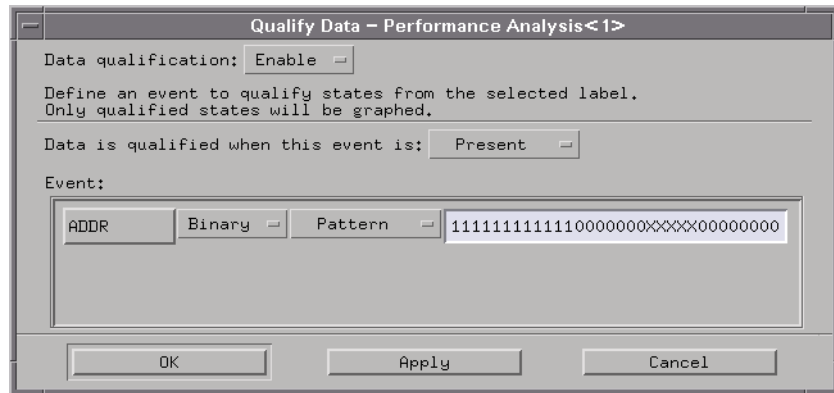
1. Set up the appropriate probing (see page 58) to capture processor execution.
2. Configure a logic analyzer tool and a SPA tool on the workspace.
3. Configure the logic analyzer.
4. Activate the SPA State Overview tool.
5. Select the address label (see page 15) in State Overview.
6. To identify and count the exceptions that occur, you can set up a state overview measurement where the X axis range is from 0xFFF00000 to 0xFFF01FFF and 32 buckets are defined.



7. To make sure only the start addresses of the exception vector service routines are counted, you can qualify the data so that only addresses in the range that end with 00 are counted.


Chapter 1: Using the System Performance Analyzer

Identifying Exceptions



8. Turn *Accumulate* ON to accumulate data (see page 31).
9. Run Repetitive (see page 23).

The result resembles:



Only one of the possible exceptions, the exception vector at 0xFFFF00500, is captured.

There are several ways to gain insight into the cause of counted

exceptions. For example:

1. Configure the state analyzer to trigger on the exception.
2. Use the Lister tool to display the acquired states.
3. Analyze the states that occurred before the fault trigger to determine what events led up to the fault.

See Also

Uses of State Overview (see page 14)

Using State Interval



Use State Interval to characterize the utilization of system resources. Concentrating optimization efforts on routines or operations that use significant system resources will result in the greatest gains in system throughput. State Interval displays a horizontal distribution of activity in code address ranges.

- Application Examples**
- “Zooming in on Potential Software Bottlenecks” on page 78
 - “Characterizing Microprocessor CPU Cycles” on page 80

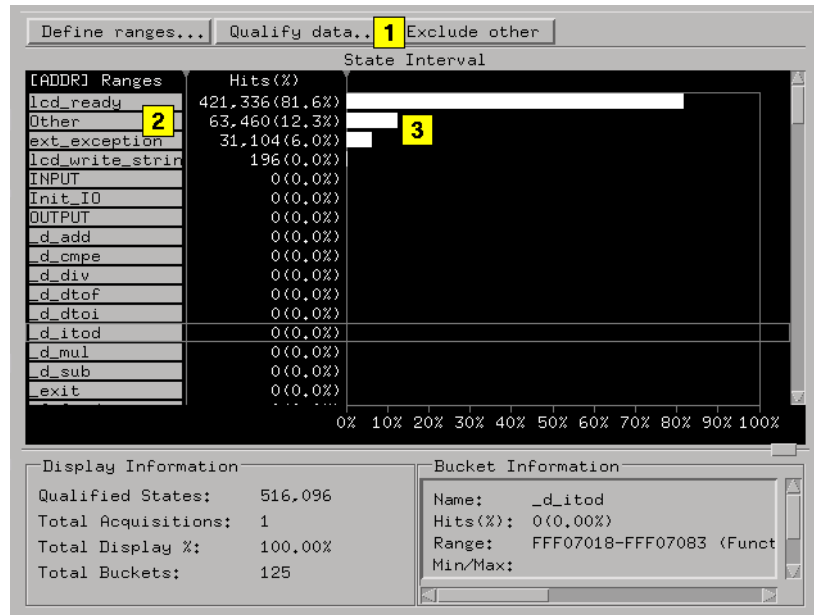
See Also

- “Defining State Interval Ranges” on page 65
- “Qualifying State Interval Data” on page 67
- Including Other Data Values (see page 63)
- “Using the Object File Symbol Selector” on page 74
- “Running and Stopping a Measurement” on page 23
- “Analyzing State Interval” on page 68
- Uses of State Interval (see page 64)
- “State Interval Information Areas” on page 71
- “Working with Graphs” on page 29
- Data Selection Criteria (see page 47)
- “Opening and Saving SPA Configurations” on page 10
- “Print the SPA Window” on page 11
- “Print SPA Data to a File” on page 12

Including Other Data Values

When you define ranges that do not cover the specified label's entire range, you can include any other occurring qualified data values in the *Other* bucket.

1. Select *Include other*.
2. A bucket named *Other* is added to the ranges list.
3. To delete the *Other* bucket, select the bucket.
4. Select *Delete*.



NOTE:

When you run a measurement, any qualified data values that are not currently defined in a bucket will be stored in the *Other* bucket. If a range is turned off, any activity in that range is included in the *Other* bucket, regardless of whether the *Other* bucket is displayed or not.

See Also

“Defining State Interval Ranges” on page 65

“Qualifying State Interval Data” on page 67

State Search Criteria (see page 48)

Uses of State Interval

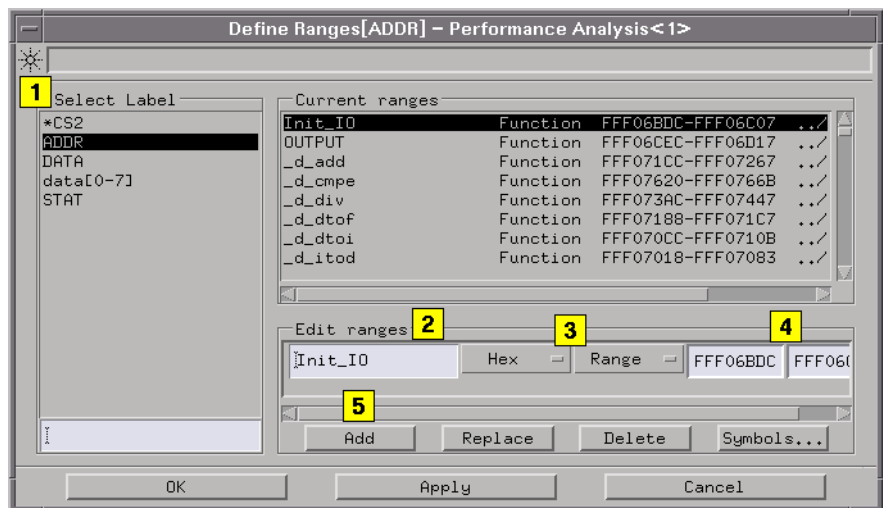
You can use State Interval to:

- Characterize the utilization of system resources.
- Find system resource utilization problems such as bottlenecks in software routines.
- Display the occurrences of system interrupt routines.
- Apply State Interval analysis to any interface where status lines define valid cycle types. Interfaces such as PCI, MCA, or EISA make good candidates for State Interval analysis.
- Identify ratios for cache hits and misses.
- Identify the percentage of time a bus is busy.
- Zoom in on potential software bottlenecks.
- Characterize microprocessor CPU Cycles.

Defining State Interval Ranges

Defining ranges qualifies states to include in the analysis.

1. Select the Label (see page 15) to display.
2. Select the *Edit ranges* text field and enter a name for the range.
3. Define the *Base* and *Range/Pattern* buttons.
4. Enter a range or pattern.
5. Select *Add* to include the defined range in the current ranges list.
6. Define and include as many additional ranges as you want.
7. To replace a range, highlight the range in the current ranges window. Edit the range in the edit ranges window, then select the *Replace* button.
8. To delete a range, highlight the range, and select the *Delete* button.
9. Select *Apply* to apply the changes or *OK* to apply and exit the dialog.



NOTE:

Selecting on the select label window displays choices to filter or sort the label list. Selecting the current ranges area provides choices to undo actions or sort and delete current ranges.

See Also

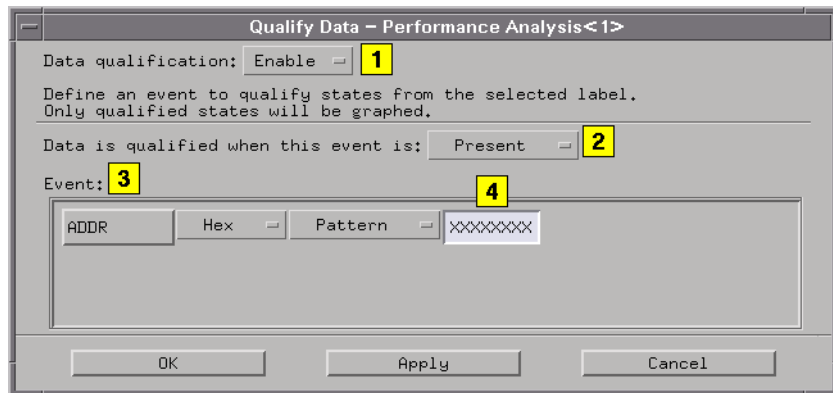
“Using the Object File Symbol Selector” on page 74

Include Other Data Values (see page 63)

State Search Criteria (see page 48)

Qualifying State Interval Data

1. Enable or disable the *Data qualification*.
2. If the data qualification is enabled, select when to qualify the data (see page 19).
3. Qualify the event by selecting the *Label*, *Base*, and *Pattern/Range* buttons.
4. Type in a pattern or range.
5. Insert as many labels as you want, and define the data qualification for each.
6. Select *Apply* to apply the changes or *OK* to apply and exit the dialog.



NOTE:

Multiple event labels are ANDed to create a single event term.

See Also


Qualifying Bus Transitions (see page 87)

Including Other Data Values (see page 63)

“Using the Object File Symbol Selector” on page 74


State Search Criteria (see page 48)

Analyzing State Interval



SPA State Interval displays states that occur within user-defined ranges of a label. State Interval is displayed on a percentage scale, and each bar represents the percentage of samples in the range as a fraction of all qualified states. State Interval is available for use with all defined labels.

How to Interpret the State Interval



- A. The maximum number of ranges is unlimited.
- B. Other states included/excluded is available and displays all states not covered by the user-defined ranges.
- C. *Hits [%]* represents the number of hits in selected bucket, and

percentage of those hits relative to the total qualified hits.

D. The highlighted bucket's information is displayed in the Bucket Information (see page 72) area.

E. The Display Information (see page 72) is updated after every run, with information based on how the data was qualified.

F. The X axis is the percentage of total qualified states.

G. The Y axis is the set of ranges defined for the selected label.

H. The total display percentage (see page 70) is the sum of the display percentages.

Effects of Defining Ranges

Defining ranges (see page 65) allows you to select the label (see page 15) of interest, define state patterns or ranges, and include OMF symbols (see page 74).

Effects of Qualifying States

Qualified states are the states used in the analysis. You can define when to qualify data (see page 19) and define a qualifying event term (see page 19) . If you define an event on which to qualify the acquired data, only those qualified states are counted. Those states are then bucketized and displayed.

See Also

Sorting the Data (see page 37)

“State Interval Information Areas” on page 71

“Setting Thresholds” on page 21

Changing Attributes (see page 36)

“Working with Graphs” on page 29

“Using the Object File Symbol Selector” on page 74

“Data Selection Criteria” on page 47


NOTE:

Choosing the select label window displays choices to filter or sort the label list. Selecting the current ranges area provides choices to undo actions or sort and delete current ranges.

Display Percentage Field

This field displays 100% or less when there are no overlapping ranges. If a graph bar reaches the 40% value, the range for that bar contains 40% of all the samples displayed. If the total of all bars is greater than 100%, ranges are overlapped and data has been counted twice.

State Interval Information Areas



State Interval information areas include messages, comments, the label area, hit area, data graph area, and Display and Bucket Information areas. When you select anywhere in the graph area, a menu displays that lets you delete the current label, delete from the current label to the end of the list, or find a label.

Message Area

The message area to the right of the *Run* icon helps you understand interactions allowed as you move through the entire dialog.

Comment Area

The comment area appears when you select *Options* and *Show Comments* in the menu bar. Use the comment area to document setup information and data results. The size of the comments file is not limited. Comments are saved when you *Print to File*.

Label Range Area

This area of the State Interval displays defined bucket (range) names for the active label. The title at the top of the column shows the label for which the buckets are defined. You can drag and drop the vertical bar to the right of the label column to resize it and the other columns.

Hit Area

This area displays the number and percentage of hits. If this area is smaller than the text to be displayed, the text is highlighted yellow to indicate the truncation. You can drag and drop the vertical bar to the right of the Hits column to resize it and the other columns.

Data Graph Area

This area displays bucket results as a percentage of hits. When you move the cursor over a bucket, the bucket information is selected and highlighted with a solid bar. You can drag and drop a bucket to reposition it in the bucket information list. You can drag and drop the vertical bar to the left of the Hits graph area to rescale it.

Find a Range Label

When you select in the graph area, you can select to find a label in the current label list. As you enter a label name, the label matching the pattern you enter moves to the top of the displayed label list. Select *Next* and *Prev* to display the next and previous occurrence of the label. If you enter a label name that does not exist, the labels do not move.

Display and Bucket Information Areas

The Display and Bucket Information Areas (see page 72) are just below the graph.


See Also

“Rescaling Graphs” on page 39

“Working with Graphs” on page 29

State Interval Display and Bucket Information

State Interval shows display and bucket information for the acquisition.

**Display Information**

Qualified States - are the number of states used in the analysis. If the acquisition is set to collect 8K states, and you define an event on which to qualify the data, only those qualified states are counted. Those

states are then bucketized and displayed.

Total Acquisitions - reflects the number of accumulated and displayed data runs in Accumulate mode. When Accumulate mode is off, the Total Acquisitions field will be 0 if no runs occurred, or 1 if the displayed data represents a single run.

Total Display % - shows the percentage of all data values displayed. If you rescale the graph to show only a portion of the data, this number reflects the new percentage of displayed and total qualified samples. If there are overlapping fields, the sum can be greater than 100%.

Total Buckets - shows the total number of buckets that are defined for the selected label.

Bucket Information

When you highlight a specific bucket (see page 17), the following information about that bucket displays in the Bucket Information area.

Name - displays the name of the bucket.

Hits (%) - displays the number of hits in the selected bucket and the percentage of those hits relative to the total qualified hits.

Range - shows the low and high data values for the selected bucket. If the bucket defines a pattern, the pattern is displayed. If you selected a symbol using the object file Symbol Selector, the symbol type (function or variable) is displayed.

Min/Max - shows the minimum and maximum data values in the selected bucket.

File - If you selected a symbol from an object file, the name of the file where the symbol is declared is displayed.

Objf - If you selected a symbol from an object file, the name of the object file where the symbol is linked is displayed.

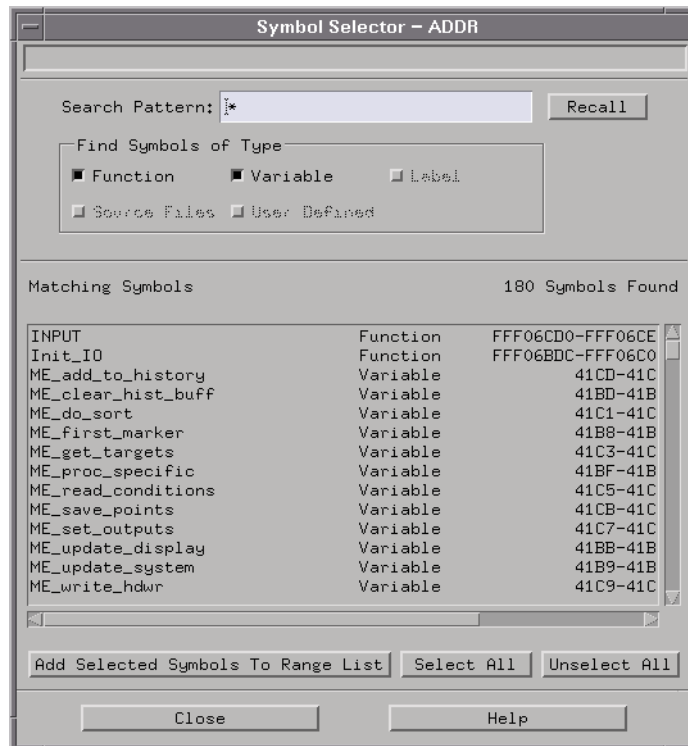
See Also

“Data Selection Criteria” on page 47

Using the Object File Symbol Selector

State Interval lets you add object file functions and variables to the list of defined ranges in the Define Ranges dialog. The Symbol Selector includes the symbol search and select symbols areas.

1. Filter (see page 75) the symbol list.
2. Select symbols (see page 75) from the symbol list (see page 75).
3. Add (see page 76) the symbols to the current ranges list.
4. Delete symbols (see page 76) that you no longer want.
5. Close the symbol selector.



After selecting *Symbols* in the State Interval Define Ranges dialog, you

can use the object file symbol selector.

See Also

“Defining State Interval Ranges” on page 65

Filtering the Symbol List

You can set the symbol search parameters so that only the symbols you want will appear in the symbol list.

1. To filter by source file, select desired source files (see page 77) either manually or by using standard Motif Extended Selection.
2. To quickly select all files in the list, select *Select All*. Or, to quickly deselect all files in the list, select *Deselect All*.
3. To filter by symbol type (functions, variables, or both), select a symbol type button.
4. To filter by regular expression (see page 76), enter an expression in the search pattern field.
5. Select *Search for Matching Symbols*. The symbol list will not be updated until you do this.

List of Symbols

The list of symbols shows the symbol name, symbol type, and the source file where it is declared. For example:

```
counter    variable  /project/main.c  
sort      function  /project/quick.c
```

The list of symbols is sorted alphabetically.

Selecting Symbols

1. Use standard Motif Extended Selection to select desired symbols. Selected symbols are highlighted in the symbol list.

2. Or, to select a symbol, select the symbol.
3. To select several contiguous symbols, hold down the Shift key and select and drag a group of symbols.
4. To select non-contiguous symbols, hold down the control key and select and drag a group of symbols.

Adding Symbols to the Defined Ranges List

- Select *Add Selected Symbols to Range List* to add only selected symbols in the current symbol list to the Defined Ranges list.
- Or, select *Add All Symbols to Range List* to add all symbols in the current symbol list to the Defined Ranges list.

To Delete Symbols from the Defined Ranges List

1. Exit the object file symbol selector.
2. Delete the symbol(s) from the Define Ranges dialog.

NOTE:

You cannot delete symbols in the object file symbol selector. So, you must exit the symbol selector, then delete symbols in the State Interval Define Ranges dialog.

Examples of Regular Expressions

- `st*` finds symbols whose names start with *st* or *ST*
- `s?t` finds symbols whose names are 3 letters long, start with *s* or *S*, and end with *t* or *T*
- `[a-z]*` finds symbols whose names start with a letter
- `*[g-h]` finds symbols whose names end with *g* or *G*, or *h* or *H*

Selecting an Object File

Source files and their associated object files are sorted alphabetically by name. To choose a file, select the file name. To select multiple files, highlight and select the file names.

Zooming in on Potential Software Bottlenecks

Use State Interval to characterize the utilization of system resources. This example measurement provides insight into which routines or operations should be optimized to realize significant improvements in system throughput.

You can use State Interval analysis to quickly find which routines are causing the biggest bottlenecks in a system. Software designers who suspect that certain routines are bottleneck candidates can test their assumptions by including these routines in the State Interval measurement.


- If you do not yet know the address ranges, you can import functions and variables from object files.
- If you do not have object files readily available, use State Overview to extract address regions that are monopolizing the CPU.

To display the routines in which most execution time is spent:

1. Set up the appropriate probing (see page 58) for the ADDRESS and STATUS lines using the Agilent Technologies 1655x Acquisition Card.
2. Configure a logic analyzer tool and a SPA tool on the workspace.
3. Create the ADDR and STAT labels in the logic analyzer.
4. If you are using an analysis probe, load the appropriate configuration file supplied with the analysis probe.
5. Configure the logic analyzer to store only valid addresses (see page 57) . You can qualify valid addresses in the logic analyzer trigger menu or in the State Interval Qualify Data (see page 67) menu.
6. Activate the State Interval tool.
7. Select the label (see page 15) named ADDR in the State Interval Define Ranges menu.
8. Define the ranges (see page 65) so that each one corresponds to the start and end address ranges of the program's routines.
9. Turn *Accumulate* on to accumulate data (see page 31).

10. Run Repetitive (see page 23).

The result resembles:



The measurement results indicate that the `lcd_ready()` routine takes the most execution time.

top 3 routines that affect system throughput are `replot()`, `convert()`, and `sort()`. These routines are responsible for 74% of all CPU resources.

The software designer should concentrate optimizing efforts on the routines that take the most execution time then test the results of the code changes with the same State Interval measurement.

See Also

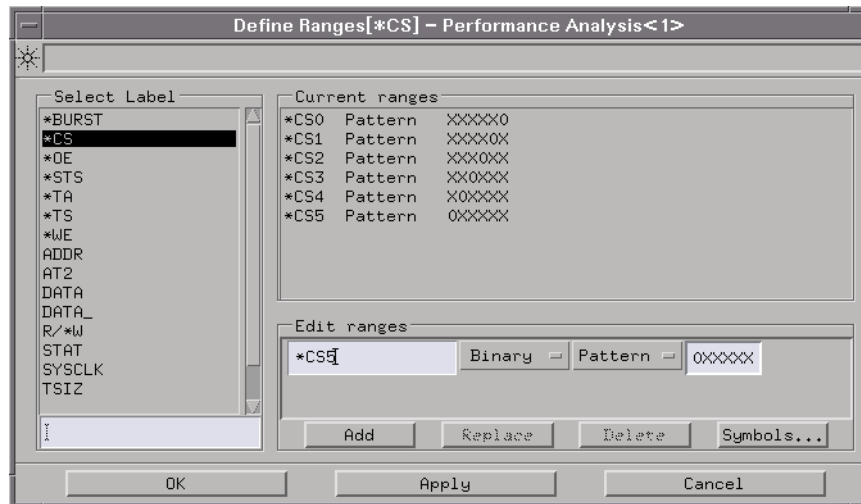
Uses of State Interval (see page 64)

Characterizing Microprocessor CPU Cycles

Microprocessor control signals tell you the type of cycles the processor is executing. You can characterize processor cycles by performing state interval measurements that compare the different cycle types.

For example, you can look at a processor's chip selects to see which devices in the processor's address space are accessed most often.


1. Set up the appropriate probing (see page 58).
2. Configure a logic analyzer tool and a SPA tool on the workspace.
3. Configure the logic analyzer.
4. Format labels for the processor control signals.
5. Activate the State Interval tool.
6. Select the label (see page 15) in the State Interval Define Ranges menu and define ranges:



7. Qualify the data (see page 67) if necessary.
8. Turn *Accumulate* ON to accumulate data (see page 31).

9. Run Repetitive (see page 23).

The result resembles:




By defining the patterns that identify each cycle, the resulting State Interval measurement displays which cycles the processor is executing most often. This gives designers an understanding of processor interface bottlenecks, and where to concentrate optimization efforts.

See Also

Uses of State Interval (see page 64)

Using Time Overview



Time Overview allows you to see where events occur in time. This is useful when you need to determine if the time between events represents a possible system specification violation. Because multiple Time Overview tools can be attached to a single instrument, you can analyze single events or a sequence of events.

Time Overview displays a vertical distribution of points in time when an event occurs, relative to the trigger. Use Markers (see page 83) to measure time intervals. Time Overview markers operate like the Listing tool markers.

See Also

“Defining Time Overview Ranges” on page 85

“Qualifying Time Overview Data” on page 86

“Defining Time Overview Buckets” on page 89

“Running and Stopping a Measurement” on page 23

“Analyzing Time Overview” on page 91

Uses of Time Overview (see page 83)

“Time Overview Information Areas” on page 93

Working with Markers (see the *Markers* help volume)

“Working with Graphs” on page 29

Data Selection Criteria (see page 47)

Time Overview Single-Event Search Criteria (see page 49)

Held Values when Using Multiple Datasets (see page 47)

“Opening and Saving SPA Configurations” on page 10

“Print the SPA Window” on page 11

“Print SPA Data to a File” on page 12

Reading a Marker Data Value

You can obtain the data value at a marker or trigger position by setting the markers to follow globally. You will then use the Listing tool to read the data value at the marker or trigger position in the Listing tool.

1. In the Time Overview graph, place a marker on the bucket of interest.
2. Configure the Listing tool on the same source used by SPA.
3. Open the Listing tool and select *Go to G1 marker*. The data value of interest is underlined with the G1 marker.
4. To calculate the time between the start and end events, expand the bucket until the maximum number of hits per bucket is 1.
5. To find the exact time between two events, place the G1 marker on the start event, and the G2 marker on the end event.
6. In the Time Overview marker window, read the G1 and G2 time values.

Incrementing a Marker Time Value

When you increment a Time Overview marker time value:

- If *there are no hits* in the bucket, the new marker value equals the low time range in the bucket.
- If *there are hits* in the bucket, the marker is placed at the minimum time value found in the bucket.

See Also

Working with Markers (see the *Markers* help volume)

Uses of Time Overview

You can use Time Overview to:

- See where events occur in time.


Using Time Overview

- Analyze asynchronous and multi-processing bus performance.
- Observe burst characteristics (utilization) of bus usage over time.

Defining Time Overview Ranges

Defining ranges allows you to select the number of buckets to graph as the X-axis start and end times.

1. Select the number of buckets (see page 89) to display.
2. Set the X-Axis Range (see page 35).
3. Select *Apply* to apply the changes or *OK* to apply and exit the dialog.



See Also


“Defining Time Overview Buckets” on page 89

Time Overview Single-Event Search Criteria (see page 49)

Qualifying Time Overview Data

You can qualify a single event or two events. These events are used in the search criteria. By default, both events are disabled, resulting in all events being qualified.

1. Enable or disable the *Data qualification* for Event 1.
2. If the data qualification is enabled, select when to qualify the data (see page 19) .
3. Define the event to qualify by selecting the *Label*, *Base*, and *Pattern/Range* buttons.
4. Type in a pattern or range.
5. Enable or disable the *Data qualification* for Event 2. If it is enabled, qualify Event 2.
6. Insert any other labels for Events 1 and 2.
7. Select *Apply* to apply the changes or *OK* to apply and exit the dialog.



NOTE:

Event labels are ANDed to create a single event term. You can use the Qualifying Bus Transitions option (see page 87) to capture any bus transitions in a single event label.

See Also

“Defining Time Overview Buckets” on page 89

Time Overview Single-Event Search Criteria (see page 49)

Time Interval and Time Overview Dual-Event Search Criteria (see page 50)

Qualifying Bus Transitions (see page 87)

Qualifying Bus Transitions

When qualifying data, it is possible to look for any transition on a label (for example, ADDR). To search for a transition on a label:

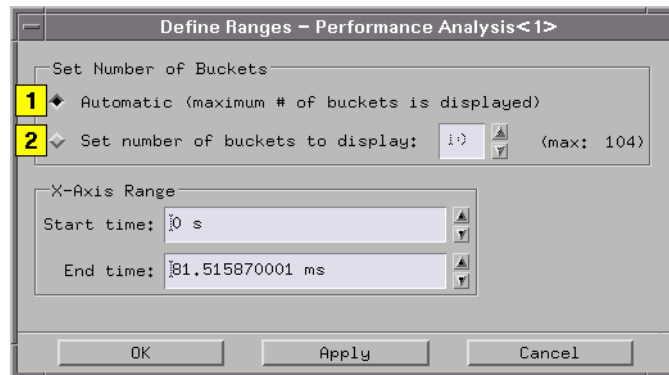
Qualifying Time Overview Data

1. Select *Options* and *Bus on*.
2. Qualify data (see page 19) based upon the Bus Mode Search Criteria (see page 52).

Defining Time Overview Buckets

You can select a specific number of buckets (see page 17) to hold and display data or have SPA automatically calculate the maximum number of buckets.

1. To display the maximum number of buckets, select *Automatic*.
2. To display another number of buckets (see page 16), select *Set number of buckets to display* and enter a number or select the up and down arrows.
3. Select *Apply* to apply the changes or *OK* to apply and exit the dialog.



See Also

Time Overview Bucket Information (see page 89)

Time Overview Single-Event Search Criteria (see page 49)

Time Interval and Time Overview Dual-Event Search Criteria (see page 50)

Time Overview Bucket Information

Bucket information is updated for the highlighted bucket in the graph. When you highlight a bucket (see page 17), you will see:

Hits (%d) - displays the number of hits in selected bucket, and percentage of those hits relative to the qualified hits shown on the

Defining Time Overview Buckets

display.

Hits (%t) - displays number of hits in selected bucket, and percentage of those hits relative to the total qualified hits.


Range - shows low and high data values for selected time bucket.

Min/Max - shows minimum and maximum qualified data values in the selected time bucket.

Mean - displays the mean time for the time values found in the selected buckets.


Standard Deviation - displays the standard deviation for the time values found in the selected buckets.

Analyzing Time Overview



Time Overview displays a bar chart of the number of times an event occurs over time.

How to Interpret the Time Overview



- A. The X axis is the defined time range for the qualified event(s).
- B. The Y axis is the relative number of occurrences of the event(s) in each time bucket. The maximum value of the Y axis is rescaled constantly (see page 92) to reflect the new bucket time values.
- C. The highlighted bucket's information is displayed in the Bucket Information area (see page 94) .
- D. The Display Information (see page 94) is updated after every run.

Information displayed is based on qualified data.

Effects of Defining Ranges

Defining ranges (see page 65) allows you to customize the display by setting the number of buckets and the X-axis start and end times.

Effects of Qualifying Events

Qualified events are the events used in the analysis. You can define when to qualify data (see page 19) and define up to two qualifying event terms (see page 19) . If you define an event on which to qualify the acquired data, only those qualified events are counted. Those events are then bucketized and displayed.

See Also

“Time Overview Information Areas” on page 93

“Working with Graphs” on page 29

“Setting Thresholds” on page 21


Changing Attributes (see page 36)

“Data Selection Criteria” on page 47

Y Axis Is Rescaled Constantly

As successive acquisitions are acquired and qualified events are filtered, the display is constantly rescaled vertically so that the upper limit of the Y axis represents the largest number of occurrences in any bucket.

Time Overview Information Areas



Time Overview work areas include messages, comments, the time graph, and Display and Bucket Information areas. When you select anywhere in the graph area, a menu displays that lets you scale the graph or place markers.

Message Area

The message area to the right of the *Run* key helps you understand interactions allowed as you move the cursor in the entire dialog.

Comment Area

The comment area appears when you select *Options* and *Show Comments* in the menu bar. Use the comment area to document setup information and data results. The size of the comments file is not limited. Comments are saved when you *Print to File*.

Data Area

The data area plots the number of hits for each qualified event. Data is plotted based on the accumulate status, number of buckets, graph attributes, thresholds, rescaling operations, the amount of time represented on the X axis.

Display and Bucket Information Areas

The Display and Bucket Information Areas (see page 94) are just below the graph.

See Also

“Rescaling Graphs” on page 39

“Working with Graphs” on page 29

Time Overview Display and Bucket Information

Time Overview shows display and bucket information for the acquisition.

Display Information	Bucket Information
Qualified Events: 226	Hits(%d): 15 (6.64%)
Total Acquisitions: 1	Hits(%t): 15 (6.64%)
Total Display %: 100%	Range: 9.135399234 ms thru 9.838122251 m
Displayed Events: 226	Min/Max: 9.166580000 ms / 9.823220000 ms
Time per Bucket: 702.723018 us	Mean: 9.492380000 ms
	StdDev: 54.832544 us

Display Information

Qualified States - are the number of states used in the distribution analysis. If the acquisition is set to collect 8K states, and you define an event on which to qualify, only those qualified events count. Those events are then bucketized and displayed.

Total Acquisitions - reflects the number of accumulated and displayed data runs in Accumulate mode. When Accumulate mode is off, the Total Acquisitions field will be 0 if no runs occurred, or 1 if the displayed data represents a single run.

Total Display % - shows the percentage of all qualified events displayed. If you rescale the graph to show only a portion of the event, this number reflects the new percentage of displayed and total qualified events.

Displayed Events - are the number of qualified events displayed on the graph.

Time per Bucket - shows the amount of time represented in each bucket.

Bucket Information

When you highlight a specific bucket (see page 17), the following information about that bucket displays in the Bucket Information area.

Hits (%d) - displays the number of hits in selected bucket, and percentage of those hits relative to the qualified hits shown on the display.

Hits (%t) - displays number of hits in selected bucket, and percentage of those hits relative to the total qualified hits.

Range - shows the low and high time values represented in the selected bucket. The resolution is 1 ps.

Min/Max - shows the minimum and maximum sampled event times found in the selected bucket.


Mean - displays the mean time for the time values found in the selected buckets.

Standard Deviation - displays the standard deviation for the time values found in the selected buckets.

See Also

“Data Selection Criteria” on page 47

Using Time Interval



Time Interval displays a horizontal distribution of execution times for a single event. Use Time Interval to view timing characteristics of routines, CPU cycles, setup and hold times, and bus transactions.

Application Examples • “Characterizing Timing of Interrupt Routines” on page 109

See Also

“Defining Time Interval Ranges” on page 98

Qualifying Time Interval Data (see page 102)

Including Other Time Ranges (see page 100)

“Running and Stopping a Measurement” on page 23

“Analyzing Time Interval” on page 104

Uses of Time Interval (see page 96)

“Time Interval Information Areas” on page 106

“Working with Graphs” on page 29

Data Selection Criteria (see page 47)

“Opening and Saving SPA Configurations” on page 10

“Print the SPA Window” on page 11

“Print SPA Data to a File” on page 12

Uses of Time Interval

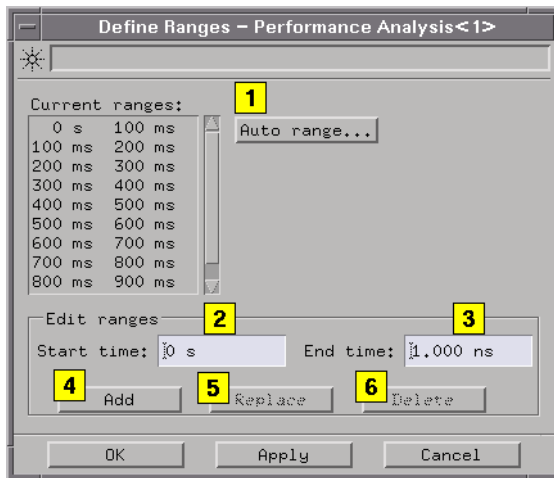
You can use Time Interval to:

- Characterize execution times of a single routine.
- Characterize timing of interrupt routines.
- Profile processor cycle times.
- Characterize critical hardware time relationships, such as setup and hold times.

Defining Time Interval Ranges

Defining ranges allows you to create timing ranges for the qualified event (see page 102).

1. To automatically set the time range (see page 99), select *Auto range*.
2. To define individual ranges, select the *Start time* option button, and enter the starting time.
3. Select the *End time* option button, and enter the ending time.
4. Select *Add* to include the time range in the current ranges list.
5. To change a range, select the range, enter new start and end times, and select *Replace*.
6. To delete a range, select the range and *Delete*.
7. Select *Apply* to apply the changes or *OK* to apply and exit the dialog.



NOTE:

Move the cursor over defined ranges, and select the desired command (undo, sort, and delete).

See Also

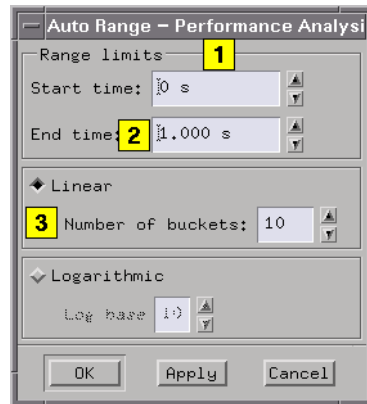
Including Other Time Ranges (see page 100)

Time Interval Search Criteria (see page 50)

Using Auto Range

The Time Interval Auto Range dialog lets you automatically define a set of time interval ranges.

1. Select the *Start time* option button, and enter the starting time.
2. Select the *End time* option button, and enter the ending time.
3. Select either the *Linear* (divisions) and define the number of buckets, or the *Logarithmic* (base) scaling type.
4. Select *Apply* to apply the changes or *OK* to apply and exit the dialog.



Consider these Auto Range settings:

Start time = 0 seconds

End time = 10 seconds

Number of divisions = 10

Range type = Linear

Each time interval range duration will be 1 second, and will span the time range of 0 to 10 seconds. If you select Logarithmic, the number of divisions is determined by the start and end times, and the base. The

maximum range values are from 0 s to 1 Ms.


See Also

“Defining Time Interval Ranges” on page 98

Including Other Time Ranges

When your defined ranges do not cover the entire range of specified times, you can include any other occurring times in the *Other* time range bucket.

1. Select *Include other*.
2. A time range named *Other* is added to the list of time ranges.
3. To delete the *Other* time range, select the bucket.
4. Select *Delete*.



NOTE:

When you run a measurement, any hit's time not currently defined in a bucket will be stored in the *Other* time range bucket. If a range is turned off, any activity in that range is included in the *Other* time range bucket, regardless of whether the *Other* time range bucket is displayed or not.

See Also

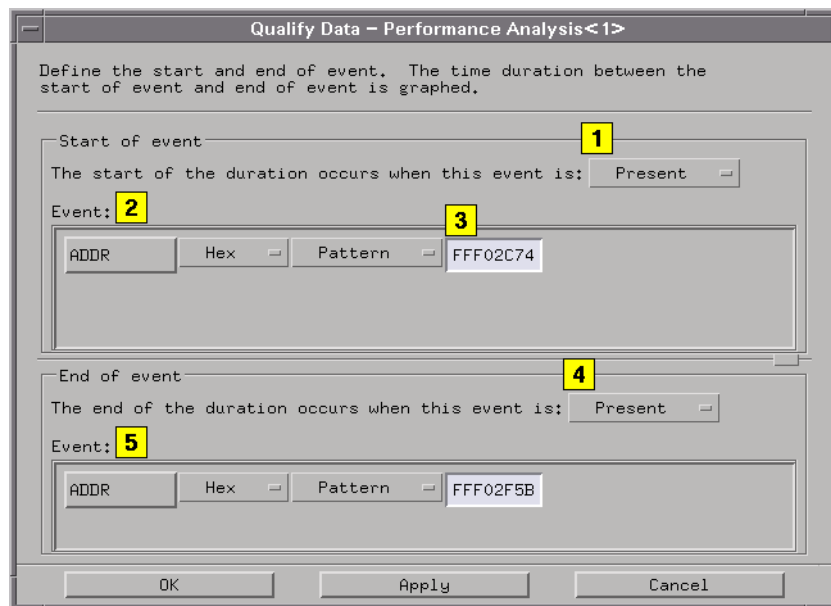
“Defining Time Interval Ranges” on page 98

Time Interval Search Criteria (see page 50)

Qualifying Time Interval Data

Qualifying data allows you to define the start and end of an event. The time duration between the start of event and end of event is used in the Time Interval search criteria.

1. Select when to qualify the data (see page 19) for the start of the event.
2. Define the start of the event by selecting the start of event *Label*, *Base*, and *Pattern/Range* buttons.
3. Type in a pattern or range.
4. Select when to qualify the data for the end of the event.
5. Define the end of the event by selecting the end of event *Label*, *Base*, and *Pattern/Range* buttons.
6. Insert any other labels necessary to qualify the start or end of the event.
7. Select *Apply* to apply the changes or *OK* to apply and exit the dialog.




NOTE: Event labels are ANDed to create a single event term. You can use the Qualifying Bus Transitions option (see page 87) to capture any bus transition on a single event label.

See Also


- Including Other Time Ranges (see page 100)
- Time Interval Search Criteria (see page 50)
- Qualifying Bus Transitions (see page 87)
- Bus Mode Search Criteria (see page 52)

Analyzing Time Interval



Time Interval shows the distribution of time intervals for the event specified with the start-of-event and end-of-event definitions.

How to Interpret the Time Interval



A. You can define start and end time ranges (see page 105) for a selected label. The minimum sample period and time interval resolution is 1 ps.

B. The maximum number of time interval ranges is unlimited.


C. Calculated statistics provide maximum time, minimum time, average time, standard deviation, and total number of time intervals sampled. and minimum times using a logarithmic or linear scale.

Effects of Defining Ranges	Defining ranges (see page 98) allows you to define time intervals for the event of interest.
Effects of Qualifying States	Qualified events are the events used in the analysis. You can define when to qualify data (see page 19) and define two qualifying event terms (see page 19) . If you define an event on which to qualify the acquired data, only those qualified events are counted. Those events are then bucketized and displayed.
See Also	<p>“Time Interval Information Areas” on page 106</p> <p>“Setting Thresholds” on page 21</p> <p>Changing Attributes (see page 36)</p> <p>“Working with Graphs” on page 29</p> <p>“Data Selection Criteria” on page 47</p> <p>Time Interval Search Criteria (see page 50)</p> <p>Bus Mode Search Criteria (see page 52)</p>

Specifying Time Event Ranges

Time event ranges do not have to be contiguous, but gaps between time ranges increase the risk of missed data. Common boundaries of adjacent ranges are not overlapping.

Time Interval Information Areas



Time Interval work areas include messages, comments, the time range area, hit area, and the graph area. When you select a range in the graph area, a menu displays that lets you delete the current label, delete from the current label to the end of the list, or find a label.

Message Area

The message area to the right of the *Run* icon helps you understand interactions allowed as you move the cursor in the entire dialog.

Comment Area

The comment area appears when you select *Options* and *Show Comments* in the menu bar. Use the comment area to document setup information and data results. The size of the comments file is not limited. Comments are saved when you *Print to File*.

Time Range Area

This area of the Time Interval displays defined time ranges for buckets. You can drag and drop the vertical bar to the right of the label column to resize it and the other columns.

Hits Area

This area displays the number percentage of hits for each bucket. If this area is smaller than the text to be displayed, the text is highlighted yellow to indicate the truncation. You can drag and drop the vertical bar to the right of the Hits column to resize it and the graph area.

Graph Area

This area displays time bucket results as a percentage of event

duration hits. When you move the cursor over a time bucket, the bucket information is selected and highlighted with a solid bar. You can drag and drop a time bucket to reposition it in the bucket information list. You can drag and drop the vertical bar to the left of the graph area to rescale the graph.

Find a Time Range Label

When you select over a time range in the graph area, you can select to find a time label in the current label list. As you enter a label name, if the label matching your pattern is displayed onscreen, it moves to the top of the displayed label list. Select *Next* and *Prev* to display the next and previous occurrence of the label.

If there is a space before the time label, you must include the space when using *Find label*. If you enter a label name that does not exist, the labels do not move.

Display and Bucket Information Areas

The Display and Bucket Information Areas (see page 107) are just below the graph.

See Also

“Rescaling Graphs” on page 39

“Working with Graphs” on page 29

Time Interval Display and Bucket Information

Time Interval shows display and bucket information for the acquisition.

Display Information	Bucket Information
Qualified Events: 113	Hits(%): 64(56.64%)
Total Acquisitions: 1	Range: 50.24 us - 58.7 us
Total Display %: 100.00%	Min/Max: 51.29 us/52.89 us
Total Buckets: 10	Mean: 51.803 us
Qualified Min/Max: 41.78 us/126.38	StdDev: 89.311 ns
Qualified Mean: 50.833 us	
Qualified StdDev: 997.8 ns	

Display Information

Qualified Events - are the number of occurrences of start and end pairs.

Total Acquisitions - reflects the number of accumulated and displayed data runs in Accumulate mode. When Accumulate mode is off, the Total Acquisitions field will be 0 if no runs occurred, or 1 if the displayed data represents a single run.

Total Display % - shows the percentage of all event pairs displayed. If you rescale the graph to show only a portion of the event pairs, this number reflects the new percentage of displayed and total qualified event pairs.

Total Buckets - shows the total number of time interval buckets that are defined.

Qualified Min/Max - shows the minimum and maximum of all qualified events.

Qualified Mean - shows the mean of all qualified events.

Qualified StdDev - shows the standard deviation of all qualified events.

Bucket Information

When you highlight a specific bucket (see page 17), the following information about that bucket displays in the Bucket Information area.

Hits (%) - displays the number of event pair hits in the selected time interval bucket and the percentage of those hits relative to the total qualified hits.

Range - shows the low and high time values for the selected bucket.

Min/Max - shows the minimum and maximum time values in the selected bucket.

Standard Deviation - displays the standard deviation for the time values found in the selected buckets.


See Also

“Data Selection Criteria” on page 47

Characterizing Timing of Interrupt Routines

Time Interval provides a distribution of execution times for a single event. This helps software designers understand timing characteristics of critical routines. It also helps hardware designers characterize events such as CPU cycle times, bus arbitration, and setup and hold times.

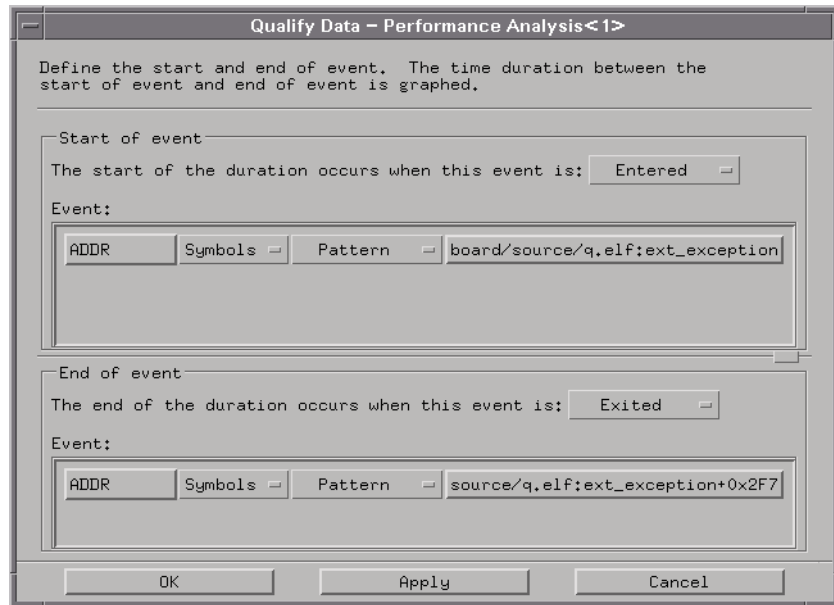
1. Set up the appropriate probing (see page 58) for the ADDRESS and STATUS lines using the Agilent Technologies 1655x Acquisition Card.
2. Configure a logic analyzer tool and a SPA tool on the workspace.
3. Create the ADDR and STAT labels in the logic analyzer.
4. If you are using an analysis probe, load the appropriate configuration file supplied with the analysis probe.
5. Activate the Time Interval tool.
6. Define Ranges (see page 98) using *AutoRange* (see page 99).



7. Qualify the data (see page 102) using the label *ADDR*. Set the Start of Event as *Entered* for the start address, and the End of Event as *Exited* for the end address.

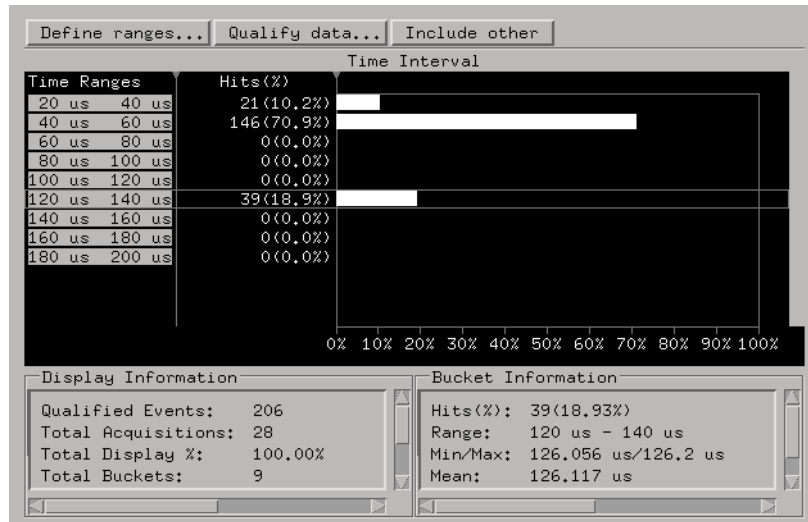
Chapter 1: Using the System Performance Analyzer

Characterizing Timing of Interrupt Routines



8. Turn *Accumulate* ON to accumulate data (see page 31).
9. Run Repetitive (see page 23).

The result resembles:



See Also

Uses of Time Interval (see page 96)

Glossary

absolute Denotes the time period or count of states between a captured state and the trigger state. An absolute count of -10 indicates the state was captured ten states before the trigger state was captured.

acquisition Denotes one complete cycle of data gathering by a measurement module. For example, if you are using an analyzer with 128K memory depth, one complete acquisition will capture and store 128K states in acquisition memory.

analysis probe A probe connected to a microprocessor or standard bus in the device under test. An analysis probe provides an interface between the signals of the microprocessor or standard bus and the inputs of the logic analyzer. Also called a *preprocessor*.

analyzer 1 In a logic analyzer with two *machines*, refers to the machine that is on by default. The default name is *Analyzer<N>*, where N is the slot letter.

analyzer 2 In a logic analyzer with two *machines*, refers to the machine that is off by default. The default name is *Analyzer<N2>*, where N is the slot letter.

arming An instrument tool must be

armed before it can search for its trigger condition. Typically, instruments are armed immediately when *Run* or *Group Run* is selected. You can set up one instrument to arm another using the *Intermodule Window*. In these setups, the second instrument cannot search for its trigger condition until it receives the arming signal from the first instrument. In some analyzer instruments, you can set up one analyzer *machine* to arm the other analyzer machine in the *Trigger Window*.

asterisk (*) See *edge terms*, *glitch*, and *labels*.

bits Bits represent the physical logic analyzer channels. A bit is a *channel* that has or can be assigned to a *label*. A bit is also a position in a label.

card This refers to a single instrument intended for use in the Agilent Technologies 16700A/B-series mainframes. One card fills one slot in the mainframe. A module may comprise a single card or multiple cards cabled together.

channel The entire signal path from the probe tip, through the cable and module, up to the label grouping.

click When using a mouse as the

pointing device, to click an item, position the cursor over the item. Then quickly press and release the *left mouse button*.

clock channel A logic analyzer *channel* that can be used to carry the clock signal. When it is not needed for clock signals, it can be used as a *data channel*, except in the Agilent Technologies 16517A.

context record A context record is a small segment of analyzer memory that stores an event of interest along with the states that immediately preceded it and the states that immediately followed it.

context store If your analyzer can perform context store measurements, you will see a button labeled *Context Store* under the Trigger tab. Typical context store measurements are used to capture writes to a variable or calls to a subroutine, along with the activity preceding and following the events. A context store measurement divides analyzer memory into a series of context records. If you have a 64K analyzer memory and select a 16-state context, the analyzer memory is divided into 4K 16-state context records. If you have a 64K analyzer memory and select a 64-state context, the analyzer memory will be

divided into 1K 64-state records.

count The count function records periods of time or numbers of state transactions between states stored in memory. You can set up the analyzer count function to count occurrences of a selected event during the trace, such as counting how many times a variable is read between each of the writes to the variable. The analyzer can also be set up to count elapsed time, such as counting the time spent executing within a particular function during a run of your target program.

cross triggering Using intermodule capabilities to have measurement modules trigger each other. For example, you can have an external instrument arm a logic analyzer, which subsequently triggers an oscilloscope when it finds the trigger state.

data channel A *channel* that carries data. Data channels cannot be used to clock logic analyzers.

data field A data field in the pattern generator is the data value associated with a single label within a particular data vector.

data set A data set is made up of all labels and data stored in memory of any single analyzer machine or

Glossary

instrument tool. Multiple data sets can be displayed together when sourced into a single display tool. The Filter tool is used to pass on partial data sets to analysis or display tools.

debug mode See *monitor*.

delay The delay function sets the horizontal position of the waveform on the screen for the oscilloscope and timing analyzer. Delay time is measured from the trigger point in seconds or states.

demo mode An emulation control session which is not connected to a real target system. All windows can be viewed, but the data displayed is simulated. To start demo mode, select *Start User Session* from the Emulation Control Interface and enter the demo name in the *Processor Probe LAN Name* field. Select the *Help* button in the *Start User Session* window for details.

deskewing To cancel or nullify the effects of differences between two different internal delay paths for a signal. Deskewing is normally done by routing a single test signal to the inputs of two different modules, then adjusting the Intermodule Skew so that both modules recognize the signal at the same time.

device under test The system under test, which contains the circuitry you are probing. Also known as a *target system*.

don't care For *terms*, a "don't care" means that the state of the signal (high or low) is not relevant to the measurement. The analyzer ignores the state of this signal when determining whether a match occurs on an input label. "Don't care" signals are still sampled and their values can be displayed with the rest of the data. Don't cares are represented by the *X* character in numeric values and the dot (.) in timing edge specifications.

dot (.) See *edge terms*, *glitch*, *labels*, and *don't care*.

double-click When using a mouse as the pointing device, to double-click an item, position the cursor over the item, and then quickly press and release the *left mouse button* twice.

drag and drop Using a Mouse: Position the cursor over the item, and then press and hold the *left mouse button*. While holding the left mouse button down, move the mouse to drag the item to a new location. When the item is positioned where you want it, release the mouse button.

Using the Touchscreen:

Position your finger over the item, then press and hold finger to the screen. While holding the finger down, slide the finger along the screen dragging the item to a new location. When the item is positioned where you want it, release your finger.

edge mode In an oscilloscope, this is the trigger mode that causes a trigger based on a single channel edge, either rising or falling.

edge terms Logic analyzer trigger resources that allow detection of transitions on a signal. An edge term can be set to detect a rising edge, falling edge, or either edge. Some logic analyzers can also detect no edge or a *glitch* on an input signal. Edges are specified by selecting arrows. The dot (.) ignores the bit. The asterisk (*) specifies a glitch on the bit.

emulation module A module within the logic analysis system mainframe that provides an emulation connection to the debug port of a microprocessor. An E5901A emulation module is used with a target interface module (TIM) or an analysis probe. An E5901B emulation module is used with an E5900A emulation probe.

emulation probe The stand-alone equivalent of an *emulation module*. Most of the tasks which can be performed using an emulation module can also be performed using an emulation probe connected to your logic analysis system via a LAN.

emulator An *emulation module* or an *emulation probe*.

Ethernet address See *link-level address*.

events Events are the things you are looking for in your target system. In the logic analyzer interface, they take a single line. Examples of events are *Label1 = XX* and *Timer 1 > 400 ns*.

filter expression The filter expression is the logical *OR* combination of all of the filter terms. States in your data that match the filter expression can be filtered out or passed through the Pattern Filter.

filter term A variable that you define in order to specify which states to filter out or pass through. Filter terms are logically *OR*'ed together to create the filter expression.

Format The selections under the logic analyzer *Format* tab tell the

Glossary

logic analyzer what data you want to collect, such as which channels represent buses (labels) and what logic threshold your signals use.

frame The Agilent Technologies or 16700A/B-series logic analysis system mainframe. See also *logic analysis system*.

gateway address An IP address entered in integer dot notation. The default gateway address is 0.0.0.0, which allows all connections on the local network or subnet. If connections are to be made across networks or subnets, this address must be set to the address of the gateway machine.

glitch A glitch occurs when two or more transitions cross the logic threshold between consecutive timing analyzer samples. You can specify glitch detection by choosing the asterisk (*) for *edge terms* under the timing analyzer Trigger tab.

grouped event A grouped event is a list of *events* that you have grouped, and optionally named. It can be reused in other trigger sequence levels. Only available in Agilent Technologies 16715A or higher logic analyzers.

held value A value that is held until

the next sample. A held value can exist in multiple data sets.

immediate mode In an oscilloscope, the trigger mode that does not require a specific trigger condition such as an edge or a pattern. Use immediate mode when the oscilloscope is armed by another instrument.

interconnect cable Short name for *module/probe interconnect cable*.

intermodule bus The intermodule bus (IMB) is a bus in the frame that allows the measurement modules to communicate with each other. Using the IMB, you can set up one instrument to *arm* another. Data acquired by instruments using the IMB is time-correlated.

intermodule Intermodule is a term used when multiple instrument tools are connected together for the purpose of one instrument arming another. In such a configuration, an arming tree is developed and the group run function is designated to start all instrument tools. Multiple instrument configurations are done in the Intermodule window.

internet address Also called Internet Protocol address or IP address. A 32-bit network address. It

Glossary

is usually represented as decimal numbers separated by periods; for example, 192.35.12.6. Ask your LAN administrator if you need an internet address.

labels Labels are used to group and identify logic analyzer channels. A label consists of a name and an associated bit or group of bits. Labels are created in the Format tab.

line numbers A line number (Line #s) is a special use of *symbols*. Line numbers represent lines in your source file, typically lines that have no unique symbols defined to represent them.

link-level address Also referred to as the Ethernet address, this is the unique address of the LAN interface. This value is set at the factory and cannot be changed. The link-level address of a particular piece of equipment is often printed on a label above the LAN connector. An example of a link-level address in hexadecimal: 0800090012AB.

local session A local session is when you run the logic analysis system using the local display connected to the product hardware.

logic analysis system The Agilent Technologies 16700A/B-series

mainframes, and all tools designed to work with it. Usually used to mean the specific system and tools you are working with right now.

machine Some logic analyzers allow you to set up two measurements at the same time. Each measurement is handled by a different machine. This is represented in the Workspace window by two icons, differentiated by a 1 and a 2 in the upper right-hand corner of the icon. Logic analyzer resources such as pods and trigger terms cannot be shared by the machines.

markers Markers are the green and yellow lines in the display that are labeled *x*, *o*, *G1*, and *G2*. Use them to measure time intervals or sample intervals. Markers are assigned to patterns in order to find patterns or track sequences of states in the data. The *x* and *o* markers are local to the immediate display, while *G1* and *G2* are global between time correlated displays.

master card In a module, the master card controls the data acquisition or output. The logic analysis system references the module by the slot in which the master card is plugged. For example, a 5-card Agilent Technologies 16555D would be referred to as *Slot C*:

Glossary

machine because the master card is in slot C of the mainframe. The other cards of the module are called *expansion cards*.

menu bar The menu bar is located at the top of all windows. Use it to select *File* operations, tool or system *Options*, and tool or system level *Help*.

message bar The message bar displays mouse button functions for the window area or field directly beneath the mouse cursor. Use the mouse and message bar together to prompt yourself to functions and shortcuts.

module/probe interconnect cable

The module/probe interconnect cable connects an E5901B emulation module to an E5900B emulation probe. It provides power and a serial connection. A LAN connection is also required to use the emulation probe.

module An instrument that uses a single timebase in its operation. Modules can have from one to five cards functioning as a single instrument. When a module has more than one card, system window will show the instrument icon in the slot of the *master card*.

monitor When using the Emulation Control Interface, running the monitor means the processor is in debug mode (that is, executing the debug exception) instead of executing the user program.

panning The action of moving the waveform along the timebase by varying the delay value in the Delay field. This action allows you to control the portion of acquisition memory that will be displayed on the screen.

pattern mode In an oscilloscope, the trigger mode that allows you to set the oscilloscope to trigger on a specified combination of input signal levels.

pattern terms Logic analyzer resources that represent single states to be found on labeled sets of bits; for example, an address on the address bus or a status on the status lines.

period (.) See *edge terms*, *glitch*, *labels*, and *don't care*.

pod pair A group of two pods containing 16 channels each, used to physically connect data and clock signals from the unit under test to the analyzer. Pods are assigned by pairs in the analyzer interface. The number of pod pairs available is determined

Glossary

by the channel width of the instrument.

pod See *pod pair*

point To point to an item, move the mouse cursor over the item, or position your finger over the item.

preprocessor See *analysis probe*.

primary branch The primary branch is indicated in the *Trigger sequence step* dialog box as either the *Then find* or *Trigger on* selection. The destination of the primary branch is always the next state in the sequence, except for the Agilent Technologies 16517A. The primary branch has an optional occurrence count field that can be used to count a number of occurrences of the branch condition. See also *secondary branch*.

probe A device to connect the various instruments of the logic analysis system to the target system. There are many types of probes and the one you should use depends on the instrument and your data requirements. As a verb, "to probe" means to attach a probe to the target system.

processor probe See *emulation probe*.

range terms Logic analyzer resources that represent ranges of values to be found on labeled sets of bits. For example, range terms could identify a range of addresses to be found on the address bus or a range of data values to be found on the data bus. In the trigger sequence, range terms are considered to be true when any value within the range occurs.

relative Denotes time period or count of states between the current state and the previous state.

remote display A remote display is a display other than the one connected to the product hardware. Remote displays must be identified to the network through an address location.

remote session A remote session is when you run the logic analyzer using a display that is located away from the product hardware.

right-click When using a mouse for a pointing device, to right-click an item, position the cursor over the item, and then quickly press and release the *right mouse button*.

sample A data sample is a portion of a *data set*, sometimes just one point. When an instrument samples the target system, it is taking a single

Glossary

measurement as part of its data acquisition cycle.

Sampling Use the selections under the logic analyzer Sampling tab to tell the logic analyzer how you want to make measurements, such as State vs. Timing.

secondary branch The secondary branch is indicated in the *Trigger sequence step* dialog box as the *Else on* selection. The destination of the secondary branch can be specified as any other active sequence state. See also *primary branch*.

session A session begins when you start a *local session* or *remote session* from the session manager, and ends when you select *Exit* from the main window. Exiting a session returns all tools to their initial configurations.

skew Skew is the difference in channel delays between measurement channels. Typically, skew between modules is caused by differences in designs of measurement channels, and differences in characteristics of the electronic components within those channels. You should adjust measurement modules to eliminate as much skew as possible so that it does not affect the accuracy of your

measurements.

state measurement In a state measurement, the logic analyzer is clocked by a signal from the system under test. Each time the clock signal becomes valid, the analyzer samples data from the system under test. Since the analyzer is clocked by the system, state measurements are *synchronous* with the test system.

store qualification Store qualification is only available in a *state measurement*, not *timing measurements*. Store qualification allows you to specify the type of information (all samples, no samples, or selected states) to be stored in memory. Use store qualification to prevent memory from being filled with unwanted activity such as no-ops or wait-loops. To set up store qualification, use the *While storing* field in a logic analyzer trigger sequence dialog.

subnet mask A subnet mask blocks out part of an IP address so that the networking software can determine whether the destination host is on a local or remote network. It is usually represented as decimal numbers separated by periods; for example, 255.255.255.0. Ask your LAN administrator if you need a the subnet mask for your network.

symbols Symbols represent patterns and ranges of values found on labeled sets of bits. Two kinds of symbols are available:

- Object file symbols - Symbols from your source code, and symbols generated by your compiler. Object file symbols may represent global variables, functions, labels, and source line numbers.
- User-defined symbols - Symbols you create.

Symbols can be used as *pattern* and *range* terms for:

- Searches in the listing display.
- Triggering in logic analyzers and in the source correlation trigger setup.
- Qualifying data in the filter tool and system performance analysis tool set.

system administrator The system administrator is a person who manages your system, taking care of such tasks as adding peripheral devices, adding new users, and doing system backup. In general, the system administrator is the person you go to with questions about implementing your software.

target system The system under test, which contains the microprocessor you are probing.

terms Terms are variables that can be used in trigger sequences. A term can be a single value on a label or set of labels, any value within a range of values on a label or set of labels, or a glitch or edge transition on bits within a label or set of labels.

TIM A TIM (Target Interface Module) makes connections between the cable from the emulation module or emulation probe and the cable to the debug port on the system under test.

time-correlated Time correlated measurements are measurements involving more than one instrument in which all instruments have a common time or trigger reference.

timer terms Logic analyzer resources that are used to measure the time the trigger sequence remains within one sequence step, or a set of sequence steps. Timers can be used to detect when a condition lasts too long or not long enough. They can be used to measure pulse duration, or duration of a wait loop. A single timer term can be used to delay trigger until a period of time after detection of a significant event.

Glossary

timing measurement In a timing measurement, the logic analyzer samples data at regular intervals according to a clock signal internal to the timing analyzer. Since the analyzer is clocked by a signal that is not related to the system under test, timing measurements capture traces of electrical activity over time. These measurements are *asynchronous* with the test system.

tool icon Tool icons that appear in the workspace are representations of the hardware and software tools selected from the toolbox. If they are placed directly over a current measurement, the tools automatically connect to that measurement. If they are placed on an open area of the main window, you must connect them to a measurement using the mouse.

toolbox The Toolbox is located on the left side of the main window. It is used to display the available hardware and software tools. As you add new tools to your system, their icons will appear in the Toolbox.

tools A tool is a stand-alone piece of functionality. A tool can be an instrument that acquires data, a display for viewing data, or a post-processing analysis helper. Tools are represented as icons in the main window of the interface.

trace See *acquisition*.

trigger sequence A trigger sequence is a sequence of events that you specify. The logic analyzer compares this sequence with the samples it is collecting to determine when to *trigger*.

trigger specification A trigger specification is a set of conditions that must be true before the instrument triggers.

trigger Trigger is an event that occurs immediately after the instrument recognizes a match between the incoming data and the trigger specification. Once trigger occurs, the instrument completes its *acquisition*, including any store qualification that may be specified.

workspace The workspace is the large area under the message bar and to the right of the toolbox. The workspace is where you place the different instrument, display, and analysis tools. Once in the workspace, the tool icons graphically represent a complete picture of the measurements.

zooming In the oscilloscope or timing analyzer, to expand and contract the waveform along the time base by varying the value in the s/Div

field. This action allows you to select specific portions of a particular waveform in acquisition memory that will be displayed on the screen. You can view any portion of the waveform record in acquisition memory.

Index

A

accumulate data mode, 27, 31, 94
acquisition time, viewing, 32
alphabetical sorting, 37
attributes, changing, 36
auto scaling, 41
axis definitions, 24, 91

B

base, 27, 34
base (number), selecting the x-axis, 22
bucket information, 17
bucket Information area, 73
bucket information area, 27, 72, 94, 107, 108
bucket information, time overview, 89
buckets, automatically calculated number of, 16
buckets, defining state overview, 16
buckets, placing cursor over, 27, 73, 94
buckets, sorting on, 37
bus mode search criteria, 52
bus transitions, 52
bus transitions, qualifying, 52

C

center about marker pairs, 43
colors for data distribution bars, 36
comment area, 26, 71, 93, 106
comments, showing, 32
CPU cycles, identifying routines taking most cycles, 56

D

data area, 93
data distribution bars, 31
data graph area, 26, 72

data or events, qualifying, 52
data selection criteria, 47
Data values in buckets, 17
data values in buckets, 89
dialog, resizing, 43
display information area, 27, 72, 94, 107
distribution of data values, 13

E

end time, 35
exceptions, identifying, 59
expand bucket, 41
expand bucket values, 42

F

filter, symbol list, 74
full scale, 41
functions, 74

G

goto marker, 43
graph area, 106
graph types and information areas, 45
graphs, clearing, 33
graphs, expanding, 39
graphs, replotting, 33
graphs, rescaling, 39
graphs, working with, 29

H

held values, 47
help, 38
hit area, 71
hits area, 106
Hits in buckets, 17
hits in buckets, 27, 73, 89, 94, 108
hits, sorting on, 37

I

information areas, 93
information areas in graphs, 45
information areas, display and bucket, 26, 72, 107
information areas, state interval, 71
interrupt routines, characterizing timing of, 109

L

label range area, 71
label, selecting, 15

M

Main System Help Page, 2
markers, placing, 43
mean time, 89
measurement, running and stopping, 23
message area, 26, 71, 93, 106
methods for searching data, 47
microprocessor CPU cycles, characterizing, 80
Min/Max bucket values, 17
min/max bucket values, 27, 42, 73, 89, 94, 108
min/max range values, 34
min/max time values, 35

N

name sort for buckets, 37

O

object file symbol selector, 74
options, rescaling graph, 41
out-of-range time values, 35
out-of-range values, 34

P

pattern, sorting, 37
place marker, 43

plotted data, 26, 93
print options, 30
print to file, 30
print windows, 30
printing SPA windows, 11

Q

qualification criteria for data, 52
qualified events, 108
qualified states, 72
qualifying data, when to, 19

R

range label, finding, 72
Range values in buckets, 17
range values in buckets, 27, 73, 89,
94, 108
range, sorting, 37
ranges, defining state interval, 65
ranges, defining state overview, 15
ranges, effects of defining, 25, 69,
92, 105
ranges, overlapping, 68, 104
repetitive run, 23
rescale display vertically, 24, 91
rules, 34
Run modes, 23

S

sash, use to expand window
portion, 41
save SPA data, 30
saving SPA configurations, 10
search criteria when pattern is
specified, 52
search parameters, symbol list
filtering, 75
single run, 23, 94
software bottlenecks (potential),
analyzing, 78
sort choices, 37
SPA data, printing to a file, 12

SPA graphs, analyzing, 46
SPA menu, closing, 31
SPA tool type, selecting, 33
standard deviation, 89
start time, 35
state interval data, qualifying, 67
state interval measurements, 78,
80
state interval, analyzing, 68
state interval, using, 62
state overview data, qualifying, 18
state overview information areas,
26
state overview measurements, 56,
59
state overview, analyzing, 24
state overview, using, 13
state search criteria, 48
states/events, effects of qualifying,
25, 69, 92, 105
stop run, 23
symbols, adding to defined ranges
list, 76
symbols, selecting, 74

T

threshold colors, 21
threshold hits, 21
thresholds, colors for, 69
thresholds, setting, 21
time difference search criteria, 50
time interval data, qualifying, 102
time interval information areas,
106
time interval measurements, 109
time interval ranges, defined, 98
time interval search criteria, 50
time interval, analyzing, 104
time interval, using, 96
time overview buckets, defining, 89
time overview data, qualifying, 86

time overview information areas,
93
time overview ranges, defining, 85
time overview search criteria, 49
time overview, analyzing, 91
time overview, using, 82
time per bucket, 94
time range area, 106
time range label, finding, 107
total acquisitions, 27, 72, 94, 108
total buckets, 72, 108
total display percentage, 27, 72,
94, 108
transition events on the bus, 52

U

undo scaling, 41

V

variables, 74

X

x-axis range, setting, 34, 35, 36

Publication Number: 5988-7604EN

