Analysis report examination with CUBE

Dr. Alexandre Otto Strube
Jülich Supercomputing Centre

CUBE

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 4.3.2 (Jun 2015)

Analysis presentation and exploration

- 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 values: for precise comparison
  - As colors: for easy identification of hotspots
  - Inclusive value when closed & exclusive value when expanded
  - Customizable via display modes

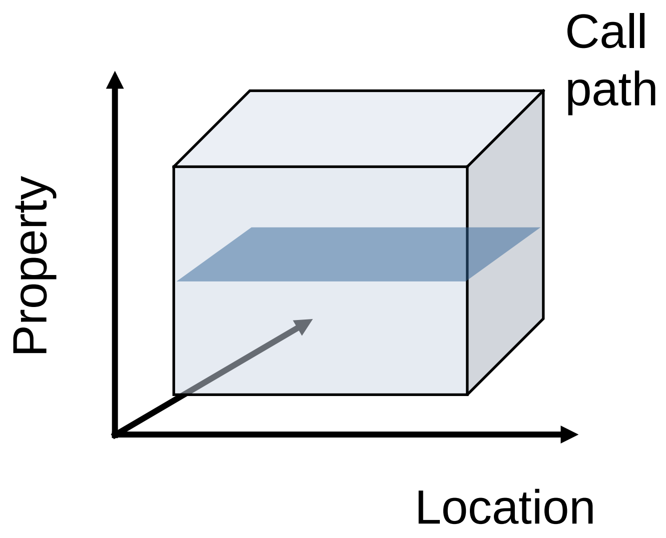

Analysis presentation

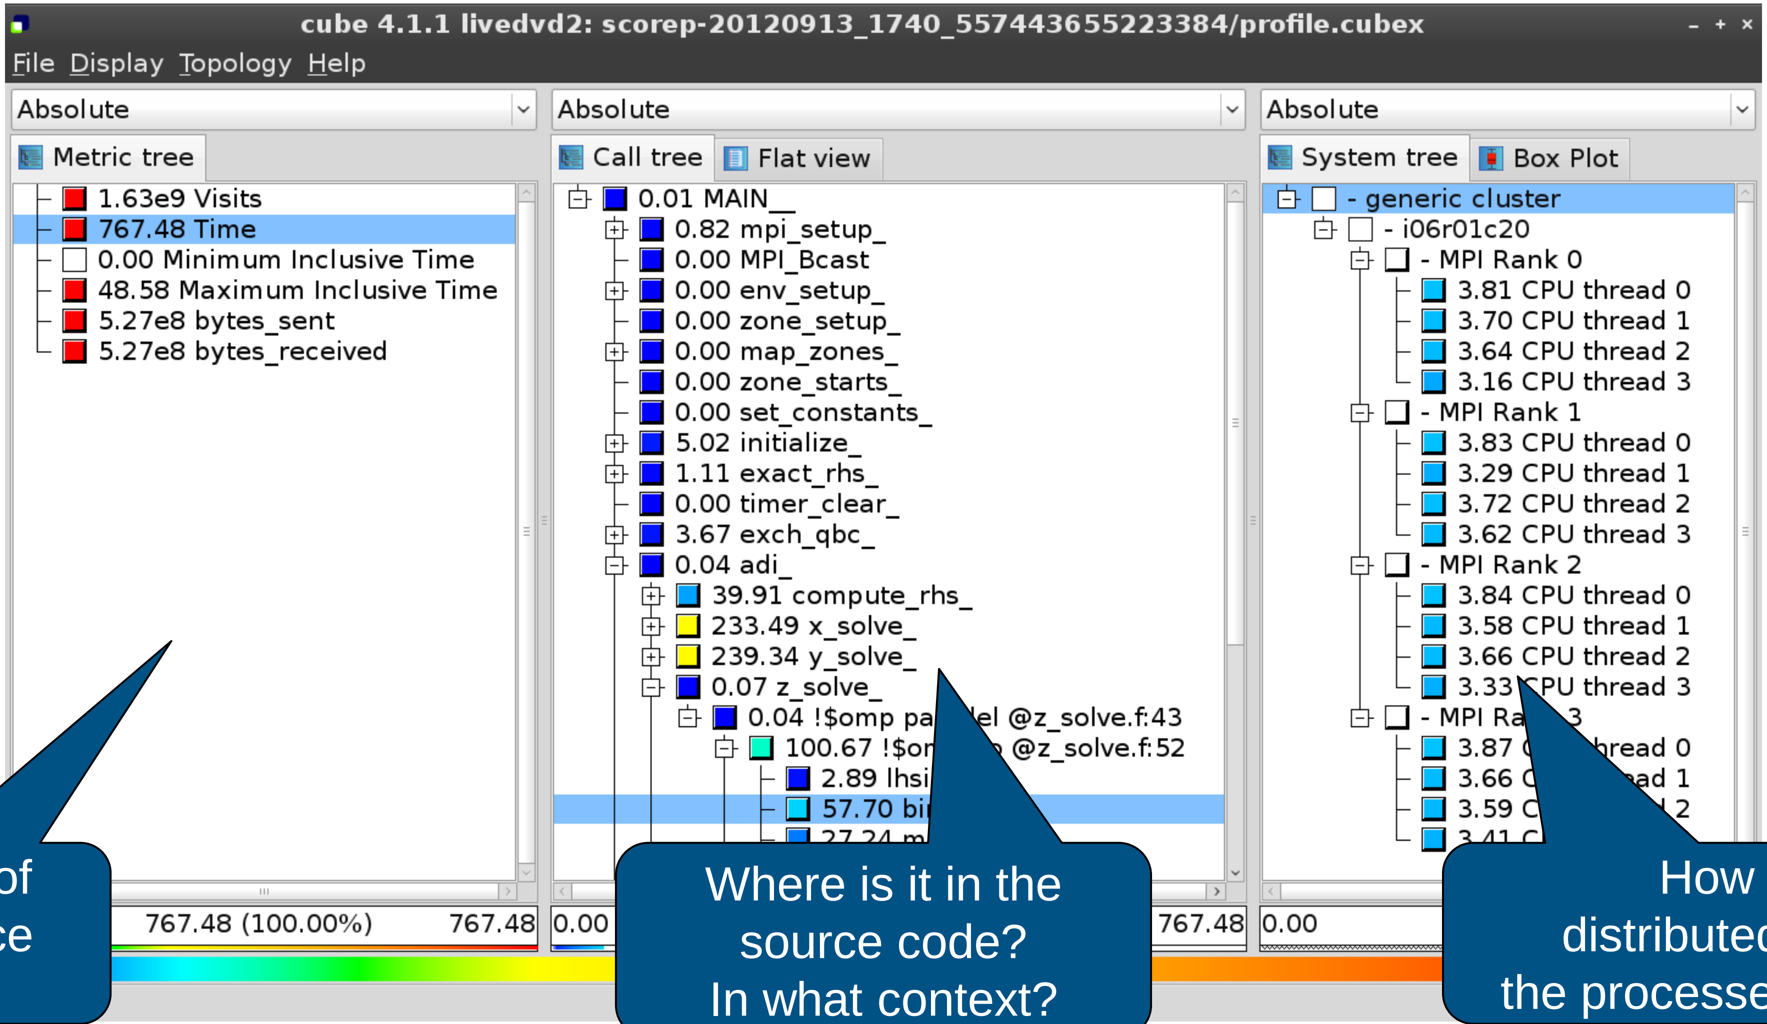

cube 4.1.1 livedvd2: scorep-20120913_1740_557443655223384/profile.cubex

```
1.63e9 Visits
767.48 Time
0.00 Minimum Inclusive Time
48.58 Maximum Inclusive Time
5.27e8 bytes_sent
5.27e8 bytes_received
```

```
0.01 MAIN__
  0.82 mpi_setup_
  0.00 MPI_Bcast
  0.00 env_setup_
  0.00 zone_setup_
  0.00 map_zones_
  0.00 zone_starts_
  0.00 set_constants_
  5.02 initialize_
  1.11 exact_rhs_
  0.00 timer_clear_
  3.67 exch_qbc_
  0.04 adi_
    39.91 compute_rhs_
    233.49 x_solve_
    239.34 y_solve_
    0.07 z_solve_
      0.04 !$omp pa    lel @z_solve.f:43
        100.67 !$or    @z_solve.f:52
          2.89 lhsi
          57.70 bi
          27.24 m
```

```
- generic cluster
  - i06r01c20
    - MPI Rank 0
      3.81 CPU thread 0
      3.70 CPU thread 1
      3.64 CPU thread 2
      3.16 CPU thread 3
    - MPI Rank 1
      3.83 CPU thread 0
      3.29 CPU thread 1
      3.72 CPU thread 2
      3.62 CPU thread 3
    - MPI Rank 2
      3.84 CPU thread 0
      3.58 CPU thread 1
      3.66 CPU thread 2
      3.33 CPU thread 3
    - MPI Ra    3
      3.87 C    hread 0
      3.66 C    ead 1
      3.59 C    2
      3.41 C
```

767.48 (100.00%) 767.48 | 0.00 767.48 | 0.00

What kind of performance metric?

Where is it in the source code? In what context?

How is it distributed across the processes/threads?

Analysis report exploration (opening view)

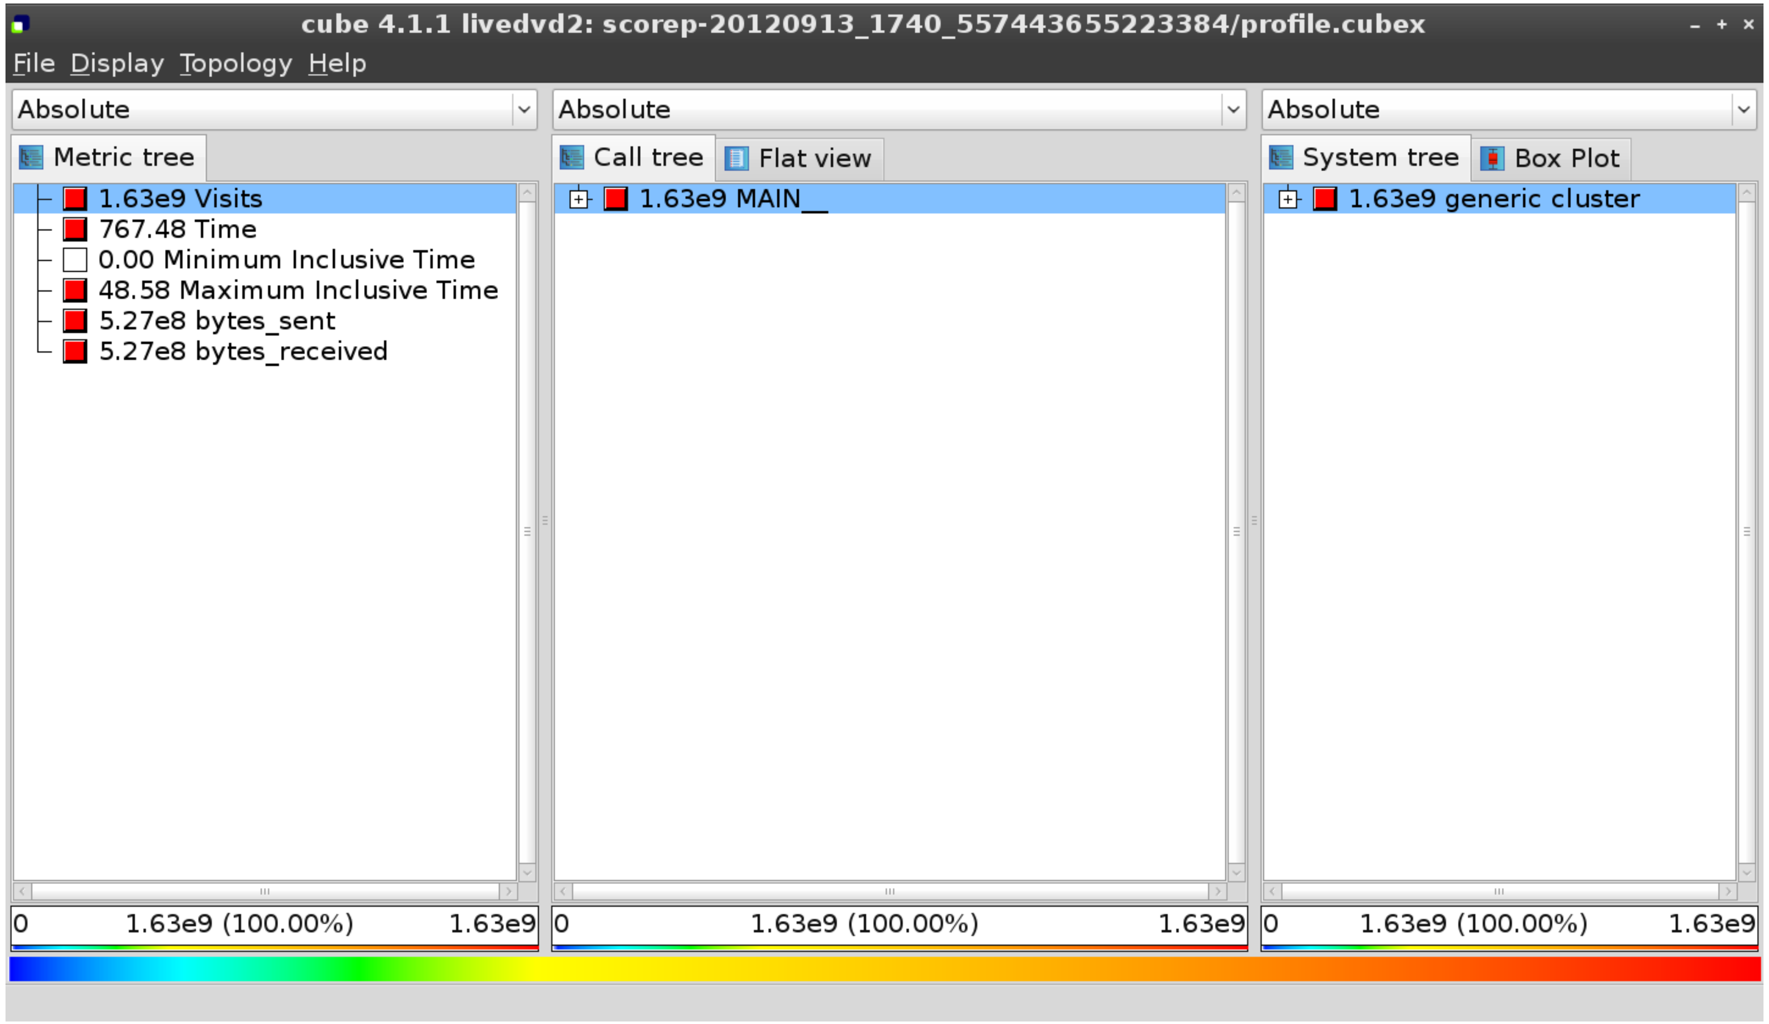
cube 4.1.1 livedvd2: scorep-20120913_1740_557443655223384/profile.cubex
File Display Topology Help
Absolute
Absolute
Absolute
Metric tree
Call tree
Flat view
System tree
Box Plot

1.63e9 Visits
767.48 Time
0.00 Minimum Inclusive Time
48.58 Maximum Inclusive Time
5.27e8 bytes_sent
5.27e8 bytes_received

1.63e9 MAIN__

1.63e9 generic cluster

0 1.63e9 (100.00%) 1.63e9

0 1.63e9 (100.00%) 1.63e9

0 1.63e9 (100.00%) 1.63e9

Metric selection

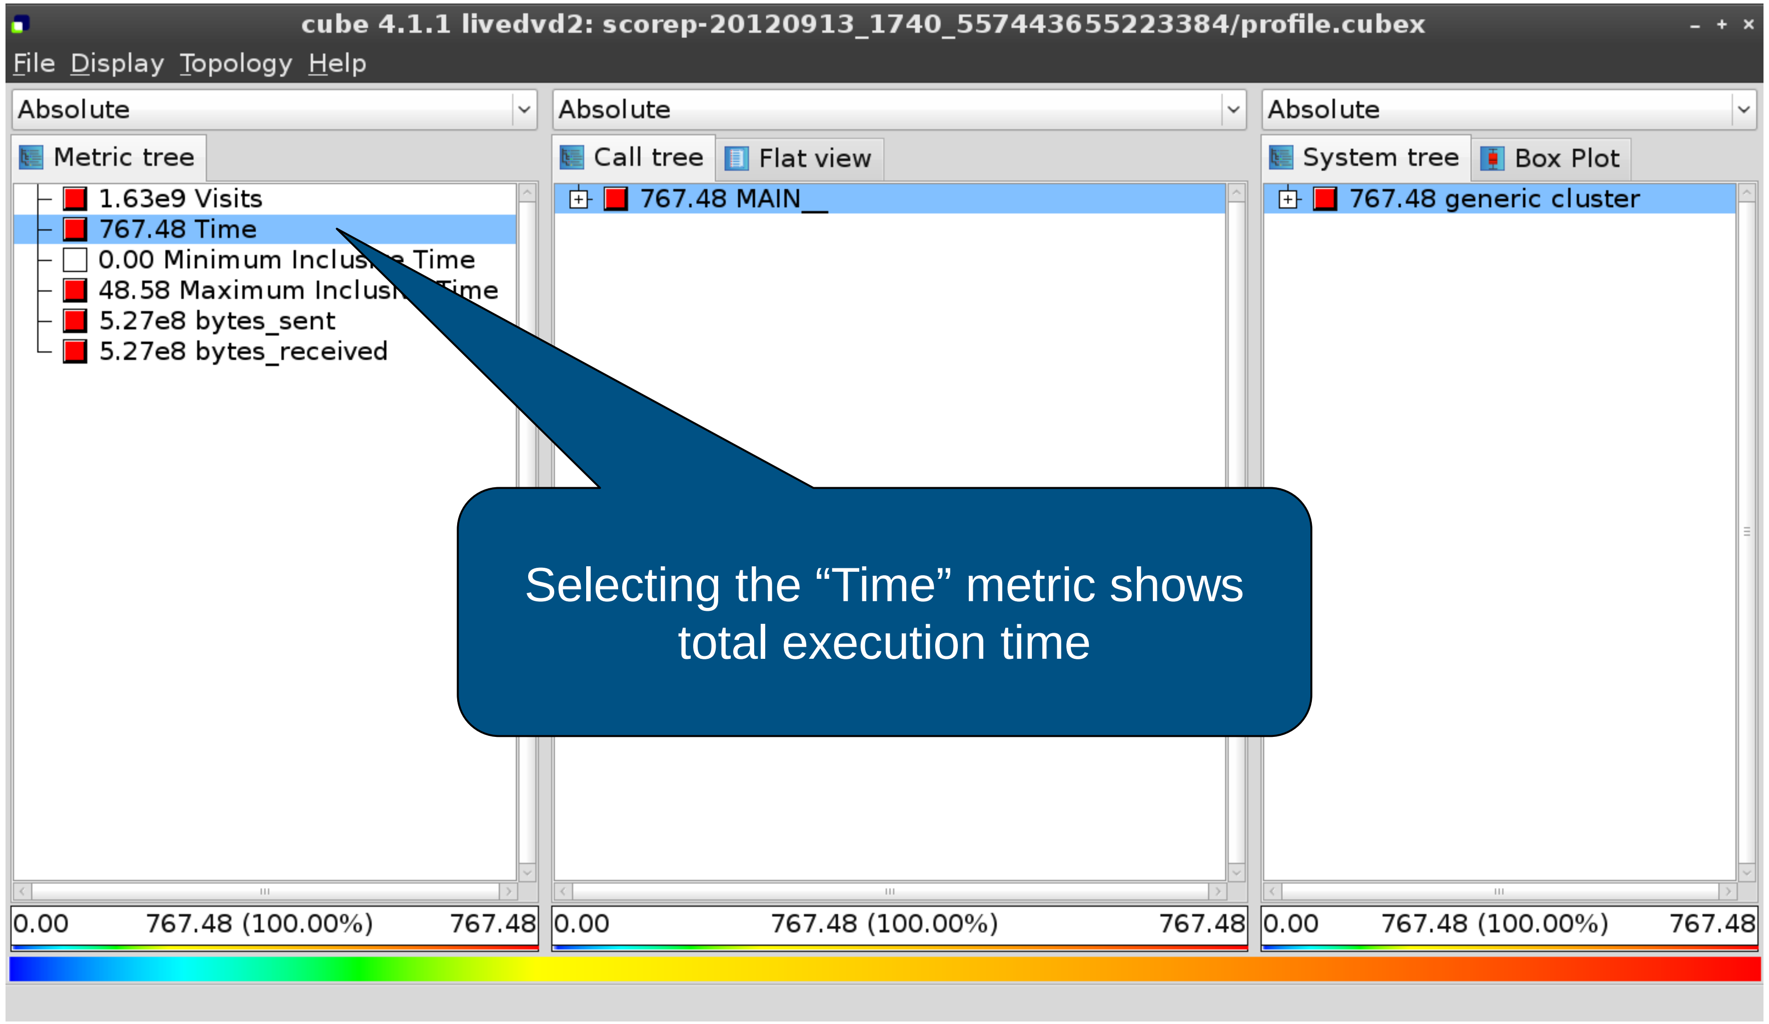
cube 4.1.1 livedvd2: scorep-20120913_1740_557443655223384/profile.cubex

File Display Topology Help

Absolute	Absolute	Absolute
Metric tree	Call tree / Flat view	System tree / Box Plot
🟥 1.63e9 Visits	🟥 767.48 MAIN__	🟥 767.48 generic cluster
🟥 767.48 Time		
☐ 0.00 Minimum Inclusive Time		
🟥 48.58 Maximum Inclusive Time		
🟥 5.27e8 bytes_sent		
🟥 5.27e8 bytes_received		
0.00 767.48 (100.00%) 767.48	0.00 767.48 (100.00%) 767.48	0.00 767.48 (100.00%) 767.48

Selecting the “Time” metric shows total execution time

Expanding the system tree

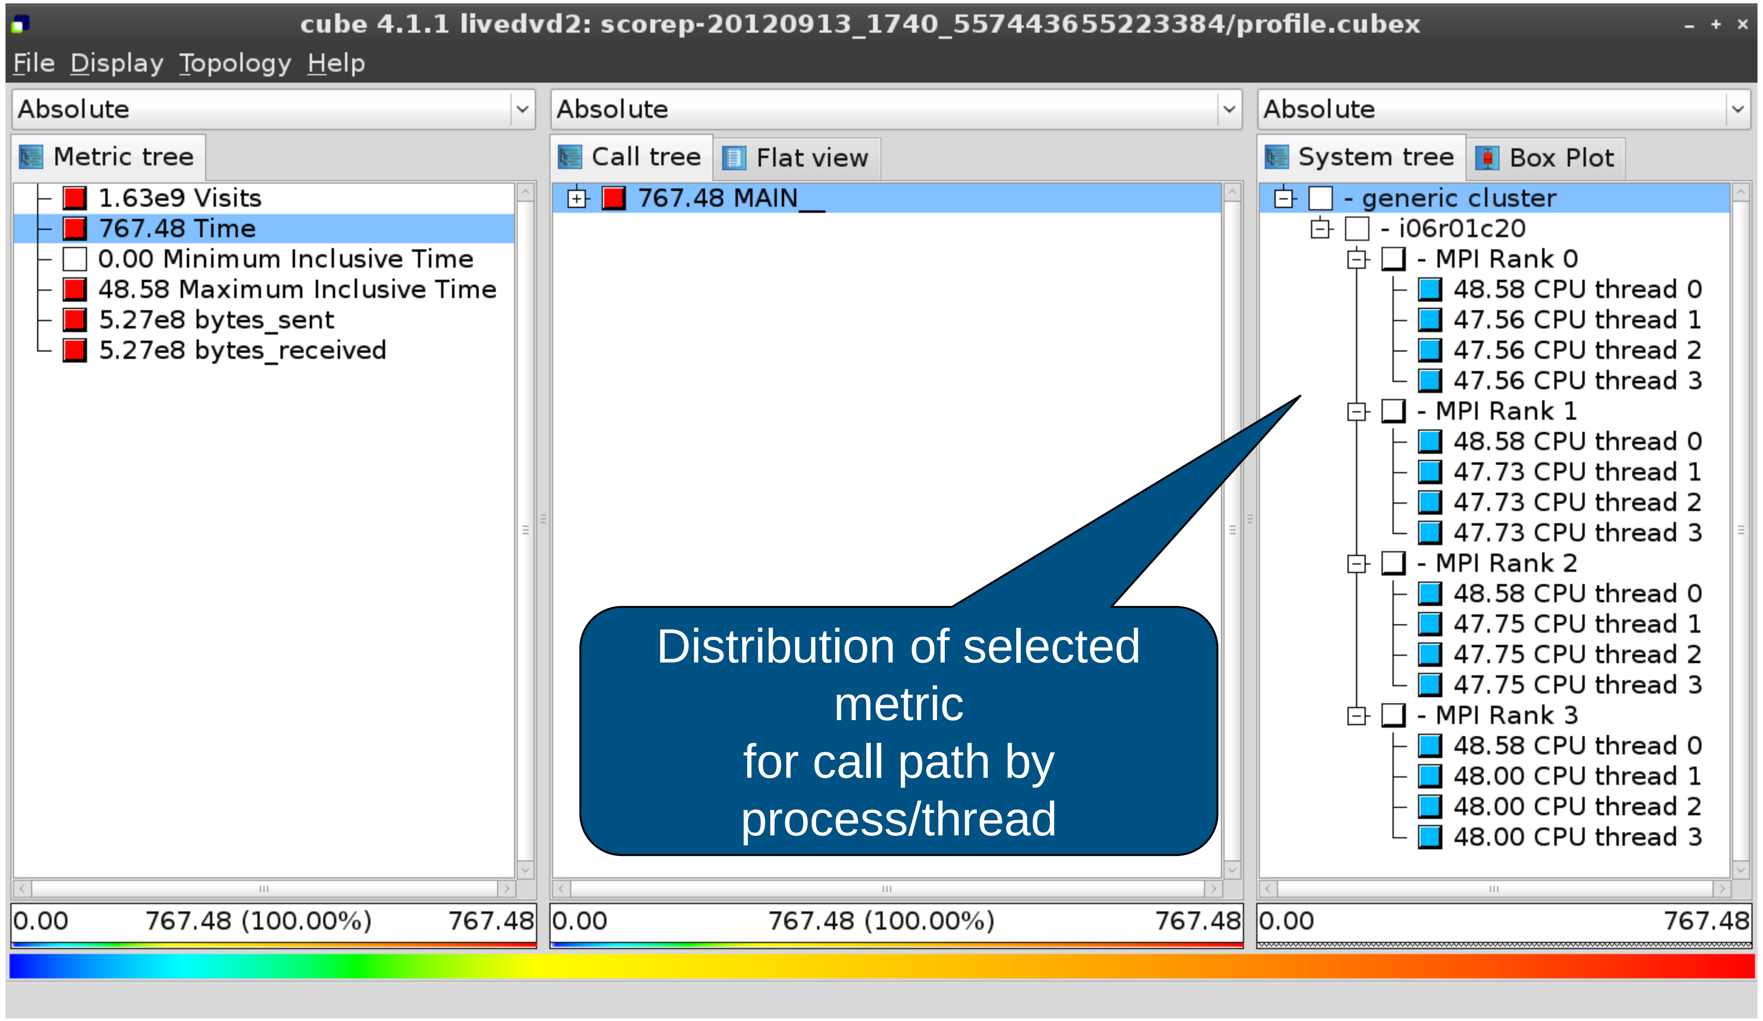

cube 4.1.1 livedvd2: scorep-20120913_1740_557443655223384/profile.cubex

File Display Topology Help

Absolute | Absolute | Absolute

Metric tree
- 1.63e9 Visits
- 767.48 Time
- 0.00 Minimum Inclusive Time
- 48.58 Maximum Inclusive Time
- 5.27e8 bytes_sent
- 5.27e8 bytes_received

Call tree | Flat view
- 767.48 MAIN__

System tree | Box Plot
- generic cluster
  - i06r01c20
    - MPI Rank 0
      - 48.58 CPU thread 0
      - 47.56 CPU thread 1
      - 47.56 CPU thread 2
      - 47.56 CPU thread 3
    - MPI Rank 1
      - 48.58 CPU thread 0
      - 47.73 CPU thread 1
      - 47.73 CPU thread 2
      - 47.73 CPU thread 3
    - MPI Rank 2
      - 48.58 CPU thread 0
      - 47.75 CPU thread 1
      - 47.75 CPU thread 2
      - 47.75 CPU thread 3
    - MPI Rank 3
      - 48.58 CPU thread 0
      - 48.00 CPU thread 1
      - 48.00 CPU thread 2
      - 48.00 CPU thread 3

0.00 767.48 (100.00%) 767.48 | 0.00 767.48 (100.00%) 767.48 | 0.00 767.48

Distribution of selected metric
for call path by
process/thread

Expanding the call tree

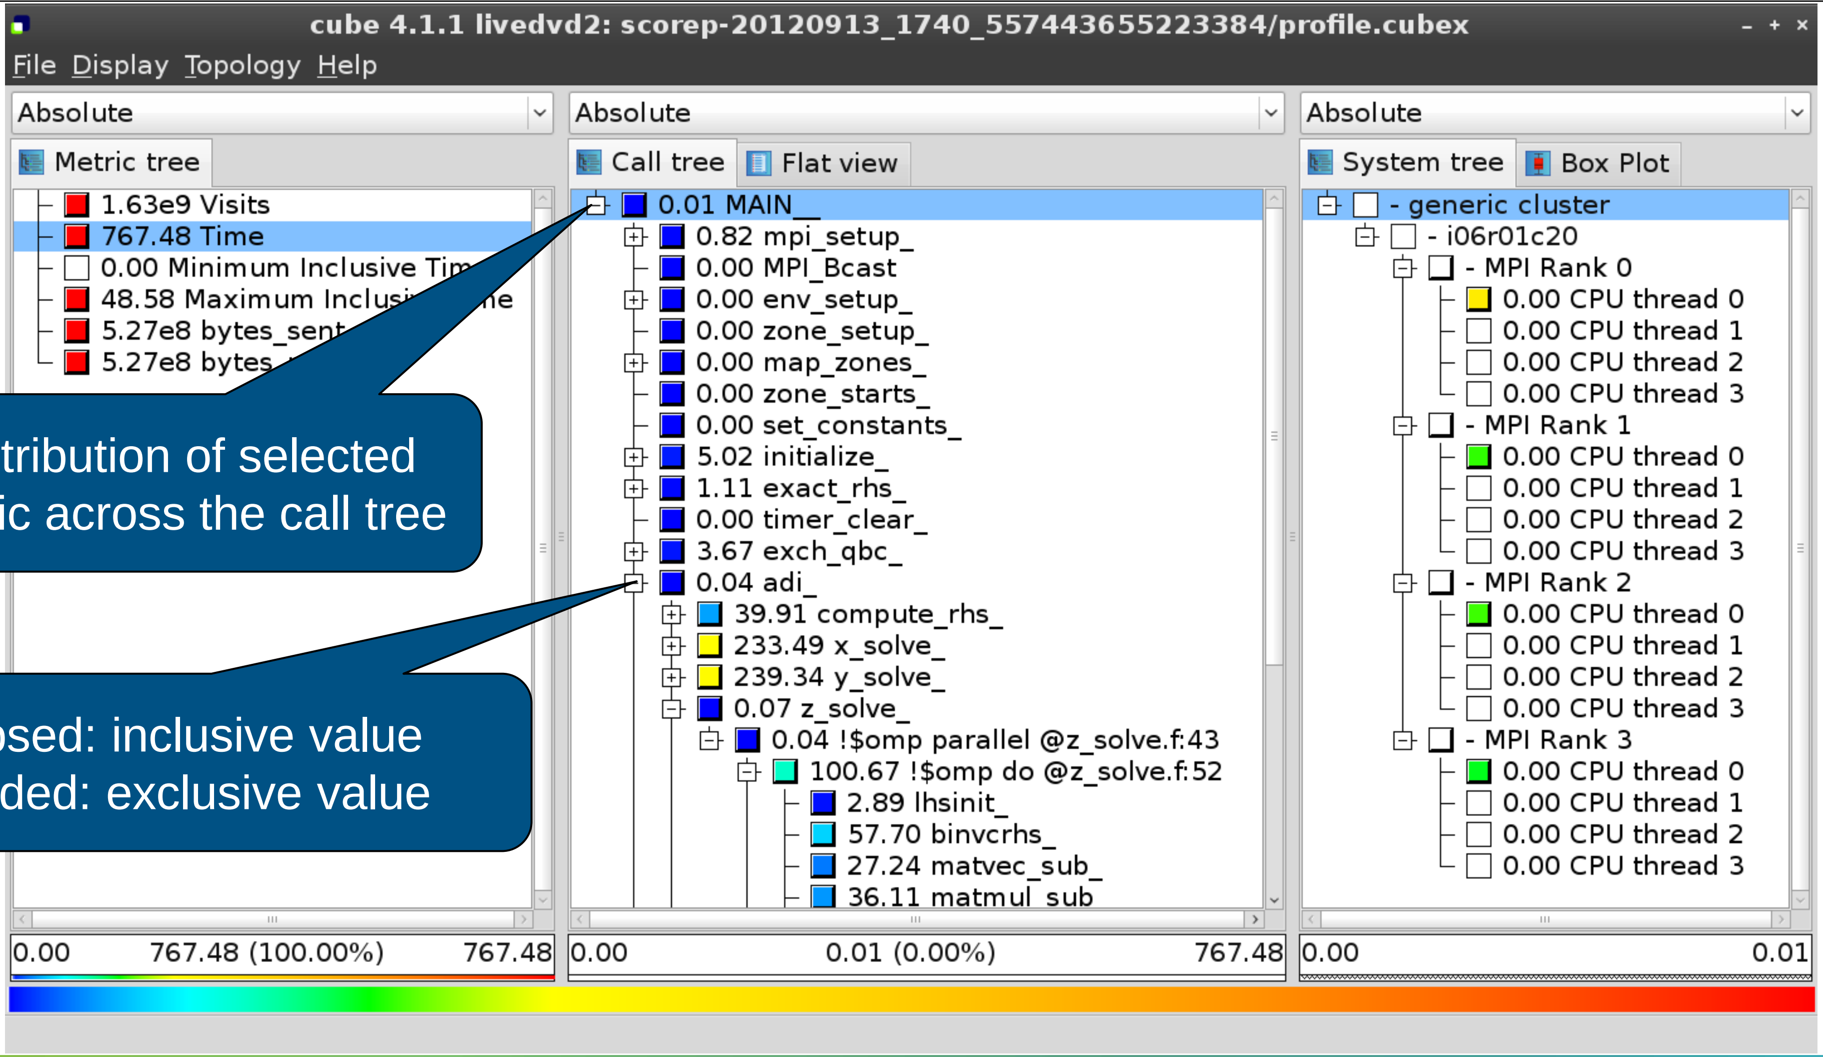
cube 4.1.1 livedvd2: scorep-20120913_1740_557443655223384/profile.cubex

Metric tree:
```
1.63e9 Visits
767.48 Time
0.00 Minimum Inclusive Tim
48.58 Maximum Inclusi      ne
5.27e8 bytes_sent
5.27e8 bytes
```

Call tree:
```
0.01 MAIN__
  0.82 mpi_setup_
  0.00 MPI_Bcast
  0.00 env_setup_
  0.00 zone_setup_
  0.00 map_zones_
  0.00 zone_starts_
  0.00 set_constants_
  5.02 initialize_
  1.11 exact_rhs_
  0.00 timer_clear_
  3.67 exch_qbc_
  0.04 adi_
    39.91 compute_rhs_
    233.49 x_solve_
    239.34 y_solve_
    0.07 z_solve_
      0.04 !$omp parallel @z_solve.f:43
        100.67 !$omp do @z_solve.f:52
          2.89 lhsinit_
          57.70 binvcrhs_
          27.24 matvec_sub_
          36.11 matmul_sub
```

System tree:
```
- generic cluster
  - i06r01c20
    - MPI Rank 0
      0.00 CPU thread 0
      0.00 CPU thread 1
      0.00 CPU thread 2
      0.00 CPU thread 3
    - MPI Rank 1
      0.00 CPU thread 0
      0.00 CPU thread 1
      0.00 CPU thread 2
      0.00 CPU thread 3
    - MPI Rank 2
      0.00 CPU thread 0
      0.00 CPU thread 1
      0.00 CPU thread 2
      0.00 CPU thread 3
    - MPI Rank 3
      0.00 CPU thread 0
      0.00 CPU thread 1
      0.00 CPU thread 2
      0.00 CPU thread 3
```

```
0.00    767.48 (100.00%)    767.48 | 0.00    0.01 (0.00%)    767.48 | 0.00    0.01
```

Distribution of selected metric across the call tree

Collapsed: inclusive value
Expanded: exclusive value

Inclusive vs. Exclusive values

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

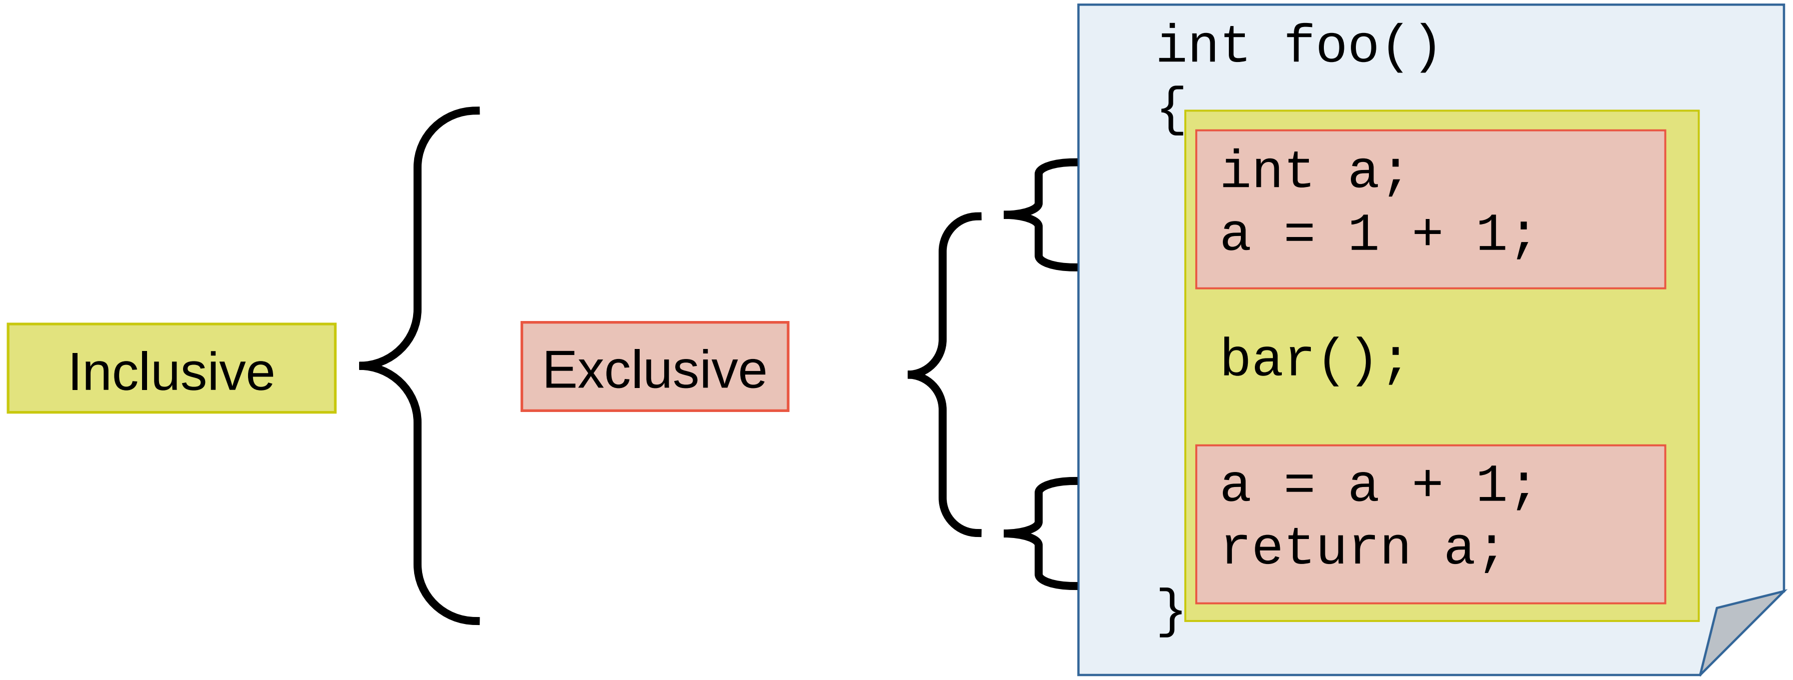

```
int foo()
{
  int a;
  a = 1 + 1;

  bar();

  a = a + 1;
  return a;
}
```

Selecting a call path

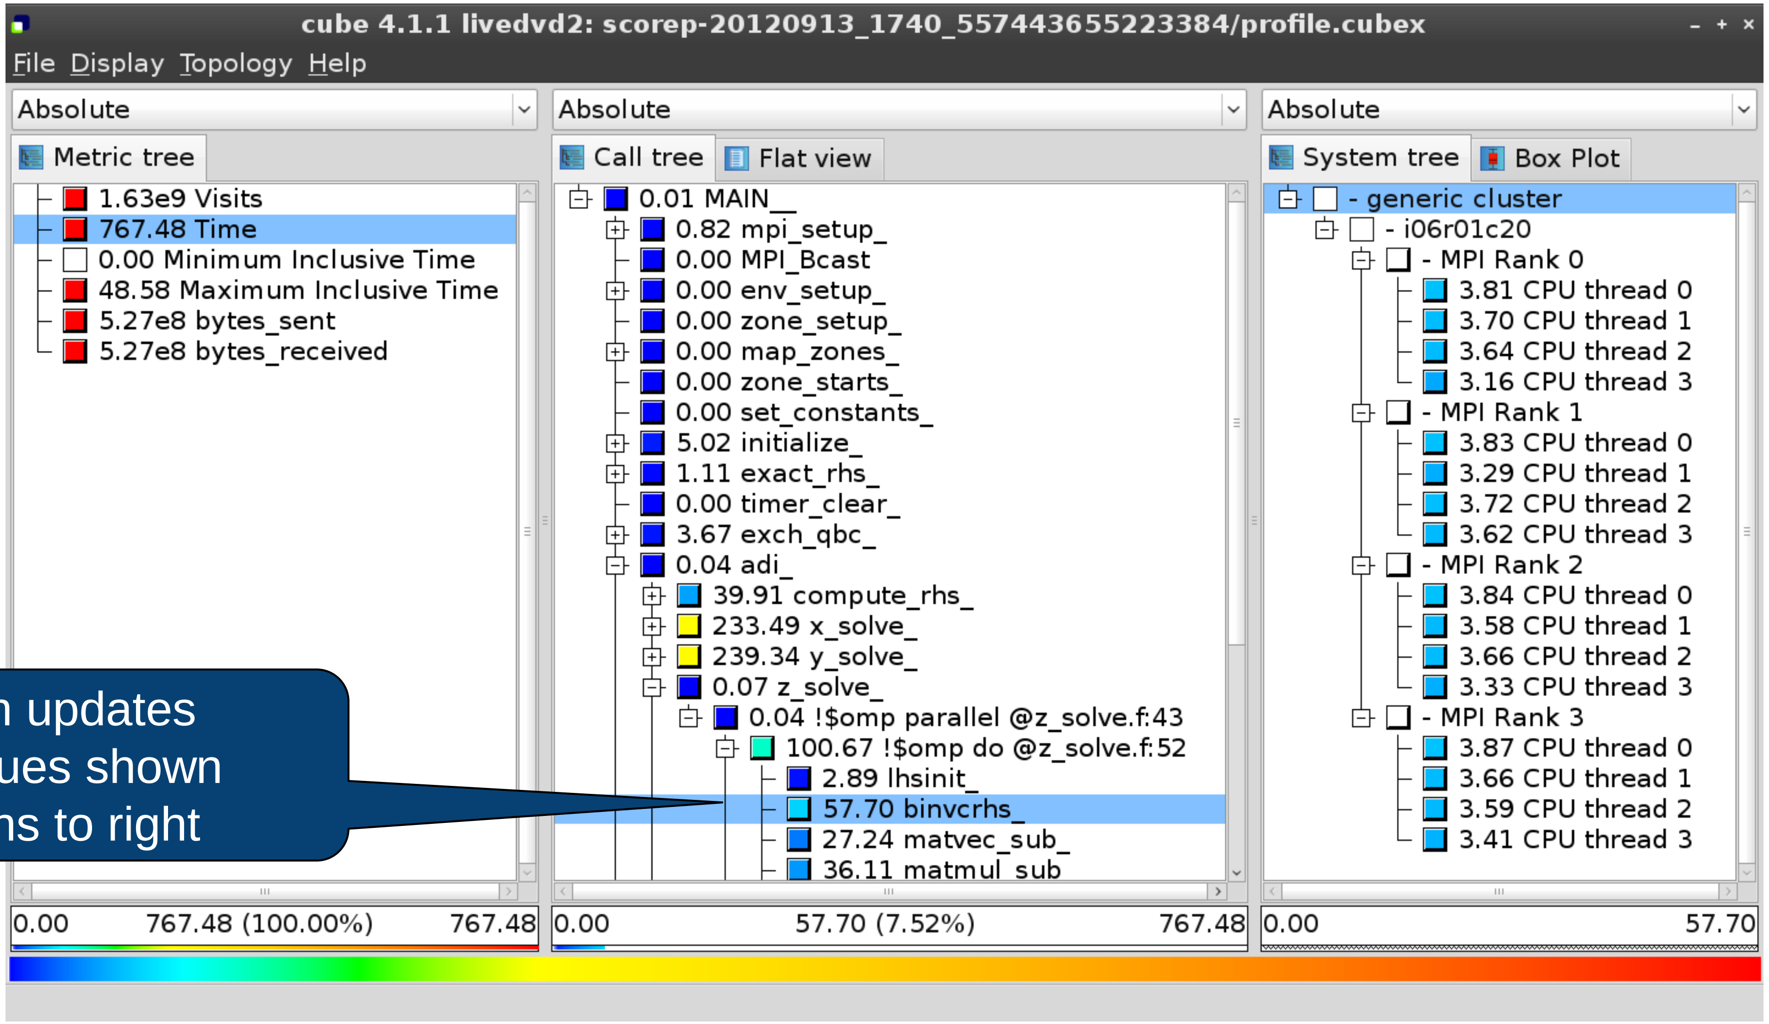
cube 4.1.1 livedvd2: scorep-20120913_1740_557443655223384/profile.cubex

File Display Topology Help

Absolute | Absolute | Absolute

Metric tree

```
├ 1.63e9 Visits
├ 767.48 Time
├ 0.00 Minimum Inclusive Time
├ 48.58 Maximum Inclusive Time
├ 5.27e8 bytes_sent
└ 5.27e8 bytes_received
```

Call tree | Flat view

```
⊟ 0.01 MAIN__
  ⊞ 0.82 mpi_setup_
    0.00 MPI_Bcast
  ⊞ 0.00 env_setup_
    0.00 zone_setup_
  ⊞ 0.00 map_zones_
    0.00 zone_starts_
    0.00 set_constants_
  ⊞ 5.02 initialize_
  ⊞ 1.11 exact_rhs_
    0.00 timer_clear_
  ⊞ 3.67 exch_qbc_
  ⊟ 0.04 adi_
    ⊞ 39.91 compute_rhs_
    ⊞ 233.49 x_solve_
    ⊞ 239.34 y_solve_
    ⊟ 0.07 z_solve_
      ⊟ 0.04 !$omp parallel @z_solve.f:43
        ⊟ 100.67 !$omp do @z_solve.f:52
            2.89 lhsinit_
            57.70 binvcrhs_
            27.24 matvec_sub_
            36.11 matmul_sub_
```

System tree | Box Plot

```
⊟ - generic cluster
  ⊟ - i06r01c20
    ⊟ - MPI Rank 0
        3.81 CPU thread 0
        3.70 CPU thread 1
        3.64 CPU thread 2
        3.16 CPU thread 3
    ⊟ - MPI Rank 1
        3.83 CPU thread 0
        3.29 CPU thread 1
        3.72 CPU thread 2
        3.62 CPU thread 3
    ⊟ - MPI Rank 2
        3.84 CPU thread 0
        3.58 CPU thread 1
        3.66 CPU thread 2
        3.33 CPU thread 3
    ⊟ - MPI Rank 3
        3.87 CPU thread 0
        3.66 CPU thread 1
        3.59 CPU thread 2
        3.41 CPU thread 3
```

0.00 767.48 (100.00%) 767.48 | 0.00 57.70 (7.52%) 767.48 | 0.00 57.70

Selection updates metric values shown in columns to right

Source-code view via context menu

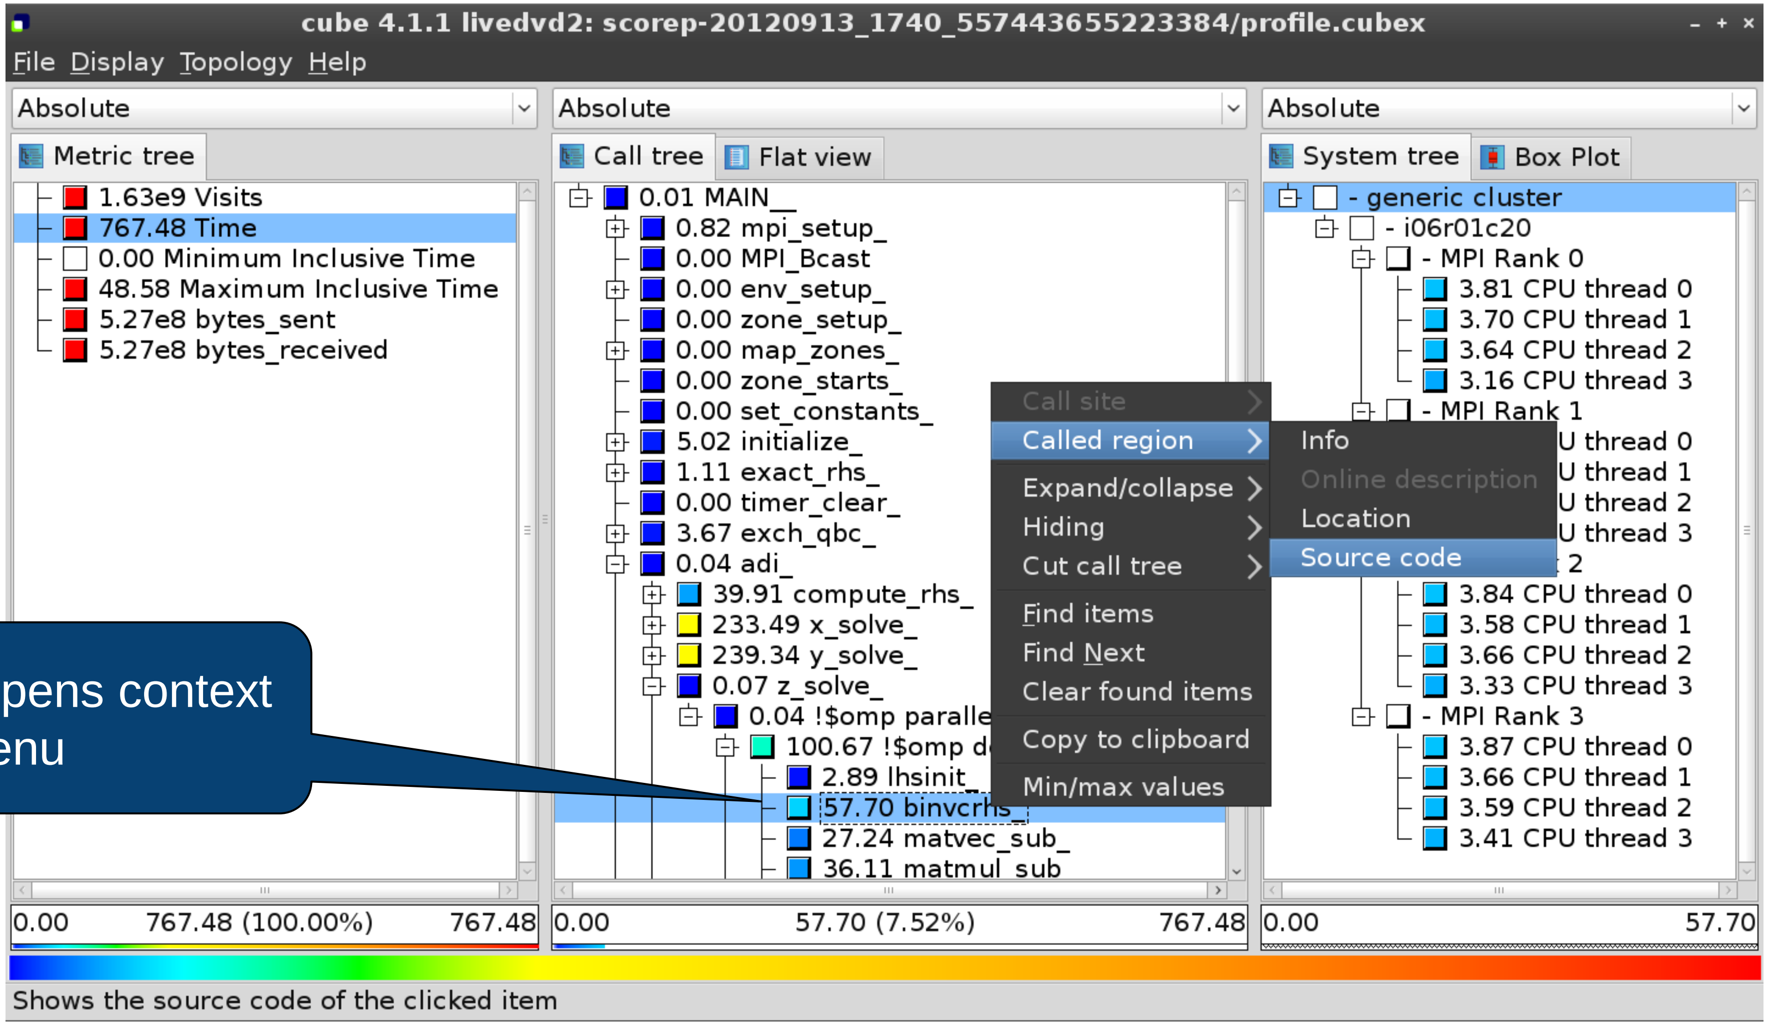

cube 4.1.1 livedvd2: scorep-20120913_1740_557443655223384/profile.cubex

```
Metric tree:
1.63e9 Visits
767.48 Time
0.00 Minimum Inclusive Time
48.58 Maximum Inclusive Time
5.27e8 bytes_sent
5.27e8 bytes_received

Call tree:
0.01 MAIN__
  0.82 mpi_setup_
  0.00 MPI_Bcast
  0.00 env_setup_
  0.00 zone_setup_
  0.00 map_zones_
  0.00 zone_starts_
  0.00 set_constants_
  5.02 initialize_
  1.11 exact_rhs_
  0.00 timer_clear_
  3.67 exch_qbc_
  0.04 adi_
    39.91 compute_rhs_
    233.49 x_solve_
    239.34 y_solve_
    0.07 z_solve_
      0.04 !$omp paralle
        100.67 !$omp d
          2.89 lhsinit_
          57.70 binvcrhs_
          27.24 matvec_sub_
          36.11 matmul_sub_

System tree:
- generic cluster
  - i06r01c20
    - MPI Rank 0
      3.81 CPU thread 0
      3.70 CPU thread 1
      3.64 CPU thread 2
      3.16 CPU thread 3
    - MPI Rank 1
      U thread 0
      U thread 1
      U thread 2
      U thread 3
    2
      3.84 CPU thread 0
      3.58 CPU thread 1
      3.66 CPU thread 2
      3.33 CPU thread 3
    - MPI Rank 3
      3.87 CPU thread 0
      3.66 CPU thread 1
      3.59 CPU thread 2
      3.41 CPU thread 3

0.00   767.48 (100.00%)   767.48
0.00   57.70 (7.52%)   767.48
0.00   57.70

Shows the source code of the clicked item
```

Right-click opens context menu

Source-code view

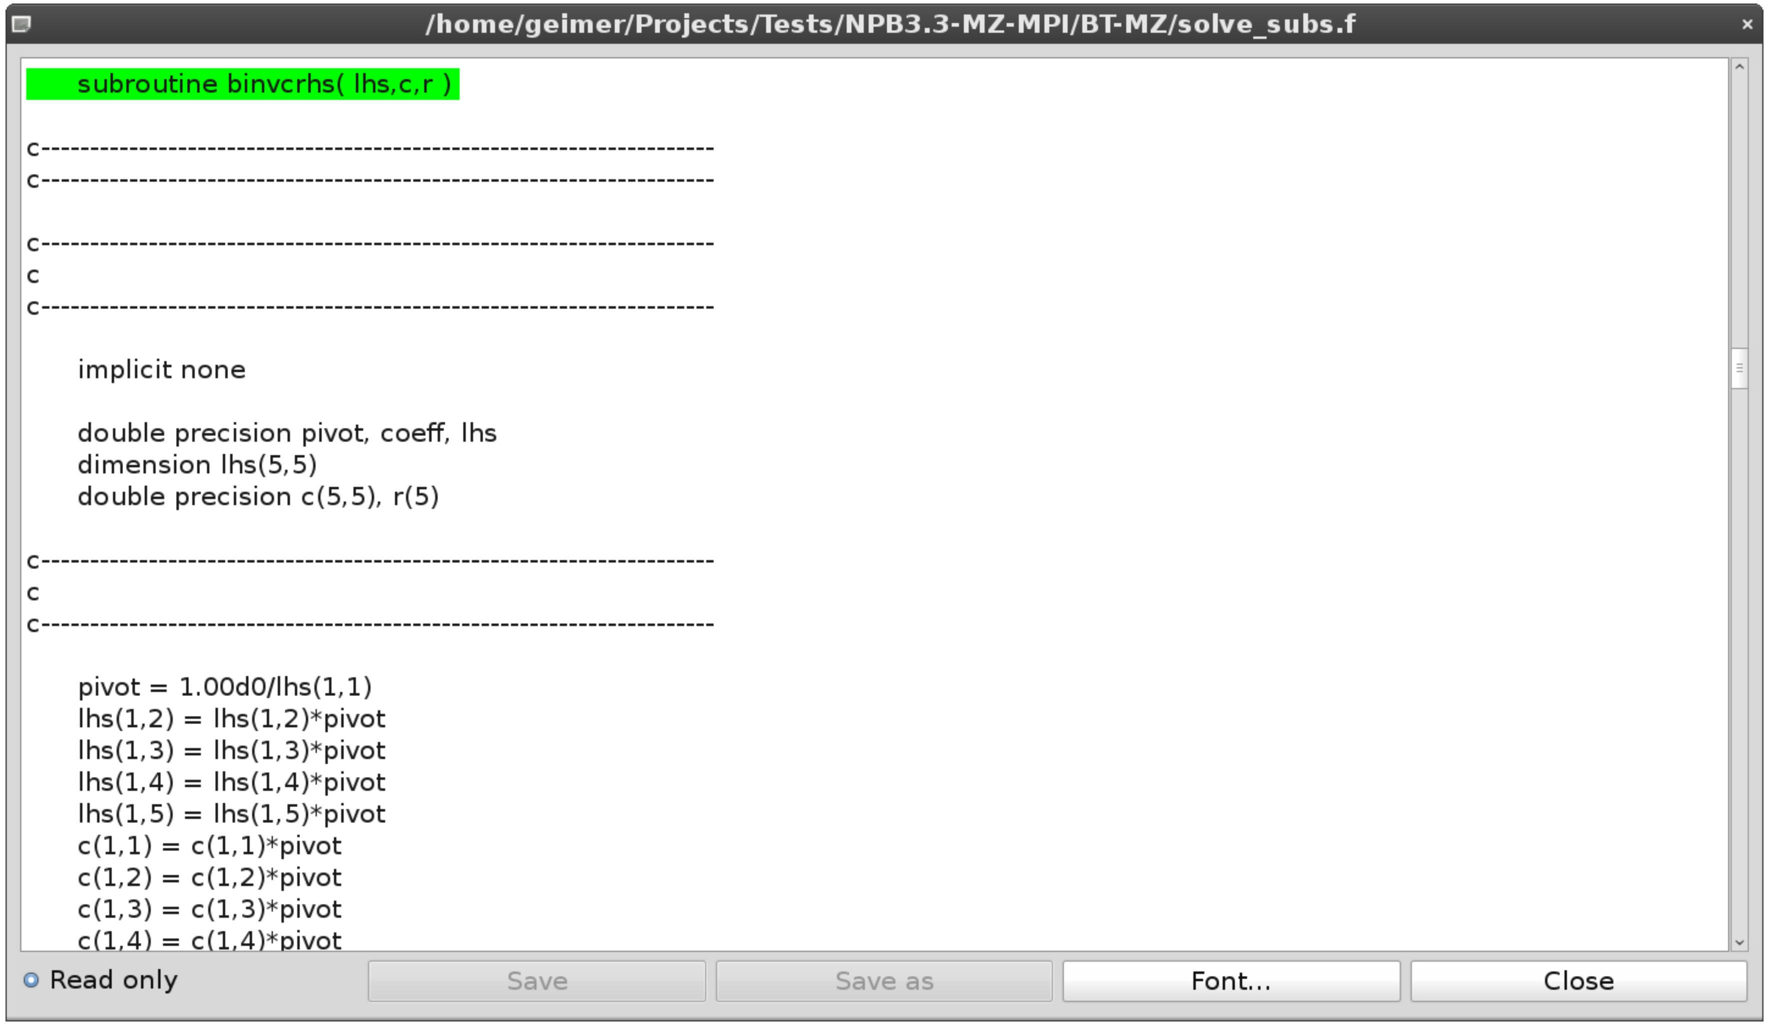

/home/geimer/Projects/Tests/NPB3.3-MZ-MPI/BT-MZ/solve_subs.f

```
      subroutine binvcrhs( 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
```

Flat profile view

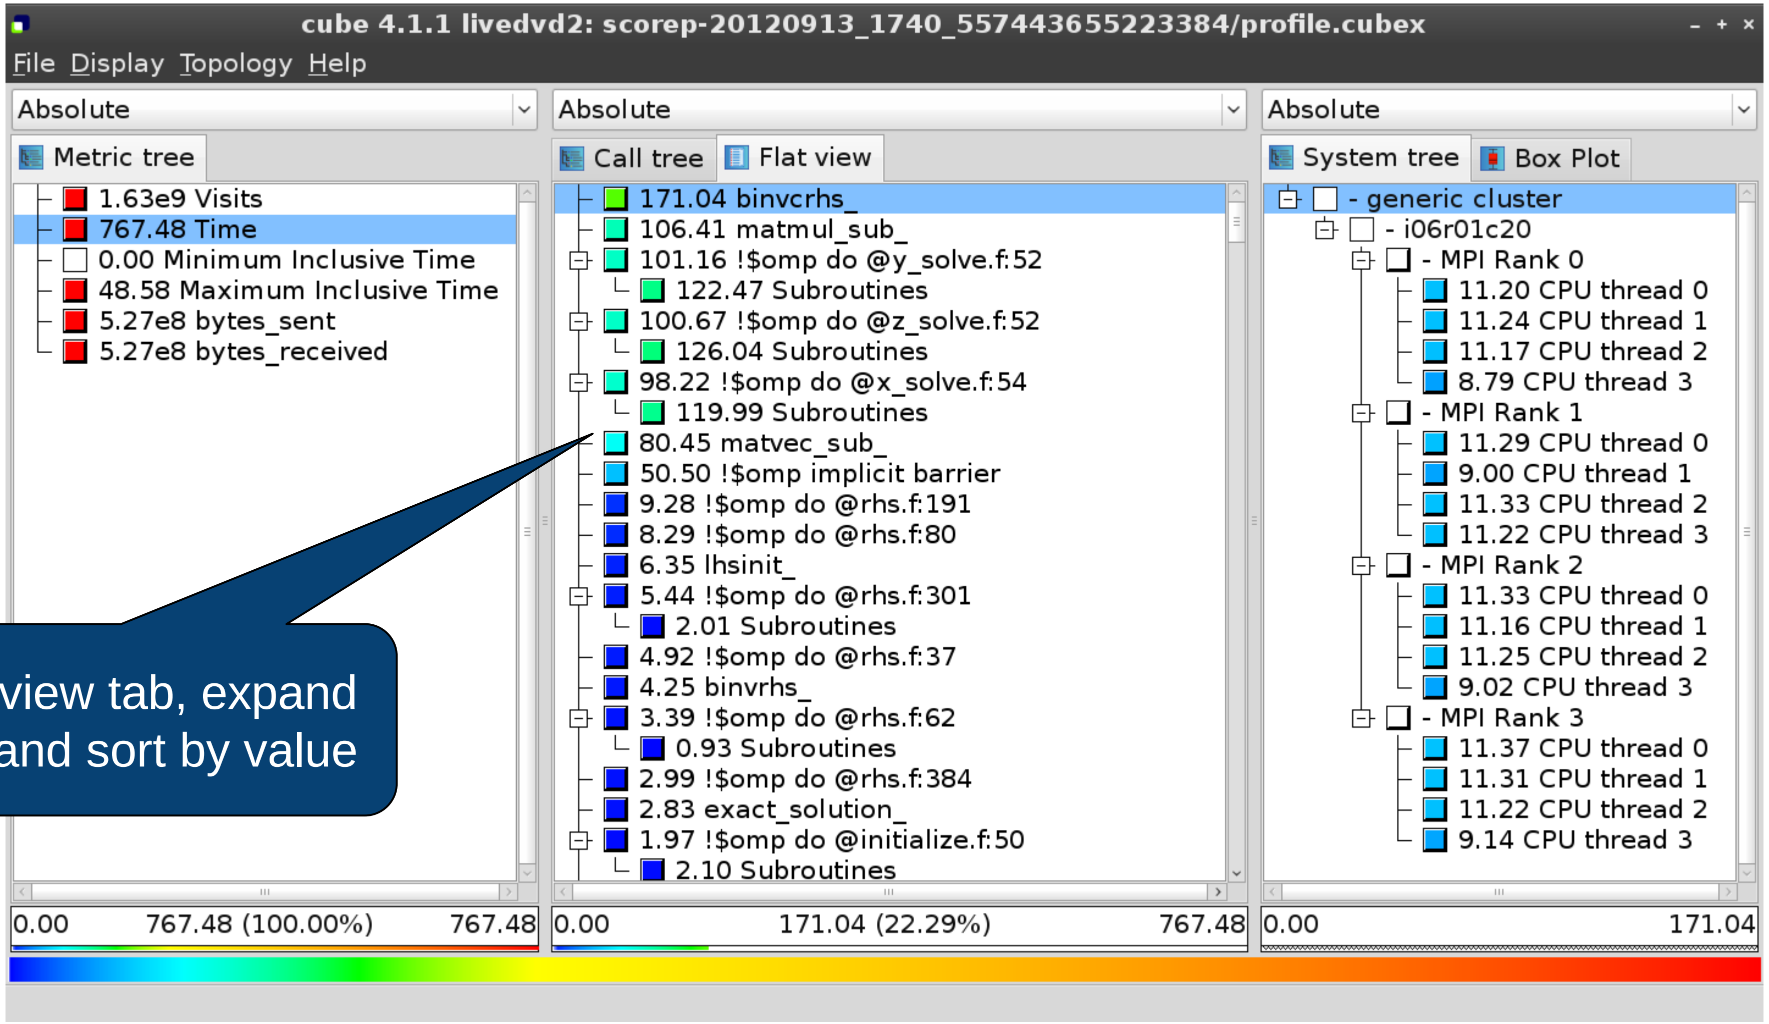
cube 4.1.1 livedvd2: scorep-20120913_1740_557443655223384/profile.cubex
File Display Topology Help
Absolute
Metric tree
Absolute
Call tree
Flat view
Absolute
System tree
Box Plot

```
Metric tree:
1.63e9 Visits
767.48 Time
0.00 Minimum Inclusive Time
48.58 Maximum Inclusive Time
5.27e8 bytes_sent
5.27e8 bytes_received

Flat view:
171.04 binvcrhs_
106.41 matmul_sub_
101.16 !$omp do @y_solve.f:52
  122.47 Subroutines
100.67 !$omp do @z_solve.f:52
  126.04 Subroutines
98.22 !$omp do @x_solve.f:54
  119.99 Subroutines
80.45 matvec_sub_
50.50 !$omp implicit barrier
9.28 !$omp do @rhs.f:191
8.29 !$omp do @rhs.f:80
6.35 lhsinit_
5.44 !$omp do @rhs.f:301
  2.01 Subroutines
4.92 !$omp do @rhs.f:37
4.25 binvrhs_
3.39 !$omp do @rhs.f:62
  0.93 Subroutines
2.99 !$omp do @rhs.f:384
2.83 exact_solution_
1.97 !$omp do @initialize.f:50
  2.10 Subroutines

System tree:
- generic cluster
  - i06r01c20
    - MPI Rank 0
      11.20 CPU thread 0
      11.24 CPU thread 1
      11.17 CPU thread 2
      8.79 CPU thread 3
    - MPI Rank 1
      11.29 CPU thread 0
      9.00 CPU thread 1
      11.33 CPU thread 2
      11.22 CPU thread 3
    - MPI Rank 2
      11.33 CPU thread 0
      11.16 CPU thread 1
      11.25 CPU thread 2
      9.02 CPU thread 3
    - MPI Rank 3
      11.37 CPU thread 0
      11.31 CPU thread 1
      11.22 CPU thread 2
      9.14 CPU thread 3

0.00    767.48 (100.00%)    767.48 | 0.00    171.04 (22.29%)    767.48 | 0.00    171.04
```

Select flat view tab, expand all nodes, and sort by value

Box plot view

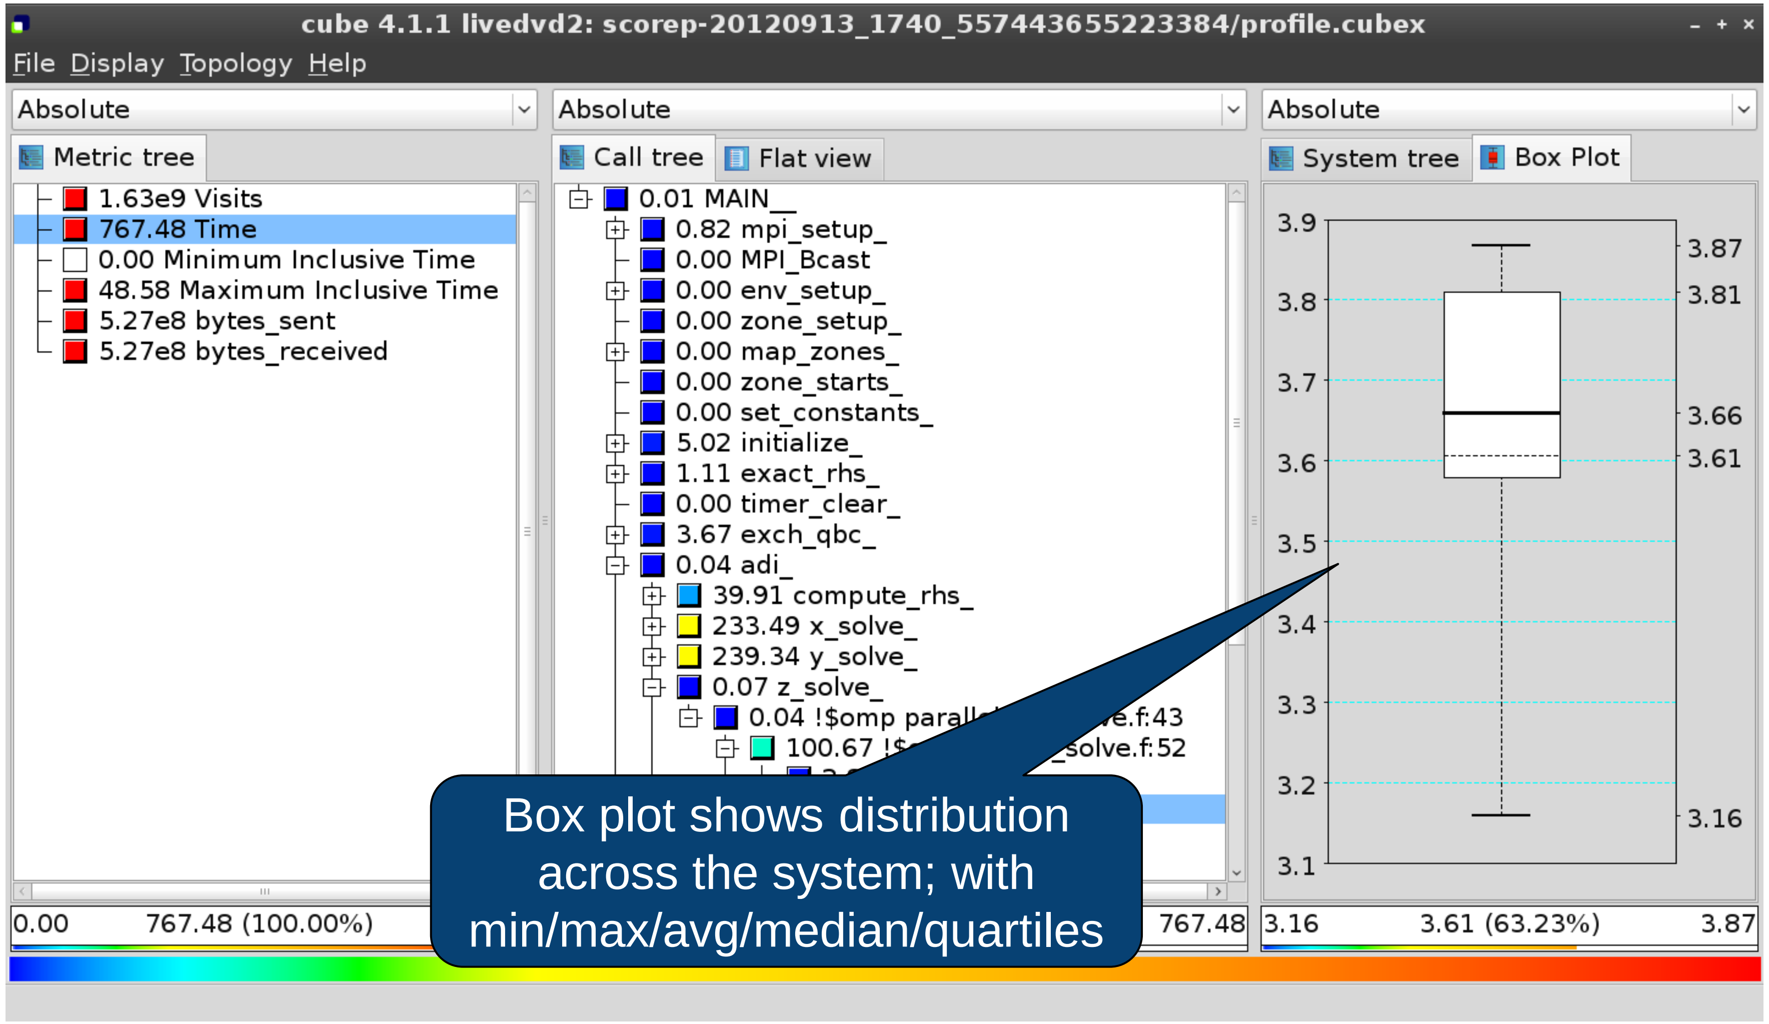

cube 4.1.1 livedvd2: scorep-20120913_1740_557443655223384/profile.cubex

```
1.63e9 Visits
767.48 Time
0.00 Minimum Inclusive Time
48.58 Maximum Inclusive Time
5.27e8 bytes_sent
5.27e8 bytes_received
```

```
0.01 MAIN__
  0.82 mpi_setup_
  0.00 MPI_Bcast
  0.00 env_setup_
  0.00 zone_setup_
  0.00 map_zones_
  0.00 zone_starts_
  0.00 set_constants_
  5.02 initialize_
  1.11 exact_rhs_
  0.00 timer_clear_
  3.67 exch_qbc_
  0.04 adi_
    39.91 compute_rhs_
    233.49 x_solve_
    239.34 y_solve_
    0.07 z_solve_
      0.04 !$omp parall…e.f:43
        100.67 !$…_solve.f:52
```

```
0.00      767.48 (100.00%)        767.48   3.16      3.61 (63.23%)      3.87
```

Box plot shows distribution across the system; with min/max/avg/median/quartiles

Alternative display modes

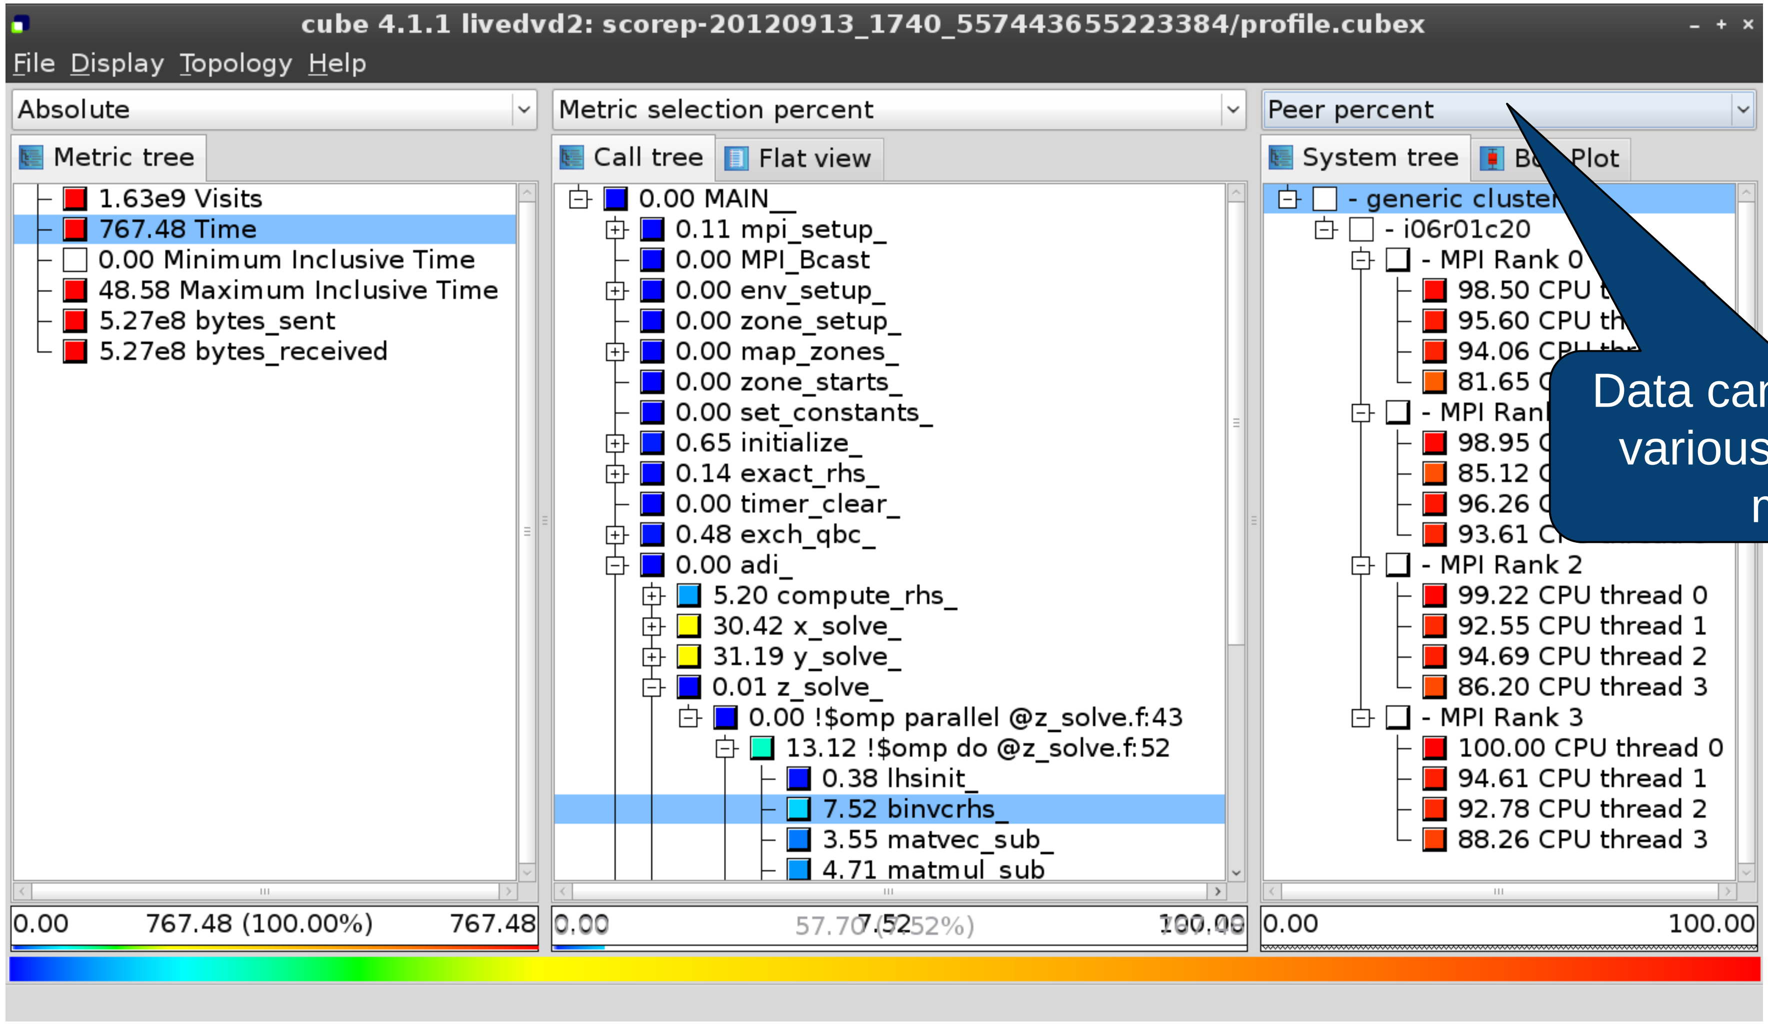

cube 4.1.1 livedvd2: scorep-20120913_1740_557443655223384/profile.cubex

```
1.63e9 Visits
767.48 Time
0.00 Minimum Inclusive Time
48.58 Maximum Inclusive Time
5.27e8 bytes_sent
5.27e8 bytes_received
```

```
0.00 MAIN__
  0.11 mpi_setup_
  0.00 MPI_Bcast
  0.00 env_setup_
  0.00 zone_setup_
  0.00 map_zones_
  0.00 zone_starts_
  0.00 set_constants_
  0.65 initialize_
  0.14 exact_rhs_
  0.00 timer_clear_
  0.48 exch_qbc_
  0.00 adi_
    5.20 compute_rhs_
    30.42 x_solve_
    31.19 y_solve_
    0.01 z_solve_
      0.00 !$omp parallel @z_solve.f:43
        13.12 !$omp do @z_solve.f:52
          0.38 lhsinit_
          7.52 binvcrhs_
          3.55 matvec_sub_
          4.71 matmul_sub
```

```
- generic cluster
  - i06r01c20
    - MPI Rank 0
      98.50 CPU t
      95.60 CPU th
      94.06 CPU thr
      81.65 C
    - MPI Rank
      98.95 C
      85.12 C
      96.26 C
      93.61 C
    - MPI Rank 2
      99.22 CPU thread 0
      92.55 CPU thread 1
      94.69 CPU thread 2
      86.20 CPU thread 3
    - MPI Rank 3
      100.00 CPU thread 0
      94.61 CPU thread 1
      92.78 CPU thread 2
      88.26 CPU thread 3
```

```
0.00    767.48 (100.00%)    767.48
0.00    7.52 (57.70 / 7.52%)    100.00
0.00    100.00
```

Data can be shown in various percentage modes

Important display modes

- 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

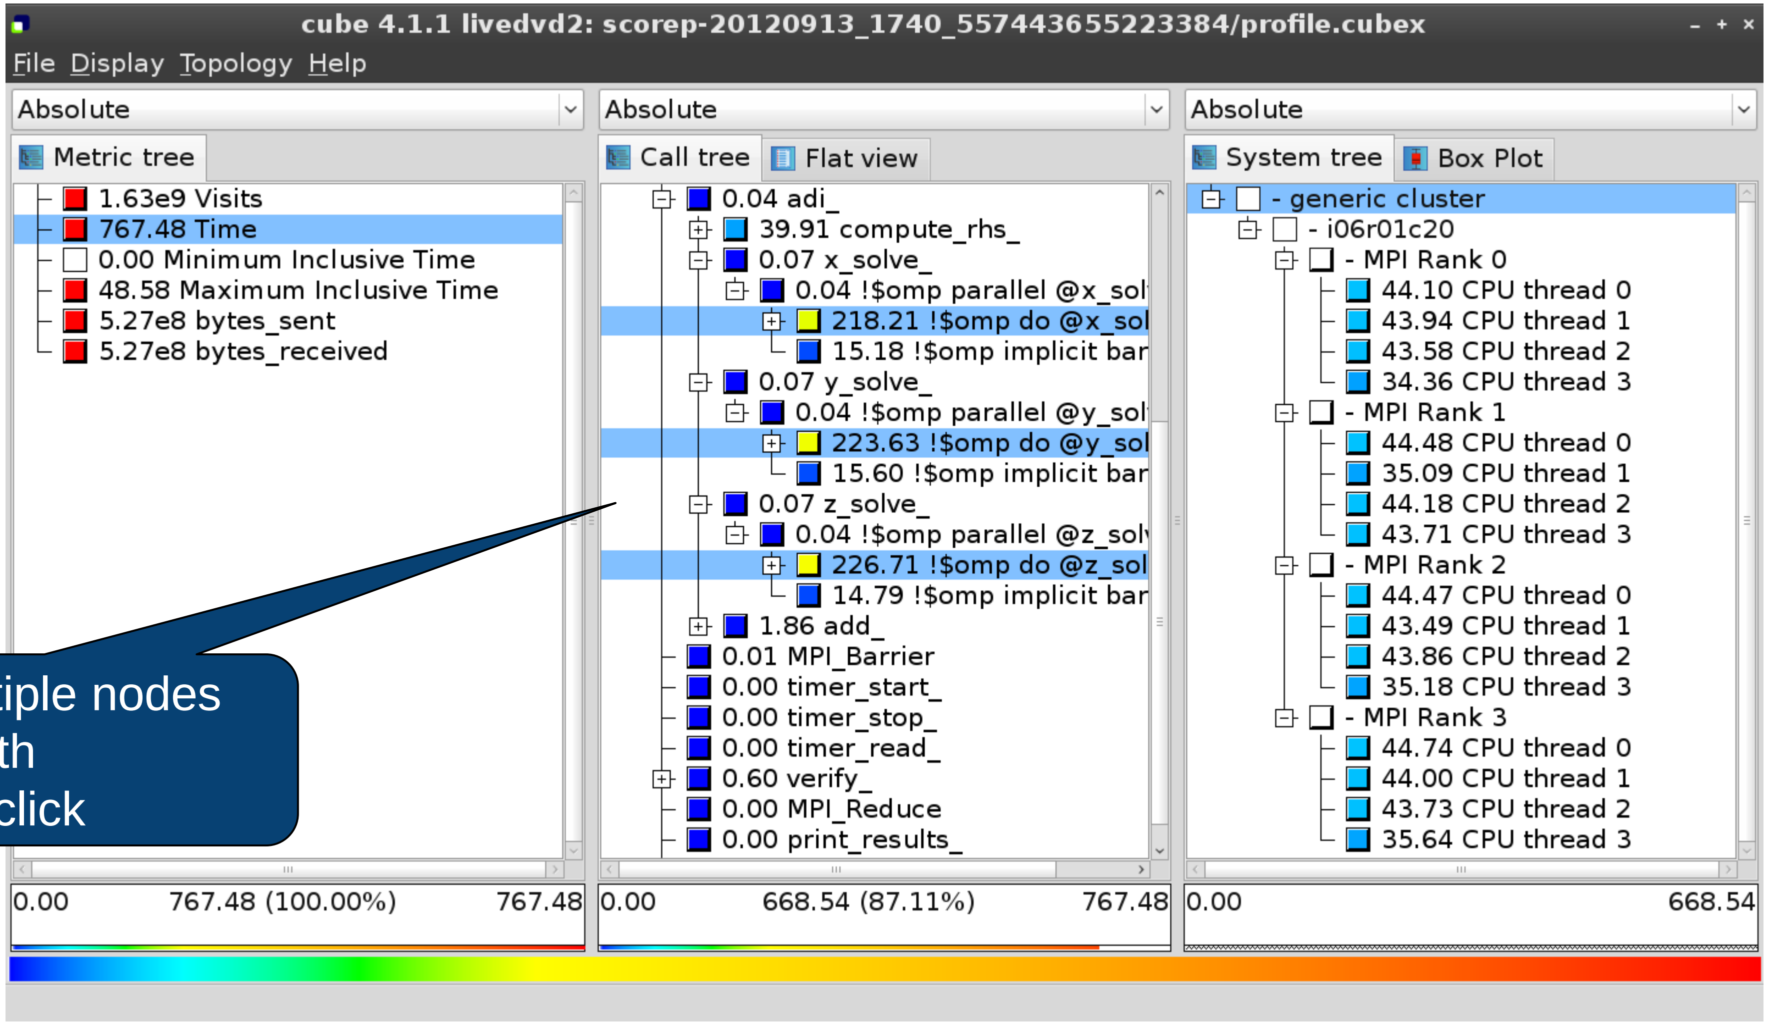

cube 4.1.1 livedvd2: scorep-20120913_1740_557443655223384/profile.cubex

File Display Topology Help

Absolute | Absolute | Absolute

Metric tree

```
├ 1.63e9 Visits
├ 767.48 Time
├ 0.00 Minimum Inclusive Time
├ 48.58 Maximum Inclusive Time
├ 5.27e8 bytes_sent
└ 5.27e8 bytes_received
```

0.00 767.48 (100.00%) 767.48

Call tree | Flat view

```
⊟ 0.04 adi_
   ⊞ 39.91 compute_rhs_
   ⊟ 0.07 x_solve_
      ⊟ 0.04 !$omp parallel @x_sol
         ⊞ 218.21 !$omp do @x_sol
         └ 15.18 !$omp implicit bar
   ⊟ 0.07 y_solve_
      ⊟ 0.04 !$omp parallel @y_sol
         ⊞ 223.63 !$omp do @y_sol
         └ 15.60 !$omp implicit bar
   ⊟ 0.07 z_solve_
      ⊟ 0.04 !$omp parallel @z_sol
         ⊞ 226.71 !$omp do @z_sol
         └ 14.79 !$omp implicit bar
   ⊞ 1.86 add_
├ 0.01 MPI_Barrier
├ 0.00 timer_start_
├ 0.00 timer_stop_
├ 0.00 timer_read_
⊞ 0.60 verify_
├ 0.00 MPI_Reduce
├ 0.00 print_results_
```

0.00 668.54 (87.11%) 767.48

System tree | Box Plot

```
⊟ - generic cluster
   ⊟ - i06r01c20
      ⊟ - MPI Rank 0
         ├ 44.10 CPU thread 0
         ├ 43.94 CPU thread 1
         ├ 43.58 CPU thread 2
         └ 34.36 CPU thread 3
      ⊟ - MPI Rank 1
         ├ 44.48 CPU thread 0
         ├ 35.09 CPU thread 1
         ├ 44.18 CPU thread 2
         └ 43.71 CPU thread 3
      ⊟ - MPI Rank 2
         ├ 44.47 CPU thread 0
         ├ 43.49 CPU thread 1
         ├ 43.86 CPU thread 2
         └ 35.18 CPU thread 3
      ⊟ - MPI Rank 3
         ├ 44.74 CPU thread 0
         ├ 44.00 CPU thread 1
         ├ 43.73 CPU thread 2
         └ 35.64 CPU thread 3
```

0.00 668.54

Select multiple nodes with Ctrl-click

Derived metrics in Cube

- Value of the derived metric is not stored, but **calculated** on-the-fly
- One defines an CubePL expression, e.g.:

  metric::time(i) / metric::visits(e)
- Types of derived metrics:
  - **Prederived**: 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()

Derived metrics in Cube GUI

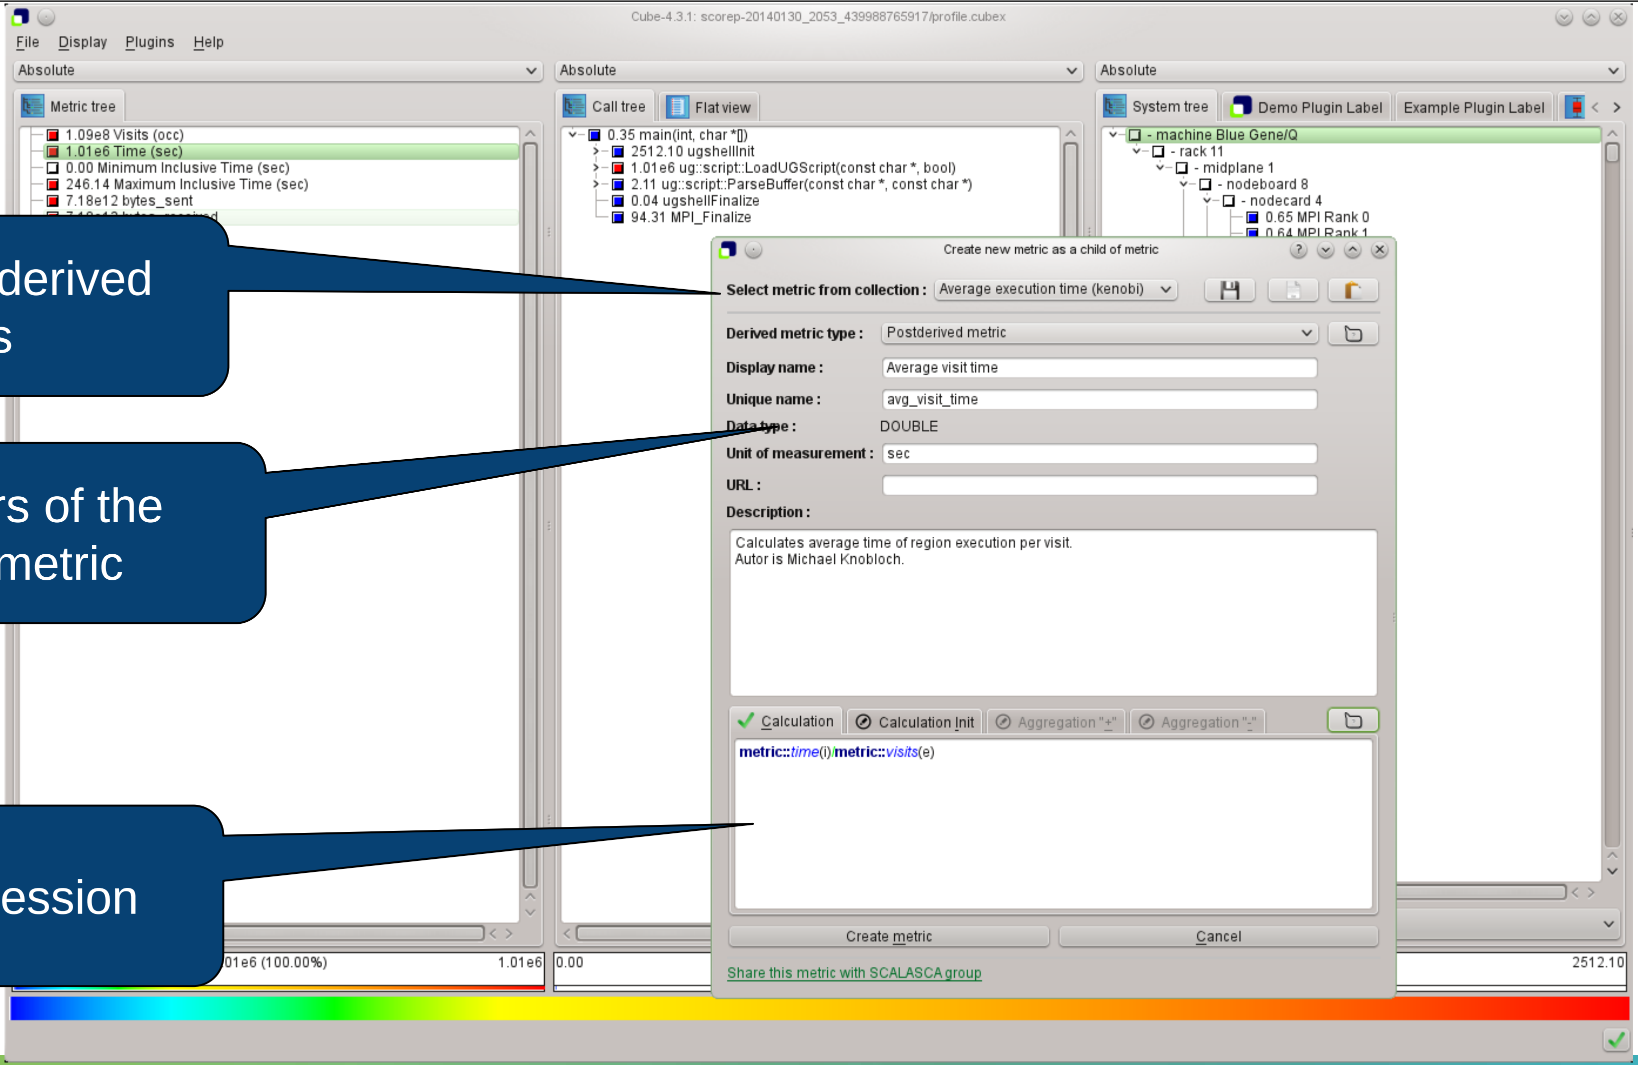

1.09e8 Visits (occ)
1.01e6 Time (sec)
0.00 Minimum Inclusive Time (sec)
246.14 Maximum Inclusive Time (sec)
7.18e12 bytes_sent

0.35 main(int, char *[])
2512.10 ugshellInit
1.01e6 ug::script::LoadUGScript(const char *, bool)
2.11 ug::script::ParseBuffer(const char *, const char *)
0.04 ugshellFinalize
94.31 MPI_Finalize

- machine Blue Gene/Q
- rack 11
- midplane 1
- nodeboard 8
- nodecard 4
0.65 MPI Rank 0

Create new metric as a child of metric

Select metric from collection :	Average execution time (kenobi)
Derived metric type :	Postderived metric
Display name :	Average visit time
Unique name :	avg_visit_time
Data type :	DOUBLE
Unit of measurement :	sec
URL :	
Description :	

Calculates average time of region execution per visit.
Autor is Michael Knobloch.

```
metric::time(i)/metric::visits(e)
```

Collection of derived metrics

Parameters of the derived metric

CubePL expression

Example derived metric FLOPS based on PAPI_FP_OPS and time

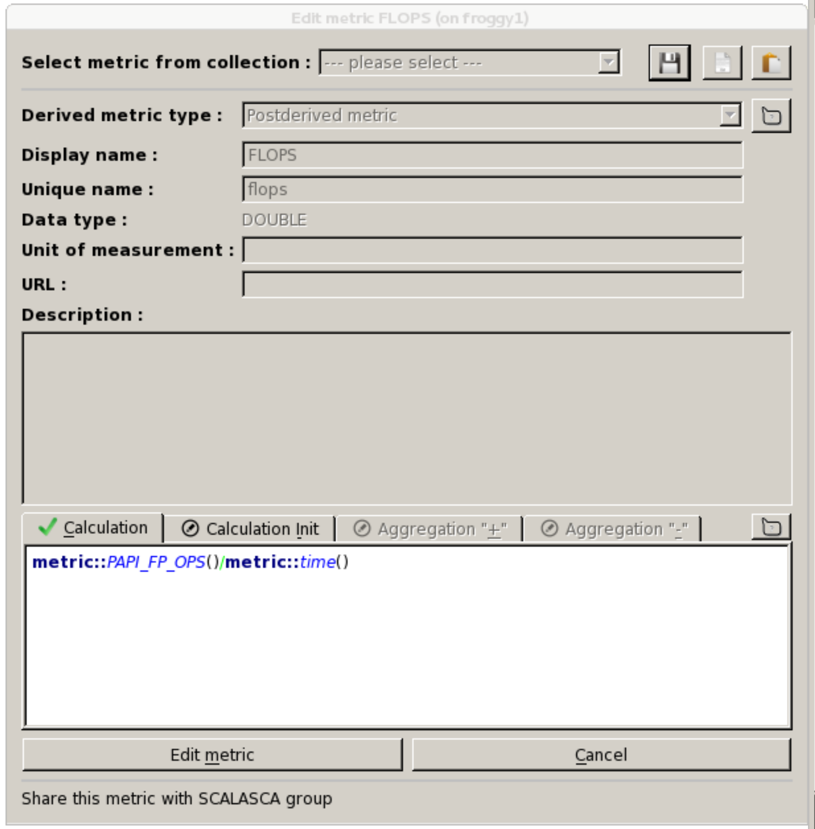
Edit metric FLOPS (on froggy1)

Select metric from collection : --- please select ---

Derived metric type : Postderived metric

Display name : FLOPS

Unique name : flops

Data type : DOUBLE

Unit of measurement :

URL :

Description :

Calculation | Calculation Init | Aggregation "+" | Aggregation "-"

```
metric::PAPI_FP_OPS()/metric::time()
```

Edit metric | Cancel

Share this metric with SCALASCA group

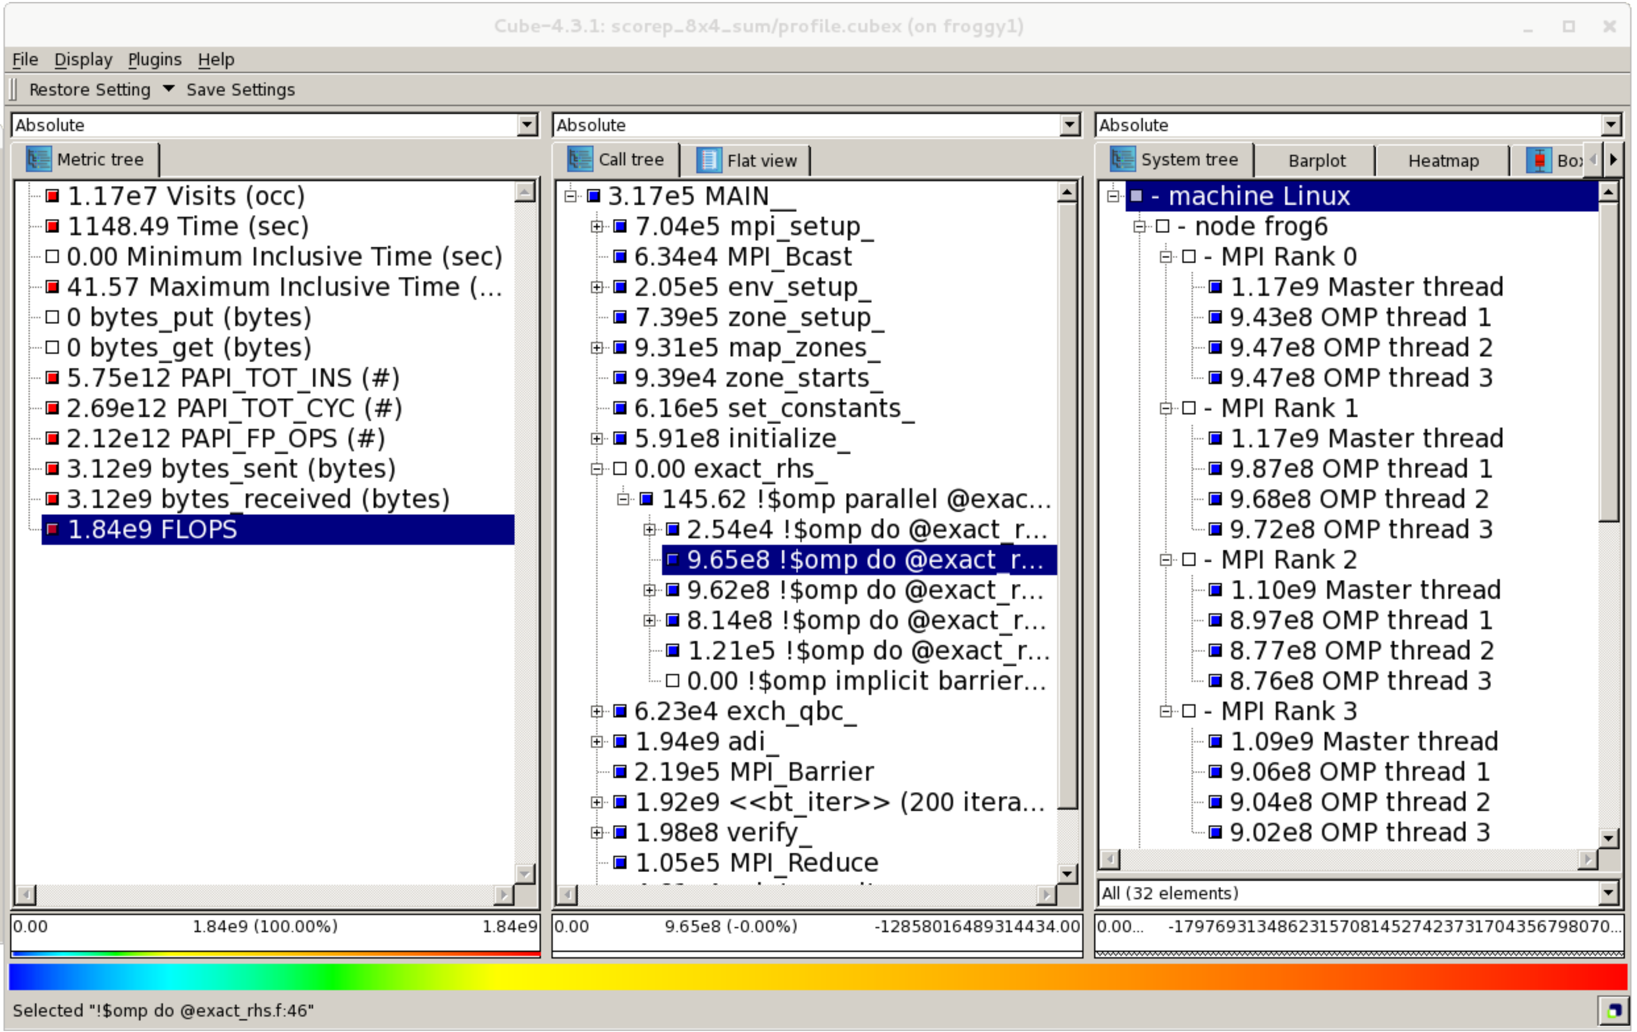
Cube-4.3.1: scorep_8x4_sum/profile.cubex (on froggy1)

File Display Plugins Help

Restore Setting ▾ Save Settings

Absolute | Absolute | Absolute

Metric tree:
- 1.17e7 Visits (occ)
- 1148.49 Time (sec)
- 0.00 Minimum Inclusive Time (sec)
- 41.57 Maximum Inclusive Time (...
- 0 bytes_put (bytes)
- 0 bytes_get (bytes)
- 5.75e12 PAPI_TOT_INS (#)
- 2.69e12 PAPI_TOT_CYC (#)
- 2.12e12 PAPI_FP_OPS (#)
- 3.12e9 bytes_sent (bytes)
- 3.12e9 bytes_received (bytes)
- 1.84e9 FLOPS

Call tree | Flat view:
- 3.17e5 MAIN__
  - 7.04e5 mpi_setup_
  - 6.34e4 MPI_Bcast
  - 2.05e5 env_setup_
  - 7.39e5 zone_setup_
  - 9.31e5 map_zones_
  - 9.39e4 zone_starts_
  - 6.16e5 set_constants_
  - 5.91e8 initialize_
  - 0.00 exact_rhs_
    - 145.62 !$omp parallel @exac...
      - 2.54e4 !$omp do @exact_r...
      - 9.65e8 !$omp do @exact_r...
      - 9.62e8 !$omp do @exact_r...
      - 8.14e8 !$omp do @exact_r...
      - 1.21e5 !$omp do @exact_r...
      - 0.00 !$omp implicit barrier...
  - 6.23e4 exch_qbc_
  - 1.94e9 adi_
  - 2.19e5 MPI_Barrier
  - 1.92e9 <<bt_iter>> (200 itera...
  - 1.98e8 verify_
  - 1.05e5 MPI_Reduce

System tree | Barplot | Heatmap | Box
- - machine Linux
  - - node frog6
    - - MPI Rank 0
      - 1.17e9 Master thread
      - 9.43e8 OMP thread 1
      - 9.47e8 OMP thread 2
      - 9.47e8 OMP thread 3
    - - MPI Rank 1
      - 1.17e9 Master thread
      - 9.87e8 OMP thread 1
      - 9.68e8 OMP thread 2
      - 9.72e8 OMP thread 3
    - - MPI Rank 2
      - 1.10e9 Master thread
      - 8.97e8 OMP thread 1
      - 8.77e8 OMP thread 2
      - 8.76e8 OMP thread 3
    - - MPI Rank 3
      - 1.09e9 Master thread
      - 9.06e8 OMP thread 1
      - 9.04e8 OMP thread 2
      - 9.02e8 OMP thread 3

All (32 elements)

0.00 1.84e9 (100.00%) 1.84e9 | 0.00 9.65e8 (-0.00%) -12858016489314434.00 | 0.00... -179769313486231570814527423731704356798070...

Selected "!$omp do @exact_rhs.f:46"

Context-sensitive help

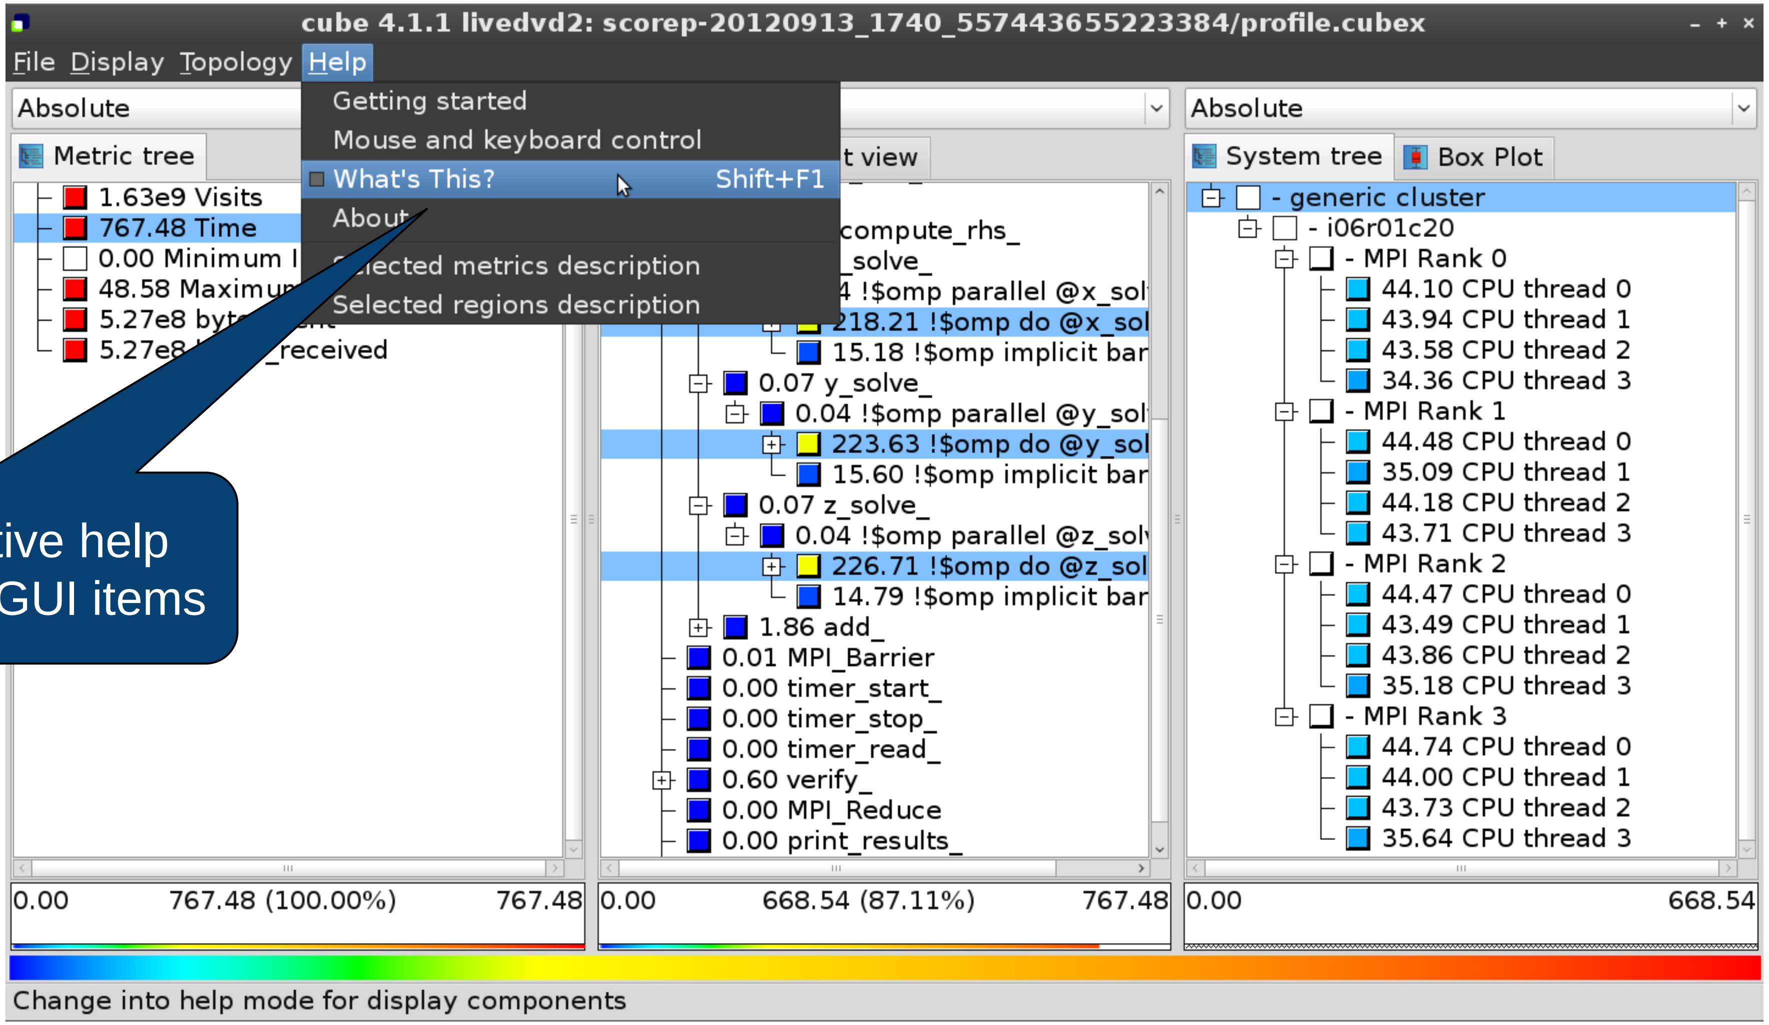

cube 4.1.1 livedvd2: scorep-20120913_1740_557443655223384/profile.cubex

File Display Topology Help
Getting started
Mouse and keyboard control
What's This? Shift+F1
About
Selected metrics description
Selected regions description
Absolute
Metric tree
System tree
Box Plot

```
Metric tree:
1.63e9 Visits
767.48 Time
0.00 Minimum I
48.58 Maximum
5.27e8 byt
5.27e8 _received

Call tree:
compute_rhs_
solve_
4 !$omp parallel @x_sol
218.21 !$omp do @x_sol
15.18 !$omp implicit bar
0.07 y_solve_
0.04 !$omp parallel @y_sol
223.63 !$omp do @y_sol
15.60 !$omp implicit bar
0.07 z_solve_
0.04 !$omp parallel @z_sol
226.71 !$omp do @z_sol
14.79 !$omp implicit bar
1.86 add_
0.01 MPI_Barrier
0.00 timer_start_
0.00 timer_stop_
0.00 timer_read_
0.60 verify_
0.00 MPI_Reduce
0.00 print_results_

System tree:
- generic cluster
  - i06r01c20
    - MPI Rank 0
      44.10 CPU thread 0
      43.94 CPU thread 1
      43.58 CPU thread 2
      34.36 CPU thread 3
    - MPI Rank 1
      44.48 CPU thread 0
      35.09 CPU thread 1
      44.18 CPU thread 2
      43.71 CPU thread 3
    - MPI Rank 2
      44.47 CPU thread 0
      43.49 CPU thread 1
      43.86 CPU thread 2
      35.18 CPU thread 3
    - MPI Rank 3
      44.74 CPU thread 0
      44.00 CPU thread 1
      43.73 CPU thread 2
      35.64 CPU thread 3

0.00   767.48 (100.00%)   767.48
0.00   668.54 (87.11%)   767.48
0.00   668.54

Change into help mode for display components
```

Context-sensitive help available for all GUI items

CUBE algebra utilities

- Extracting solver sub-tree from analysis report

```
% cube_cut  -r '<<ITERATION>>'  scorep_bt-mz_B_8x8_sum/profile.cubex
Writing cut.cubex... done.
```

- Calculating difference of two reports

```
% cube_diff  scorep_bt-mz_B_8x8_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

Loop Unrolling

- Show time dependent behavior by unrolling iterations

- Preparations:
  - Mark loops by using Score-P user instrumentation in your source code

```
SCOREP_USER_REGION_BEGIN( scorep_bt_loop, "<<bt_iter>>", SCOREP_USER_REGION_TYPE_DYNAMIC )
```

- Result in the CUBE profile:
  - Iterations shown as separate call trees
  - ➢Useful for checking results for specific iterations

or

  - Select your user instrumented region and mark it as loop
  - Choose hide iterations
  - ➢View the Barplot statistics or the (thread x iterations) Heatmap

Loop Unrolling - Barplot

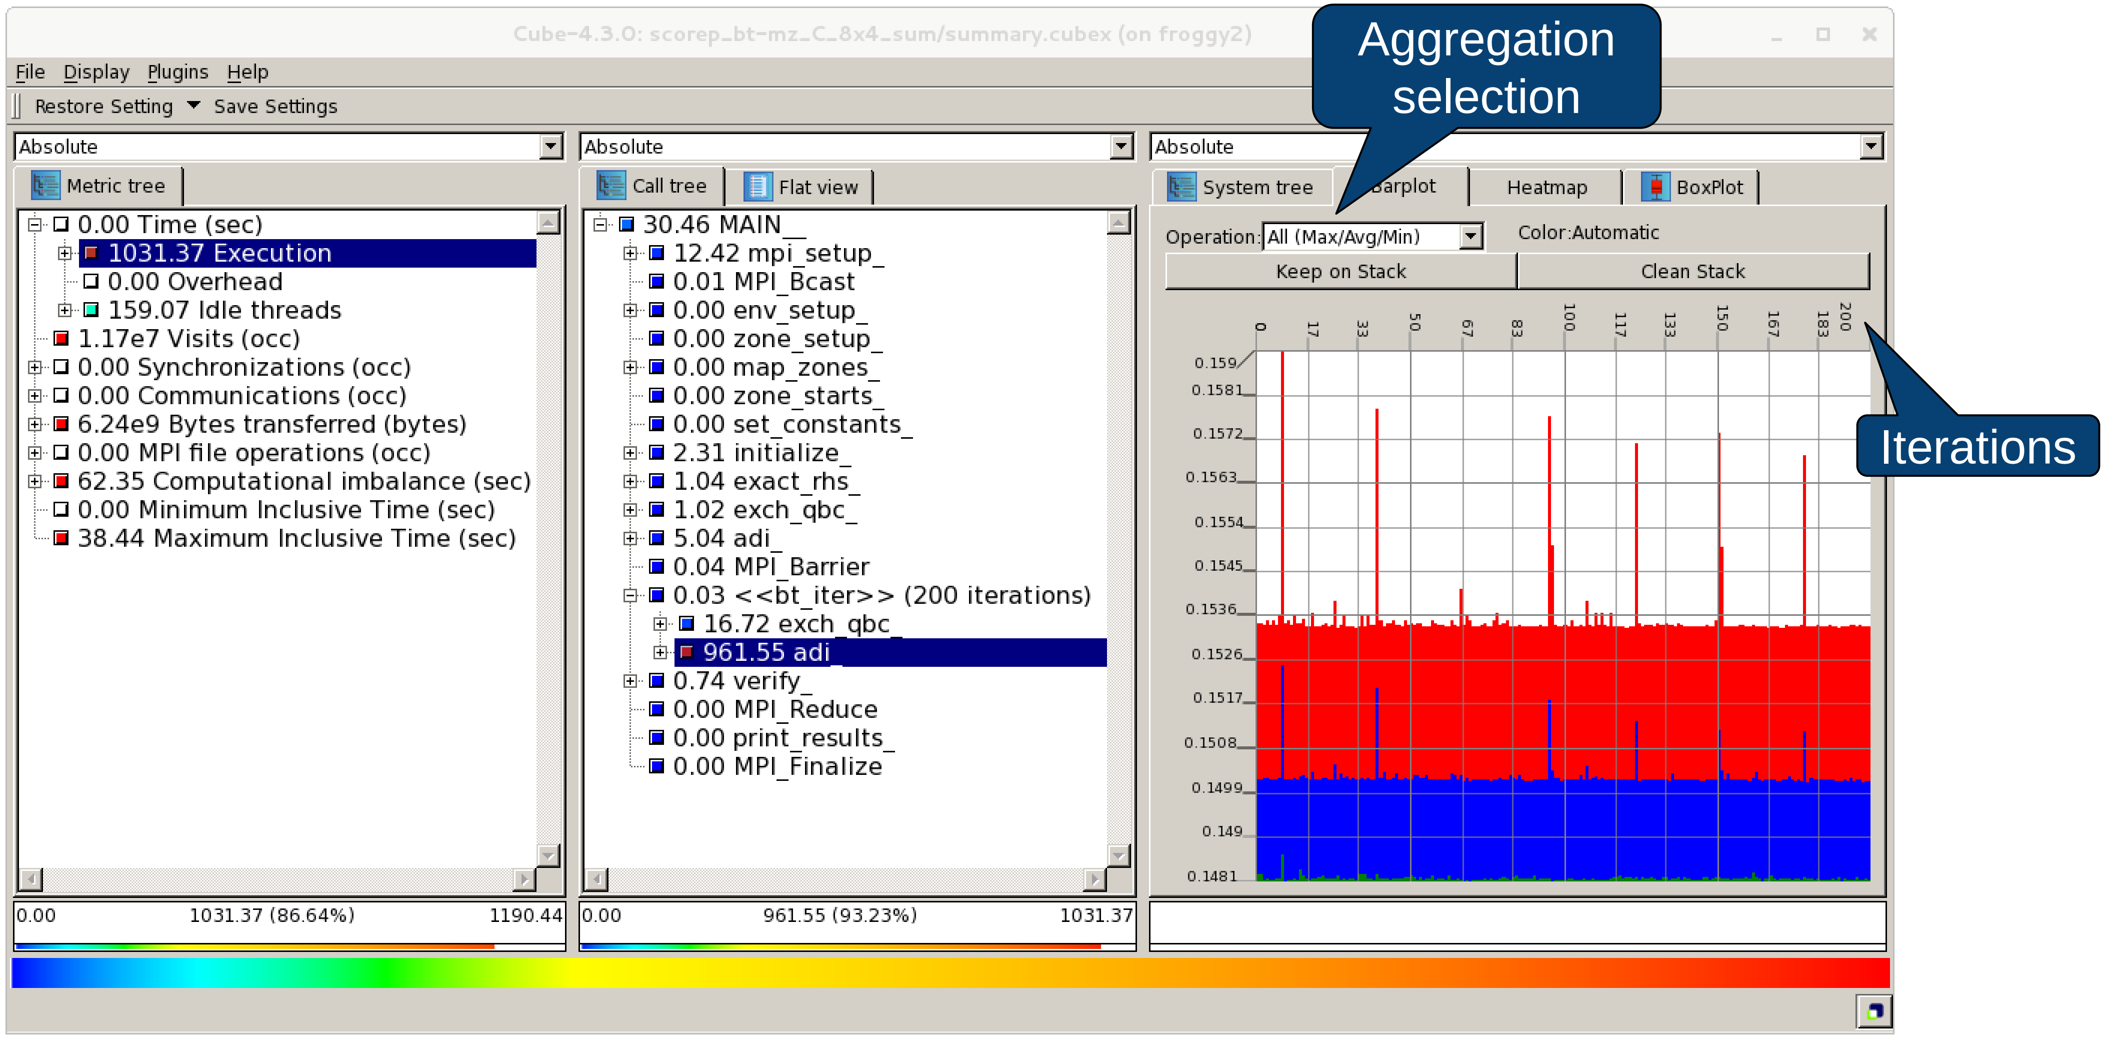
Cube-4.3.0: scorep_bt-mz_C_8x4_sum/summary.cubex (on froggy2)

File Display Plugins Help

Restore Setting Save Settings

Absolute | Absolute | Absolute

Metric tree:
```
0.00 Time (sec)
  1031.37 Execution
  0.00 Overhead
  159.07 Idle threads
1.17e7 Visits (occ)
0.00 Synchronizations (occ)
0.00 Communications (occ)
6.24e9 Bytes transferred (bytes)
0.00 MPI file operations (occ)
62.35 Computational imbalance (sec)
0.00 Minimum Inclusive Time (sec)
38.44 Maximum Inclusive Time (sec)
```

Call tree:
```
30.46 MAIN__
  12.42 mpi_setup_
  0.01 MPI_Bcast
  0.00 env_setup_
  0.00 zone_setup_
  0.00 map_zones_
  0.00 zone_starts_
  0.00 set_constants_
  2.31 initialize_
  1.04 exact_rhs_
  1.02 exch_qbc_
  5.04 adi_
  0.04 MPI_Barrier
  0.03 <<bt_iter>> (200 iterations)
    16.72 exch_qbc_
    961.55 adi_
  0.74 verify_
  0.00 MPI_Reduce
  0.00 print_results_
  0.00 MPI_Finalize
```

0.00 1031.37 (86.64%) 1190.44

0.00 961.55 (93.23%) 1031.37

Loop Unrolling – Heatmap

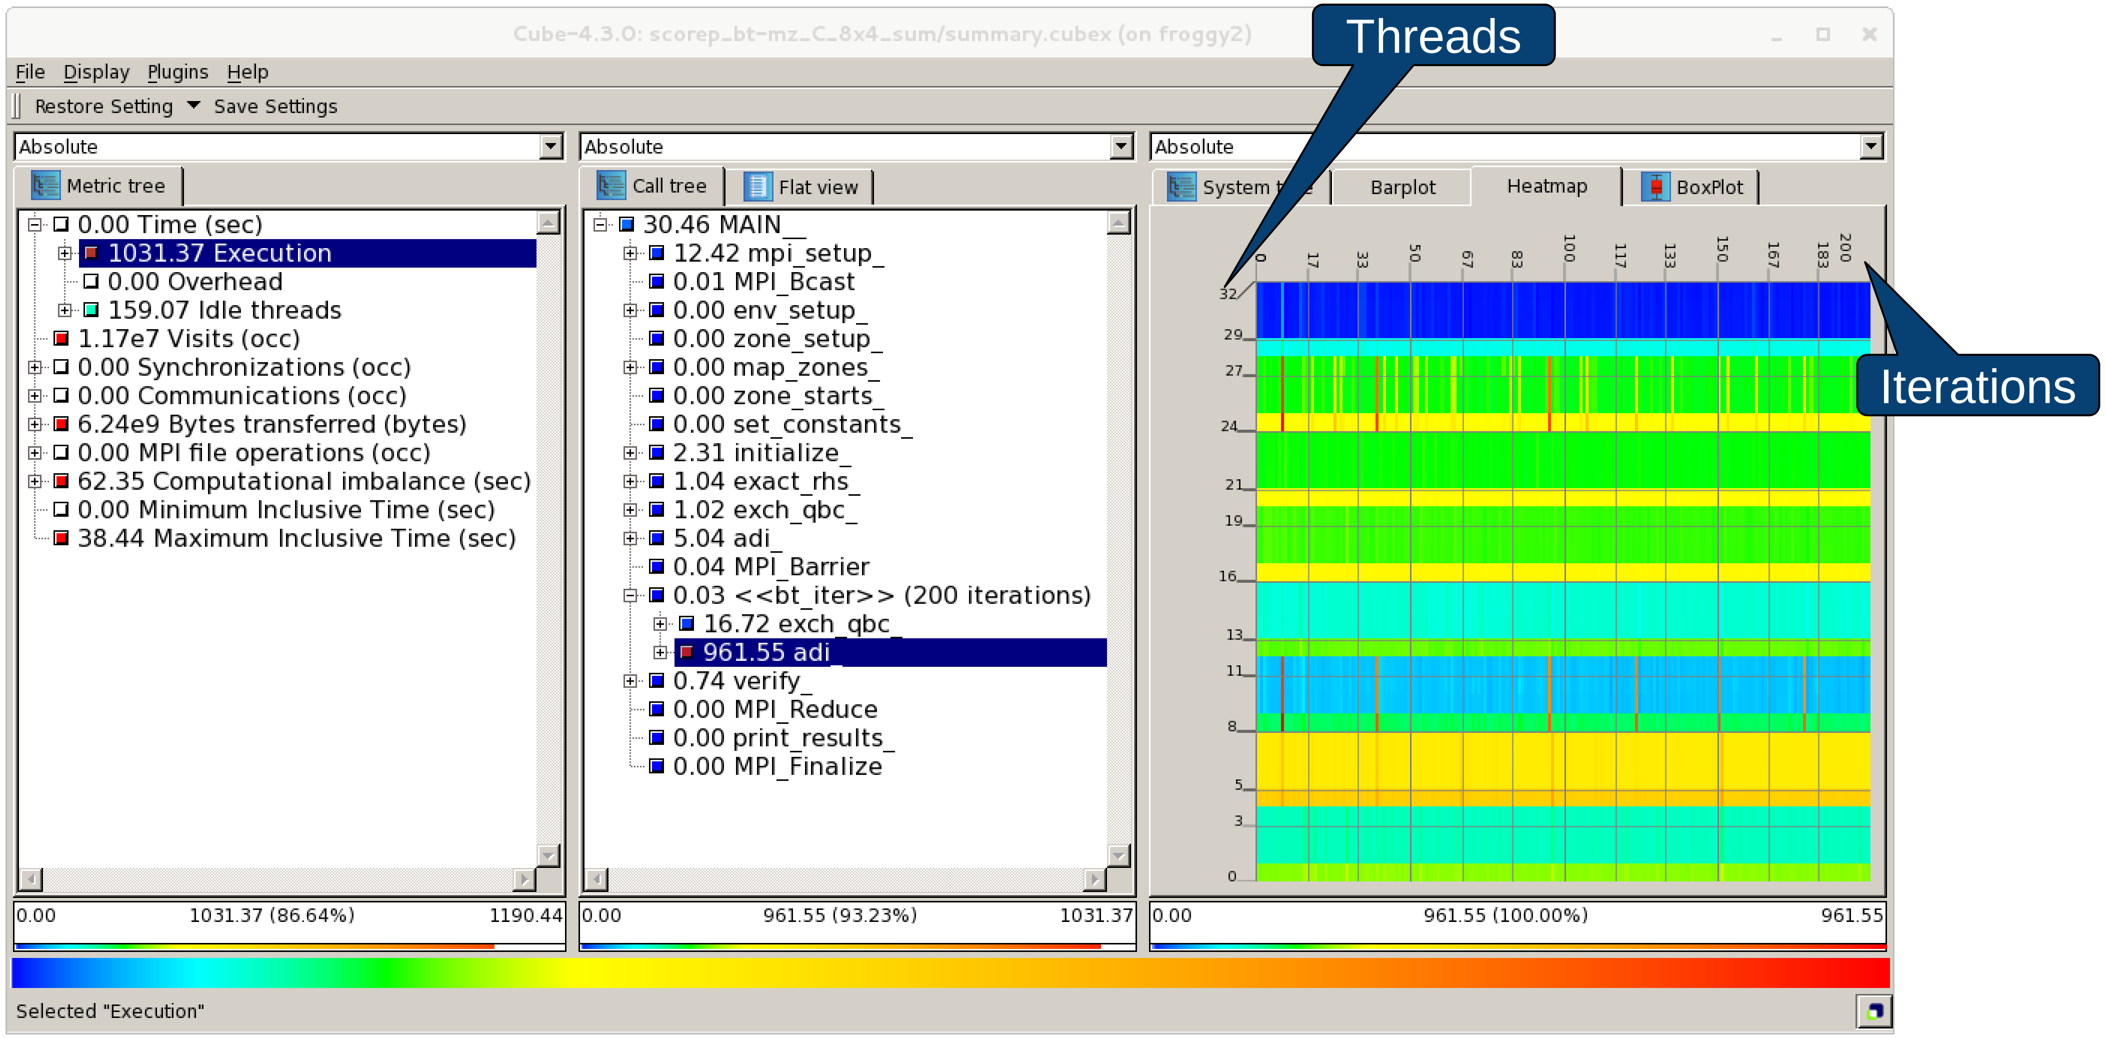

```
0.00 Time (sec)
  1031.37 Execution
  0.00 Overhead
  159.07 Idle threads
1.17e7 Visits (occ)
0.00 Synchronizations (occ)
0.00 Communications (occ)
6.24e9 Bytes transferred (bytes)
0.00 MPI file operations (occ)
62.35 Computational imbalance (sec)
0.00 Minimum Inclusive Time (sec)
38.44 Maximum Inclusive Time (sec)
```

```
30.46 MAIN__
  12.42 mpi_setup_
  0.01 MPI_Bcast
  0.00 env_setup_
  0.00 zone_setup_
  0.00 map_zones_
  0.00 zone_starts_
  0.00 set_constants_
  2.31 initialize_
  1.04 exact_rhs_
  1.02 exch_qbc_
  5.04 adi_
  0.04 MPI_Barrier
  0.03 <<bt_iter>> (200 iterations)
    16.72 exch_qbc_
    961.55 adi_
  0.74 verify_
  0.00 MPI_Reduce
  0.00 print_results_
  0.00 MPI_Finalize
```

```
0.00    1031.37 (86.64%)    1190.44
0.00    961.55 (93.23%)     1031.37
0.00    961.55 (100.00%)    961.55
```

Selected "Execution"

Further information

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

Obtaining a profile from BT-MZ

- module load scorep cube

- You need the scorep-enabled version of BT-MZ
- Run the command “cube” on the command line. You will be shown a “blank” screen
- File->Open
- Go to the bin.scorep/ directory of BT. Find the profile execution “scorep-20151028_XXXX_XXXXXXXXX/profile.cubex”

- You now have the results of your execution!