

# VI-HPS



## Analysis report examination with CUBE

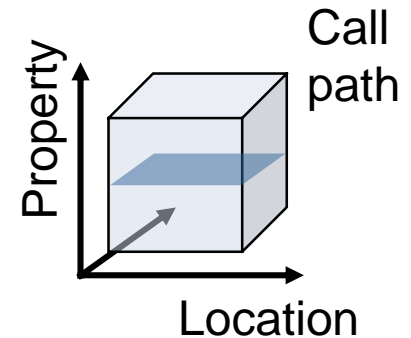
Ilya Zhukov

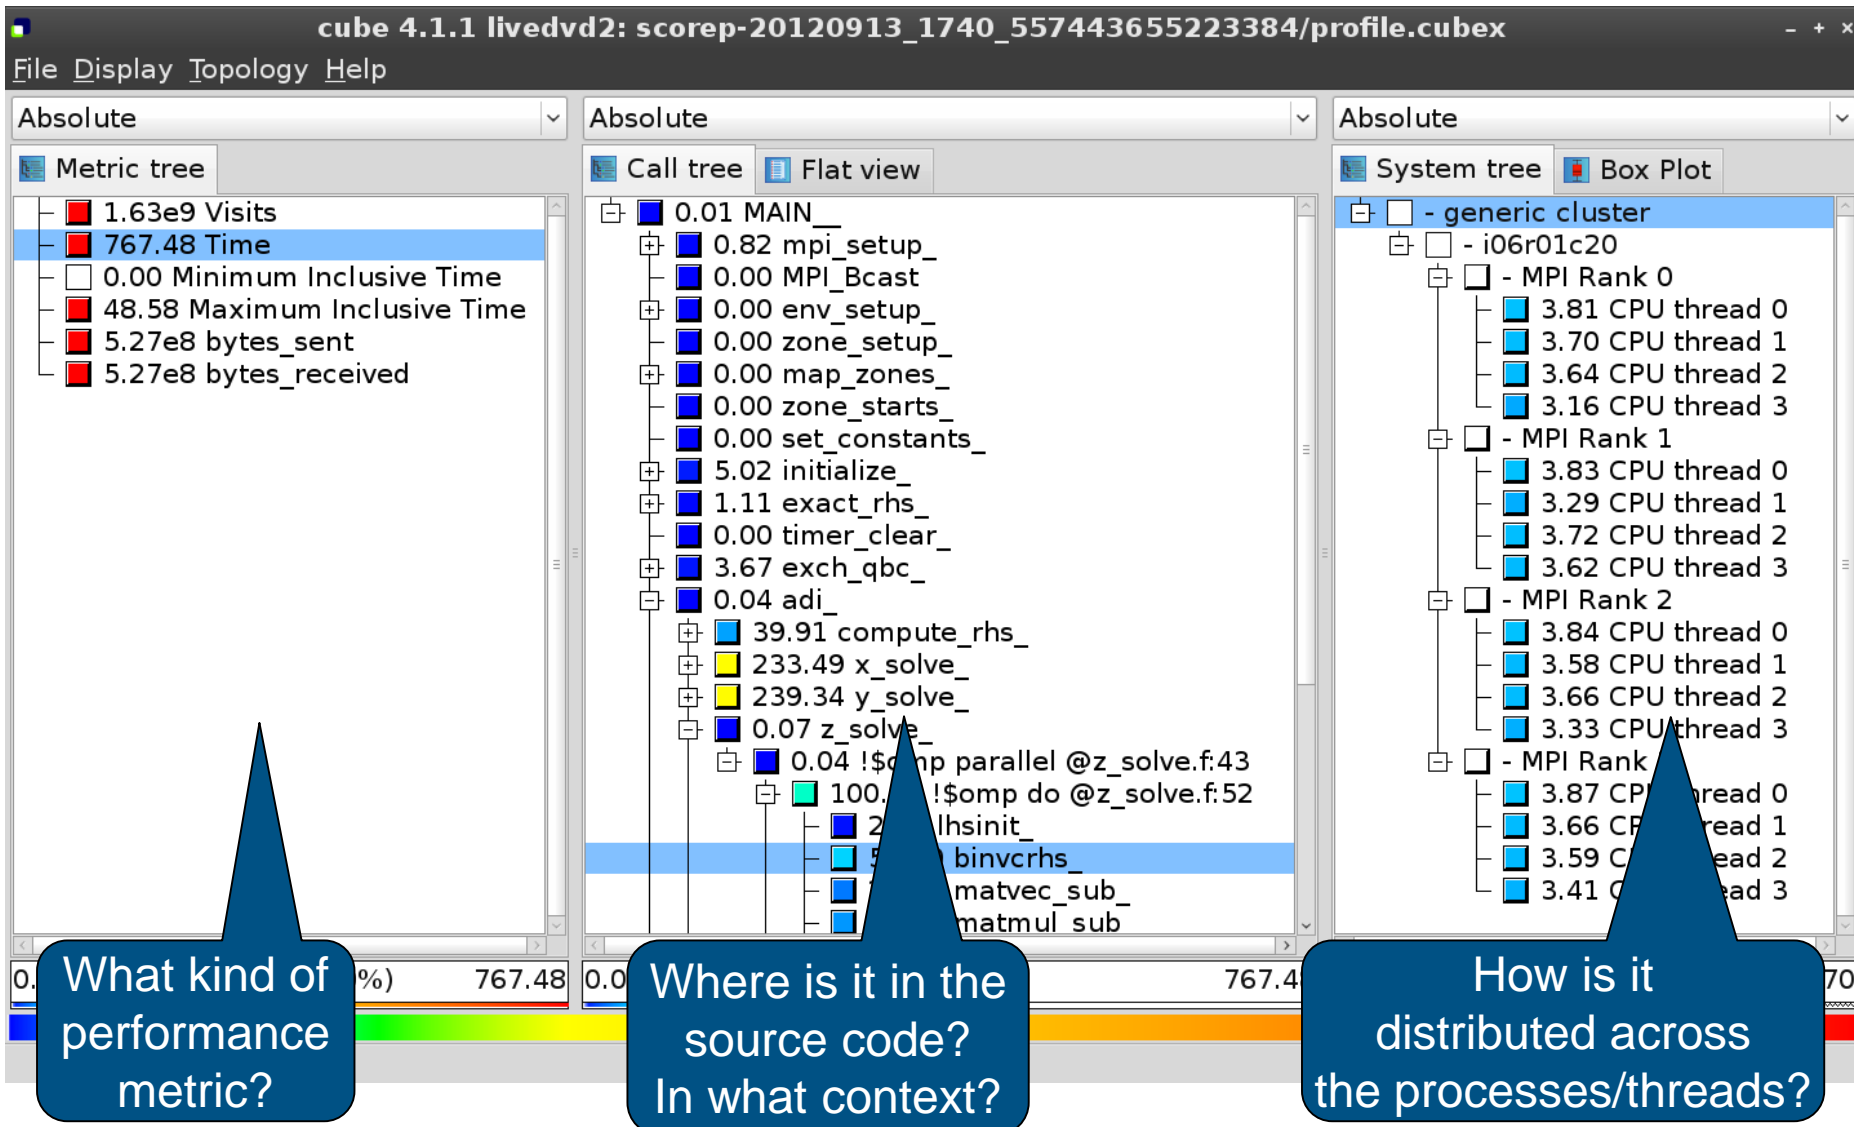
Jülich Supercomputing Centre



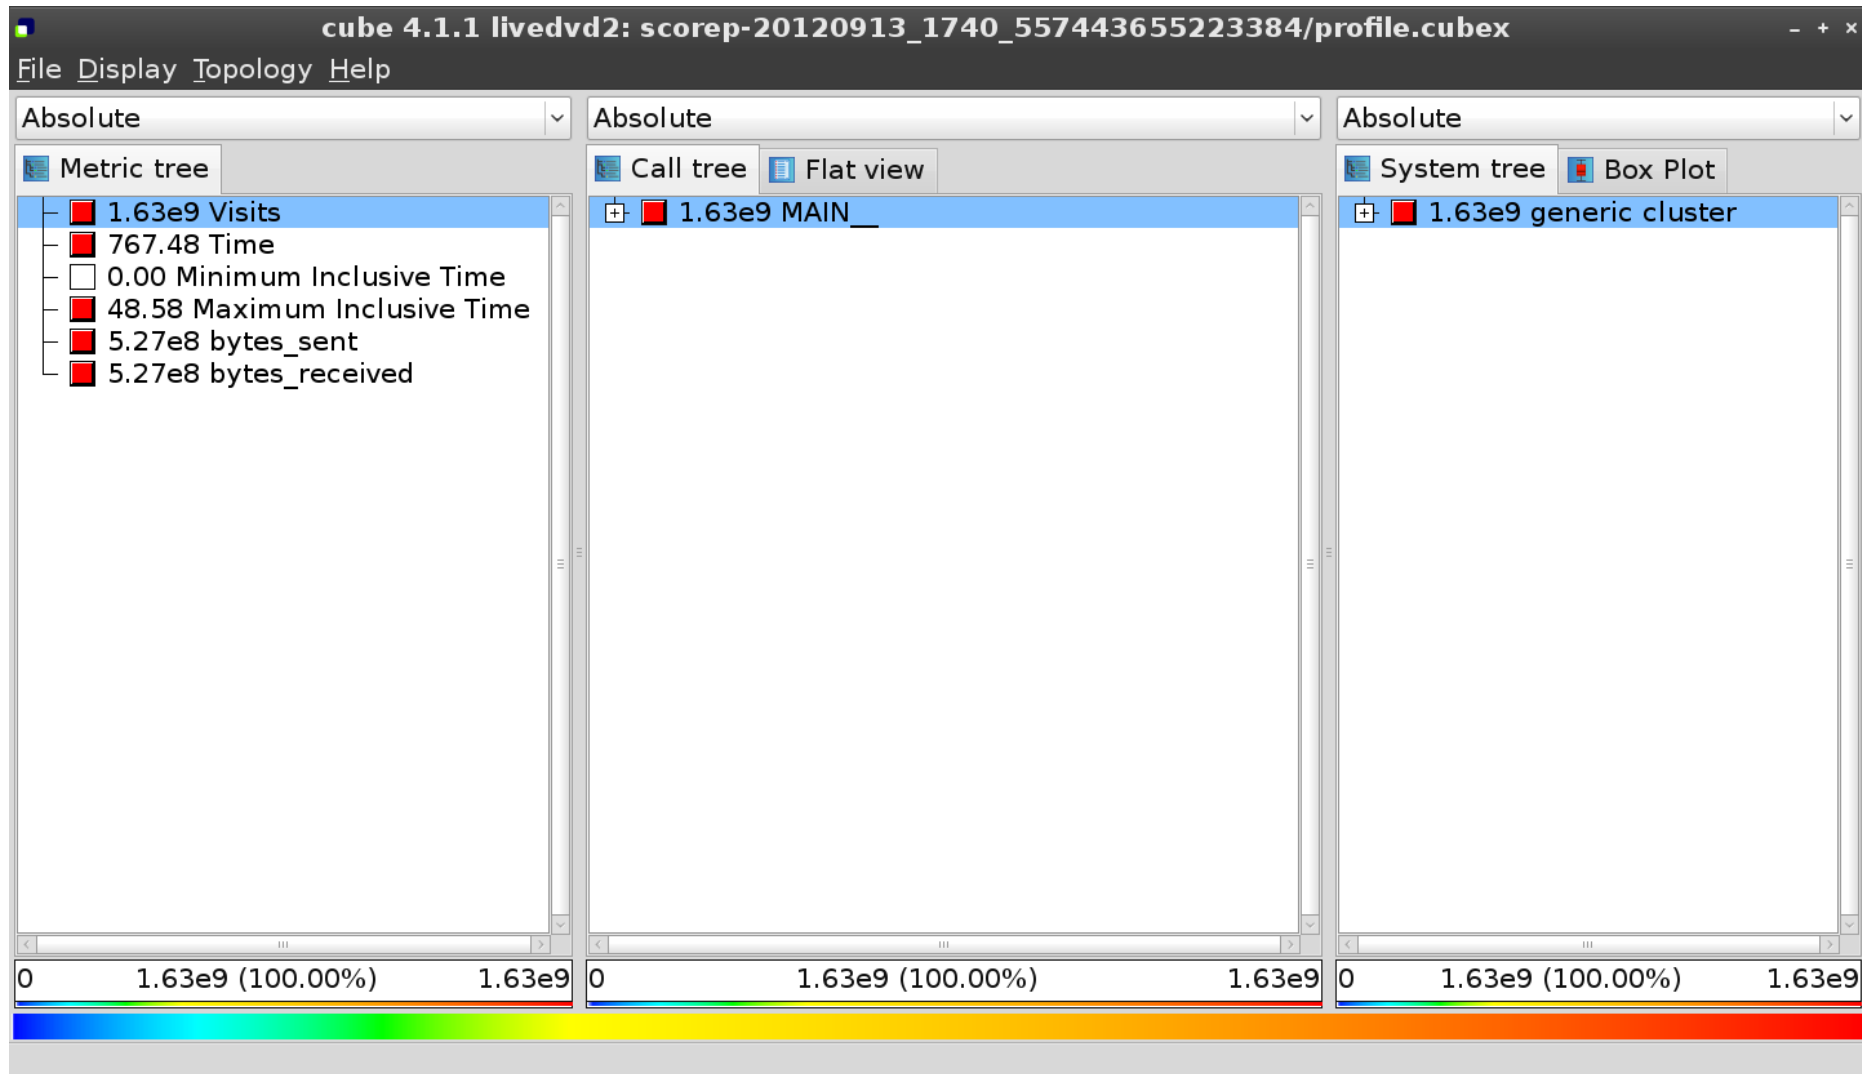
- Parallel program analysis report exploration tools
  - Libraries for XML report reading & writing
  - Algebra utilities for report processing
  - GUI for interactive analysis exploration
    - requires Qt4
- Originally developed as part of Scalasca toolset
- Now available as a separate component
  - Can be installed independently of Score-P, e.g., on laptop or desktop
  - Latest release: CUBE 4.3.0 (January 2015)

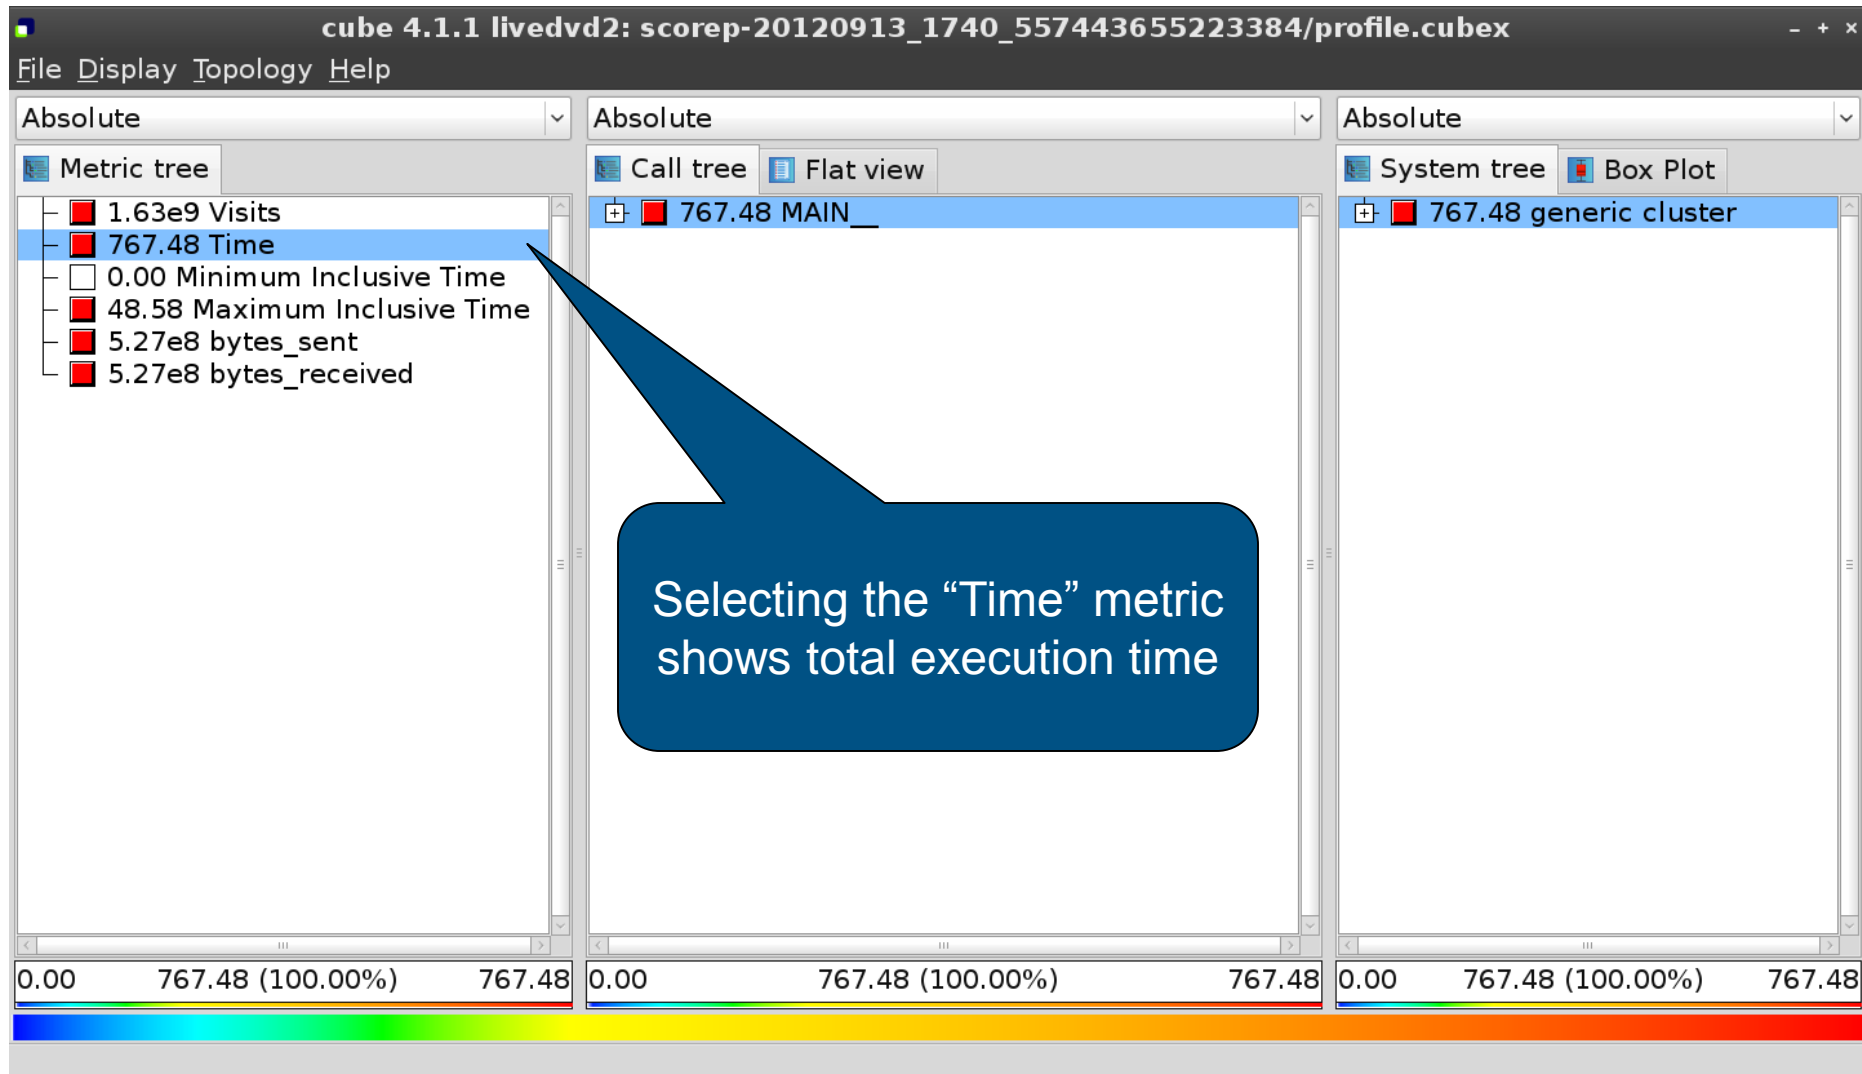
- Representation of values (severity matrix) on three hierarchical axes
  - Performance property (metric)
  - Call path (program location)
  - System location (process/thread)
- Three coupled tree browsers
- CUBE displays severities
  - As value: for precise comparison
  - As colour: for easy identification of hotspots
  - Inclusive value when closed & exclusive value when expanded
  - Customizable via display modes



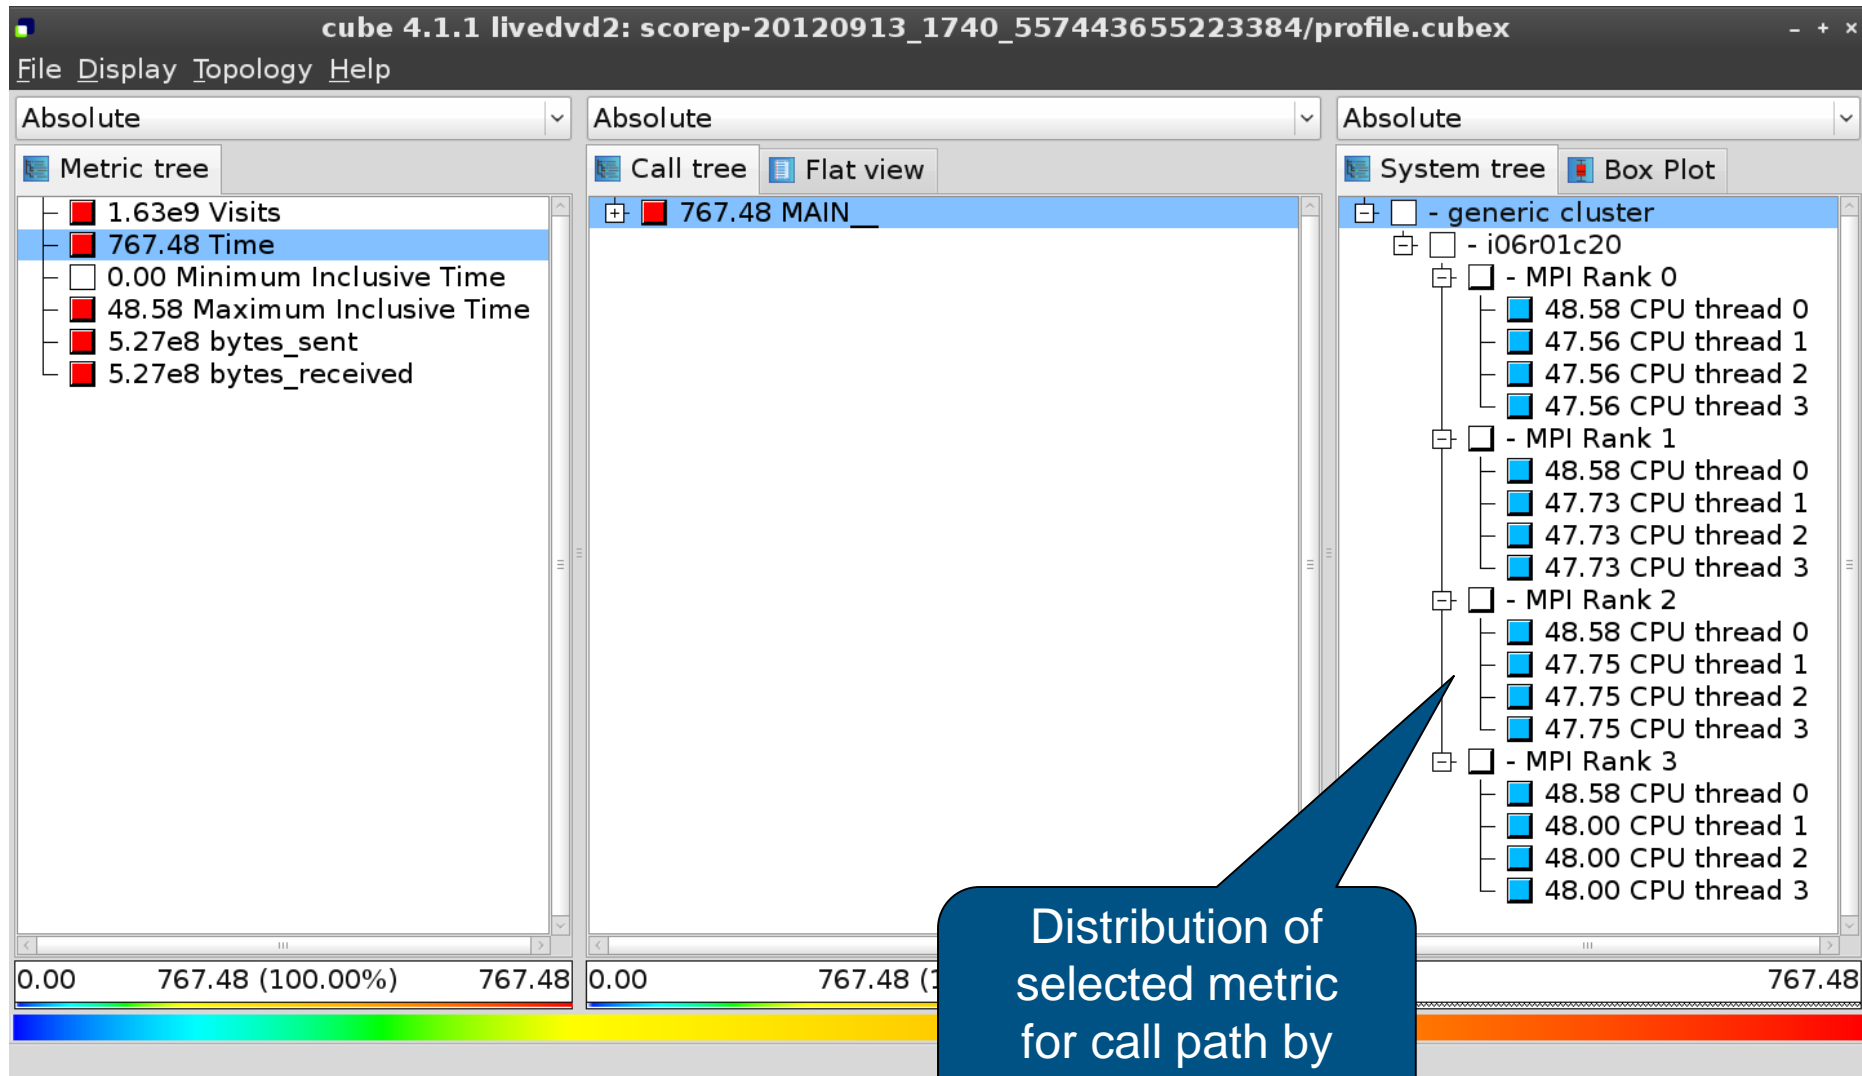


# Analysis report exploration (opening view)

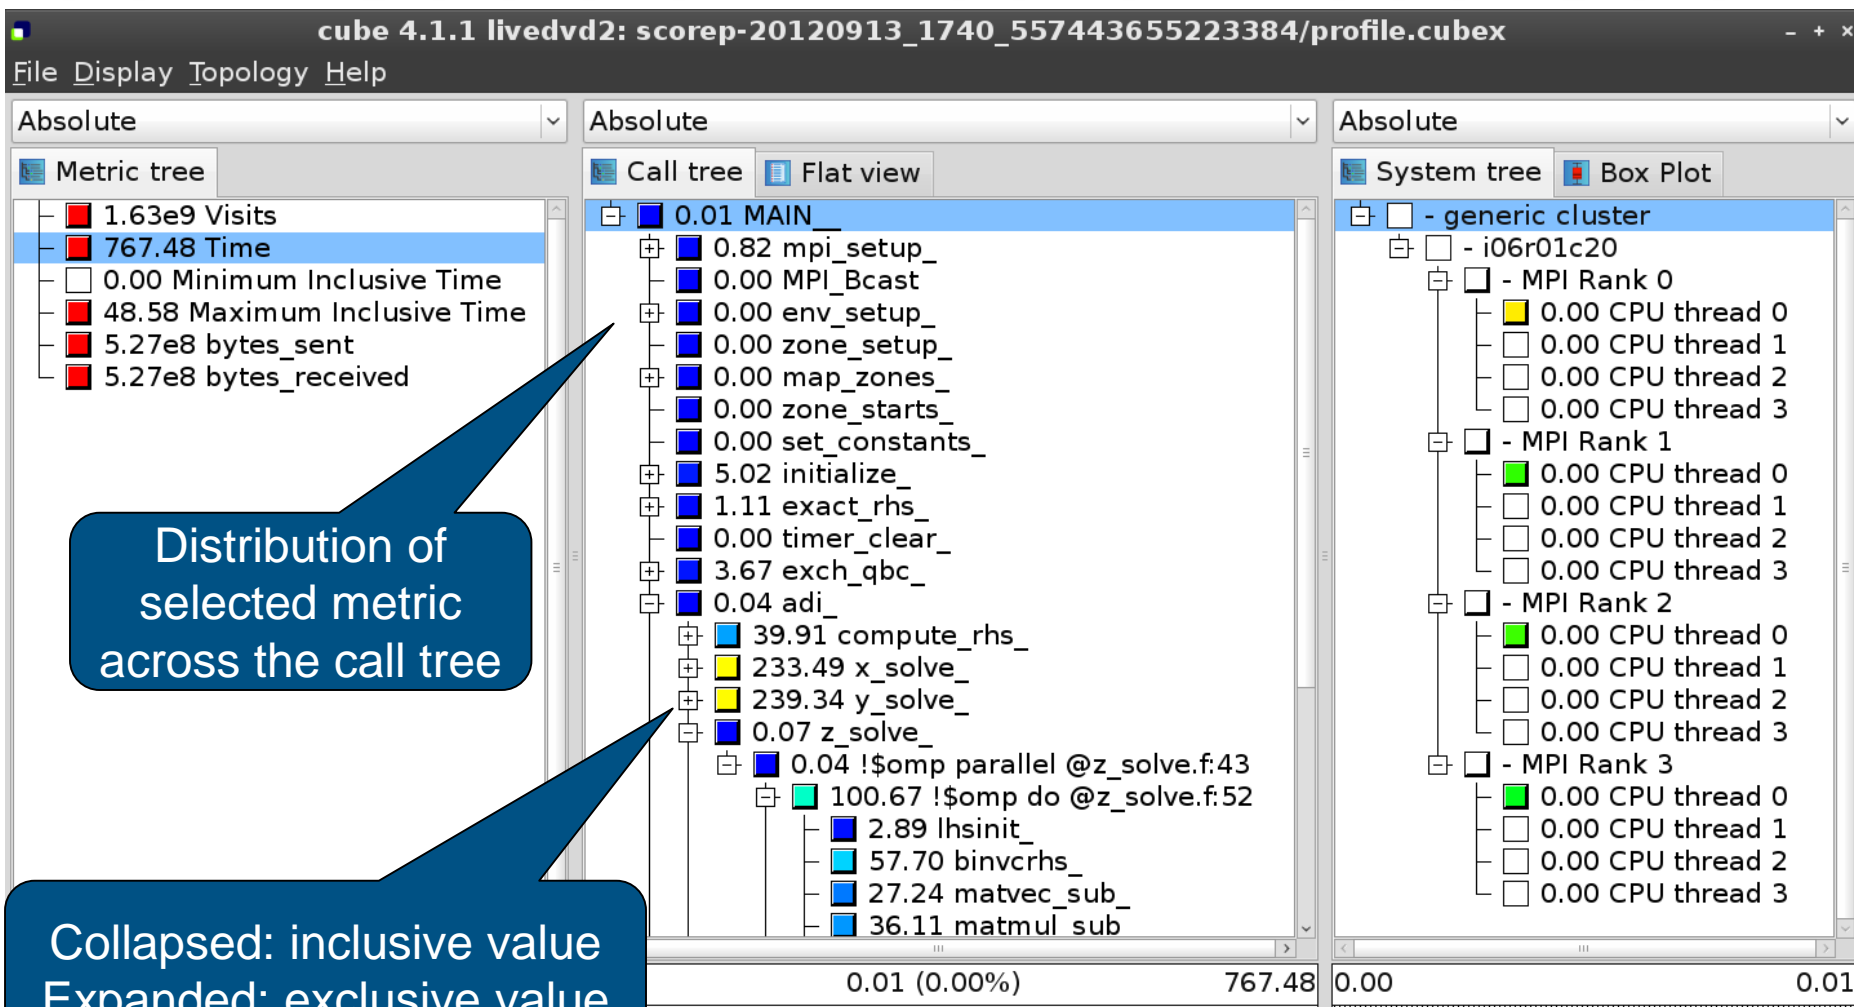




# Expanding the system tree

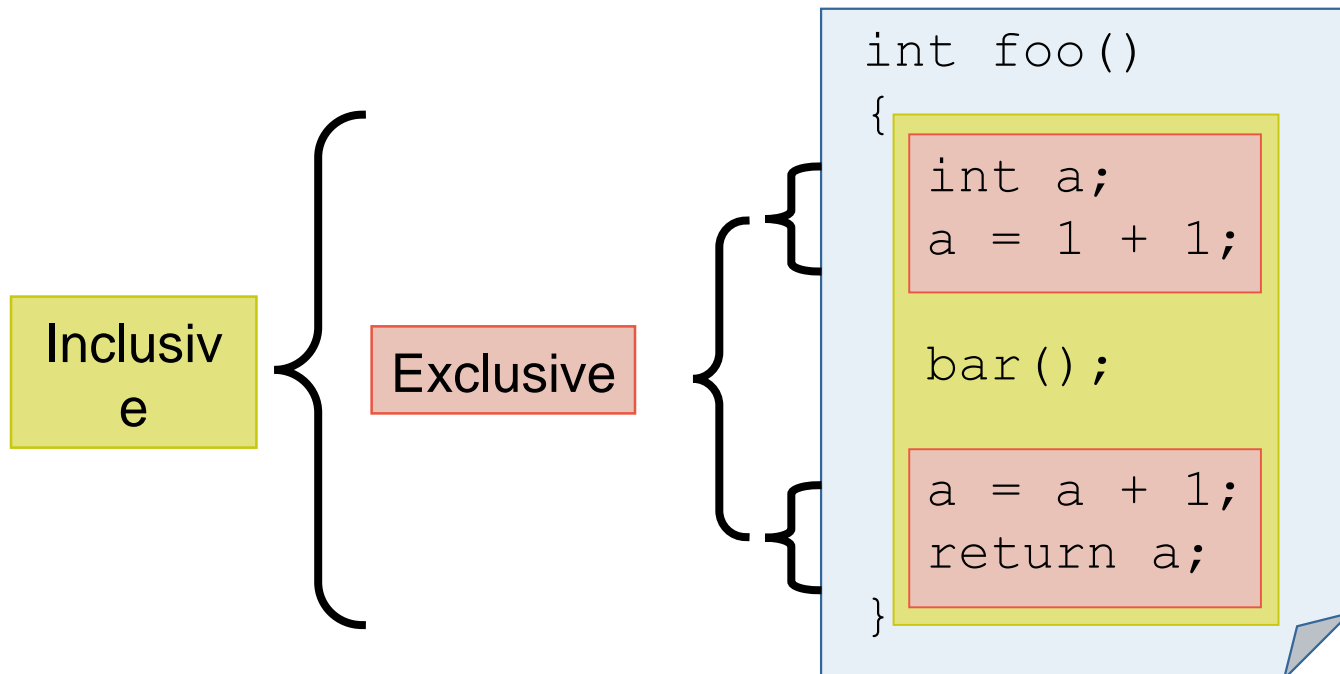


# Expanding the call tree

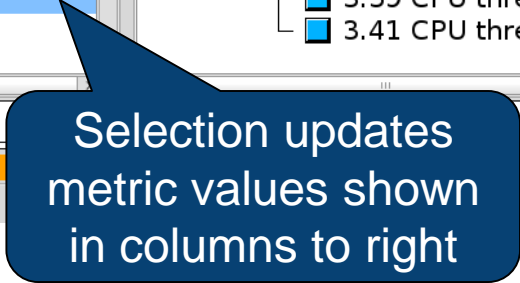




- Inclusive
  - Information of all sub-elements aggregated into single value
- Exclusive
  - Information cannot be subdivided further



# VI-HPS

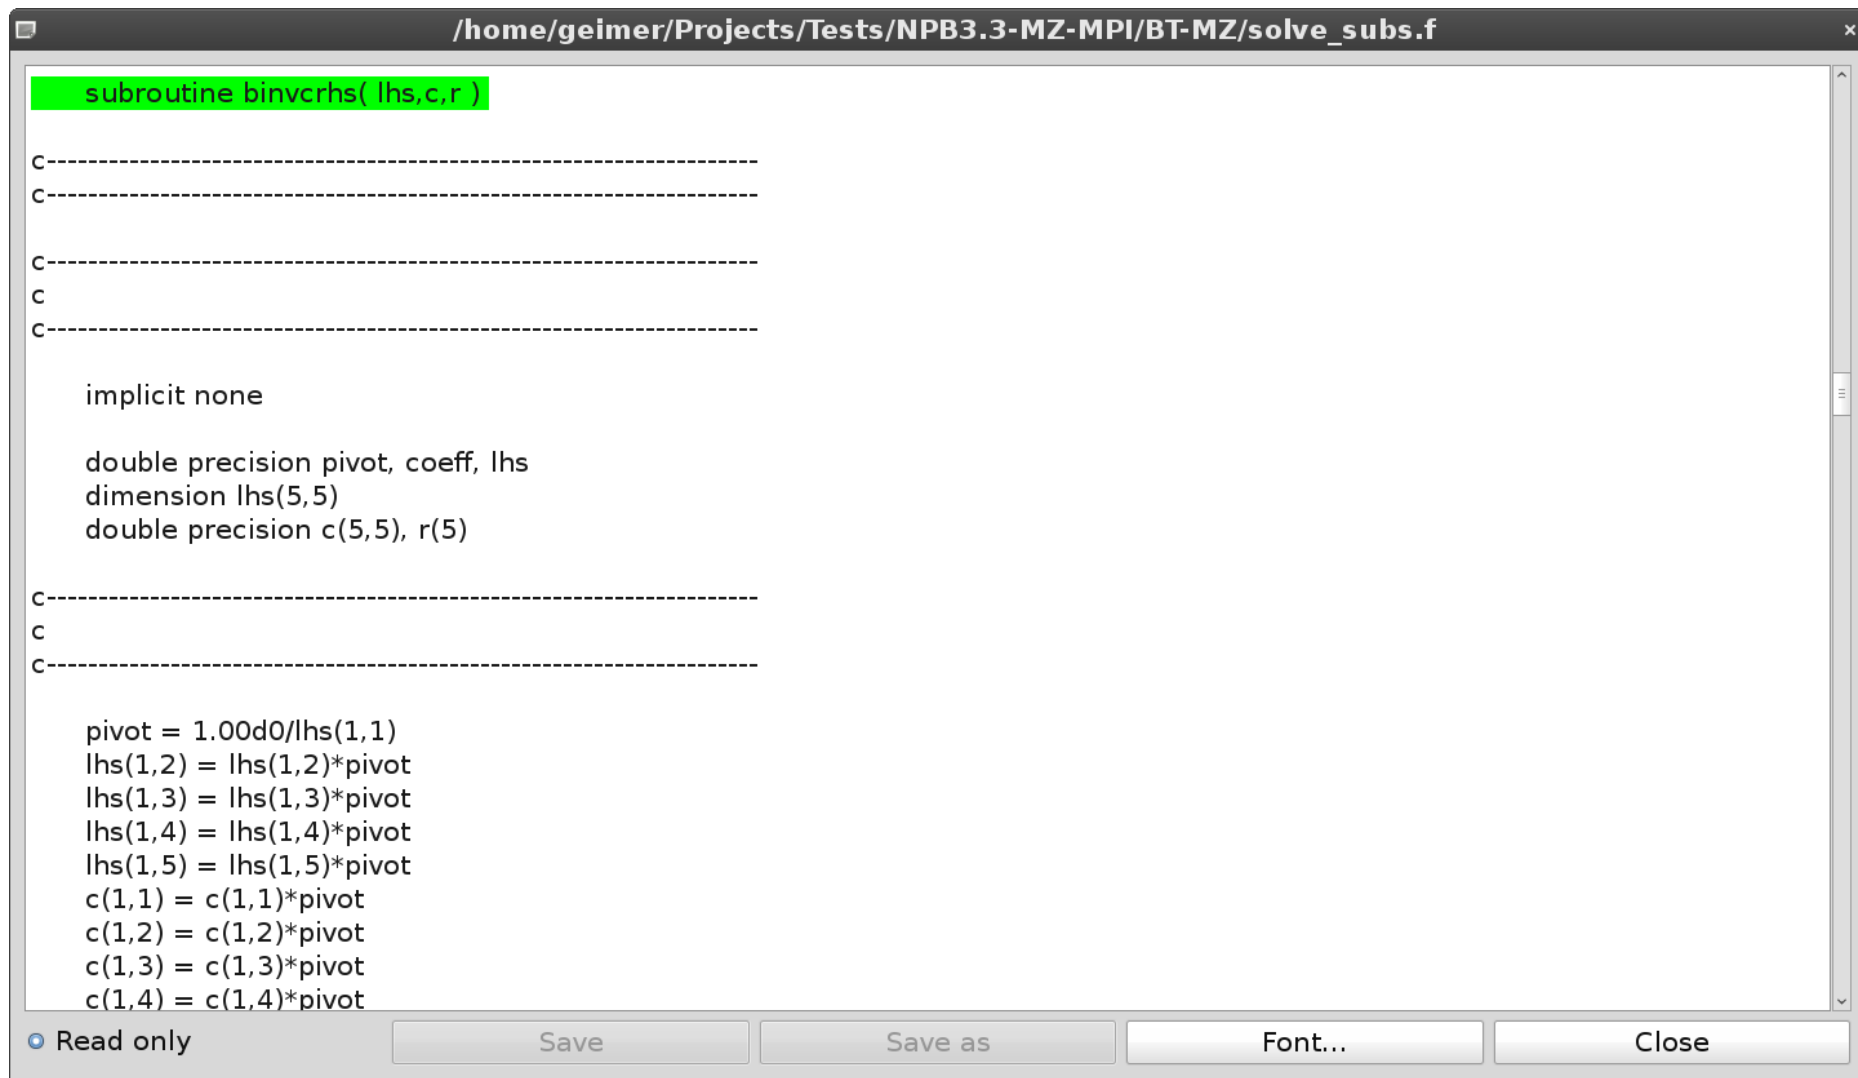


# Source-code view via context menu

The screenshot displays the VI-HPS interface with three main panels:

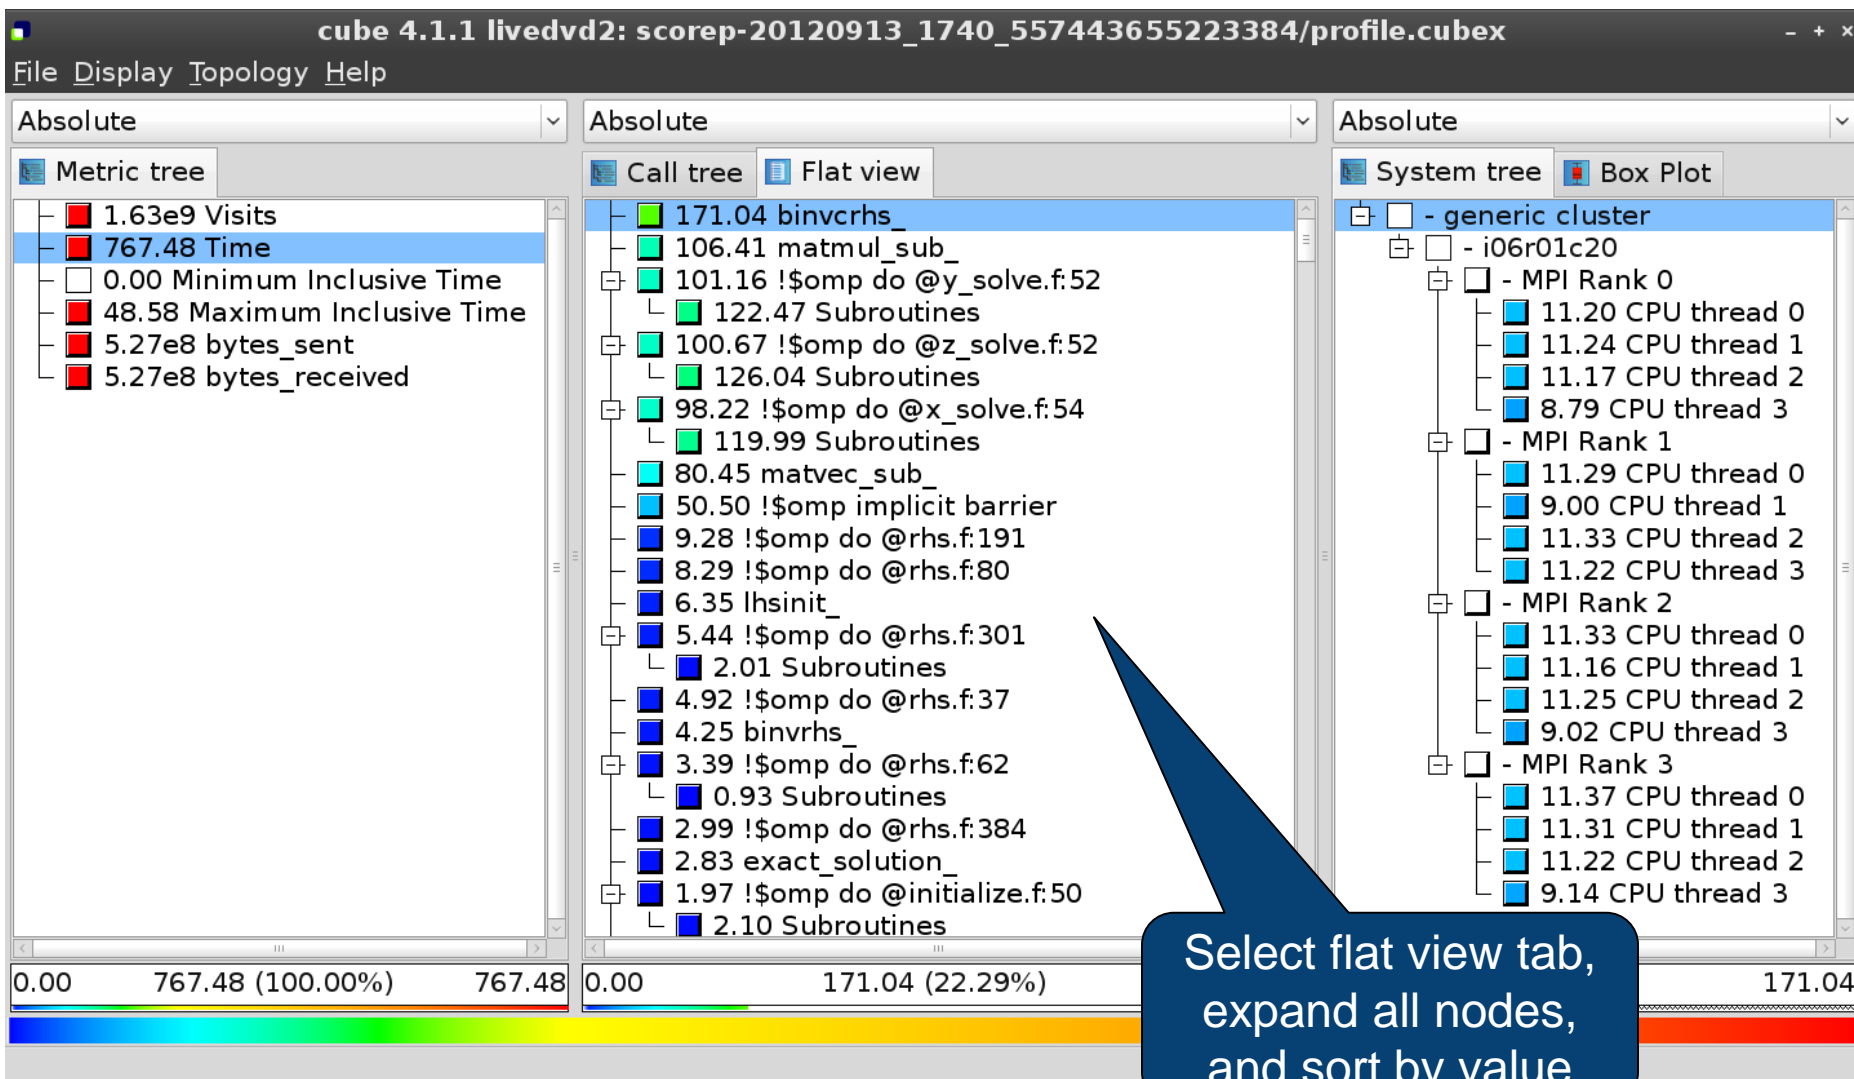
- Metric tree:** Shows various performance metrics such as Visits, Time, Minimum Inclusive Time, Maximum Inclusive Time, bytes\_sent, and bytes\_received.
- Call tree:** Displays a hierarchical view of function calls. The '57.70 binvcrhs' node is highlighted, and a context menu is open over it. The menu options include: Call site, Called region, Expand/collapse, Hiding, Cut call tree, Find items, Find Next, Clear found items, Copy to clipboard, and Min/max values. The 'Source code' option is selected.
- System tree:** Shows the system configuration, including the generic cluster, i06r01c20, and MPI Ranks 0, 1, 2, and 3, each with its own set of CPU threads.

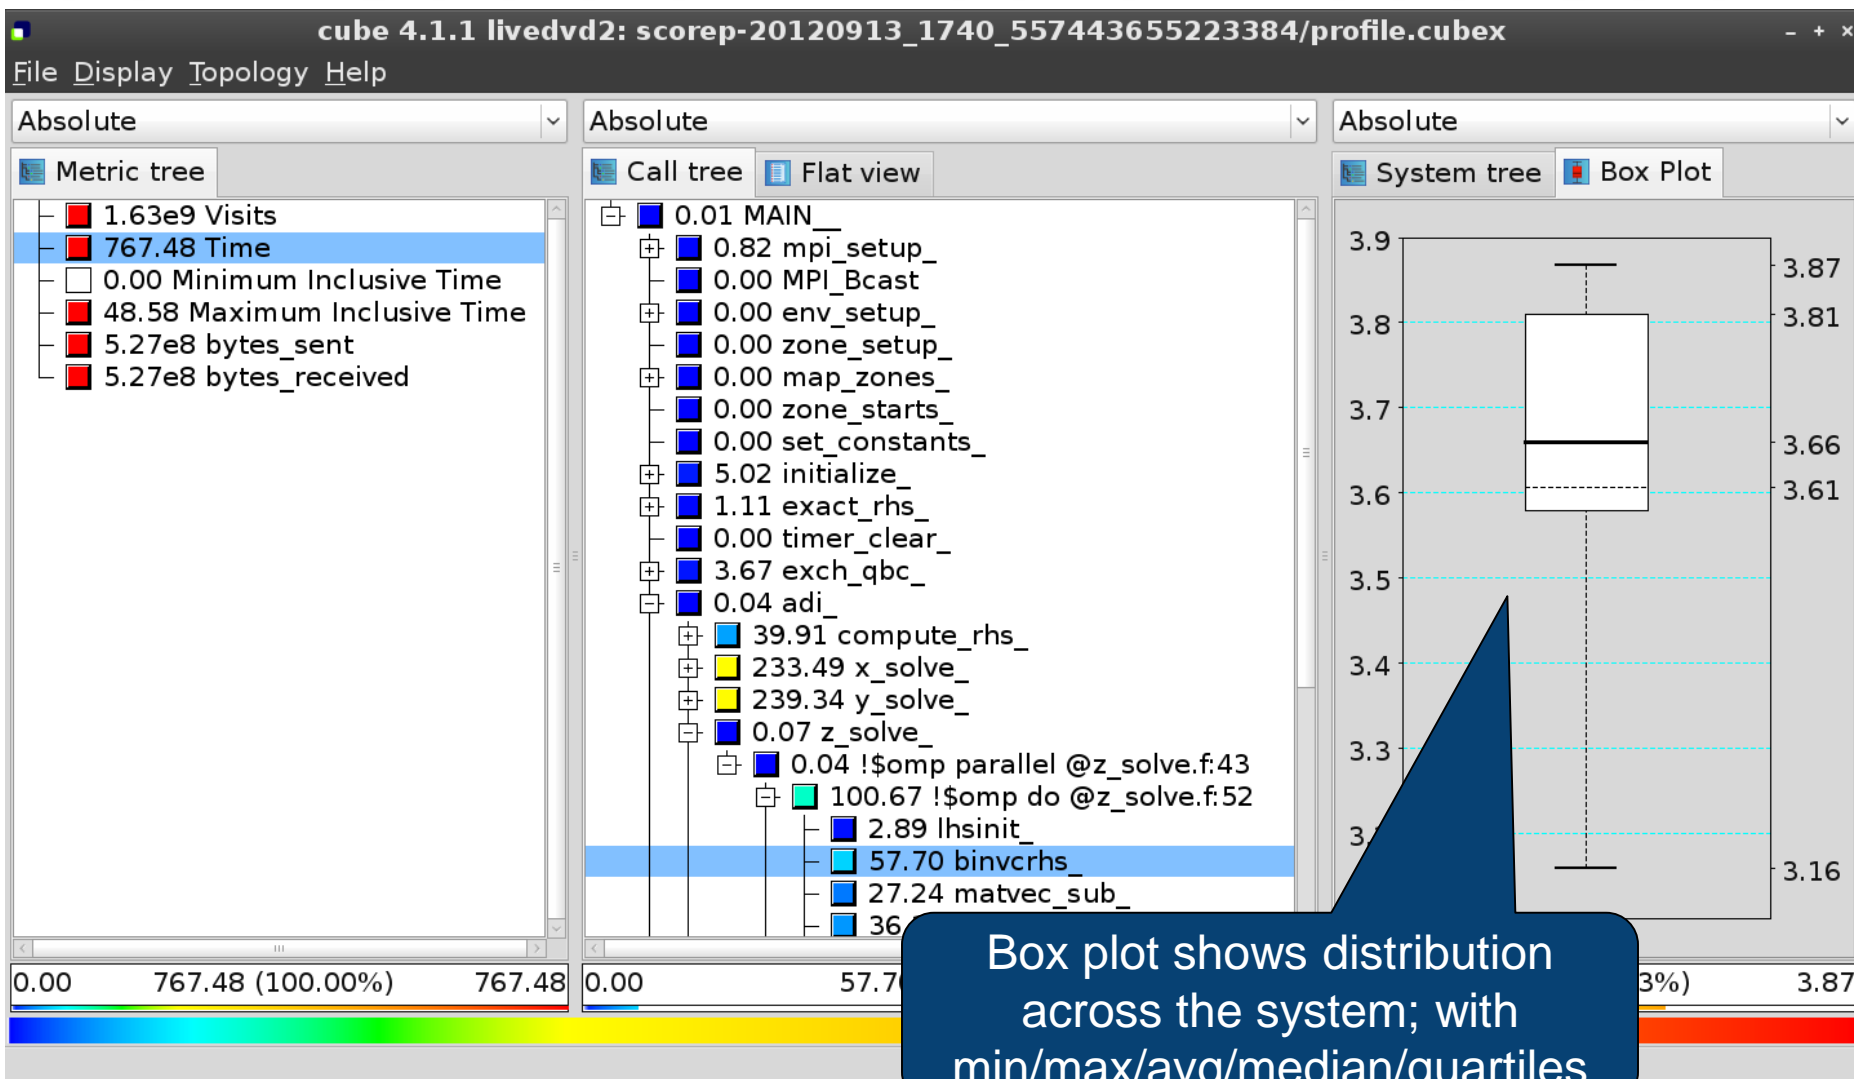
A blue callout bubble points to the context menu with the text: "Right-click opens context menu".

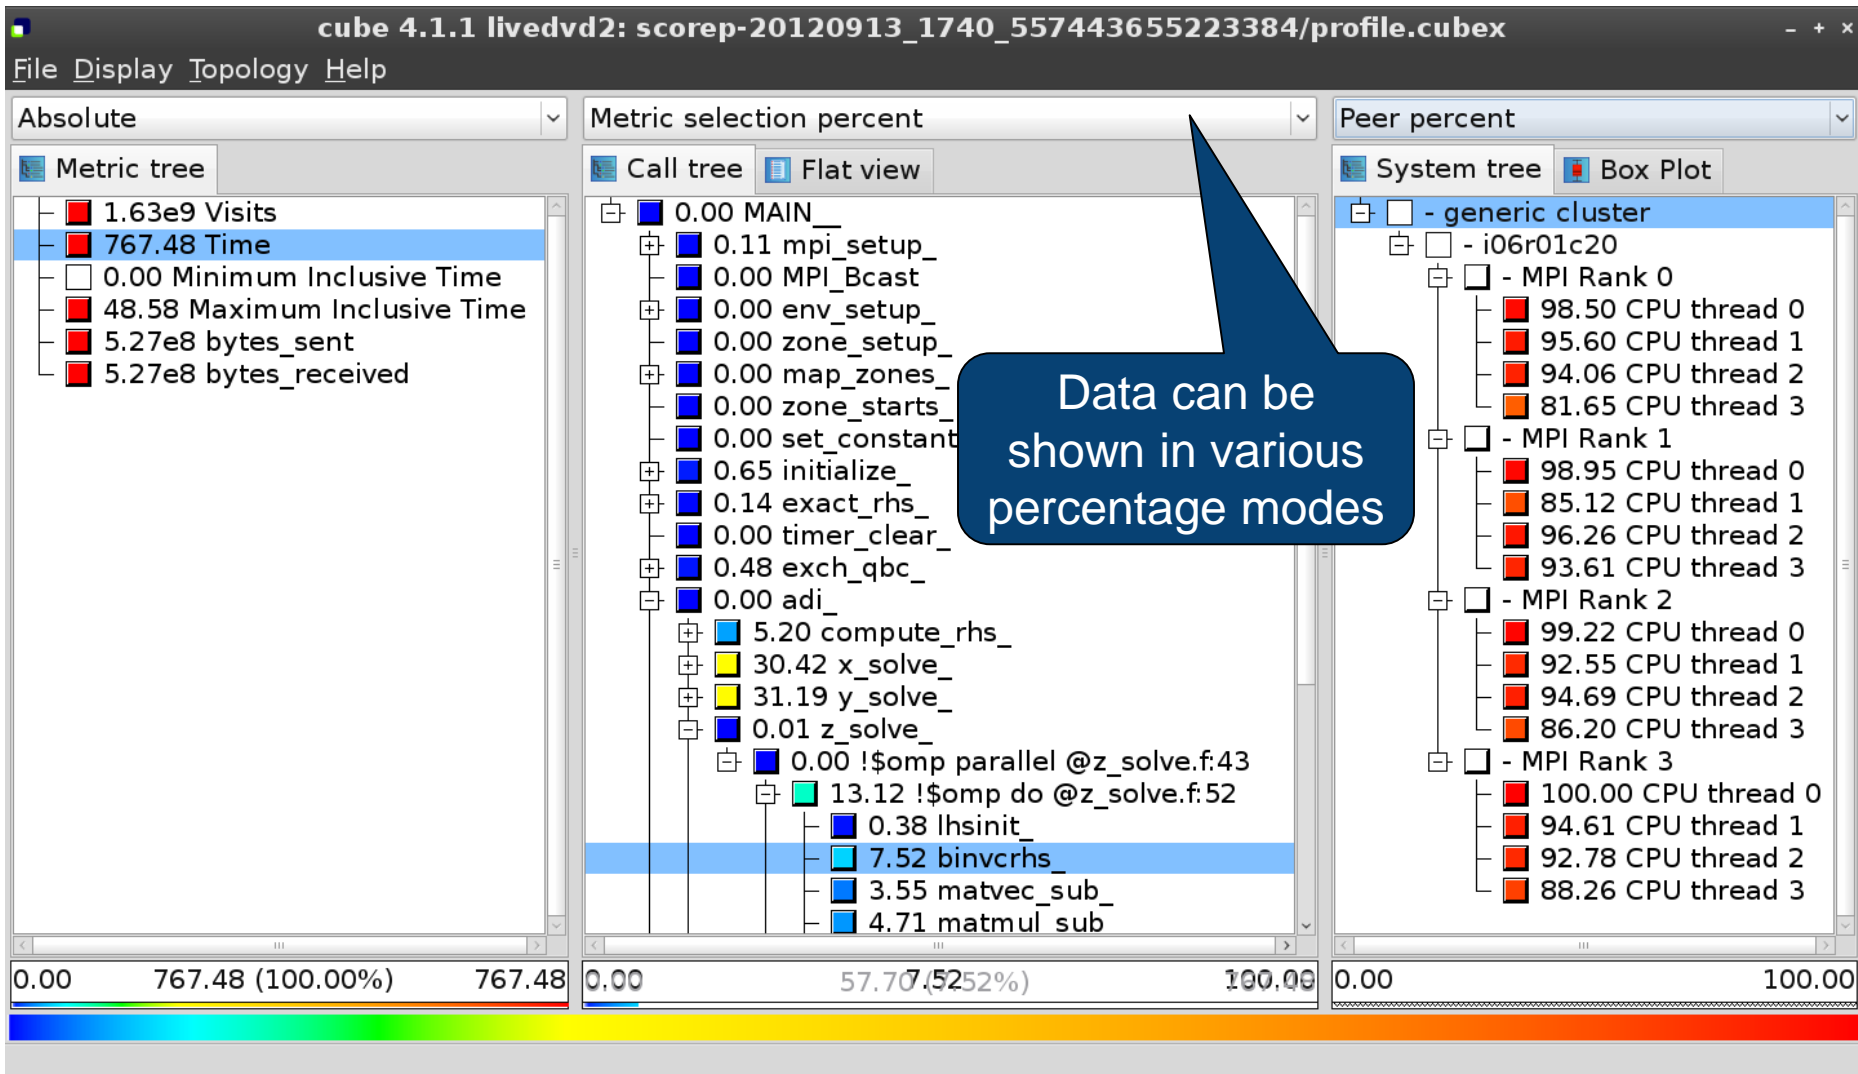


```
subroutine binvrchs( lhs,c,r )  
  
C-----  
C-----  
  
C-----  
C  
C-----  
  
implicit none  
  
double precision pivot, coeff, lhs  
dimension lhs(5,5)  
double precision c(5,5), r(5)  
  
C-----  
C  
C-----  
  
pivot = 1.00d0/lhs(1,1)  
lhs(1,2) = lhs(1,2)*pivot  
lhs(1,3) = lhs(1,3)*pivot  
lhs(1,4) = lhs(1,4)*pivot  
lhs(1,5) = lhs(1,5)*pivot  
c(1,1) = c(1,1)*pivot  
c(1,2) = c(1,2)*pivot  
c(1,3) = c(1,3)*pivot  
c(1,4) = c(1,4)*pivot
```

Read only Save Save as Font... Close



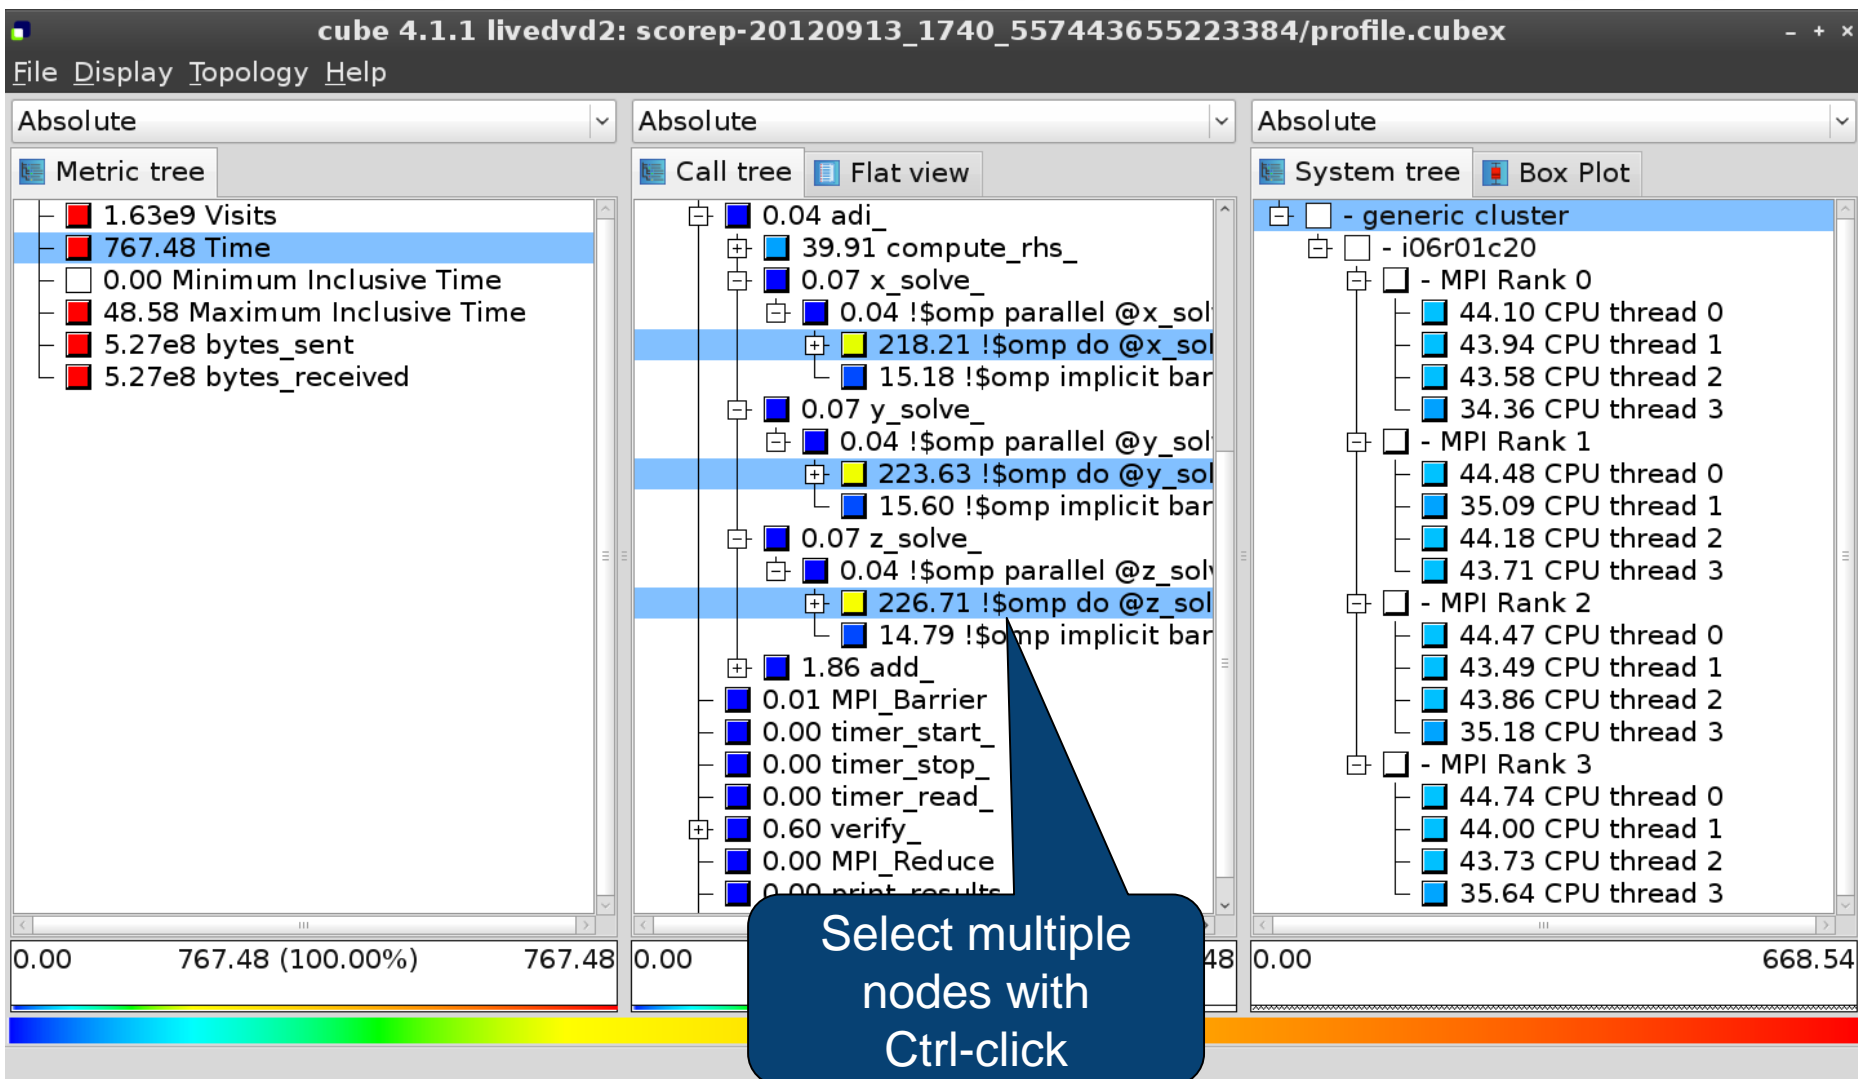




- Absolute
  - Absolute value shown in seconds/bytes/counts
- Selection percent
  - Value shown as percentage w.r.t. the selected node “on the left” (metric/call path)
- Peer percent (system tree only)
  - Value shown as percentage relative to the maximum peer value



# Multiple selection



- Value of the derived metric is not stored, but **calculated** on-the-fly
- One defines an *CubePL* expression, like  
`metric::time(i)/metric::visits(e)`
- Types of derived metrics:
  - **Predetermined**: evaluation of the *CubePL* expression is done before the aggregation
  - **Postderived**: evaluation of the *CubePL* expression is performed after the aggregation
- Examples:
  - “Average execution time”  
Postderived metric with an expression:  
`metric::time(i)/metric::visits(e)`
  - “Number of FLOP per second”  
Postderived metric with an expression:  
`metric::FLOP()/metric::time()`



The screenshot displays the Cube GUI interface with three main panels: 'Metric tree', 'Call tree', and 'System tree'. The 'Metric tree' panel on the left lists various metrics, including '1.09e8 Visits (occ)', '1.01e6 Time (sec)', '0.00 Minimum Inclusive Time (sec)', '246.14 Maximum Inclusive Time (sec)', '7.18e12 bytes\_sent', and '7.18e12 bytes\_received'. The 'Call tree' panel in the center shows a hierarchical view of system calls, such as '0.35 main(int, char \*[])', '2512.10 ugshellInit', and '1.01e6 ug::script:LoadUGScript(const char \*, bool)'. The 'System tree' panel on the right shows a hierarchical view of system components, including 'machine Blue Gene/Q', 'rack 11', 'midplane 1', 'nodecard 8', 'nodecard 4', and 'MPI Rank 0'.

Three blue callout boxes highlight key features:

- Collection of derived metrics**: Points to the 'Metric tree' panel.
- Parameters of the derived metric**: Points to the 'Create new metric as a child of metric' dialog box.
- CubePL expression**: Points to the 'Calculation' field in the dialog box.

The 'Create new metric as a child of metric' dialog box is open, showing the following fields:

- Select metric from collection**: Average execution time (kenobi)
- Derived metric type**: Postderived metric
- Display name**: Average visit time
- Unique name**: avg\_visit\_time
- Type**: DOUBLE
- Unit of measurement**: sec
- URL**: (empty)
- Description**: Calculates average time of region execution per visit. Autor is Michael Knobloch.
- Calculation**: `metric::time(l)/metric::visits(e)`

The dialog box also includes buttons for 'Calculation', 'Calculation Init', 'Aggregation "+"', and 'Aggregation "-"', as well as 'Create metric' and 'Cancel' buttons at the bottom.



The screenshot displays the VI-HPS application window titled "cube 4.1.1 livedvd2: scorep-20120913\_1740\_557443655223384/profile.cubex". The "Help" menu is open, showing options: "Getting started", "Mouse and keyboard control", "What's This? (Shift+F1)", "About", "Selected metrics description", and "Selected regions description". A blue callout bubble points to the "What's This?" option with the text: "Context-sensitive help available for all GUI items".

The GUI is divided into three main panels:

- Metric tree (Left):** Displays a list of metrics with color-coded bars. The "767.48 Time" metric is highlighted in blue. Below the list is a progress bar showing "0.00 767.48 (100.00%) 767.48".
- System tree (Right):** Displays a hierarchical tree of system components. The "generic cluster" is expanded, showing "i06r01c20" and four MPI Ranks. Each rank lists CPU threads and their times. Below the tree is a progress bar showing "0.00 668.54 (87.11%) 767.48".
- Bottom Panel:** A color-coded progress bar and a status bar that reads "Change into help mode for display components".

- Extracting solver sub-tree from analysis report

```
% cube_cut -r '<<ITERATION>>' scorep_bt-mz_W_4x4_sum/profile.cubex  
Writing cut.cubex... done.
```

- Calculating difference of two reports

```
% cube_diff scorep_bt-mz_W_4x4_sum/profile.cubex cut.cubex  
Writing diff.cubex... done.
```

- Additional utilities for merging, calculating mean, etc.
  - Default output of `cube_utility` is a new report `utility.cubex`
- Further utilities for report scoring & statistics
- Run utility with “-h” (or no arguments) for brief usage info

- CUBE
  - Parallel program analysis report exploration tools
    - Libraries for XML report reading & writing
    - Algebra utilities for report processing
    - GUI for interactive analysis exploration
  - Available under New BSD open-source license
  - Documentation & sources:
    - <http://www.scalasca.org>
  - User guide also part of installation:
    - ``cube-config --cube-dir`/share/doc/CubeGuide.pdf`
  - Contact:
    - `mailto:scalasca@fz-juelich.de`

