NEC

NEC Ftrace Viewer User's Guide (G2AT01E)

SX-Aurora TSUBASA

Proprietary Notice

Introduction

The NEC Ftrace Viewer User's Guide describes the performance analysis tool "NEC Ftrace Viewer" for the SX-Aurora TSUBASA system.

Targeted readers in this document

This document is written mainly for general users and programmers and it assumes that the reader knows Fortran Compiler (nfort), C Compiler (ncc), C++ Compiler (nc++) and NEC MPI.

How to read this document

This document consists of the following chapters. Please refer to the target reader on the right side of the table and read on.

Chapter	Title	Contents	Targeted Readers
Chapter 1	Overview	This chapter explains an overview of NEC Ftrace Viewer.	general users programmers
Chapter 2	Getting Started	This chapter explains a tuning procedure using NEC Ftrace Viewer	general users programmers
Chapter 3	Basic Appearance	This chapter explains the basics such as window layout and arrangement	general users programmers
Chapter 4	Performance Data Visualization	This chapter explains the details of charts and tables	general users programmers
Chapter 5	Troubleshooting	This chapter explains how to solve problems that occur when using NEC Ftrace Viewer.	general users programmers

Related Documents

The following documents have additional information related to NEC Ftrace Viewer.

- Fortran Compiler (nfort)
  Fortran User's Guide
- C Compiler (ncc) and C++ Compiler (nc++)
  C/C++ User's Guide
- NEC MPI
  NEC MPI User's Guide
- FTRACE
  PROGINF/FTRACE User's Guide
- VEO
  The Tutorial and API Reference of Alternative VE Offloading

Conventions

The following conventions are used throughout this document.

- Bold text that starts with hyphen (-)
  It indicates a command option.
- Dollar sign ($) in an example of executing command
  It indicates **bash** (Bourne-Again Shell) prompt.
- **ncc/nc++/nfort/mpincc/mpinc++/mpinfort**
  In this document, these commands to compile and link a program are called "compiler".
- Text surrounded by square brackets ([ ])
  It indicates a menu or a button. For example, [File] indicates the "File" menu or the "File" button. [File]→[Exit] indicates the "Exit" sub-menu in the "File" menu.

Contents

Chapter 1 Overview 1

1.1 Performance Tuning 1
1.2 NEC Ftrace Viewer 1

Chapter 2 Getting Started 3

2.1 Tuning Procedure 3
2.2 Creating an Executable to Measure Performance 3
2.3 Running Program and Collecting Analysis Information Files 4
2.4 Checking Analysis Information Files 5

Chapter 3 Basic Appearance 7

3.1 Window Layout 7
3.2 Checking Loaded Analysis Information Files 7
3.3 Opening a New Chart or Table 8
3.4 Chart and Table Arrangement 9
3.4.1 Moving a View 10
3.4.2 Minimize and Maximize a Part 11
3.4.3 Displaying as Fast View 12
3.4.4 Closing a View 12
3.5 Loading New Analysis Information Files 12
3.5.1 Loading specified files 13
3.5.2 Loading a specified directory 13
3.6 Loading Analysis Information Files on Startup 14

Chapter 4 Performance Data Visualization 15

4.1 Visualized Elements 15
4.2 Metrics 16
4.3 Chart 19
4.3.1 Chart Basics 19
4.3.2 Function Metric Chart 21
4.3.3 Function Breakdown Chart 22
4.3.4 Function Statistics Chart 23
4.3.5 Process Metrics Chart 25
4.3.6 Process Breakdown Chart 25
4.3.7 MPI Communication Chart 27

4.3.8 Process Histogram 28
4.3.9 Selecting Metric 28
4.3.10 Selecting Process 29
4.3.11 Selecting Function 31
4.3.12 Selecting Colors Assigned to Metrics and Functions 33
4.3.13 Exporting as Image 33
4.4 Table 34
4.4.1 Table Basics 35
4.4.2 Function Table 37
4.4.3 User Region Table 38
4.4.4 Selecting and Sorting Metrics 38
4.4.5 Selecting a method of drill down 39
Chapter 5 Troubleshooting 41
Appendix A Revision History 43
Appendix B Trademarks and Licenses 45
B.1 Trademarks 45
B.2 Licenses 45

Chapter 1 Overview

1.1 Performance Tuning

The improvement of a program to be executed faster and more effectively is called "Performance Tuning". Typically, a program can be executed much faster by improving the point where a large amount of execution time is spent. But it becomes difficult to find the point and understand why the point spends a large amount of execution time as a program become complex. It is ineffective to modify the inappropriate point by the inappropriate method. Therefore, for performance tuning, you need to know performance information of a program accurately by support tools.

1.2 NEC Ftrace Viewer

NEC Ftrace Viewer is a performance analysis GUI tool for Fortran, C, and C++ programs. This tool has the following features.

- This tool graphically displays a result of performance measurement by "Simple Performance Analysis Function" (ftrace function).
- This tool supports programs using automatic, OpenMP and MPI parallelization and displays exclusive execution time of functions/subroutines for each process/thread, and MPI communication time, in order to make it easier to identify performance bottlenecks and load imbalances.
- This tool supports performance analysis from a single thread program to a large scale parallel program by narrowing down the functions, threads, or processes based on performance information and a range of charts and tables.

Chapter 2 Getting Started

2.1 Tuning Procedure

The tuning procedure by NEC Ftrace Viewer is as follows. Use NEC Ftrace Viewer to check analysis information files. If you find performance issues in the program, fix the issues and repeat the tuning procedure. The tuning is finished when a performance of the program achieves a goal.

(1) Creating an executable to measure performance

(2) Running program and collecting analysis information files

(3) Checking analysis information files

2.2 Creating an Executable to Measure Performance

Execute a compiler with **-ftrace** as follows. An executable file to measure performance is created.

Example In the case of using Fortran Compiler (nfort)

```
$ nfort -ftrace a.f90
```

Example In the case of using C Compiler (ncc)

```
$ ncc -ftrace a.c
```

Example In the case of using C++ Compiler (nc++)

```
$ nc++ -ftrace a.cc
```

Example In the case of using MPI shell script for Fortran compiler (nfort)

```
$ mpinfort -ftrace a.f90
```

Example In the case of using MPI shell script for C Compiler (ncc)

```
$ mpincc -ftrace a.c
```

Example In the case of using MPI shell script for C++ Compiler (nc++)

```
$ mpinc++ -ftrace a.cc
```

2.3 Running Program and Collecting Analysis Information Files

Execute the executable file to measure performance and collect analysis information files. You can execute the executable file in the same way as a normal executable file. At the end of execution, the executable file outputs one or more analysis information files in a working directory where the executable file is executed. In the case of non-MPI program, a single analysis information file is created. In the case of MPI program, analysis information files are created for each MPI process. The file name is given as follows.

- In the case of non-MPI Program
  ftrace.out
- In the case of MPI Program
  ftrace.out.*univ.rank*
  where *univ* is the value of MPIUNIVERSE and *rank* is the value of MPIRANK, which are environment variables given by NEC MPI.

Example In the case of non-MPI Program

```
$ /opt/nec/ve/bin/ve_exec ./a.out        (Execute an executable file)
$ ls ftrace.out
ftrace.out
```

Example In the case of MPI program

```
$ mpirun -np 4 /opt/nec/ve/bin/ve_exec ./a.out
(Execute an executable file)
$ ls ftrace.out.*
ftrace.out.0.0 ftrace.out.0.1 ftrace.out.0.2 ftrace.out.0.3
```

2.4 Checking Analysis Information Files

Start NEC Ftrace Viewer as follows.

```
$ /opt/nec/ve/ftraceviewer/ftraceviewer
$ /opt/nec/ve/ftraceviewer/ftraceviewer-c
(start with C locale (LANG=C)
```

When NEC Ftrace Viewer is started, the following window appears.

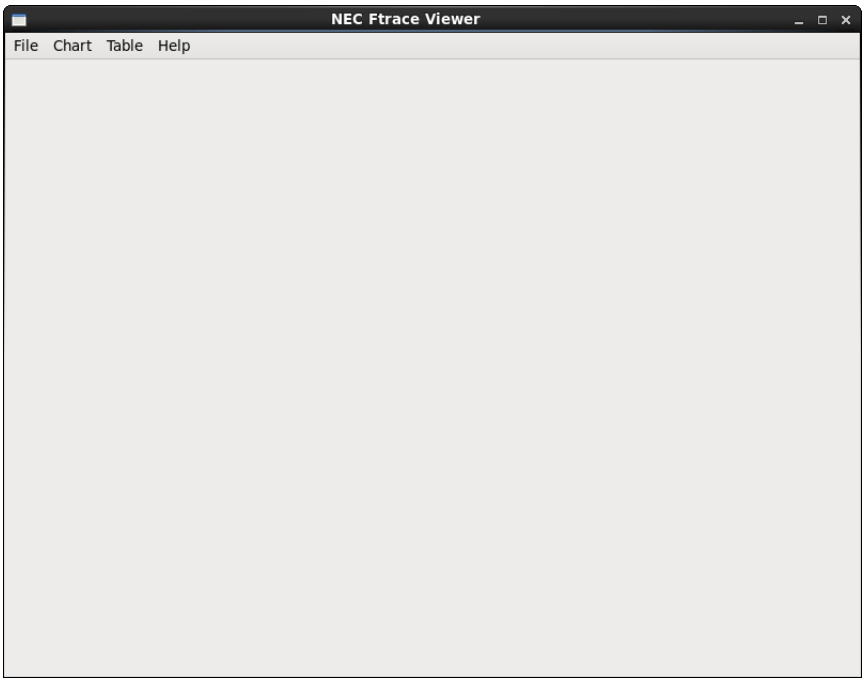

Open and Load the analysis information file by the following menu at the top of the window.

- In the case of non-MPI program
  Choose [File]→[Open File]
- In the case of MPI program
  Choose [File]→[Open Directory]

Note

In the case of VEO program, choose [File]→[Open File] and load single analysis information file same as non-MPI program.

When file selection dialog is garbled, you need to install your locale font packages or start NEC Ftrace Viewer with C locale (LANG=C) using /opt/nec/ve/ftraceviewer/ftraceviewer-c command.

After loading, a chart view appears in the upper part of the window and a table view appears in the lower part of the window. The basic appearance such as window layout and arrangement are explained in Chapter 3. The details of charts and tables are explained in Chapter 4. In order to terminate NEC Ftrace Viewer, choose [File]→[Exit].

Chapter 3 Basic Appearance

3.1 Window Layout

After loading analysis information files, the window consists of the following elements.

(1) The window title

(2) The window menu bar

(3) The area of chart views

(4) The area of table views

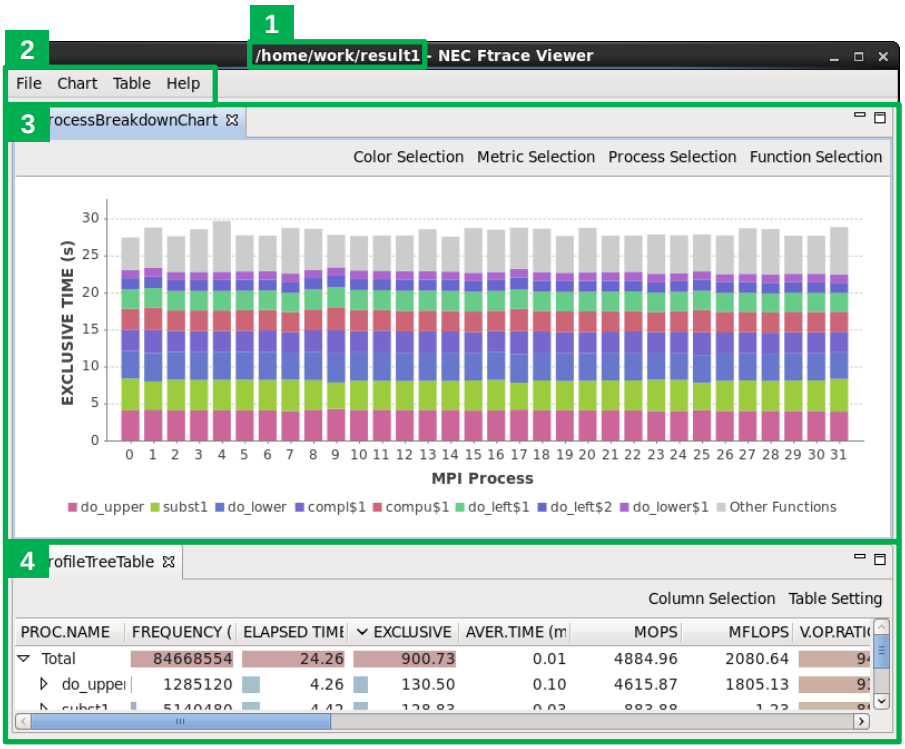

3.2 Checking Loaded Analysis Information Files

The window title displays the full path name of one or more loaded analysis information files or a directory name from where the analysis information files are loaded and the number of the loaded analysis information files.

- In the case of loading an analysis information file by [File]→[Open File]

  The full path name of the specified file is displayed.

- In the case of loading analysis information files by [File]→[Open File]

  The concatenated string with "~" character between the full path name of the first specified file and the name of the last one is displayed.

- In the case of loading analysis information files by [File]→[Open Directory]

  The full path name of the specified directory is displayed.

3.3 Opening a New Chart or Table

In order to open a new chart, choose a sub-menu in [Chart]. A tab for the chart is added to the area of chart views. The details of charts are explained in “4.3 Chart”.

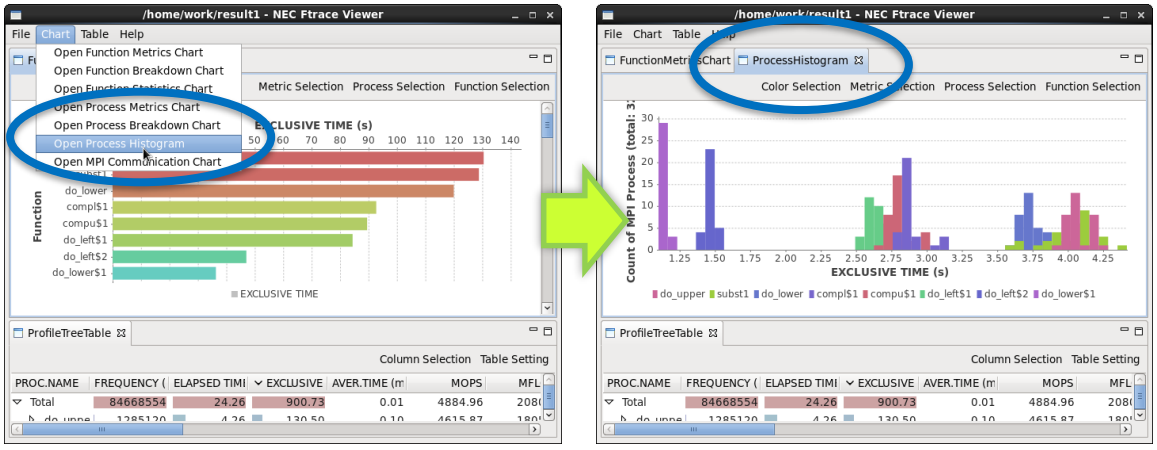

In order to open a new table, choose [Table]→[Open Profile Tree Table]. A tab for new table is added in the area of table views. The details of tables are explained in "4.4 Table".

3.4 Chart and Table Arrangement

In NEC Ftrace Viewer, the window has one or more pages, and each page has one or more views that display a chart or table. You can arrange the charts and tables by manipulating the pages or views. The page and view consist of the following elements.

(1) The page

(2) The tab

(3) The close button

(4) The minimize button

(5) The maximize button

(6) The restore button

(7) The fast view button

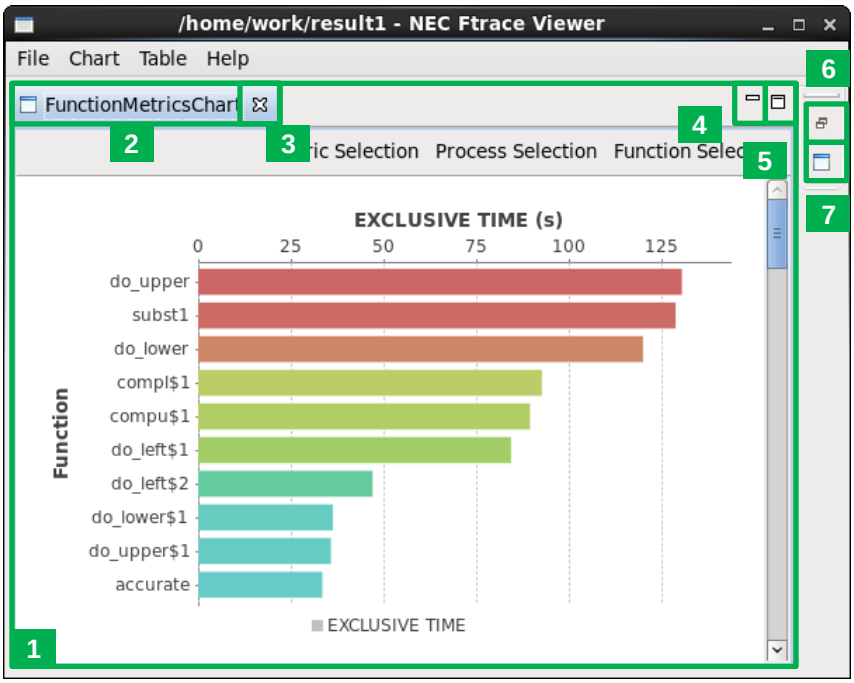

3.4.1 Moving a View

A view can be moved between pages. If you drag and drop the tab of the view to the edge of the page, the page is split and the multiple views are displayed at the same time.

If you want to turn the view back to the original page, drag and drop the view to after or before the tab in the original page. After that, if the other page has no view, the page disappeares.

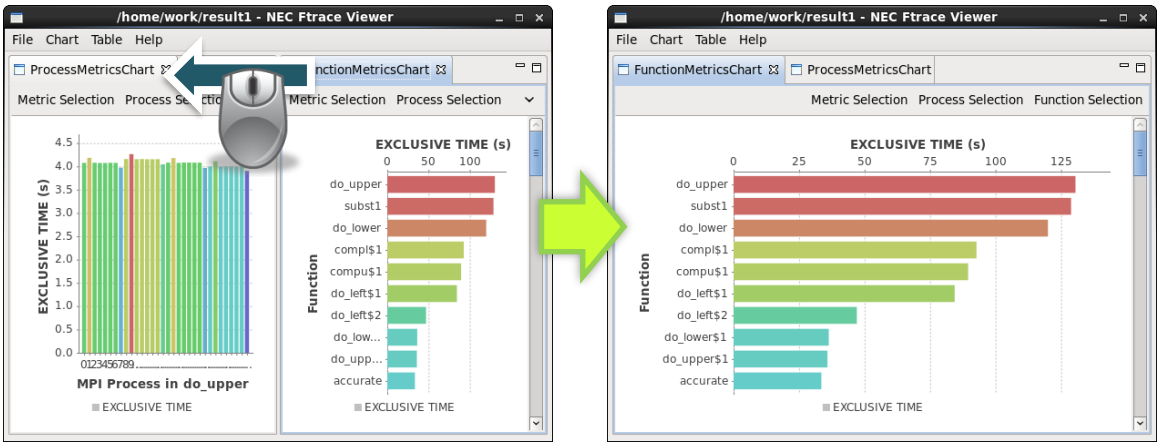

If you drag and drop the tab out of the window, the view is split off from the original window. If you want to turn the view back to the original window, drag and drop the tab of the view to the tab in the original page.

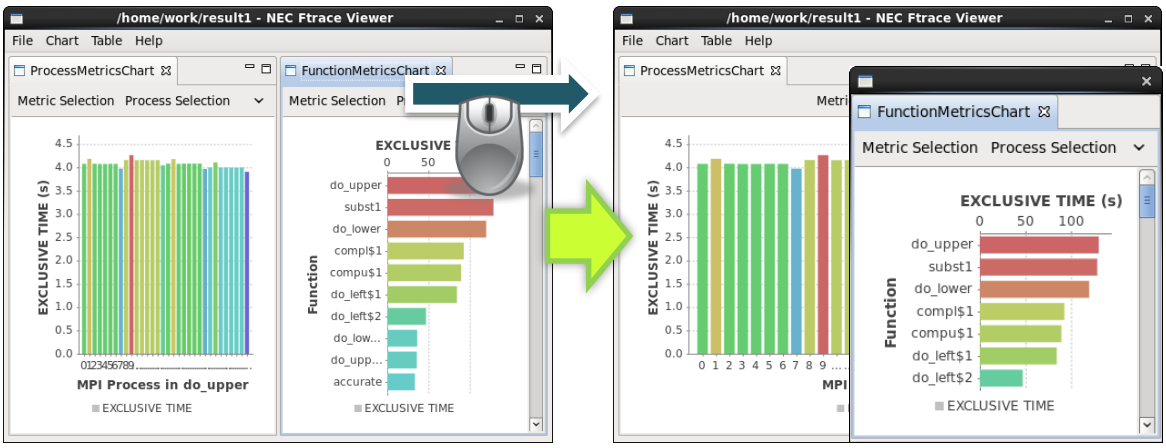

3.4.2 Minimize and Maximize a Part

When the minimize button is clicked in a page, the page is stored in the right or bottom of the window. Click the restore button to return to the stored page. When the maximize button is clicked, all other pages are stored. Double-click the tab in the page to minimize and maximize the page.

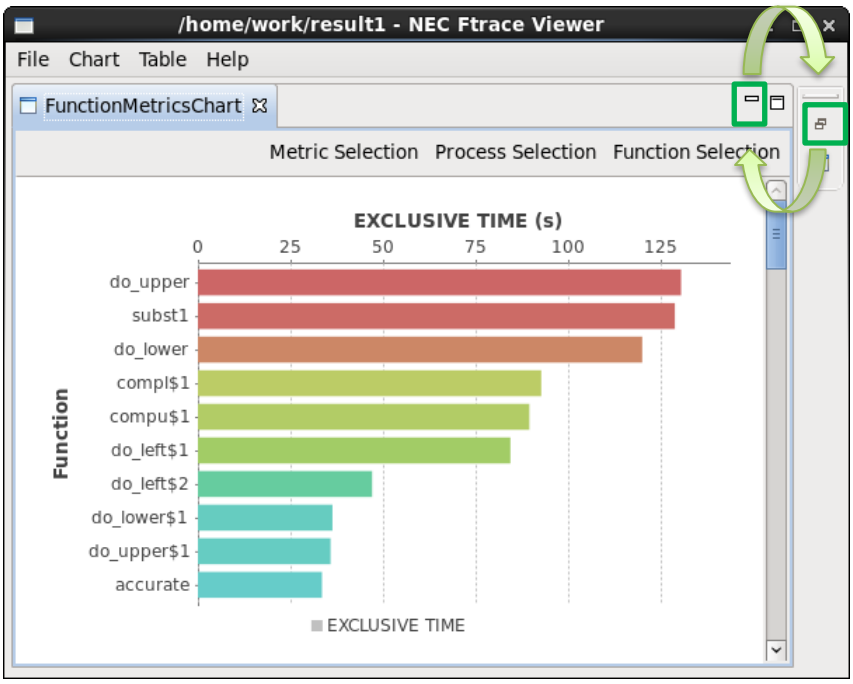

3.4.3 Displaying as Fast View

In a minimized page, the fast view button appears in each view. Click the fast view button to resume the view onto the current page.

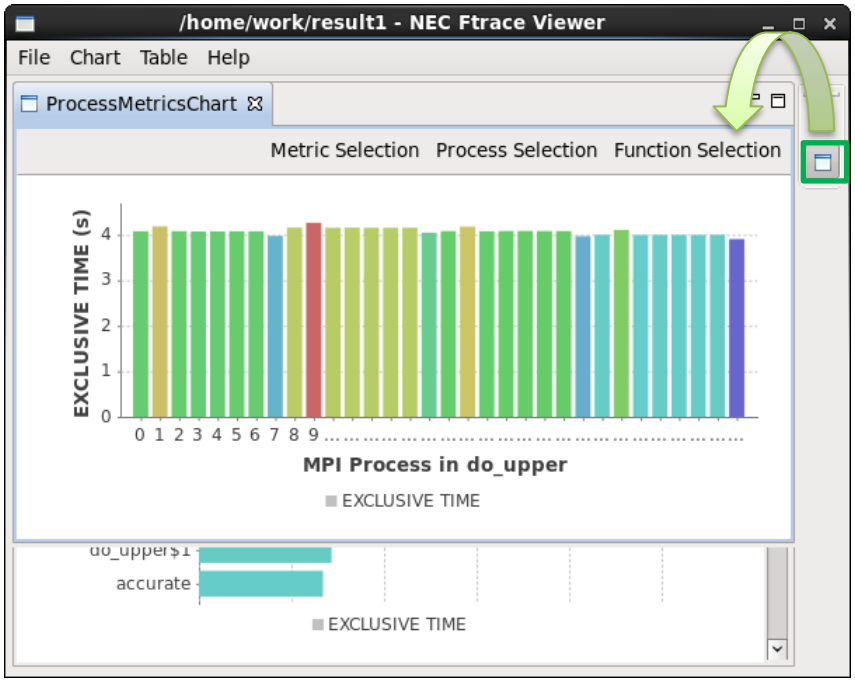

3.4.4 Closing a View

Click the close button. You can open the closed view again from the menu.

3.5 Loading New Analysis Information Files

When [File]→[Open File] or [File]→[Open Directory] is chosen, The previous analysis information files are closed and new analysis information files are loaded.

Note

If you want to leave the previous analysis information files open, start another NEC Ftrace Viewer and load analysis information files over there.

3.5.1 Loading specified files

Choose [File]→[Open File] to open the dialog and specify one or more analysis information files. Only "ftrace.out*" filtered files are shown by default in the dialog. If you change file name, change the filter to "*.*".

Note

Clicking files holding Ctrl key allow you to select multiple files that are anywhere on your file list and Clicking files holding Shift key allow you to select a group of files that are contiguous on your file list.

The loading would be failed if the analysis information file of non-MPI program and that of MPI program are selected at once. Also, the incorrect result may be shown if multiple of analysis information files of VEO program are loaded at the same time.

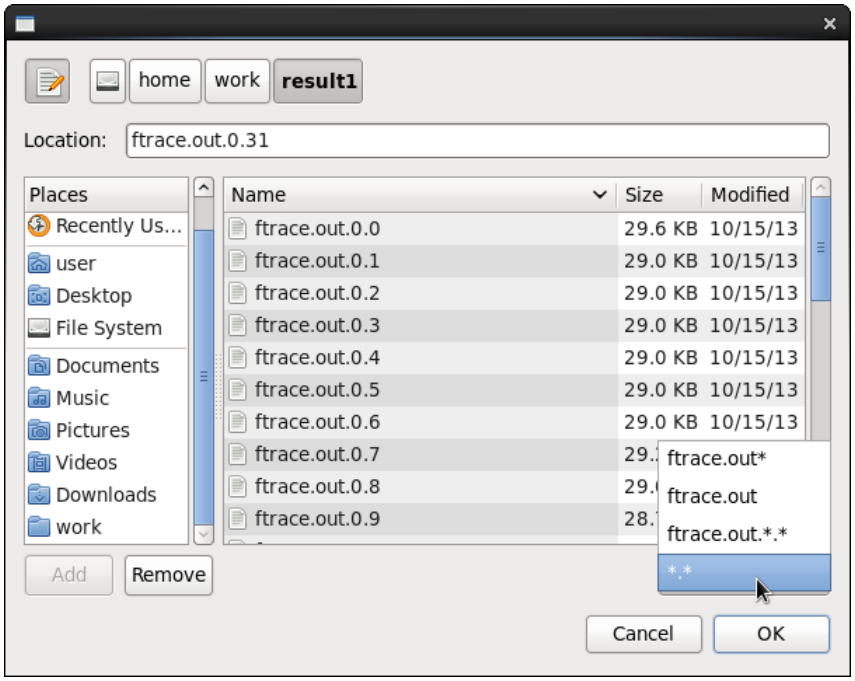

3.5.2 Loading a specified directory

Choose [File]→[Open Directory] to open the dialog and specify a directory. All analysis information files under the specified directory except the sub-directories are loaded.

Note

The loading would be failed if there are the analysis information file of non-MPI program and that of MPI program in the specified directory at once.

3.6 Loading Analysis Information Files on Startup

NEC Ftrace Viewer loads the specified analysis information files on startup by executing the **fv** command with -f option as follows

Example In the case of non-MPI Program

```
$ /opt/nec/ve/bin/ftraceviewer/ftraceviewer -f ftrace.out
```

Example In the case of MPI Program

```
$ /opt/nec/ve/bin/ftraceviewer/ftraceviewer -f ftrace.out.*
```

Chapter 4 Performance Data Visualization

4.1 Visualized Elements

In NEC Ftrace Viewer, the following elements by ftrace function are visualized.

- Execution method
  - Single-thread execution
  - Shared memory parallel execution by automatic-parallelized program or OpenMP-parallelized program
  - Distributed memory parallel execution by MPI-parallelized program
  - Hybrid parallel execution, that is means the combination of shared and distributed memory parallel executions.
- Execution context
  - Thread by shared memory parallel execution
  - MPI process by distributed memory parallel execution
  - MPI processes in MPI universe by distributed memory parallel execution
  - Thread by hybrid parallel execution
  - MPI process by hybrid parallel execution
  - MPI processes in MPI universe by hybrid parallel execution
- Execution code
  - Function, procedure, or subroutine
  - Parallelized-region by automatic or OpenMP parallelization
  - User-specified region by ftrace function (e.g. user region)

4.2 Metrics

In NEC Ftrace Viewer, visualizations of the following metrics are available. The metrics are the same as ftrace function.

(*1) the item is output when specifying the analysis information files that the program executed on VE30

Metric	Unit	Description
FREQUENCY	-	Calling count of a function
EXCLUSIVE TIME	second	Exclusive CPU time to execute one function
AVER.TIME	millisecond	Average CPU time required to execute one function
MOPS	-	MOPS value
MFLOPS	-	MFLOPS value
V.OP.RATIO	percent	Vector operation ratio
AVER.V.LEN	-	Average vector length
VECTOR TIME	second	Vector instruction execution time
L1CACHE MISS	second	L1 cache miss time
CPU PORT CONF	second	CPU port conflict time at memory access
LD L3 HIT ELEM.%	percent	Ratio of the elements hitting L3 cache by load instructions (*1)
VLD LLC HIT ELEM.%	percent	Ratio of the elements hitting LLC by vector load instructions
ELAPSED TIME	second	Elapsed time
MPI COMM.TIME	second	Send/receive elapsed time in MPI communication. This also includes the send/receive time of MPI functions.
MPI IDLE TIME	second	The idle time before MPI communication, and elapsed time of synchronization

Metric	Unit	Description
MPI COMM.TIME/ELAPSED	percent	The ratio of MPI COMM.TIME to ELAPSED TIME in a function
MPI IDLE TIME/ELAPSED	percent	The ratio of MPI IDLE TIME to ELAPSED TIME in a function
MPI AVER.LEN	kilobyte (using base 1024)	The average communication amount per MPI communication (MPI TOTAL LEN / MPI COUNT).
MPI COUNT	-	The MPI communication count
MPI TOTAL LEN	megabyte (using base 1024)	The total MPI communication amount

In the case of analysis information files collected with VE_PERF_MODE=VECTOR-MEM, the following metrics are available instead of L1CACHE MISS, CPU PORT CONF., VLD LLC HIT E.%.
(*1) the value over 100 is truncated.
(*2) the item is output when specifying the analysis information files that the program executed on VE30
(*3) the item is output when specifying the analysis information files that the program was executed on VE10, VE10E or VE20

Metric	Unit	Description
L1ICACHE MISS	second	L1 instruction cache miss time
L1OCACHE MISS	second	L1 operand cache miss time
L2CACHE MISS	second	L2 cache miss time
LD L3 HIT ELEM.%	percent	Ratio of the elements hitting L3 cache by load instructions (*2)
VLD LLC HIT ELEM.%	percent	Ratio of the elements hitting LLC by vector load instructions (*2)
REQ. B/F	-	B/F calculated from bytes specified by load and store instructions (*1)
REQ. ST B/F	-	B/F calculated from bytes specified by store instructions (*1)
REQ. LD B/F	-	B/F calculated from bytes specified by load instructions (*1)
ACT. LD B/F	-	B/F calculated from bytes of actual memory access by load instructions (*1) (*2)
ACT. VLD B/F	-	B/F calculated from bytes of actual memory access by vector load instructions (*1) (*3)
FLOP COUNT	-	Number of floating-point data execution elements
FMA ELEM.	-	Number of FMA execution elements

4.3 Chart

Using charts, performance data for each execution code or execution context can be visualized in various formats. The following charts are available.

- Charts where performance data are represented for each execution code to find the code where a large portion of execution time is spent.
  - Function Metric Chart
  - Function Breakdown Chart (automatic or OpenMP-parallelized program only)
  - Function Statistics Chart
- Charts where performance data are represented for each execution context to find load imbalance between MPI processes or threads.
  - Process Metric Chart
  - Process Breakdown Chart
  - MPI Communication Chart (MPI program only)
- Chart where a performance distribution of MPI processes is represented to find MPI processes with bad performance in a large amount of MPI processes.
  - Process Histogram (MPI program only)

In charts, unless otherwise specified, the execution code is also called "function" and the execution context is also called "process".

4.3.1 Chart Basics

In charts where performance data are represented for each process, you can zoom in or out the bar chart. When mouse wheel is scrolled up pressing "Ctrl" key on the bar chart, the bar chart is zoomed

in. When mouse wheel is scrolled down pressing "Ctrl" key on the bar chart, the bar chart is zoomed out.

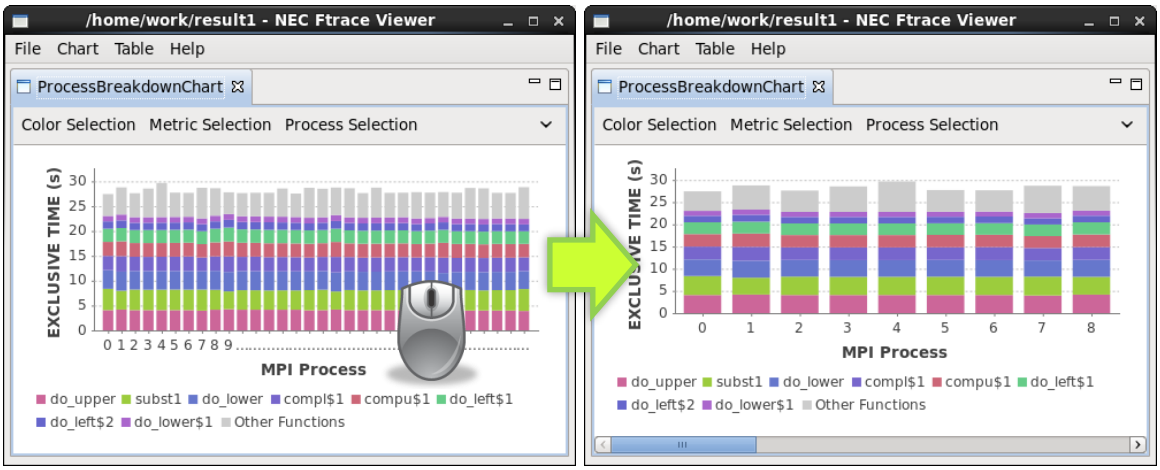

You can select metrics, processes and functions on a chart by the chart menu. The selection is explained in "4.3.9 Selecting Metric", "4.3.10 Selecting Process" and "4.3.11 Selecting Function".

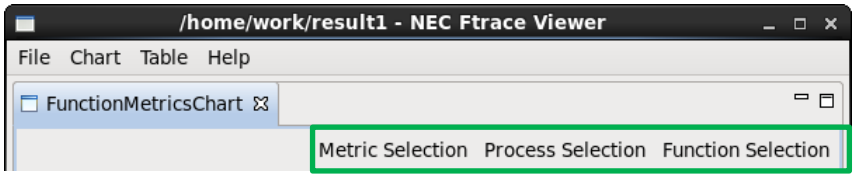

4.3.2 Function Metric Chart

This chart represents performance data for each function with a colored bar and lines. The axis of a bar is displayed in the upper hand. The axes of a line are displayed in lower hand. You can select multiple metrics and each metric is assienged to a line individually.

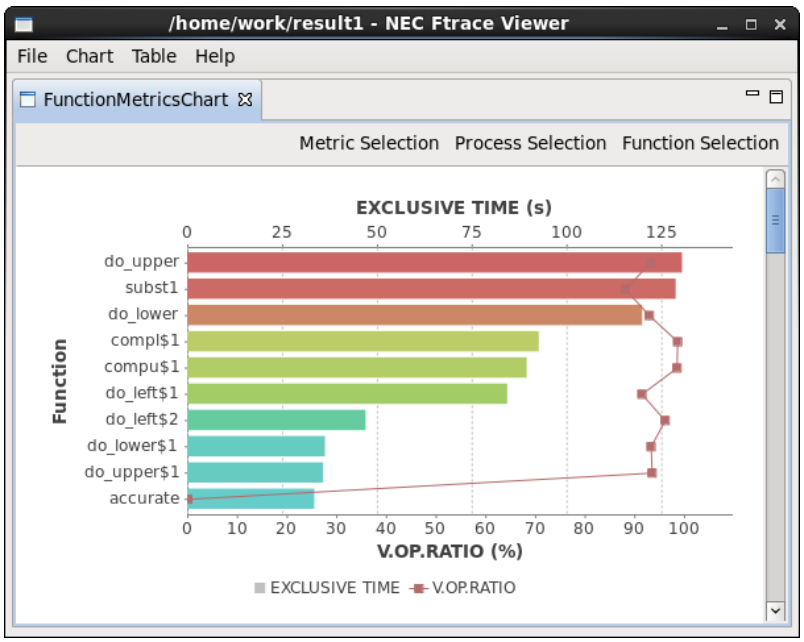

A color of bar represents the ratio of a performance data in selected funcions. In the case of red, the performance data of the function is the largest. In the case of green, the performance data of the function is in the middle. In the case of blue, the performance data of the function is the smallest. In the following chart, the performance data of the function "subst1" is smallest and the bar is painted in blue because "do_upper" and "subst1" are selected.

Otherwise, in the following chart, the performance data of the parallelized-region "compl$1" is smallest and the bar is painted in blue because the functions from "do_upper" to "compl$1" are selected.

4.3.3 Function Breakdown Chart

This chart represents performance data for each function with a bar which is painted on a per-thread. For example, in the following chart, the function "subst1" which is executed in 1 thread is painted in 1 color, and the parallelized-region "compl$1" which is executed in 4 threads is painted in 4 colors.

You can select multiple metrics in the chart as well as Function Metrics Chart.

Note

In a bar, the metric which is possible to be accumulated is only available.

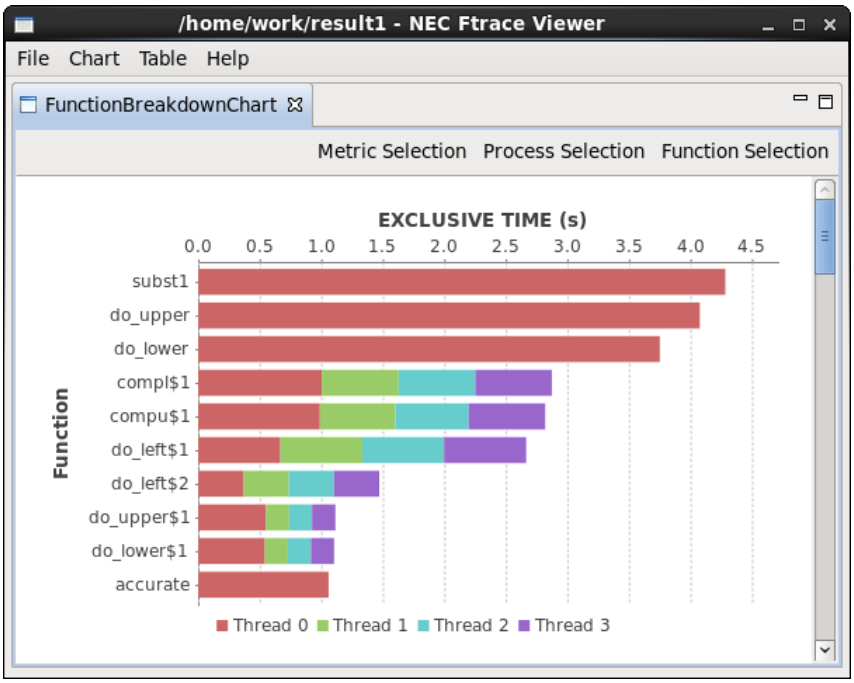

4.3.4 Function Statistics Chart

This chart represents the maximum, minimum, average, and standard deviation of a performance data the in processes for each function.

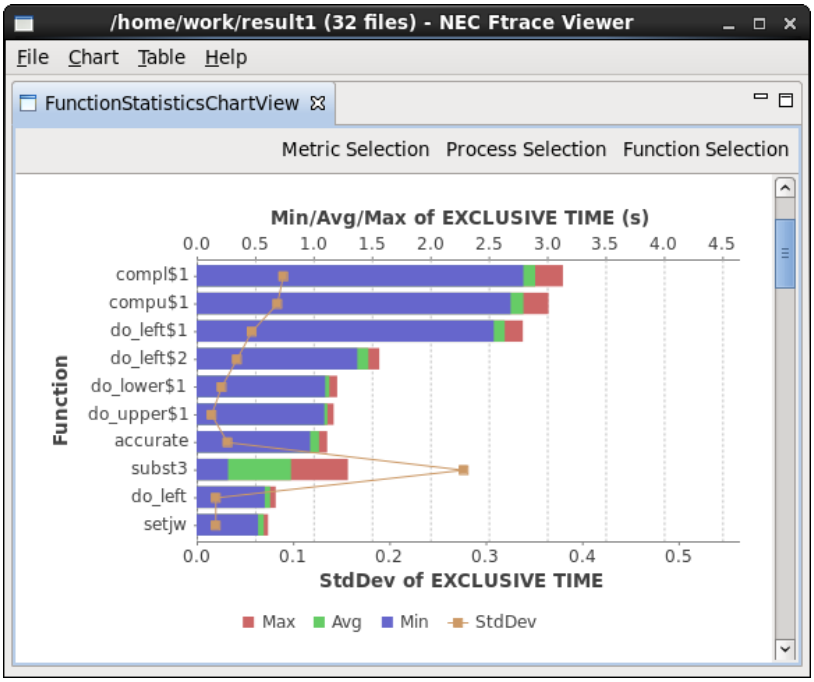

The bar represents maximum, minimum and average as follows. With respect to the left end, the boundary between blue and green is minimum (Min), the boundary between green and red is average (Avg) and the right end is maximum (Max).

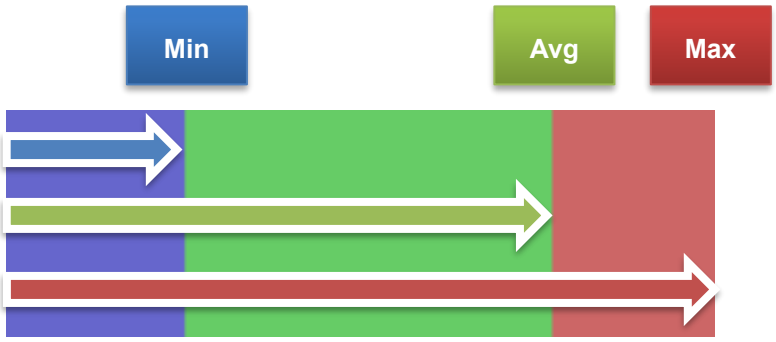

If the red part is large, the performance data of some process or threads are much larger than one of the others. For example, in the following chart, the red part of the parallelized-region "compl$1" is large because the execlusive time of 0-th thread is much largeer than the other threads.

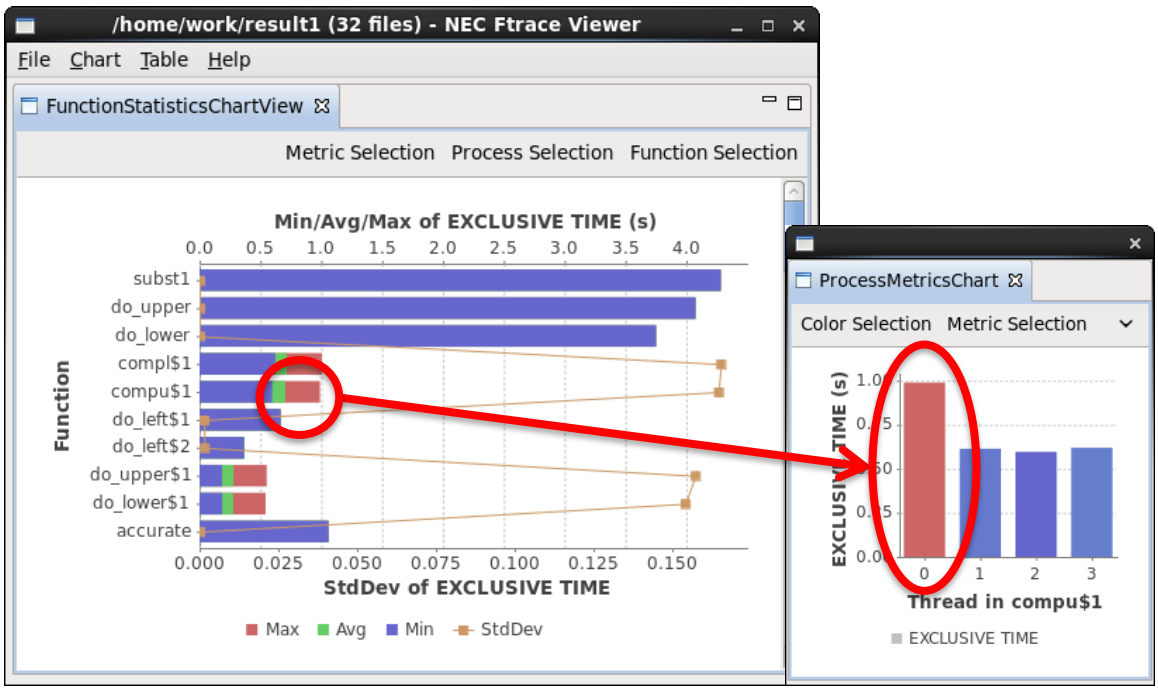

The line is represent standard deviation of a performance data. If the value is large, the performance data between processes or threas are vary widely.

Note

In the chart as well as Function Metrics Chart, functions are sorted in the order of a performance data of whole selected processes or threads. So, it may seem that the functions where the performance data of some process or threads are much larger than one of the others are not sorted as follows.

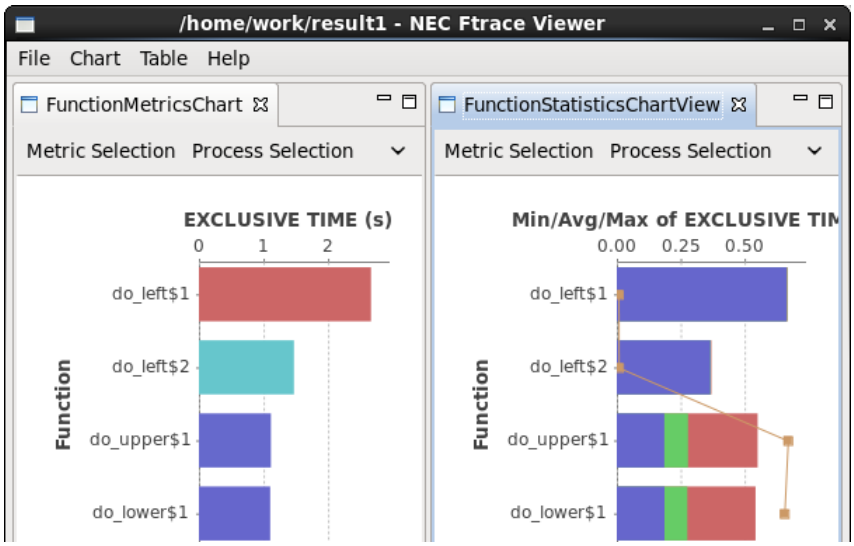

4.3.5 Process Metrics Chart

This chart represents performance data for each processes with a colored bar and lines in the specified function. The axis of a bar is displayed in the left hand. The axes of a line are displayed in right hand. You can select multiple metrics and each metric is assienged to a line individually.

A color of bar represents the ratio of a performance data in selected processes. In the case of red, the performance data of the process is the largest. In the case of green, the performance data of the process is in the middle. In the case of blue, the performance data of the process is the smallest. In the following chart, the performance data of 9-th MPI process is largest and that of 31-th MPI process is smallest.

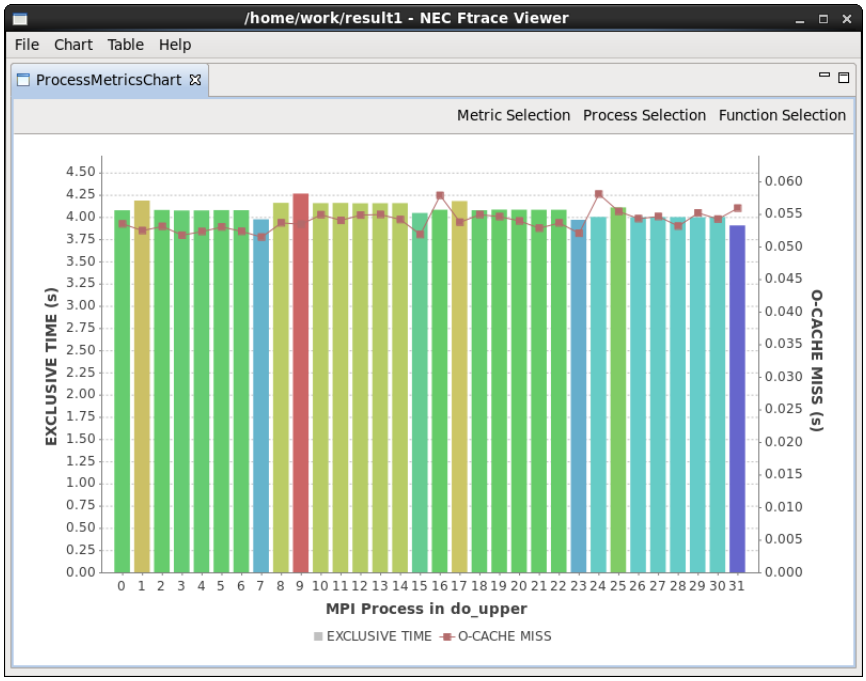

4.3.6 Process Breakdown Chart

This chart represents a performance data for each MPI process with a bar which is painted on a per-function in the order of the performance data. The function that the performance data is the

largest is painted in the lowest place. The functions that the performance data is small are gathered into one and painted in gray in the upper place.

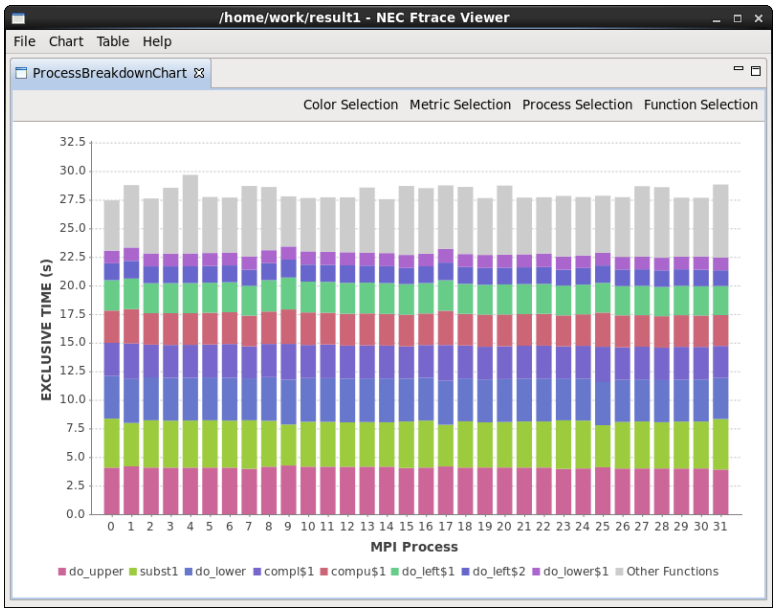

You can select the functions which is painted indivisually and change the order of the functions by [Function Selection]. The checked functions are selected from top to bottom and painted indivisually from botom to top. The non-checked functions are gathered into one as "Other Function" and painted in gray in the upper place. The details of [Function Selection] are explained in "4.3.11 Selecting Function".

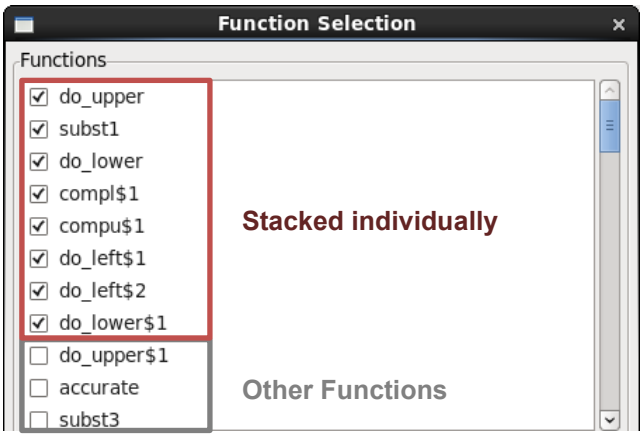

Note

In the chart, the metric which is possible to be accumulated is only available.

4.3.7 MPI Communication Chart

This chart represents elapsed time, MPI communication time and MPI communication idle time for each MPI process in the specified function.

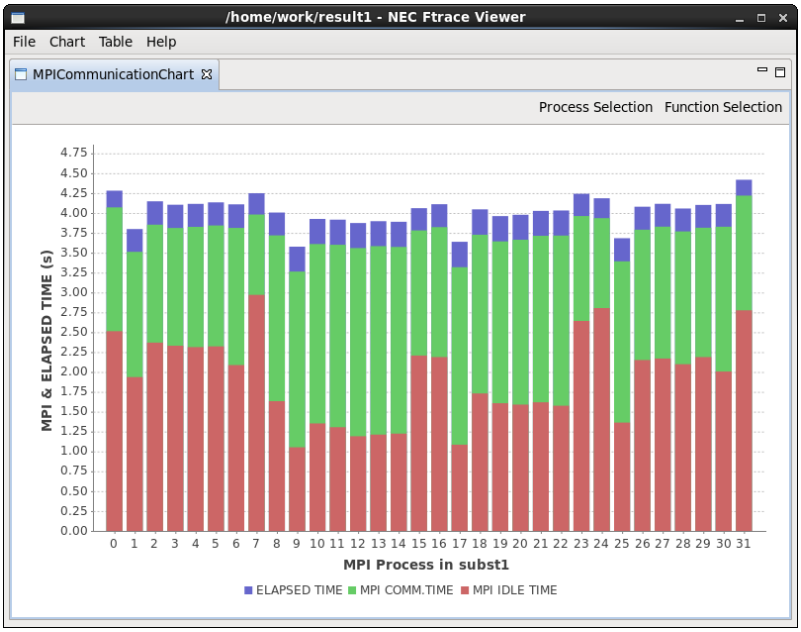

The bar represents elapsed time, MPI comunication time and MPI communication idle time as follows. With respect to the bottom end, the boundary between red and green is MPI communication idle time (MPI IDLE TIME), the boundary between green and blue is MPI communication time (MPI COMM.TIME) and the top end is elapsed time (ELAPSED TIME).

Note

MPI communication time without idle time is a green part because MPI COMM.TIME contains MPI IDLE TIME. Similarly, elapsed time except MPI communication such as the execution of user code is a blue part because ELAPSED TIME contains MPI COMM.TIME.

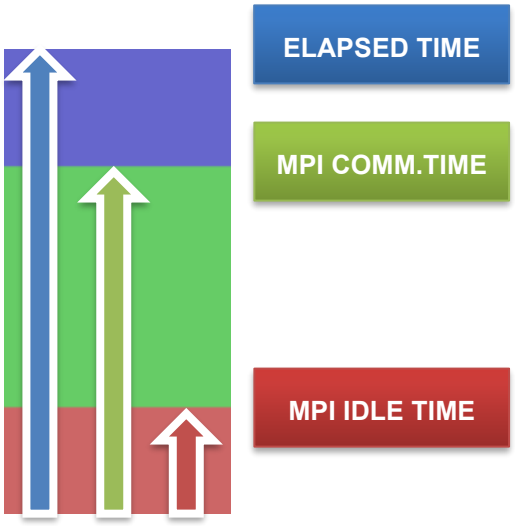

4.3.8 Process Histogram

This chart represents a performance distribution of MPI processes in the specified function. The performance data is on the horizontal axis and the number of MPI processes within specified values is on the vertical axis. When a bar of histogram gets focused by mouse, MPI RANK of MPI process in the bar is popped up.

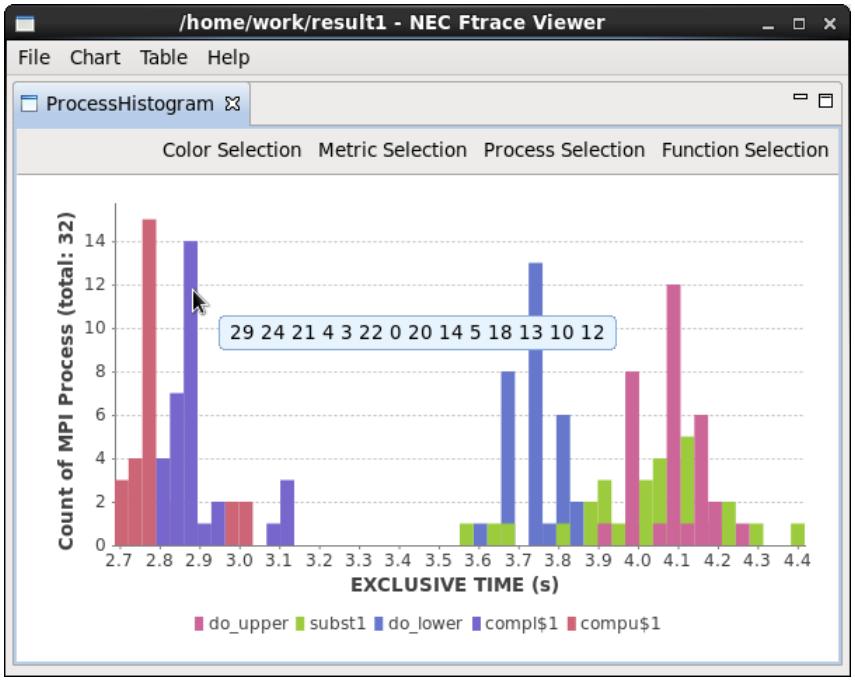

4.3.9 Selecting Metric

In order to select metrics in the chart, open a selection dialog by [Metric Selection] and select the metrics which the bar in the chart represents from "Main Metric" and the metrics the lines in the chart represent from "Sub Metrics"

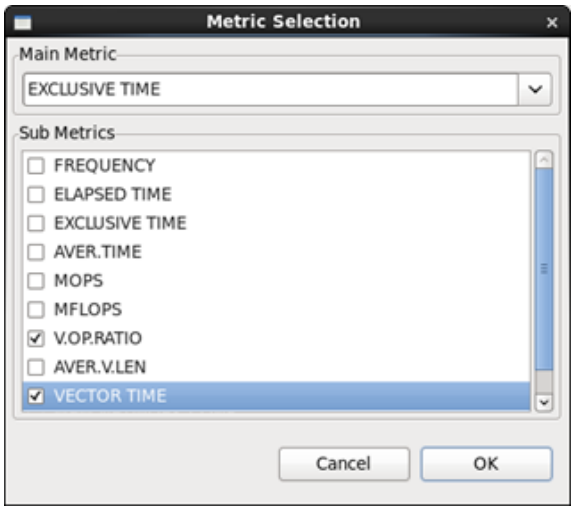

In lines of the chart, the checked metrics are selected from top to bottom and drawn in the order. In order to change the order, arrange the metrics in the list by drag-and-drop.

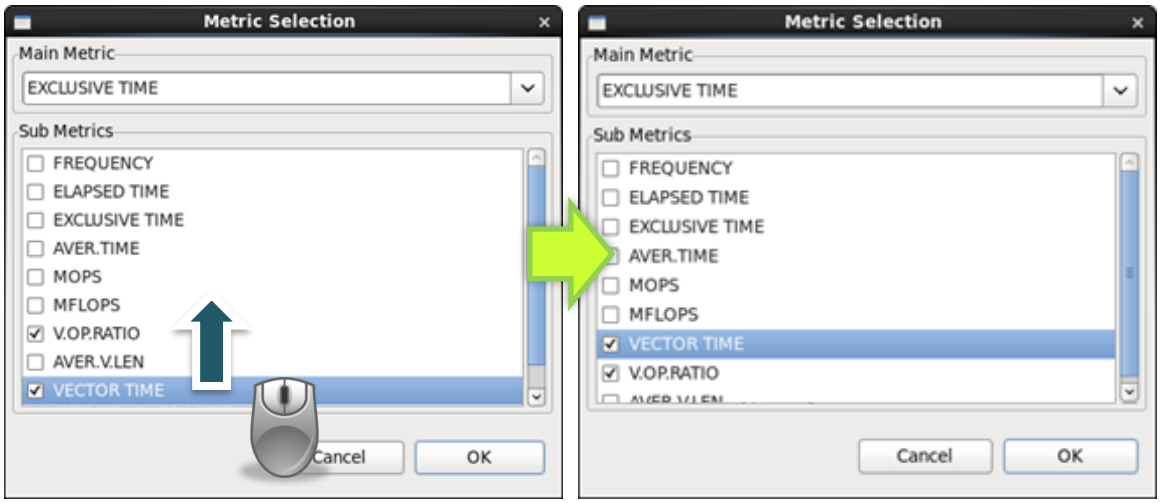

4.3.10 Selecting Process

In order to select processes in the chart, open a selection dialog by [Selection Selection].

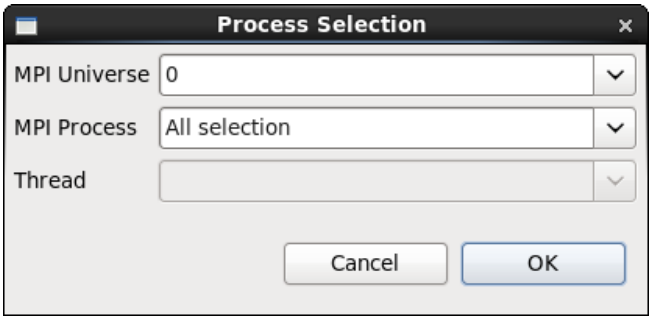

The selection is narrowed down in the order of MPI universe, MPI process and thread. The following methods of narrowing down are available.

Note

In the case of non-MPI program, "MPI Universe" and "MPI Process" is displayed in narrowing target. In this case, 0 is set in "MPI Universe" and "MPI Process".

The number of thread is a 0-based value to identify the thread.

Narrowing Target	Narrowing Method	Selected Processes
MPI Universe	All selection	All MPI universes
	The number of MPIUNIVERSE	The MPI universe that the environment variable "MPIUNIVERSE" is equals the specified number
MPI Process	All selection	All MPI processes in the specified MPI universe
	Custom selection	Any MPI process in the specified MPI universe
	The number of MPIRANK	In the specified MPI universe, The MPI process that the environment variable "MPIRANK" is equals the specified number
Thread	All selection	All threads in the specified MPI process
	Custom selection	Any thread in the specified MPI process
	The number of Thread	In the specified MPI process, the thread identified by the specified number

In order to select any processes, use "Custom Selection". In "Custom Selection", the MPI process or thread which identified a number input in the text area in lower part is selected.

You can use "Start", "End" and "Stride" to input numbers within a range.

- Start Specify the minimum number input in the text area
- End Specify the maximum number input in the text area.
- Stride The numbers are input within from "Start" to "End" skipping "Stride".

For example, when you set "Start" to 8, "End" to 8 and "Stride" to 2, the even numbers from 8 to 24 are input in the text area.

Using the slider, you can specify "Start" and "End". The left end of slider represents the value of "Start" and the right end of one represents the value of "End". When the slider is moved back or forth, "Start" and "End" is increased and decreased. If you drag the blue part in the slider, "Start" and "End" is increased and decreased at the same time.

Using "Sort by metric", you can sort MPI process or thread by a metric. When a metric is specified in pull-down menu, the numbers in the text area are sorted in decreasing order from the largest

performance data by the specified metric. When PROC.NAME is specified, the numbers are sorted by name in numeric order. In order to sort in the reverse order, turn the "Reverse" checkbox on.

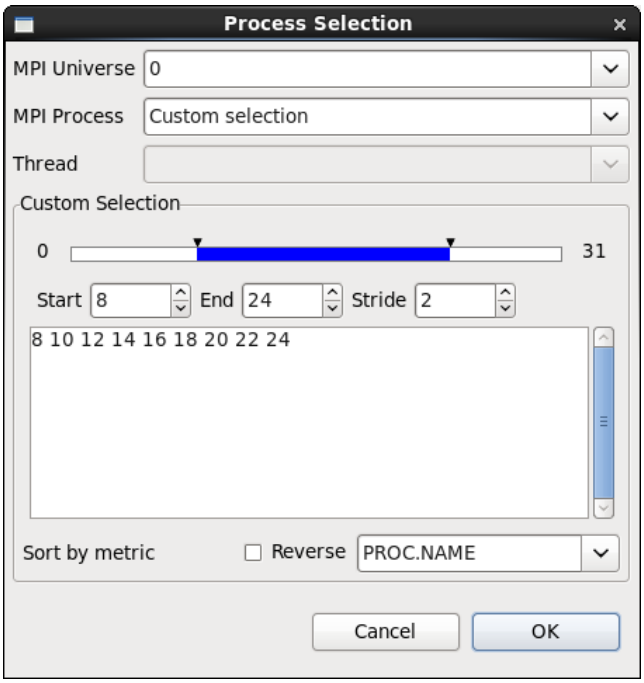

4.3.11 Selecting Function

In order to select functions in the chart, open a selection dialog by [Function Selection].

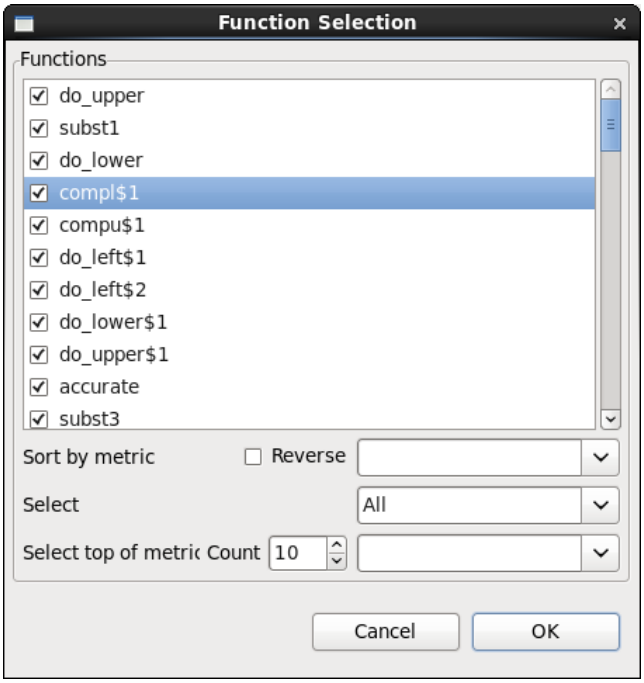

The checked function is selected from top to bottom and displayed in the chart in the order. The functions by automatic parallelization or OpenMP are output with '$number' appended. In order to change the order, arrange the functions in the list by drag-and-drop.

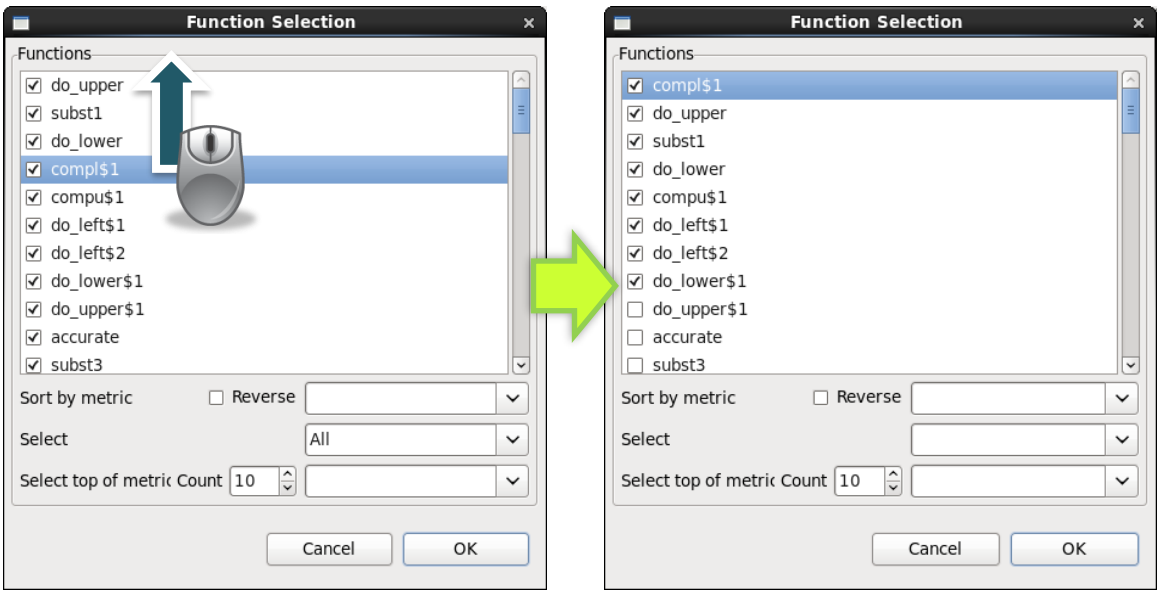

Using "Sort by metric", you can sort functions by a metric. When a metric is specified in pull-down menu, the functions in the list are sorted in decreasing order from the largest performance data by the specified metric. When PROC.NAME is specified, the functions are sorted by name in dictionary order. In order to sort in the reverse order, turn the "Reverse" checkbox on.

You can select the functions which fulfill a condition using "Select" or "Select top of metric". In the case of "Select", when you select the following condition in the pull-down menu, the functions which fulfill the condition are selected. In the case of "Select top of metric", when you specify the metric and the number of the selection, about the specified metrics, the top N functions are selected; N is the specified number.

Condition	Selected function
All	All functions
Executed one or more times	Functions that is executed one or more times
Executed by some threads	Functions that is executed in multiple threads
Executed by some MPI process	Functions that is executed in multiple MPI processes
None	No function is selected

4.3.12 Selecting Colors Assigned to Metrics and Functions

In the following charts, colors assigned to metrics and functions can changed.

- Colors assigned to line charts of metrics:
  Function Metrics Chart, Function Breakdown Chart and Process Metrics Chart
- Colors assigned to functions:
  Process Breakdown Chart and Process Histogram

Open a selection dialog by [Color Selection] and select a new color. The selections are applied to all charts.

Note

The change in coloring are applied at the timing of re-drawing a chart. If a chart is not changed in coloring, resize the chart to re-draw.

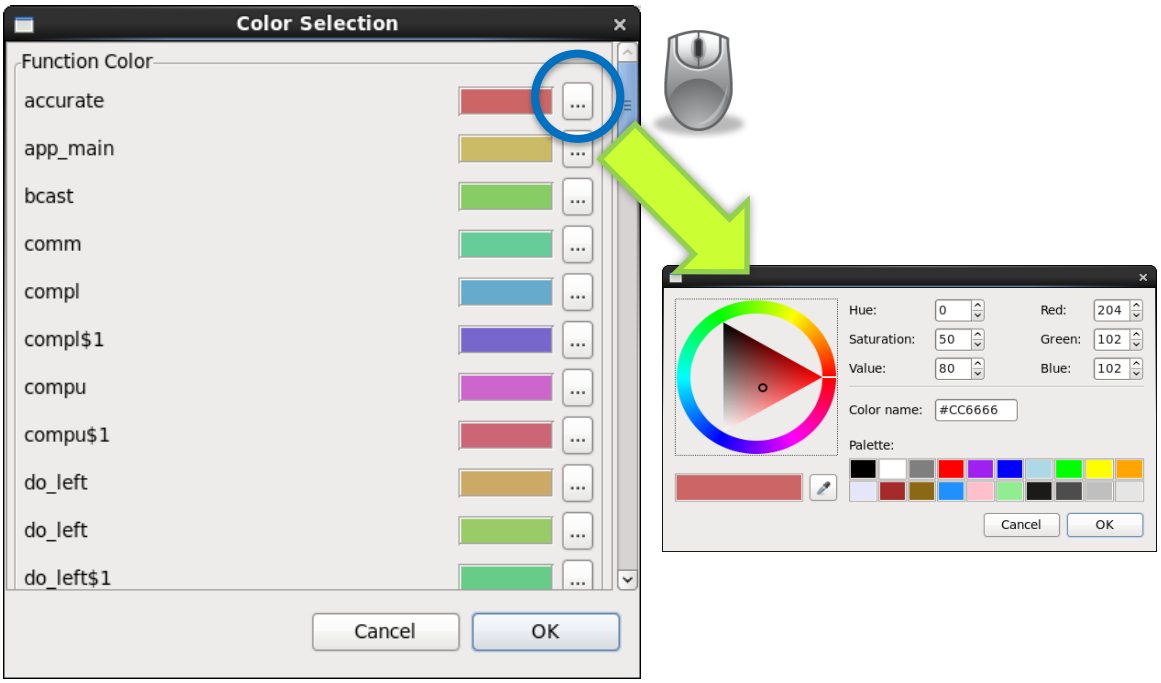

4.3.13 Exporting as Image

When you right-click the mouse on the chart and choose [Save as] in the context menu, the chart image is saved in PNG format.

Note

The part that displayed on the view is only included In the image. The part that can not be displayed without scrolling is not included.

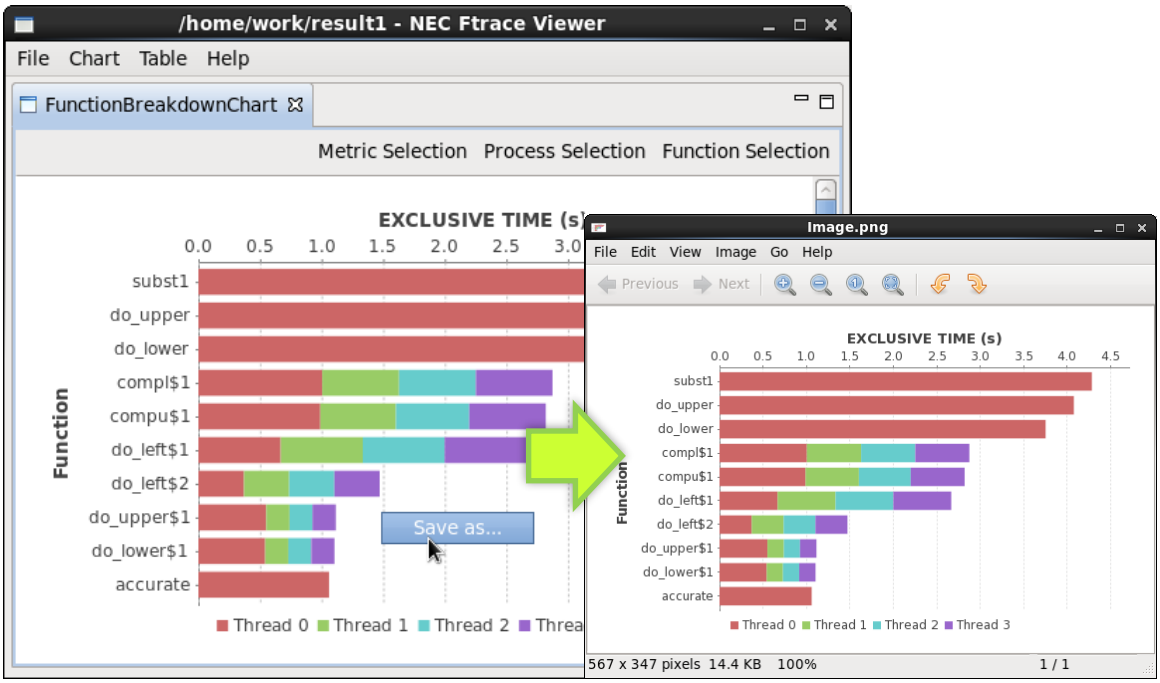

4.4 Table

Using table, performance data of all metrics can be listed for each function, subroutine, parallelized-region, user region, MPI universe, MPI process and thread. You can drill down for each. For example, in the case of hybrid parallel execution, you can drill down each function and confirm performance data for each the MPI process which executes the function. In addition, you can drill down each MPI process and confirm each performance data for each thread in the MPI process

The following tables are available.

- Function Table

  This table lists performance data for each function, subroutine and parallelized-region.

- User Region Table

  This table lists performance data for each user region.

4.4.1 Table Basics

In order to drill down, click a expand icon in PROC.NAME column. In the following table, the function "subst1" is drilled down.

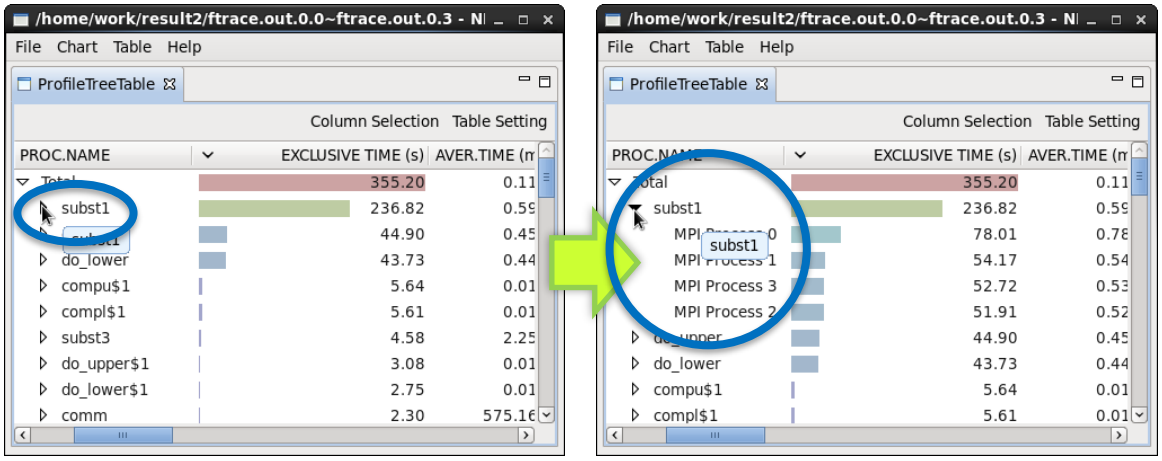

In order to sort the table by a metric, click the column header of the metric.

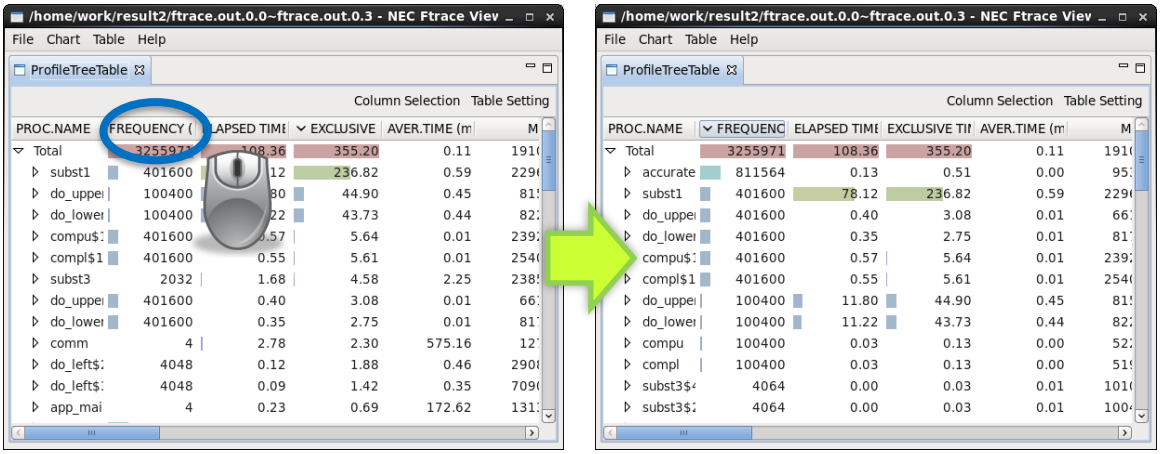

In order to move a column, drag and drop the column header to the destination.

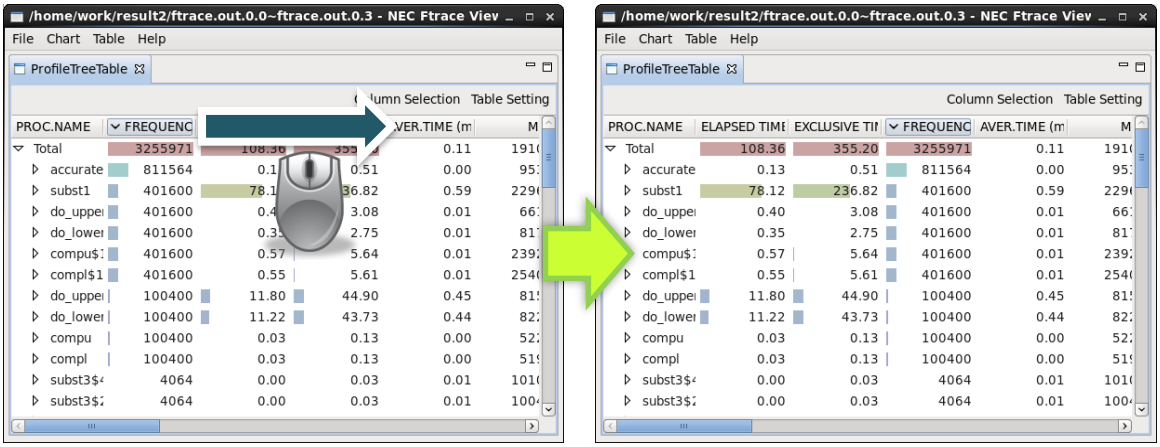

You can select metrics to show in the table and the method of drill down by the table menu. The details are explained in "4.4.4 Selecting and Sorting Metrics" and "4.4.5 Selecting a method of drill down".

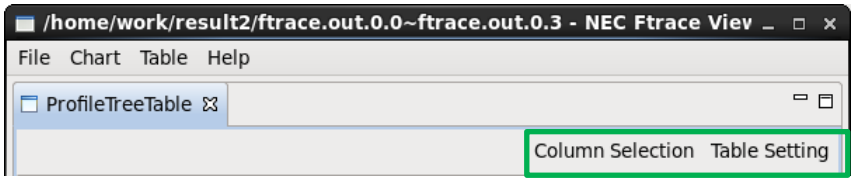

4.4.2 Function Table

This table lists performance data for each function, subroutine and parallelized-region. You can drill down in the order of function/subroutine/parallelized-region, MPI universe, MPI process and thread by default. You can change the order. The change is explained in "4.4.5 Selecting a method of drill down".

In some metrics, the performance data is represented by a bar chart in the cell.

- In the case of the metric which is possible to be accumulated, the ratio to the "Total" is painted from blue (0%) to red (100%)
- In the case of the metric which of unit is percent, the value is painted from blue (0%) to red (100%).

Note

In the metrics except above such as MOPS and MFLOPS, the bar chart is not available.

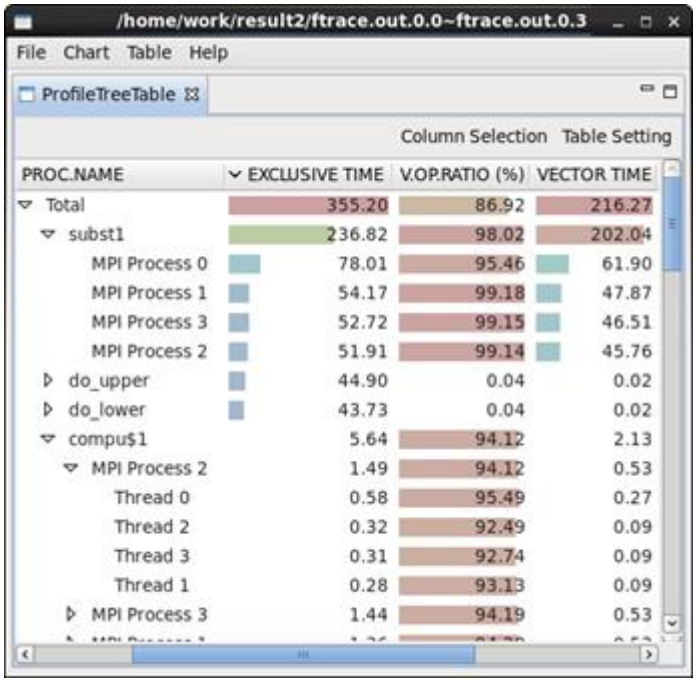

4.4.3 User Region Table

This table lists performance data for each user region. You can drill down in the order of user region, MPI universe, MPI process and thread.

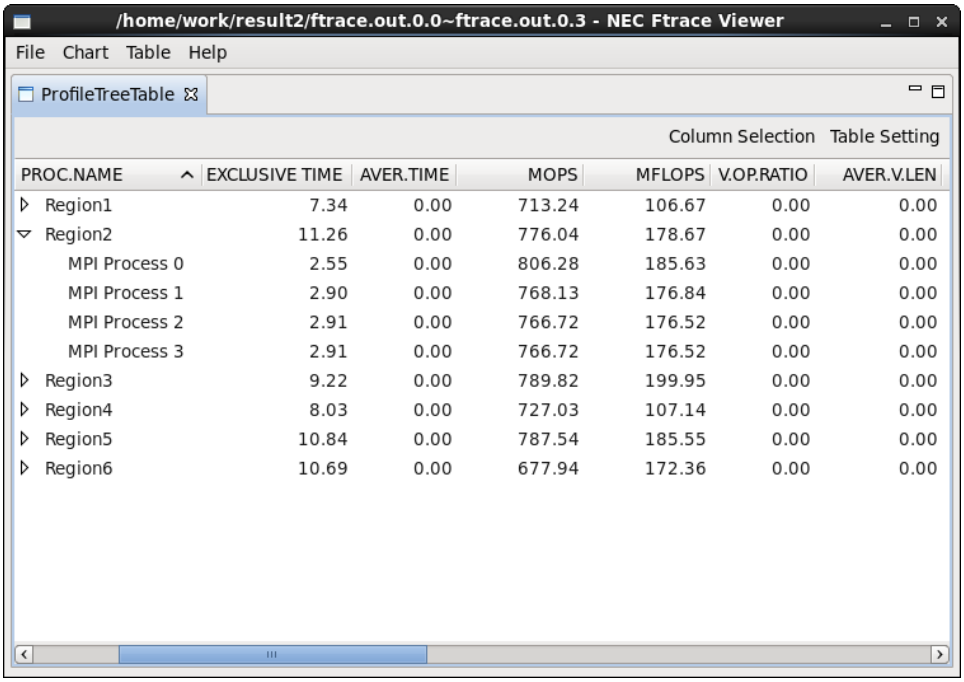

PROC.NAME	EXCLUSIVE TIME	AVER.TIME	MOPS	MFLOPS	V.OP.RATIO	AVER.V.LEN
Region1	7.34	0.00	713.24	106.67	0.00	0.00
Region2	11.26	0.00	776.04	178.67	0.00	0.00
MPI Process 0	2.55	0.00	806.28	185.63	0.00	0.00
MPI Process 1	2.90	0.00	768.13	176.84	0.00	0.00
MPI Process 2	2.91	0.00	766.72	176.52	0.00	0.00
MPI Process 3	2.91	0.00	766.72	176.52	0.00	0.00
Region3	9.22	0.00	789.82	199.95	0.00	0.00
Region4	8.03	0.00	727.03	107.14	0.00	0.00
Region5	10.84	0.00	787.54	185.55	0.00	0.00
Region6	10.69	0.00	677.94	172.36	0.00	0.00

4.4.4 Selecting and Sorting Metrics

Open a selection dialog by [Table Setting]. When you want to hide metrics, turn the checkboxes of the metric off. You can arrange the metrics in the list by drag-and-drop.

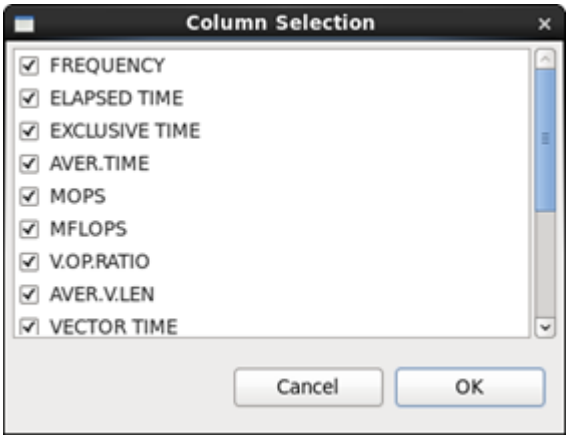

4.4.5 Selecting a method of drill down

Open a setting dialog by [Table Setting]. When "Function Table" is specified, the table in the view is changed to "Function Table". You can select the method of drill down in the follows.

- Function > MPI Process > Thread order
  You can drill down in the order of function, MPI universe, MPI process, thread.
- MPI Process > Function > Thread order
  You can drill down in the order of MPI universe, MPI process, function, and, thread.
- MPI Process > Thread > Function order
  You can drill down in the order of MPI universe, MPI process, thread, and function.

When "User Region Table" is specified, the table in the view is changed to "User Region Table".

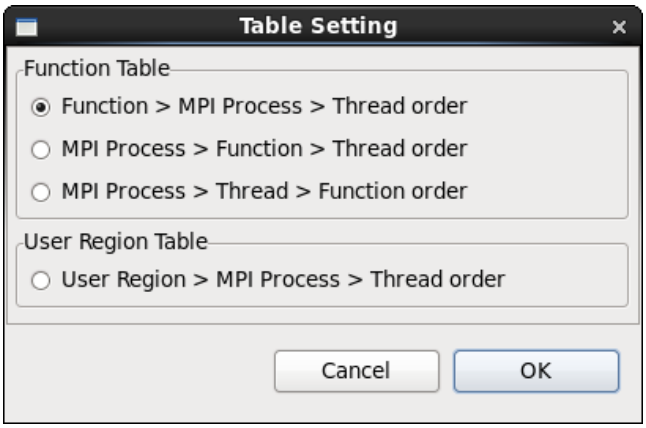

Chapter 5 Troubleshooting

This chapter explains how to solve problems that occur when using NEC Ftrace Viewer.

1. NEC Ftrace Viewer does not start

If the following error is output to the log, an attempt to start NEC Ftrace Viewer with X11 forwarding has failed. Make sure you can use X11 forwarding by other GUI applications.

```
!MESSAGE Application error
!STACK 1
org.eclipse.swt.SWTError: No more handles [gtk_init_check() failed]
```

2. File selection dialog and directory selection dialog are garbled

Fonts for your locale are not installed. you need to install your locale font packages or start NEC Ftrace Viewer with C locale (LANG=C) using /opt/nec/ve/ftraceviewer/ftraceviewer-c command.

3. Failure to load the analysis information files, or the expected result is not displayed

NEC Ftrace Viewer supports loading one analysis information file output by one non-MPI program execution, or one or more analysis information files output by one MPI program execution. Also, loading the following analysis information files is not supported. In these cases, divide the analysis information files within the supporting and load them separately.

- Simultaneous loading of analysis information files for non-MPI and MPI programs
- Loading multiple analysis information files for programs using VEO
- Simultaneous loading of analysis information files output by MPI processes executed in different types of VEs
- Simultaneous loading of analysis information files output by MPI processes when the value of the environment variable VE_PERF_MODE is different for each MPI process

4. NEC Ftrace Viewer is slow

NEC Ftrace Viewer limits the default maximum memory usage to 2GiB. Therefore, if a large number of analysis information files are loaded, the memory may run out and the operation may slow down. The memory usage of NEC Ftrace Viewer can be specified by the command line option at startup. For example, to increase the maximum memory usage to 10GiB, add "-vmargs -Xmx10G" to the end of the command line options.

Note

Be sure to add "-vmargs" and subsequent options to the end of the command line. Otherwise, command line options will be misinterpreted, and launch will fail.

```
$ /opt/nec/ve/ftraceviewer/ftraceviewer -f ftrace.out.* -vmargs -Xmx10G
```

Appendix A Revision History

Edition	Issue	Changes
1	February 2018	Original
2	January 2020	• Added 4.2 Metrics
3	December 2021	• Updated Appendix B
4	November 2022	• Updated 2.4 Checking Analysis Information Files • Updated 3.5 Loading New Analysis Information Files
5	March 2023	• Updated 4.2 Metrics • Added 5 Troubleshooting • Updated Appendix B

Appendix B Trademarks and Licenses

B.1 Trademarks

- Eclipse is a registered trademark of Eclipse Foundation.
- All other product, brand, or trade names in this publication are the trademarks or registered trademarks of their respective owners.

B.2 Licenses

The following list is software and the license included in NEC Ftrace Viewer.

Software	License
JFreeChart http://www.jfree.org/index.html	LGPL Version 2.1
JCommon http://www.jfree.org/index.html	LGPL Version 2.1
PngEncoder http://catcode.com/pngencoder/	LGPL Version 2.1
Guava https://code.google.com/p/guava-libraries/	Apache License 2.0
Eclipse RCP https://www.eclipse.org/	Eclipse Public License Version 1.0
OSGi Materials https://www.osgi.org	Apache License 2.0
Jakarta Commons Collections https://commons.apache.org/proper/commons-collections/	Apache License 2.0
ICU4J http://site.icu-project.org/	ICU License

Software	License
Independent JPEG Group's JPEG software release 6b http://ijg.org/	Original License
PuTTY 0.58 http://www.chiark.greenend.org.uk/~sgtatham/putty/	MIT License
Gnome (contained in “Gnome Binding” of Eclipse RCP) http://www.gnome.org/	LGPL Version 2.1
GTK+ (contained in “GTK+ Binding” of Eclipse RCP) http://www.gtk.org/	LGPL Version 2.1
WebKitGTK+ (contained in “WebKitGTK+ Binding” of Eclipse RCP) http://webkitgtk.org/	LGPL Version 2.1
libsoup (contained in “WebKitGTK+ Binding” of Eclipse RCP) http://www.gnome.org/	LGPL Version 2.1
WebKit JavaScriptCore (contained in “WebKitGTK+ Binding” of Eclipse RCP) http://webkit.org/	BSD License
Mozilla (contained in “Mozilla Binding” of Eclipse RCP) http://www.mozilla.org/	Mozilla Public License Version 1.1
XULRunner 1.9 (contained in “Mozilla Binding” of Eclipse RCP) http://www.mozilla.org/	Mozilla Public License Version 1.1
Cairo (contained in “Cairo Binding” of Eclipse RCP) http://cairographics.org/	Mozilla Public License Version 1.1
Cairo for Linux (contained in “Cairo Library” of Eclipse RCP) http://cairographics.org/	Mozilla Public License Version 1.1

Software	License
pixman 0.1.6 (contained in "Cairo Library" of Eclipse RCP) http://www.pixman.org/	Original License
W3C DTDs http://www.w3.org/	W3C Software Notice and License

NEC Ftrace Viewer consists of Eclipse plug-ins by Eclipse RCP, and each plug-in includes any or all of the above software or derivative works. The license information for each plug-in (Eclipse RCP) and the included software can be viewed in a web browser (firefox) by the following steps: open NEC Ftrace Viewer, select [Help]→[About]→[Installation Details]→[Plug-ins] tab, select a plug-in and click [Legal Info] Button. Also, you can view the license information from about.html located in /opt/nec/ve/ftraceviewer/plugins/<plug-in id>_<version>.jar file or /opt/nec/ve/ftraceviewer/plugins/<plug-in id>_<version> directory. The source codes related to these software are located in /opt/nec/ve/ftraceviewer/plugins/<plug-in id>.source_<version>.jar file or /opt/nec/ve/ftraceviewer/plugins/<plug-in id>.source_<version> directory according to the license. Please refer to the license information of each plug-in for details about the licenses and source codes.

Note

If the new license information page does not appear, please close your web browser and try again.

SX-Aurora TSUBASA system software

NEC Ftrace Viewer User's Guide

February	2018	Edition	1
March	2023	Edition	5

NEC Corporation