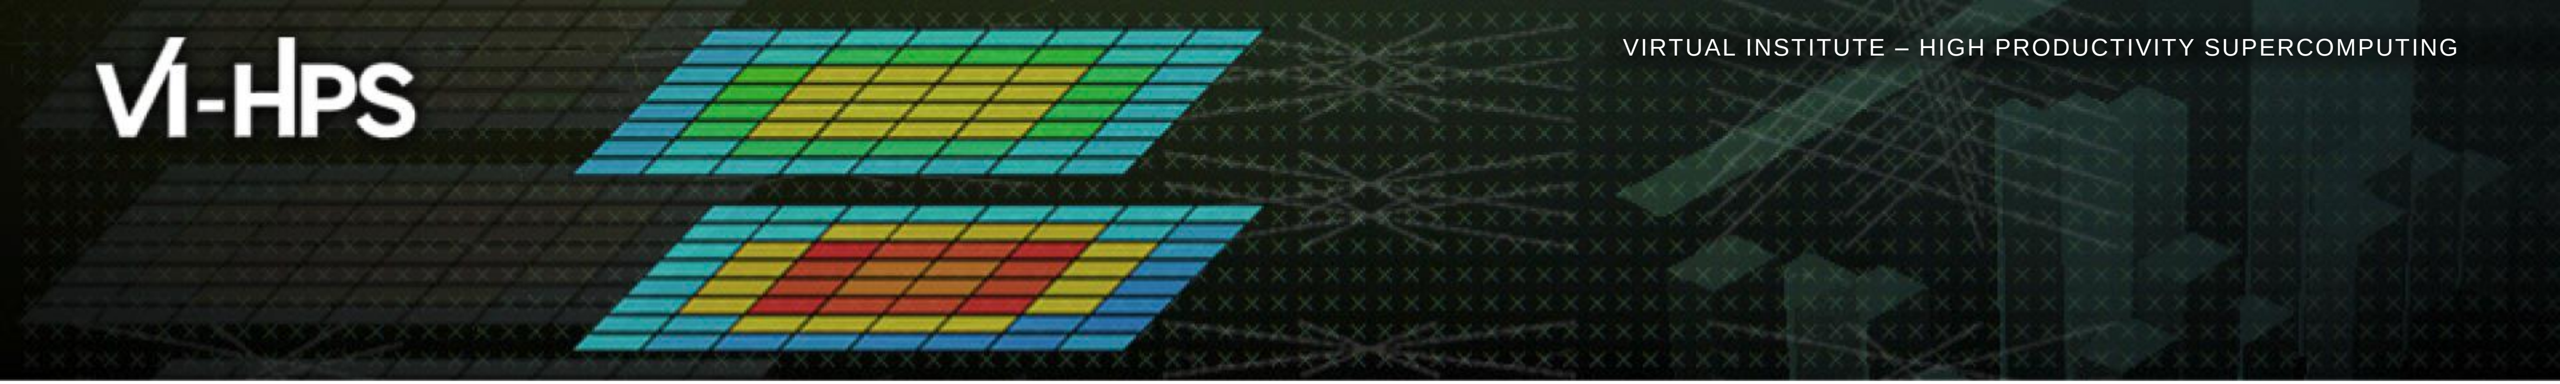

Analysis report examination with Cube

Markus Geimer
Jülich Supercomputing Centre

Cube

- Parallel program analysis report exploration tools
  - Libraries for XML+binary report reading & writing
  - Algebra utilities for report processing
  - GUI for interactive analysis exploration
    - Requires Qt4 ≥4.6 or Qt 5
- Originally developed as part of the Scalasca toolset
- Now available as a separate components
  - Can be installed independently of Score-P, e.g., on laptop or desktop
  - Latest release: Cube v4.4.x (March 2019)

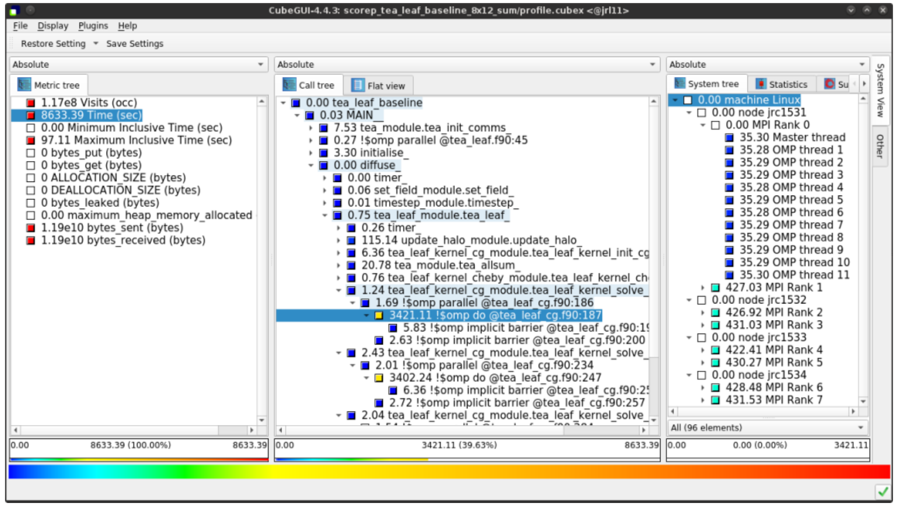

```
1.17e8 Visits (occ)
8633.39 Time (sec)
0.00 Minimum Inclusive Time (sec)
97.11 Maximum Inclusive Time (sec)
0 bytes_put (bytes)
0 bytes_get (bytes)
0 ALLOCATION_SIZE (bytes)
0 DEALLOCATION_SIZE (bytes)
0 bytes_leaked (bytes)
0.00 maximum_heap_memory_allocated
1.19e10 bytes_sent (bytes)
1.19e10 bytes_received (bytes)
```

```
0.00 tea_leaf_baseline
  0.03 MAIN__
    7.53 tea_module.tea_init_comms_
    0.27 !$omp parallel @tea_leaf.f90:45
    3.30 initialise_
    0.00 diffuse_
      0.00 timer_
      0.06 set_field_module.set_field_
      0.01 timestep_module.timestep_
      0.75 tea_leaf_module.tea_leaf_
        0.26 timer_
        115.14 update_halo_module.update_halo_
        6.36 tea_leaf_kernel_cg_module.tea_leaf_kernel_init_cg
        20.78 tea_module.tea_allsum_
        0.76 tea_leaf_kernel_cheby_module.tea_leaf_kernel_ch
        1.24 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_
          1.69 !$omp parallel @tea_leaf_cg.f90:186
            3421.11 !$omp do @tea_leaf_cg.f90:187
              5.83 !$omp implicit barrier @tea_leaf_cg.f90:1
            2.63 !$omp implicit barrier @tea_leaf_cg.f90:200
        2.43 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_
          2.01 !$omp parallel @tea_leaf_cg.f90:234
            3402.24 !$omp do @tea_leaf_cg.f90:247
              6.36 !$omp implicit barrier @tea_leaf_cg.f90:2
            2.72 !$omp implicit barrier @tea_leaf_cg.f90:257
        2.04 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_
```

```
0.00 machine Linux
  0.00 node jrc1531
    0.00 MPI Rank 0
      35.30 Master thread
      35.28 OMP thread 1
      35.29 OMP thread 2
      35.29 OMP thread 3
      35.28 OMP thread 4
      35.29 OMP thread 5
      35.28 OMP thread 6
      35.29 OMP thread 7
      35.29 OMP thread 8
      35.29 OMP thread 9
      35.29 OMP thread 10
      35.30 OMP thread 11
    427.03 MPI Rank 1
  0.00 node jrc1532
    426.92 MPI Rank 2
    431.03 MPI Rank 3
  0.00 node jrc1533
    422.41 MPI Rank 4
    430.27 MPI Rank 5
  0.00 node jrc1534
    428.48 MPI Rank 6
    431.53 MPI Rank 7
All (96 elements)
```

```
0.00    8633.39 (100.00%)    8633.39
0.00    3421.11 (39.63%)     8633.39
0.00    0.00 (0.00%)         3421.11
```

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

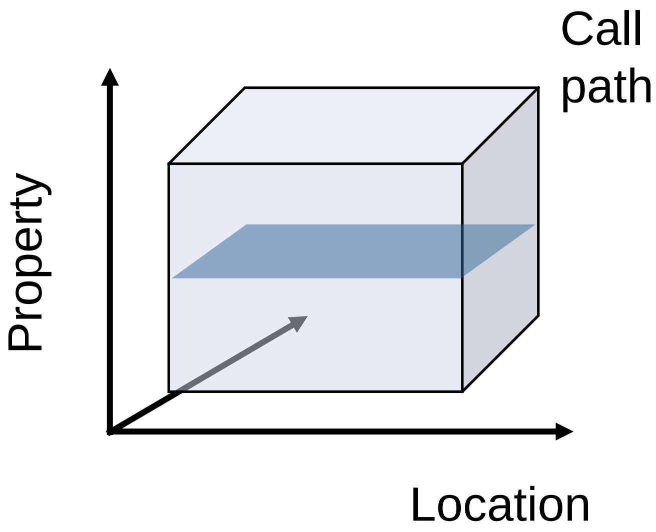

Inclusive vs. exclusive values

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

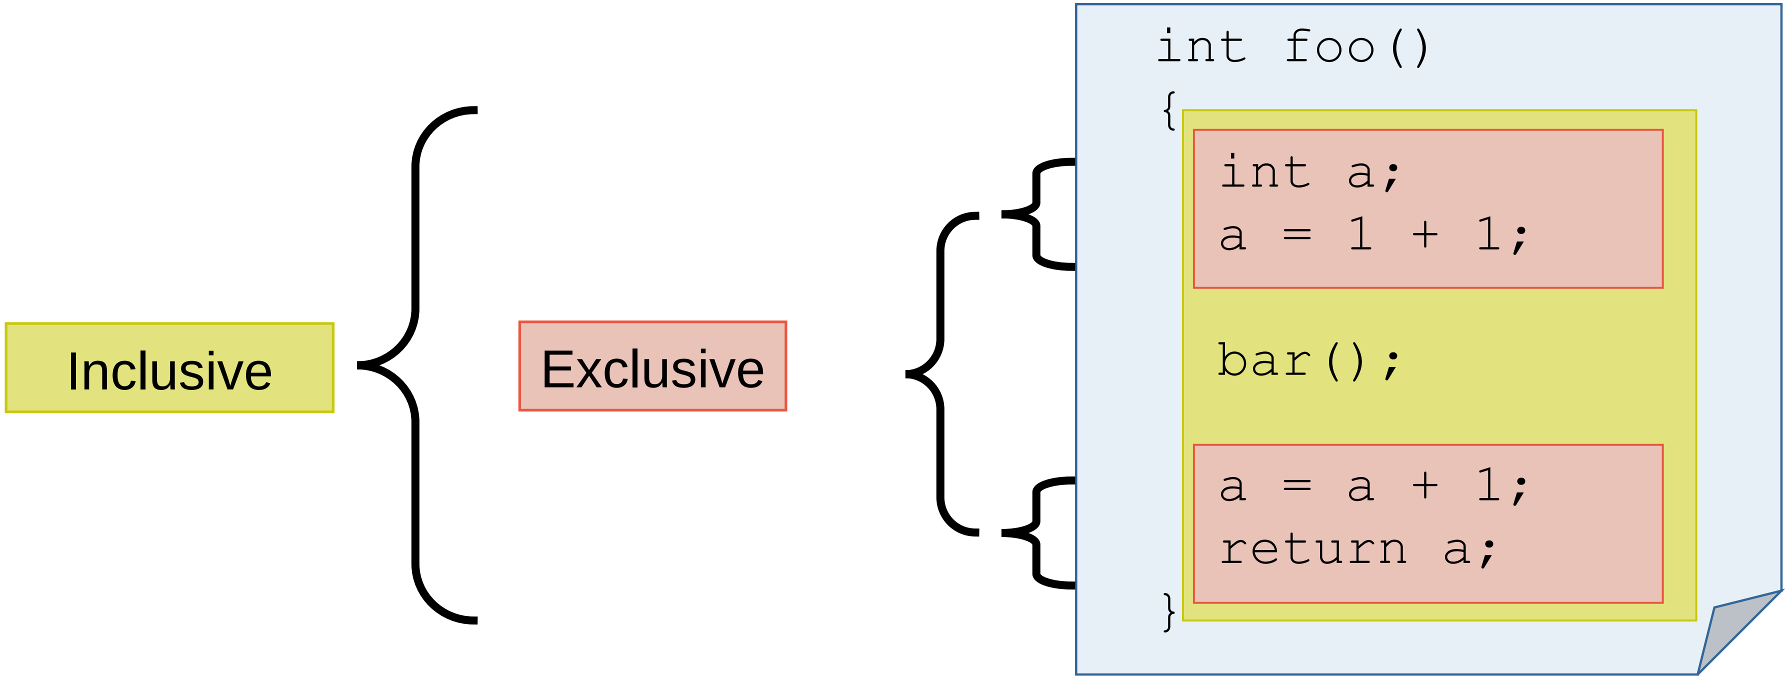

```
int foo()
{
  int a;
  a = 1 + 1;

  bar();

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

Case study:
TeaLeaf

Case study: TeaLeaf

- HPC mini-app developed by the UK Mini-App Consortium
  - Solves the linear 2D heat conduction equation on a spatially decomposed regular grid using a 5 point stencil with implicit solvers
  - Part of the Mantevo 3.0 suite
  - Available on GitHub: http://uk-mac.github.io/TeaLeaf/

- Measurements of TeaLeaf reference v1.0 taken on Jureca cluster @ JSC
  - Using Intel 19.0.3 compilers, Intel MPI 2019.3, and Score-P 5.0
  - Run configuration
    - 8 MPI ranks with 12 OpenMP threads each
    - Distributed across 4 compute nodes (2 ranks per node)
    - Test problem "5": 4000 × 4000 cells, CG solver

```
% cp -r /p/scratch/share/VI-HPS/examples/TeaLeaf .
% cd TeaLeaf
% cube scorep_tea_leaf_baseline_8x12_sum/profile.cubex
                              [GUI showing summary analysis report]
```

Score-P analysis report exploration (opening view)

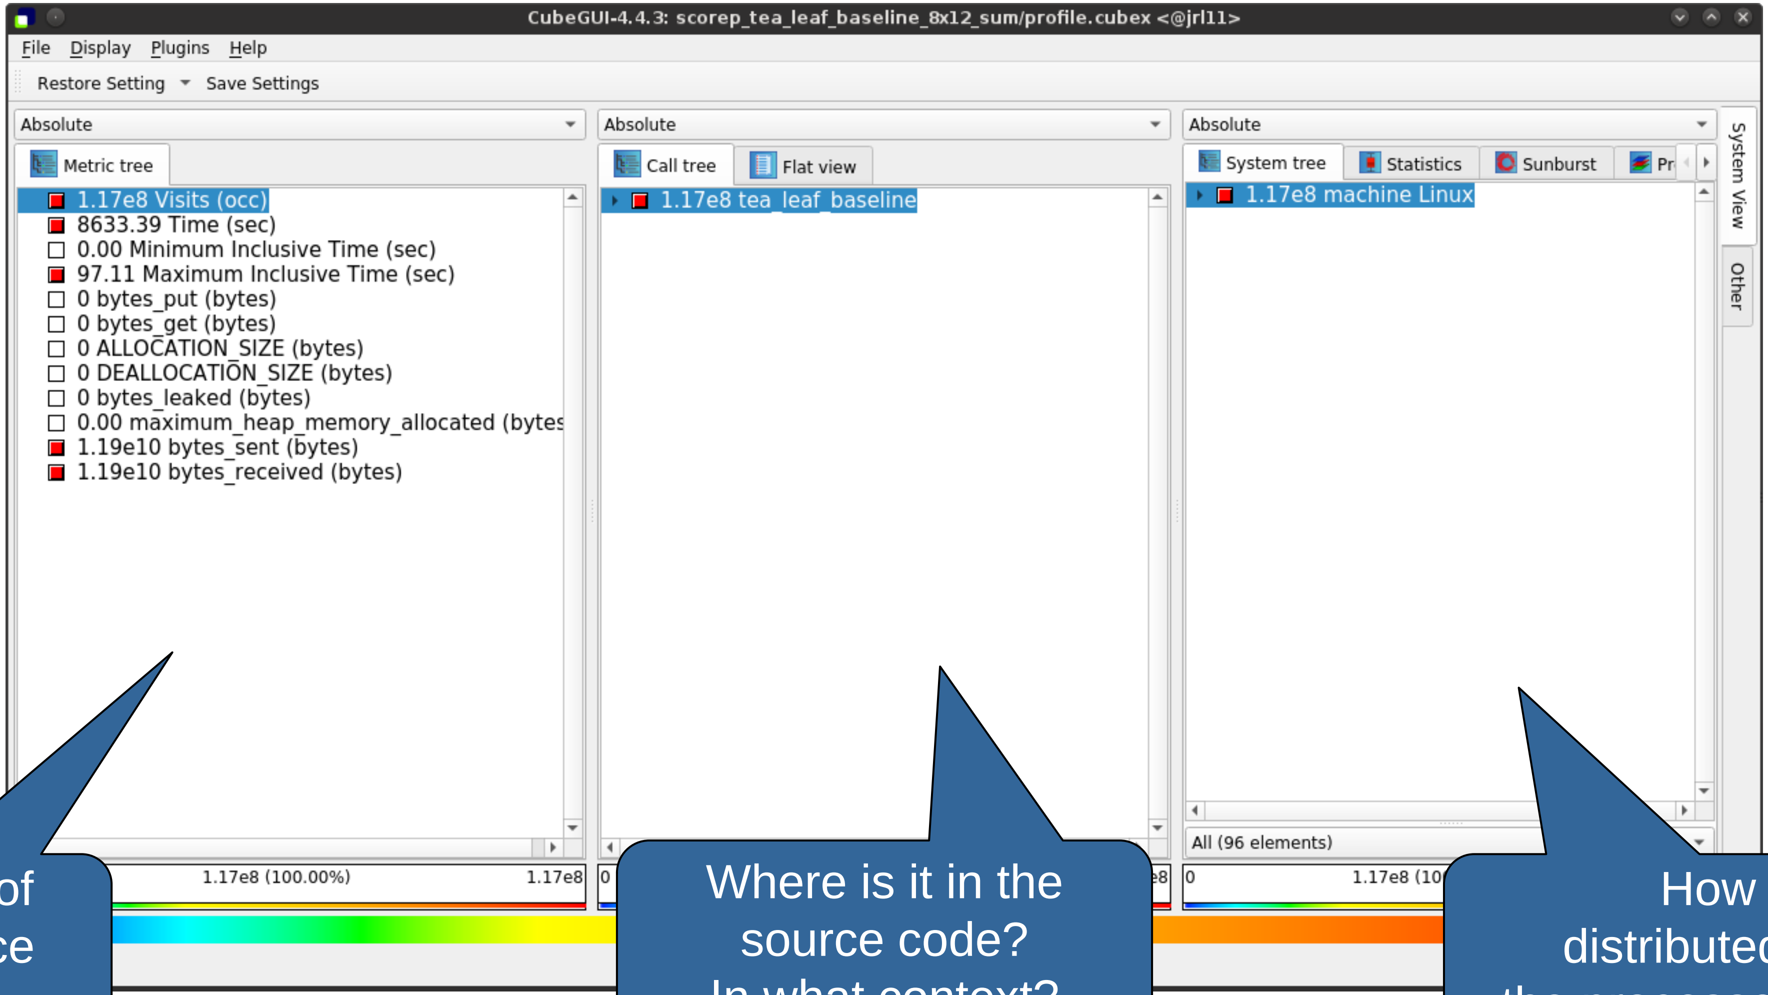

CubeGUI-4.4.3: scorep_tea_leaf_baseline_8x12_sum/profile.cubex <@jrl11>

- 🟥 1.17e8 Visits (occ)
- 🟥 8633.39 Time (sec)
- □ 0.00 Minimum Inclusive Time (sec)
- 🟥 97.11 Maximum Inclusive Time (sec)
- □ 0 bytes_put (bytes)
- □ 0 bytes_get (bytes)
- □ 0 ALLOCATION_SIZE (bytes)
- □ 0 DEALLOCATION_SIZE (bytes)
- □ 0 bytes_leaked (bytes)
- □ 0.00 maximum_heap_memory_allocated (bytes
- 🟥 1.19e10 bytes_sent (bytes)
- 🟥 1.19e10 bytes_received (bytes)

🟥 1.17e8 tea_leaf_baseline

🟥 1.17e8 machine Linux

1.17e8 (100.00%) 1.17e8

0 1.17e8 (10

What kind of performance metric?

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

How is it distributed across the processes/threads?

Metric selection

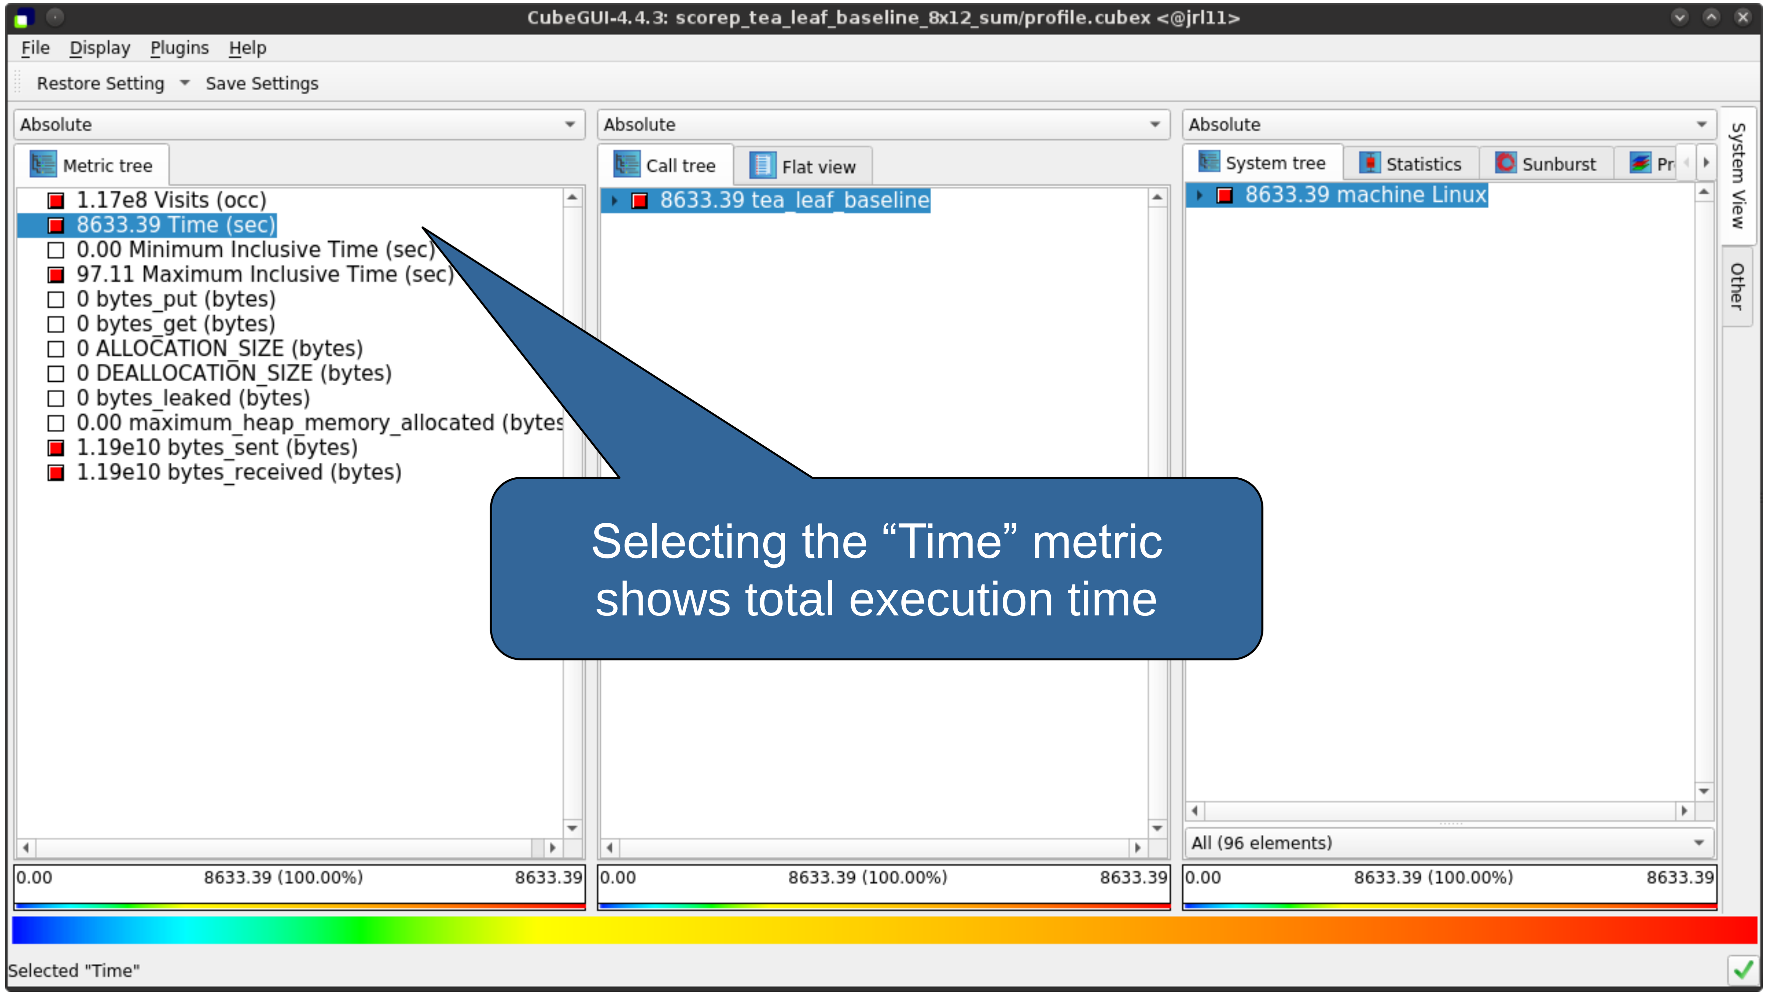

🟥 1.17e8 Visits (occ)
🟥 8633.39 Time (sec)
□ 0.00 Minimum Inclusive Time (sec)
🟥 97.11 Maximum Inclusive Time (sec)
□ 0 bytes_put (bytes)
□ 0 bytes_get (bytes)
□ 0 ALLOCATION_SIZE (bytes)
□ 0 DEALLOCATION_SIZE (bytes)
□ 0 bytes_leaked (bytes)
□ 0.00 maximum_heap_memory_allocated (bytes
🟥 1.19e10 bytes_sent (bytes)
🟥 1.19e10 bytes_received (bytes)

▸ 🟥 8633.39 tea_leaf_baseline

▸ 🟥 8633.39 machine Linux

0.00 8633.39 (100.00%) 8633.39

0.00 8633.39 (100.00%) 8633.39

0.00 8633.39 (100.00%) 8633.39

Selected "Time"

Selecting the “Time” metric shows total execution time

Expanding the system tree

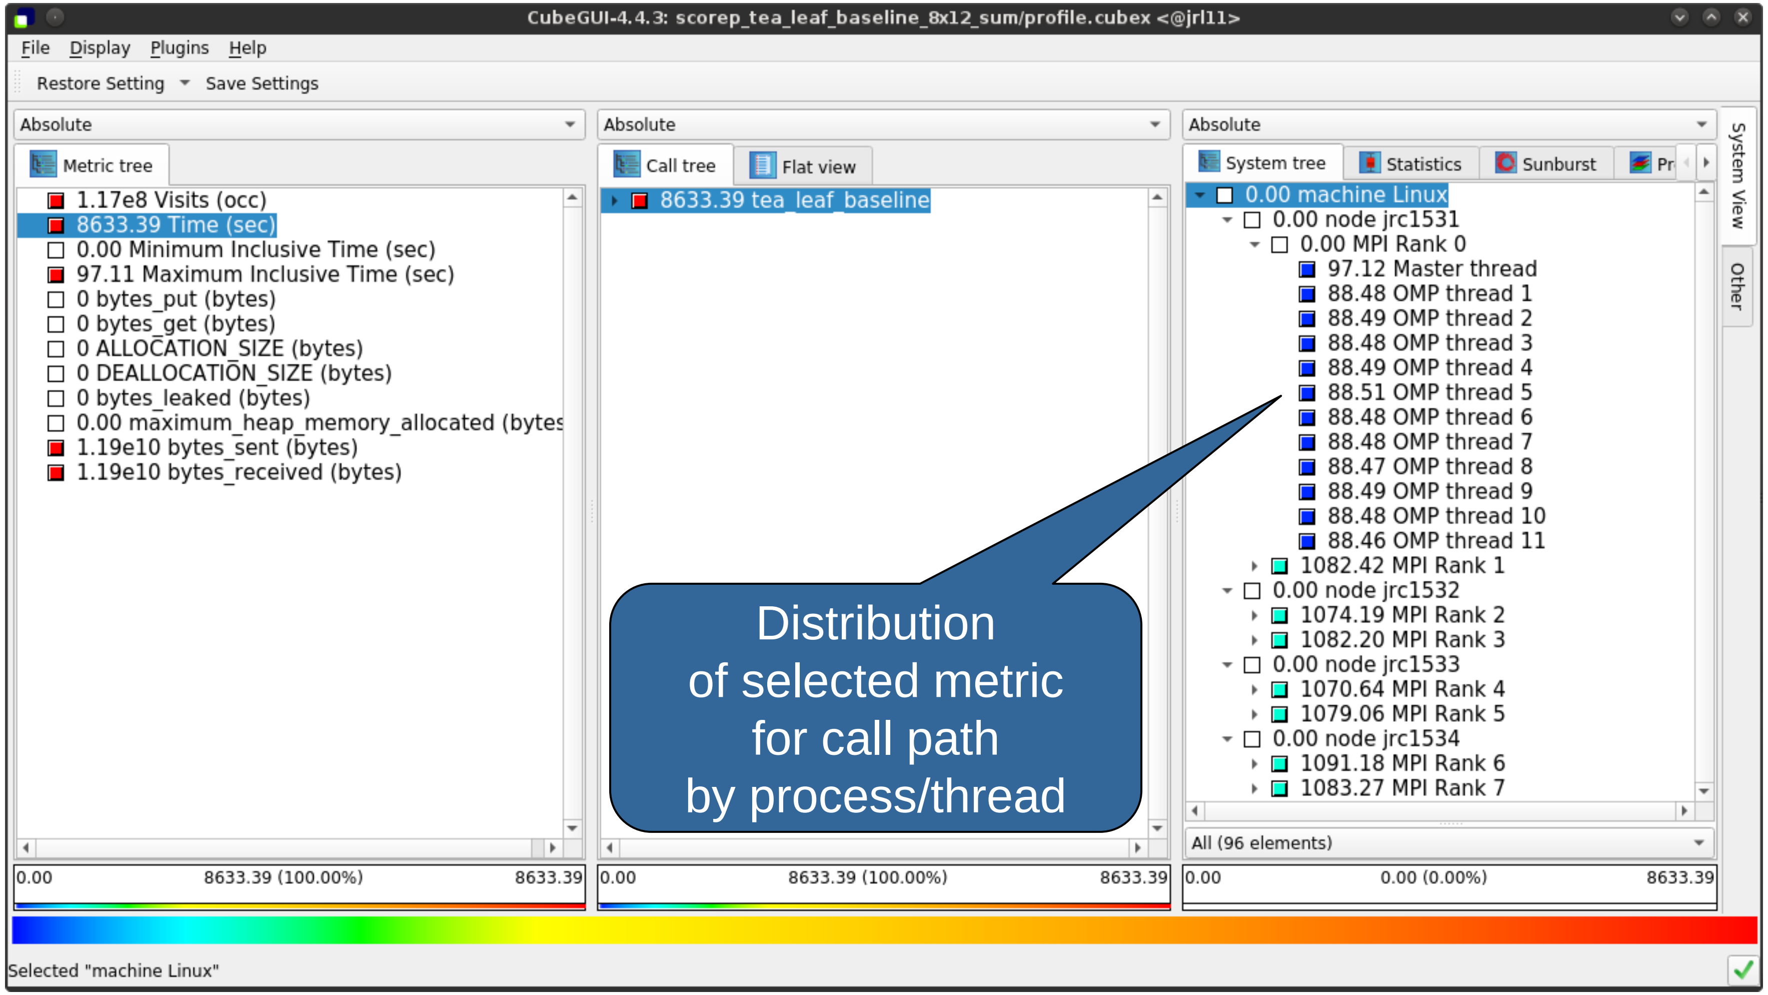

CubeGUI-4.4.3: scorep_tea_leaf_baseline_8x12_sum/profile.cubex <@jrl11>

File Display Plugins Help

Restore Setting Save Settings

Absolute

Metric tree

- 🟥 1.17e8 Visits (occ)
- 🟥 8633.39 Time (sec)
- □ 0.00 Minimum Inclusive Time (sec)
- 🟥 97.11 Maximum Inclusive Time (sec)
- □ 0 bytes_put (bytes)
- □ 0 bytes_get (bytes)
- □ 0 ALLOCATION_SIZE (bytes)
- □ 0 DEALLOCATION_SIZE (bytes)
- □ 0 bytes_leaked (bytes)
- □ 0.00 maximum_heap_memory_allocated (bytes
- 🟥 1.19e10 bytes_sent (bytes)
- 🟥 1.19e10 bytes_received (bytes)

0.00 8633.39 (100.00%) 8633.39

Absolute

Call tree Flat view

- 🟥 8633.39 tea_leaf_baseline

0.00 8633.39 (100.00%) 8633.39

Absolute

System tree Statistics Sunburst Pr

- □ 0.00 machine Linux
  - □ 0.00 node jrc1531
    - □ 0.00 MPI Rank 0
      - 🟦 97.12 Master thread
      - 🟦 88.48 OMP thread 1
      - 🟦 88.49 OMP thread 2
      - 🟦 88.48 OMP thread 3
      - 🟦 88.49 OMP thread 4
      - 🟦 88.51 OMP thread 5
      - 🟦 88.48 OMP thread 6
      - 🟦 88.48 OMP thread 7
      - 🟦 88.47 OMP thread 8
      - 🟦 88.49 OMP thread 9
      - 🟦 88.48 OMP thread 10
      - 🟦 88.46 OMP thread 11
    - 🟩 1082.42 MPI Rank 1
  - □ 0.00 node jrc1532
    - 🟩 1074.19 MPI Rank 2
    - 🟩 1082.20 MPI Rank 3
  - □ 0.00 node jrc1533
    - 🟩 1070.64 MPI Rank 4
    - 🟩 1079.06 MPI Rank 5
  - □ 0.00 node jrc1534
    - 🟩 1091.18 MPI Rank 6
    - 🟩 1083.27 MPI Rank 7

All (96 elements)

0.00 0.00 (0.00%) 8633.39

System View Other

Selected "machine Linux"

Distribution of selected metric for call path by process/thread

Expanding the call tree

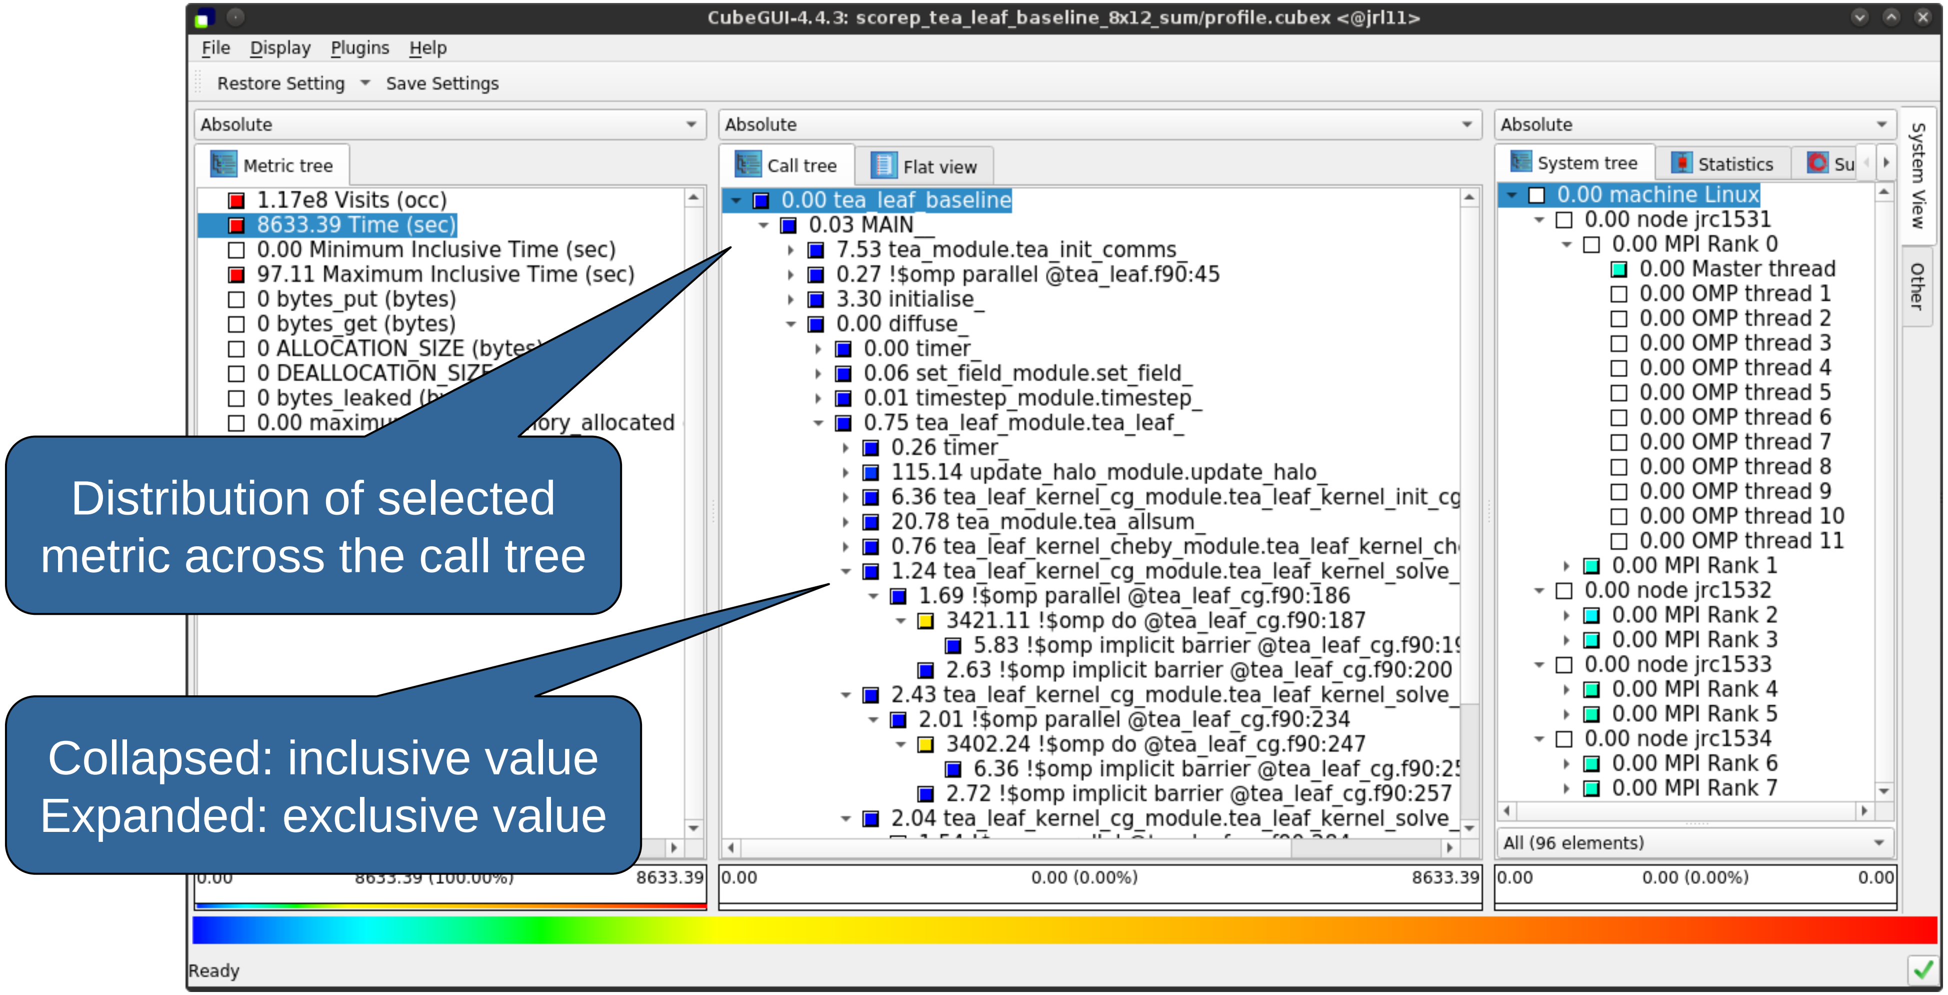

CubeGUI-4.4.3: scorep_tea_leaf_baseline_8x12_sum/profile.cubex <@jrl11>

```
■ 1.17e8 Visits (occ)
■ 8633.39 Time (sec)
□ 0.00 Minimum Inclusive Time (sec)
■ 97.11 Maximum Inclusive Time (sec)
□ 0 bytes_put (bytes)
□ 0 bytes_get (bytes)
□ 0 ALLOCATION_SIZE (bytes
□ 0 DEALLOCATION_SIZE
□ 0 bytes_leaked (b
□ 0.00 maximu                ory_allocated
```

```
■ 0.00 tea_leaf_baseline
  ■ 0.03 MAIN__
    ■ 7.53 tea_module.tea_init_comms_
    ■ 0.27 !$omp parallel @tea_leaf.f90:45
    ■ 3.30 initialise_
    ■ 0.00 diffuse_
      ■ 0.00 timer_
      ■ 0.06 set_field_module.set_field_
      ■ 0.01 timestep_module.timestep_
      ■ 0.75 tea_leaf_module.tea_leaf_
        ■ 0.26 timer_
        ■ 115.14 update_halo_module.update_halo_
        ■ 6.36 tea_leaf_kernel_cg_module.tea_leaf_kernel_init_cg
        ■ 20.78 tea_module.tea_allsum_
        ■ 0.76 tea_leaf_kernel_cheby_module.tea_leaf_kernel_ch
        ■ 1.24 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_
          ■ 1.69 !$omp parallel @tea_leaf_cg.f90:186
            ■ 3421.11 !$omp do @tea_leaf_cg.f90:187
              ■ 5.83 !$omp implicit barrier @tea_leaf_cg.f90:19
            ■ 2.63 !$omp implicit barrier @tea_leaf_cg.f90:200
        ■ 2.43 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_
          ■ 2.01 !$omp parallel @tea_leaf_cg.f90:234
            ■ 3402.24 !$omp do @tea_leaf_cg.f90:247
              ■ 6.36 !$omp implicit barrier @tea_leaf_cg.f90:25
            ■ 2.72 !$omp implicit barrier @tea_leaf_cg.f90:257
        ■ 2.04 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_
```

```
□ 0.00 machine Linux
  □ 0.00 node jrc1531
    □ 0.00 MPI Rank 0
      ■ 0.00 Master thread
      □ 0.00 OMP thread 1
      □ 0.00 OMP thread 2
      □ 0.00 OMP thread 3
      □ 0.00 OMP thread 4
      □ 0.00 OMP thread 5
      □ 0.00 OMP thread 6
      □ 0.00 OMP thread 7
      □ 0.00 OMP thread 8
      □ 0.00 OMP thread 9
      □ 0.00 OMP thread 10
      □ 0.00 OMP thread 11
    ■ 0.00 MPI Rank 1
  □ 0.00 node jrc1532
    ■ 0.00 MPI Rank 2
    ■ 0.00 MPI Rank 3
  □ 0.00 node jrc1533
    ■ 0.00 MPI Rank 4
    ■ 0.00 MPI Rank 5
  □ 0.00 node jrc1534
    ■ 0.00 MPI Rank 6
    ■ 0.00 MPI Rank 7
```

```
0.00    8633.39 (100.00%)    8633.39
0.00    0.00 (0.00%)         8633.39
0.00    0.00 (0.00%)         0.00
```

Distribution of selected metric across the call tree

Collapsed: inclusive value
Expanded: exclusive value

Selecting a call path

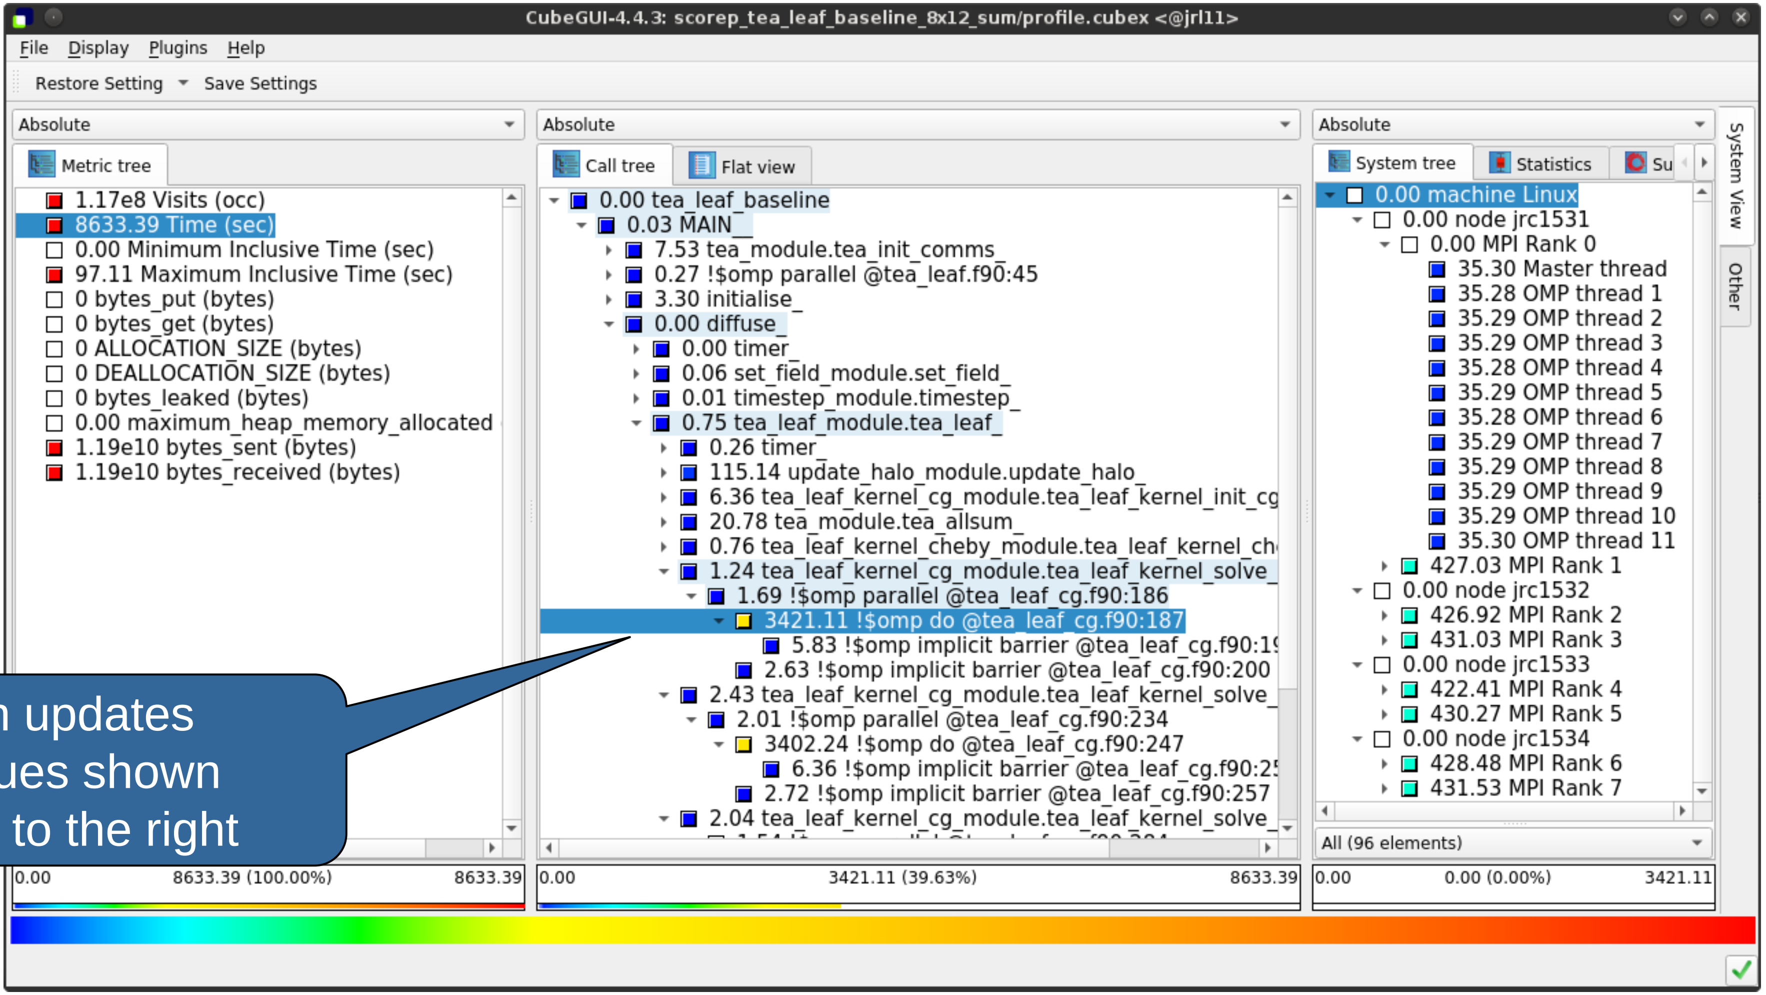

```
🟥 1.17e8 Visits (occ)
🟥 8633.39 Time (sec)
□ 0.00 Minimum Inclusive Time (sec)
🟥 97.11 Maximum Inclusive Time (sec)
□ 0 bytes_put (bytes)
□ 0 bytes_get (bytes)
□ 0 ALLOCATION_SIZE (bytes)
□ 0 DEALLOCATION_SIZE (bytes)
□ 0 bytes_leaked (bytes)
□ 0.00 maximum_heap_memory_allocated
🟥 1.19e10 bytes_sent (bytes)
🟥 1.19e10 bytes_received (bytes)
```

```
0.00                8633.39 (100.00%)                8633.39
```

```
🟦 0.00 tea_leaf_baseline
  🟦 0.03 MAIN__
    🟦 7.53 tea_module.tea_init_comms_
    🟦 0.27 !$omp parallel @tea_leaf.f90:45
    🟦 3.30 initialise_
    🟦 0.00 diffuse_
      🟦 0.00 timer_
      🟦 0.06 set_field_module.set_field_
      🟦 0.01 timestep_module.timestep_
      🟦 0.75 tea_leaf_module.tea_leaf_
        🟦 0.26 timer_
        🟦 115.14 update_halo_module.update_halo_
        🟦 6.36 tea_leaf_kernel_cg_module.tea_leaf_kernel_init_cg
        🟦 20.78 tea_module.tea_allsum_
        🟦 0.76 tea_leaf_kernel_cheby_module.tea_leaf_kernel_ch
        🟦 1.24 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_
          🟦 1.69 !$omp parallel @tea_leaf_cg.f90:186
            🟨 3421.11 !$omp do @tea_leaf_cg.f90:187
              🟦 5.83 !$omp implicit barrier @tea_leaf_cg.f90:1
            🟦 2.63 !$omp implicit barrier @tea_leaf_cg.f90:200
        🟦 2.43 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_
          🟦 2.01 !$omp parallel @tea_leaf_cg.f90:234
            🟨 3402.24 !$omp do @tea_leaf_cg.f90:247
              🟦 6.36 !$omp implicit barrier @tea_leaf_cg.f90:2
            🟦 2.72 !$omp implicit barrier @tea_leaf_cg.f90:257
        🟦 2.04 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_
```

```
0.00                3421.11 (39.63%)                8633.39
```

```
□ 0.00 machine Linux
  □ 0.00 node jrc1531
    □ 0.00 MPI Rank 0
      🟦 35.30 Master thread
      🟦 35.28 OMP thread 1
      🟦 35.29 OMP thread 2
      🟦 35.29 OMP thread 3
      🟦 35.28 OMP thread 4
      🟦 35.29 OMP thread 5
      🟦 35.28 OMP thread 6
      🟦 35.29 OMP thread 7
      🟦 35.29 OMP thread 8
      🟦 35.29 OMP thread 9
      🟦 35.29 OMP thread 10
      🟦 35.30 OMP thread 11
    🟩 427.03 MPI Rank 1
  □ 0.00 node jrc1532
    🟩 426.92 MPI Rank 2
    🟩 431.03 MPI Rank 3
  □ 0.00 node jrc1533
    🟩 422.41 MPI Rank 4
    🟩 430.27 MPI Rank 5
  □ 0.00 node jrc1534
    🟩 428.48 MPI Rank 6
    🟩 431.53 MPI Rank 7
```

```
0.00                0.00 (0.00%)                3421.11
```

Selection updates metric values shown in columns to the right

Multiple selection

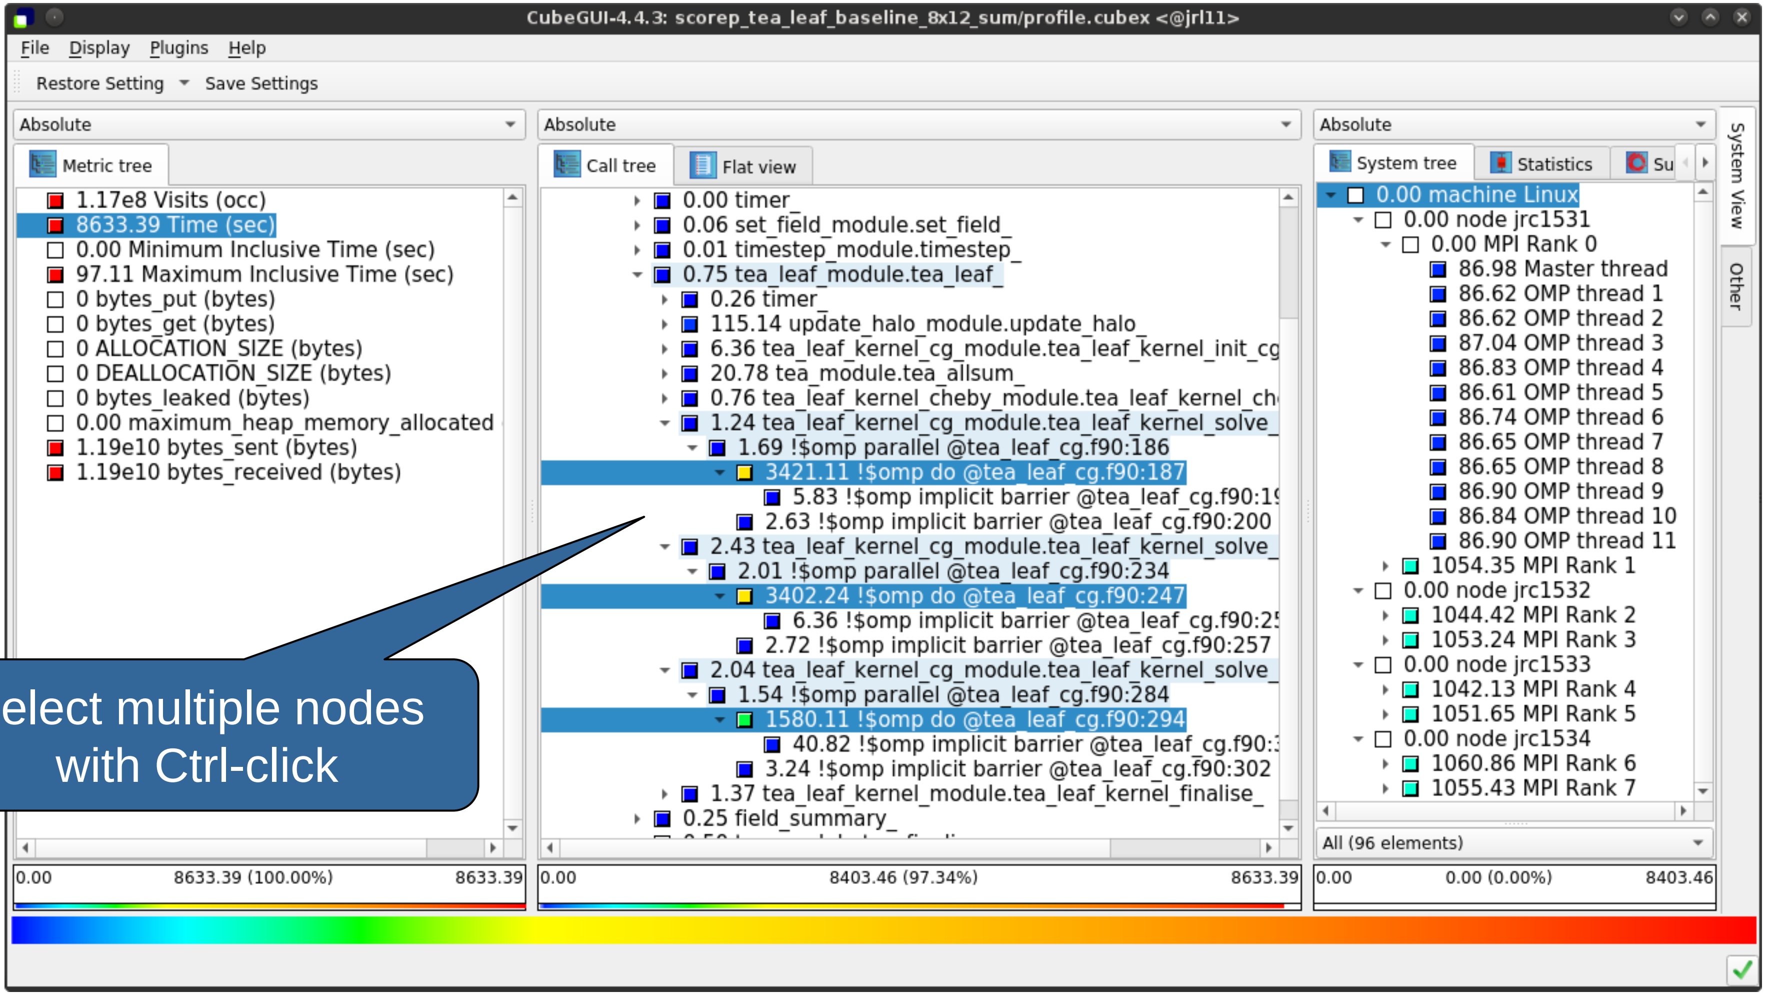
CubeGUI-4.4.3: scorep_tea_leaf_baseline_8x12_sum/profile.cubex <@jrl11>
File Display Plugins Help
Restore Setting Save Settings
Absolute
Metric tree
Call tree
Flat view
System tree
Statistics
Su
System View
Other
All (96 elements)

```
1.17e8 Visits (occ)
8633.39 Time (sec)
0.00 Minimum Inclusive Time (sec)
97.11 Maximum Inclusive Time (sec)
0 bytes_put (bytes)
0 bytes_get (bytes)
0 ALLOCATION_SIZE (bytes)
0 DEALLOCATION_SIZE (bytes)
0 bytes_leaked (bytes)
0.00 maximum_heap_memory_allocated
1.19e10 bytes_sent (bytes)
1.19e10 bytes_received (bytes)
```

```
0.00 timer_
0.06 set_field_module.set_field_
0.01 timestep_module.timestep_
0.75 tea_leaf_module.tea_leaf_
  0.26 timer_
  115.14 update_halo_module.update_halo
  6.36 tea_leaf_kernel_cg_module.tea_leaf_kernel_init_cg
  20.78 tea_module.tea_allsum_
  0.76 tea_leaf_kernel_cheby_module.tea_leaf_kernel_ch
  1.24 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_
    1.69 !$omp parallel @tea_leaf_cg.f90:186
      3421.11 !$omp do @tea_leaf_cg.f90:187
        5.83 !$omp implicit barrier @tea_leaf_cg.f90:1
      2.63 !$omp implicit barrier @tea_leaf_cg.f90:200
  2.43 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_
    2.01 !$omp parallel @tea_leaf_cg.f90:234
      3402.24 !$omp do @tea_leaf_cg.f90:247
        6.36 !$omp implicit barrier @tea_leaf_cg.f90:2
      2.72 !$omp implicit barrier @tea_leaf_cg.f90:257
  2.04 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_
    1.54 !$omp parallel @tea_leaf_cg.f90:284
      1580.11 !$omp do @tea_leaf_cg.f90:294
        40.82 !$omp implicit barrier @tea_leaf_cg.f90:
      3.24 !$omp implicit barrier @tea_leaf_cg.f90:302
  1.37 tea_leaf_kernel_module.tea_leaf_kernel_finalise_
0.25 field_summary_
```

```
0.00 machine Linux
  0.00 node jrc1531
    0.00 MPI Rank 0
      86.98 Master thread
      86.62 OMP thread 1
      86.62 OMP thread 2
      87.04 OMP thread 3
      86.83 OMP thread 4
      86.61 OMP thread 5
      86.74 OMP thread 6
      86.65 OMP thread 7
      86.65 OMP thread 8
      86.90 OMP thread 9
      86.84 OMP thread 10
      86.90 OMP thread 11
    1054.35 MPI Rank 1
  0.00 node jrc1532
    1044.42 MPI Rank 2
    1053.24 MPI Rank 3
  0.00 node jrc1533
    1042.13 MPI Rank 4
    1051.65 MPI Rank 5
  0.00 node jrc1534
    1060.86 MPI Rank 6
    1055.43 MPI Rank 7
```

0.00	8633.39 (100.00%)	8633.39
0.00	8403.46 (97.34%)	8633.39
0.00	0.00 (0.00%)	8403.46

Select multiple nodes with Ctrl-click

Box plot view

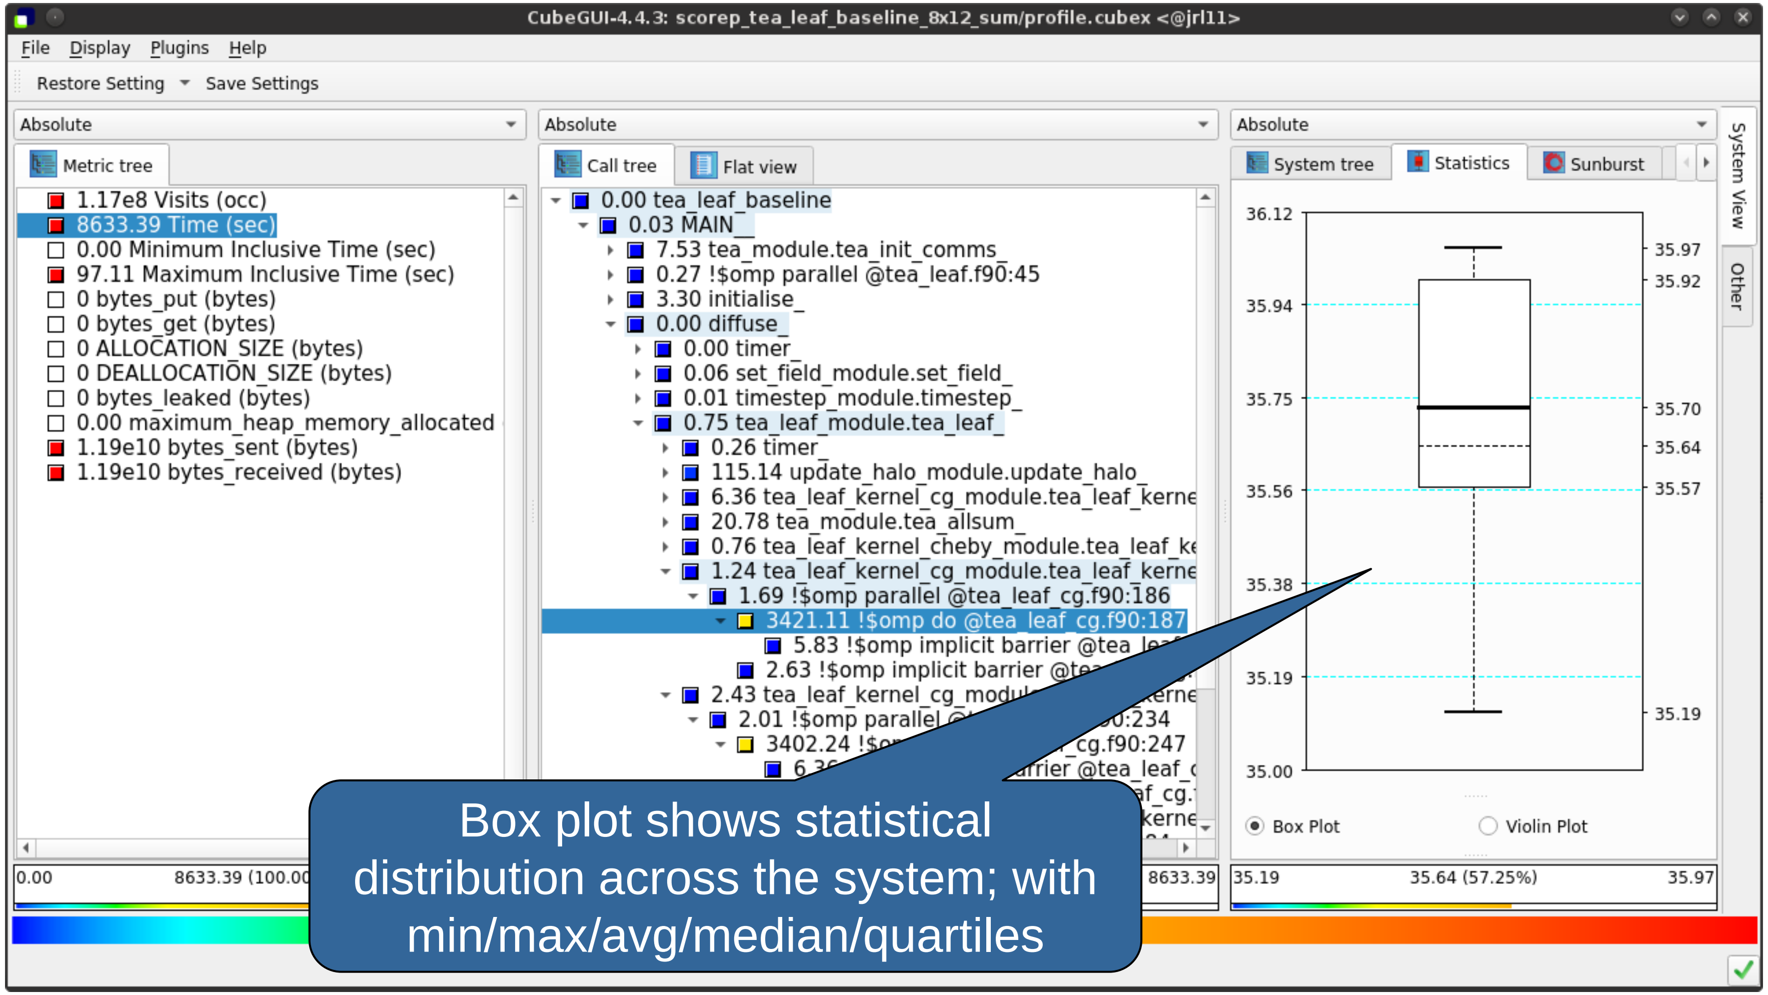

CubeGUI-4.4.3: scorep_tea_leaf_baseline_8x12_sum/profile.cubex <@jrl11>

- 🟥 1.17e8 Visits (occ)
- 🟥 8633.39 Time (sec)
- □ 0.00 Minimum Inclusive Time (sec)
- 🟥 97.11 Maximum Inclusive Time (sec)
- □ 0 bytes_put (bytes)
- □ 0 bytes_get (bytes)
- □ 0 ALLOCATION_SIZE (bytes)
- □ 0 DEALLOCATION_SIZE (bytes)
- □ 0 bytes_leaked (bytes)
- □ 0.00 maximum_heap_memory_allocated
- 🟥 1.19e10 bytes_sent (bytes)
- 🟥 1.19e10 bytes_received (bytes)

```
0.00 tea_leaf_baseline
  0.03 MAIN__
    7.53 tea_module.tea_init_comms_
    0.27 !$omp parallel @tea_leaf.f90:45
    3.30 initialise_
    0.00 diffuse_
      0.00 timer_
      0.06 set_field_module.set_field_
      0.01 timestep_module.timestep_
      0.75 tea_leaf_module.tea_leaf_
        0.26 timer_
        115.14 update_halo_module.update_halo_
        6.36 tea_leaf_kernel_cg_module.tea_leaf_kerne
        20.78 tea_module.tea_allsum_
        0.76 tea_leaf_kernel_cheby_module.tea_leaf_ke
        1.24 tea_leaf_kernel_cg_module.tea_leaf_kerne
          1.69 !$omp parallel @tea_leaf_cg.f90:186
            3421.11 !$omp do @tea_leaf_cg.f90:187
              5.83 !$omp implicit barrier @tea_lea
            2.63 !$omp implicit barrier @te
        2.43 tea_leaf_kernel_cg_modul                    kerne
          2.01 !$omp parallel @                     0:234
            3402.24 !$o                        _cg.f90:247
              6.3                          rrier @tea_leaf_c
                                            af_cg.
                                            kerne
```

0.00	8633.39 (100.00		8633.39	35.19	35.64 (57.25%)	35.97

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

Violin plot view

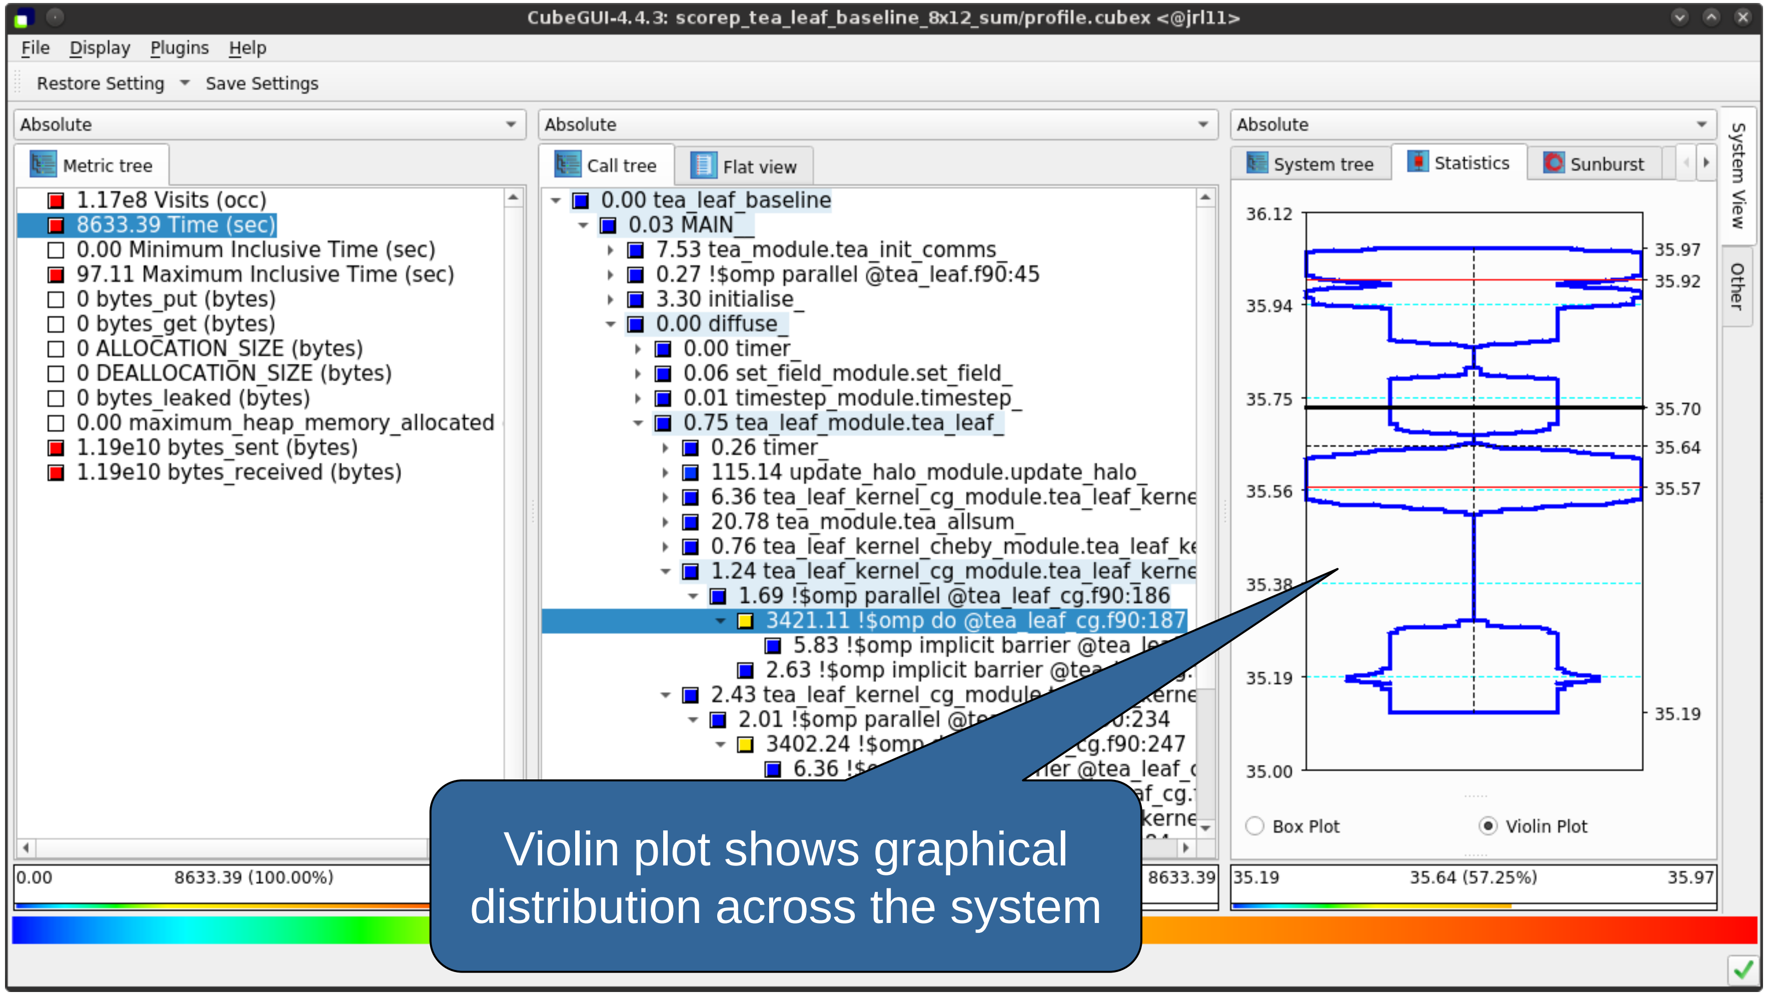

CubeGUI-4.4.3: scorep_tea_leaf_baseline_8x12_sum/profile.cubex <@jrl11>

Metric tree:

- 🟥 1.17e8 Visits (occ)
- 🟥 8633.39 Time (sec)
- □ 0.00 Minimum Inclusive Time (sec)
- 🟥 97.11 Maximum Inclusive Time (sec)
- □ 0 bytes_put (bytes)
- □ 0 bytes_get (bytes)
- □ 0 ALLOCATION_SIZE (bytes)
- □ 0 DEALLOCATION_SIZE (bytes)
- □ 0 bytes_leaked (bytes)
- □ 0.00 maximum_heap_memory_allocated
- 🟥 1.19e10 bytes_sent (bytes)
- 🟥 1.19e10 bytes_received (bytes)

Call tree:

```
▾ 0.00 tea_leaf_baseline
  ▾ 0.03 MAIN__
    ▸ 7.53 tea_module.tea_init_comms_
    ▸ 0.27 !$omp parallel @tea_leaf.f90:45
    ▸ 3.30 initialise_
    ▾ 0.00 diffuse_
      ▸ 0.00 timer_
      ▸ 0.06 set_field_module.set_field_
      ▸ 0.01 timestep_module.timestep_
      ▾ 0.75 tea_leaf_module.tea_leaf_
        ▸ 0.26 timer_
        ▸ 115.14 update_halo_module.update_halo_
        ▸ 6.36 tea_leaf_kernel_cg_module.tea_leaf_kerne
        ▸ 20.78 tea_module.tea_allsum_
        ▸ 0.76 tea_leaf_kernel_cheby_module.tea_leaf_ke
        ▾ 1.24 tea_leaf_kernel_cg_module.tea_leaf_kerne
          ▾ 1.69 !$omp parallel @tea_leaf_cg.f90:186
            ▾ 3421.11 !$omp do @tea_leaf_cg.f90:187
                5.83 !$omp implicit barrier @tea_le
              2.63 !$omp implicit barrier @tea
        ▾ 2.43 tea_leaf_kernel_cg_module          kerne
          ▾ 2.01 !$omp parallel @te            0:234
            ▾ 3402.24 !$omp                cg.f90:247
                6.36 !$                 ier @tea_leaf_c
                                        af_cg.f
                                        kerne
```

0.00 8633.39 (100.00%)

8633.39

35.19 35.64 (57.25%) 35.97

Violin plot shows graphical distribution across the system

Topology view

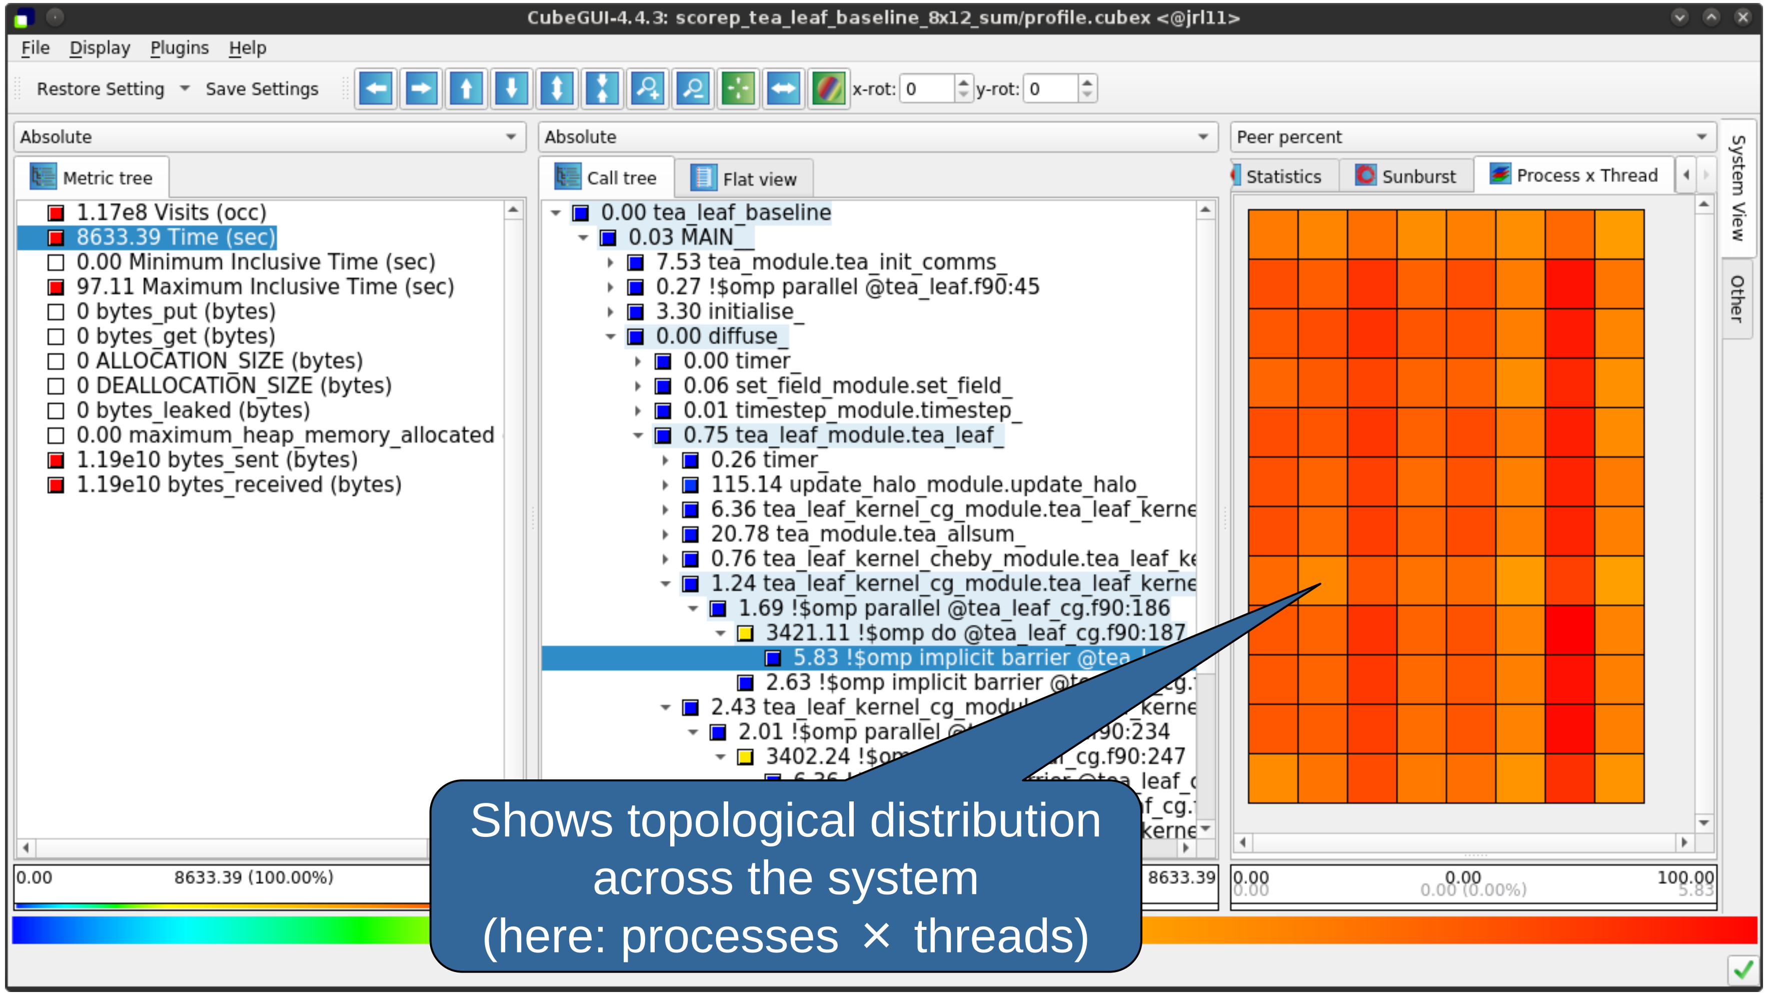

CubeGUI-4.4.3: scorep_tea_leaf_baseline_8x12_sum/profile.cubex <@jrl11>

Absolute

```
☐ 1.17e8 Visits (occ)
☐ 8633.39 Time (sec)
☐ 0.00 Minimum Inclusive Time (sec)
☐ 97.11 Maximum Inclusive Time (sec)
☐ 0 bytes_put (bytes)
☐ 0 bytes_get (bytes)
☐ 0 ALLOCATION_SIZE (bytes)
☐ 0 DEALLOCATION_SIZE (bytes)
☐ 0 bytes_leaked (bytes)
☐ 0.00 maximum_heap_memory_allocated
☐ 1.19e10 bytes_sent (bytes)
☐ 1.19e10 bytes_received (bytes)
```

Absolute

```
0.00 tea_leaf_baseline
  0.03 MAIN__
    7.53 tea_module.tea_init_comms_
    0.27 !$omp parallel @tea_leaf.f90:45
    3.30 initialise_
    0.00 diffuse_
      0.00 timer_
      0.06 set_field_module.set_field_
      0.01 timestep_module.timestep_
      0.75 tea_leaf_module.tea_leaf_
        0.26 timer_
        115.14 update_halo_module.update_halo_
        6.36 tea_leaf_kernel_cg_module.tea_leaf_kerne
        20.78 tea_module.tea_allsum_
        0.76 tea_leaf_kernel_cheby_module.tea_leaf_ke
        1.24 tea_leaf_kernel_cg_module.tea_leaf_kerne
          1.69 !$omp parallel @tea_leaf_cg.f90:186
            3421.11 !$omp do @tea_leaf_cg.f90:187
              5.83 !$omp implicit barrier @tea_l
            2.63 !$omp implicit barrier @te                cg.f
        2.43 tea_leaf_kernel_cg_modu                       kerne
          2.01 !$omp parallel @                      f90:234
            3402.24 !$om                             _cg.f90:247
              6.36                               rrier @tea_leaf_c
                                                          f_cg.f
                                                          kerne
```

Peer percent

0.00	8633.39 (100.00%)	8633.39

0.00	0.00	100.00
0.00	0.00 (0.00%)	5.83

Shows topological distribution across the system (here: processes × threads)

Topology view (cont.)

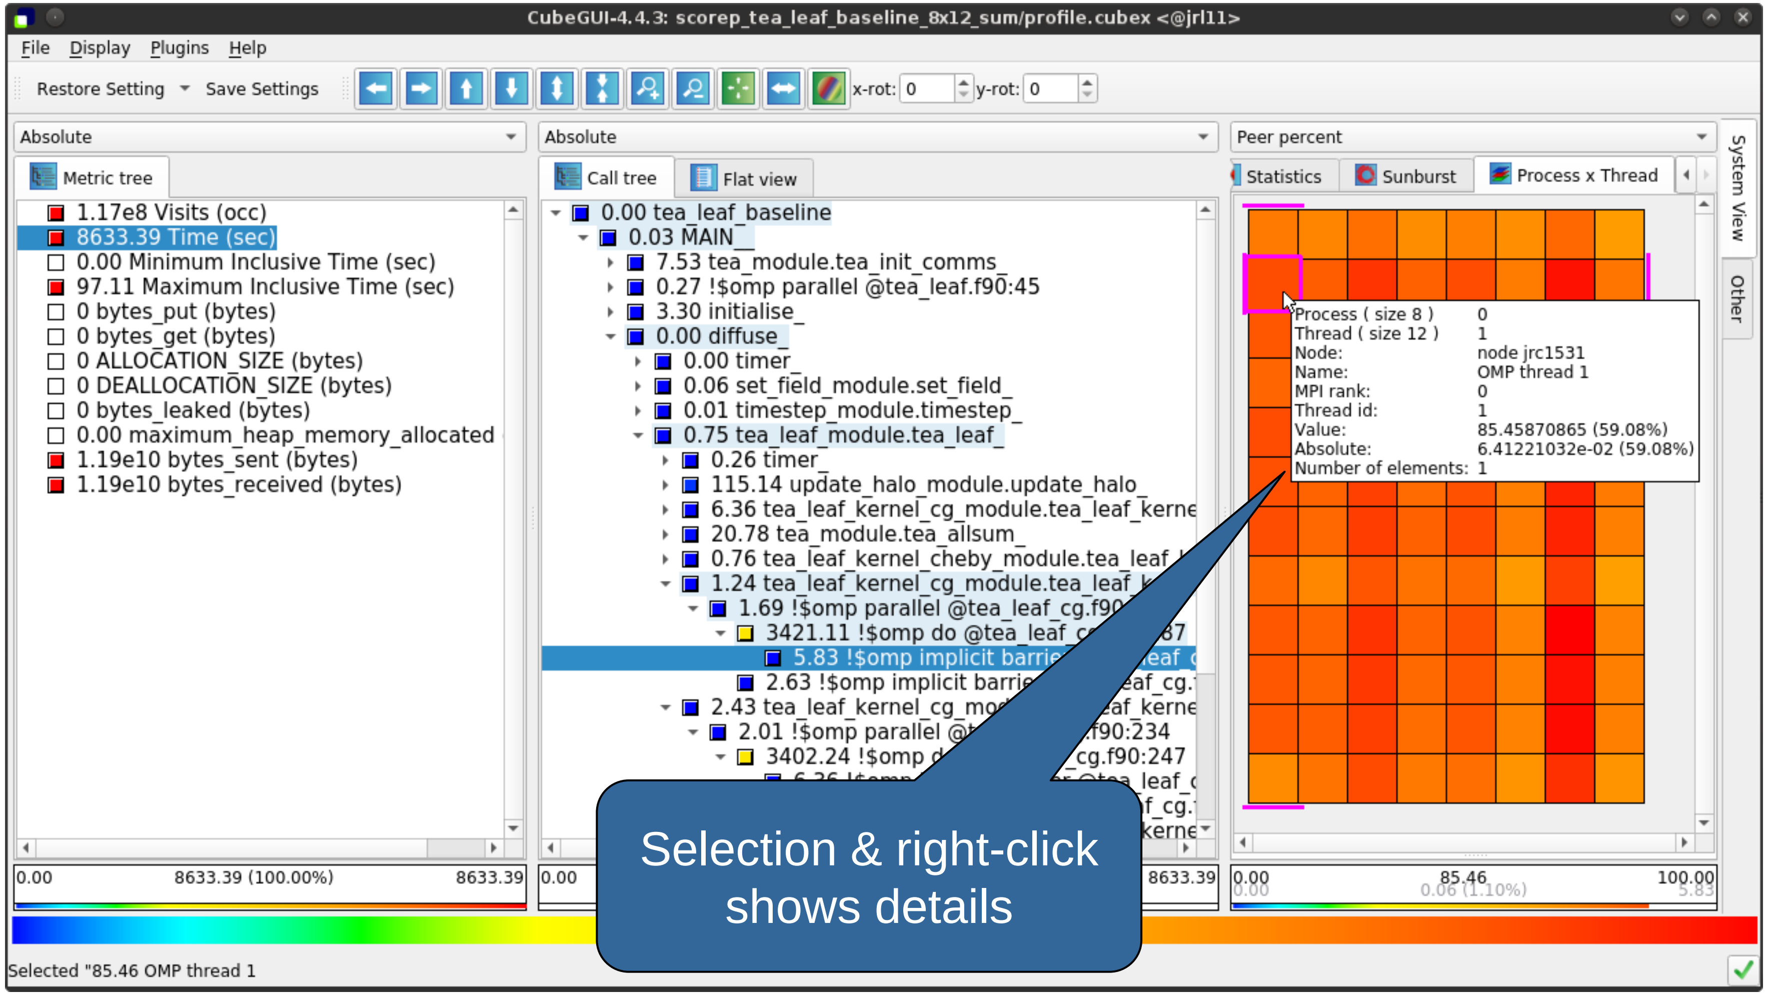

```
1.17e8 Visits (occ)
8633.39 Time (sec)
0.00 Minimum Inclusive Time (sec)
97.11 Maximum Inclusive Time (sec)
0 bytes_put (bytes)
0 bytes_get (bytes)
0 ALLOCATION_SIZE (bytes)
0 DEALLOCATION_SIZE (bytes)
0 bytes_leaked (bytes)
0.00 maximum_heap_memory_allocated
1.19e10 bytes_sent (bytes)
1.19e10 bytes_received (bytes)
```

```
0.00 tea_leaf_baseline
  0.03 MAIN__
    7.53 tea_module.tea_init_comms_
    0.27 !$omp parallel @tea_leaf.f90:45
    3.30 initialise_
    0.00 diffuse_
      0.00 timer_
      0.06 set_field_module.set_field_
      0.01 timestep_module.timestep_
      0.75 tea_leaf_module.tea_leaf_
        0.26 timer_
        115.14 update_halo_module.update_halo_
        6.36 tea_leaf_kernel_cg_module.tea_leaf_kerne
        20.78 tea_module.tea_allsum_
        0.76 tea_leaf_kernel_cheby_module.tea_leaf_
        1.24 tea_leaf_kernel_cg_module.tea_leaf_
          1.69 !$omp parallel @tea_leaf_cg.f90
            3421.11 !$omp do @tea_leaf_c        87
              5.83 !$omp implicit barrie        eaf_c
            2.63 !$omp implicit barrie        eaf_cg.
        2.43 tea_leaf_kernel_cg_mo        af_kerne
          2.01 !$omp parallel @        f90:234
            3402.24 !$omp d        _cg.f90:247
              6.36 !$omp        @tea_leaf_c
                                f_cg.
                                kerne
```

```
Process ( size 8 )      0
Thread ( size 12 )      1
Node:                   node jrc1531
Name:                   OMP thread 1
MPI rank:               0
Thread id:              1
Value:                  85.45870865 (59.08%)
Absolute:               6.41221032e-02 (59.08%)
Number of elements:     1
```

```
0.00        8633.39 (100.00%)        8633.39
0.00                                 8633.39
0.00        85.46                    100.00
0.00        0.06 (1.10%)             5.83
```

Selected "85.46 OMP thread 1

Selection & right-click shows details

Alternative display modes

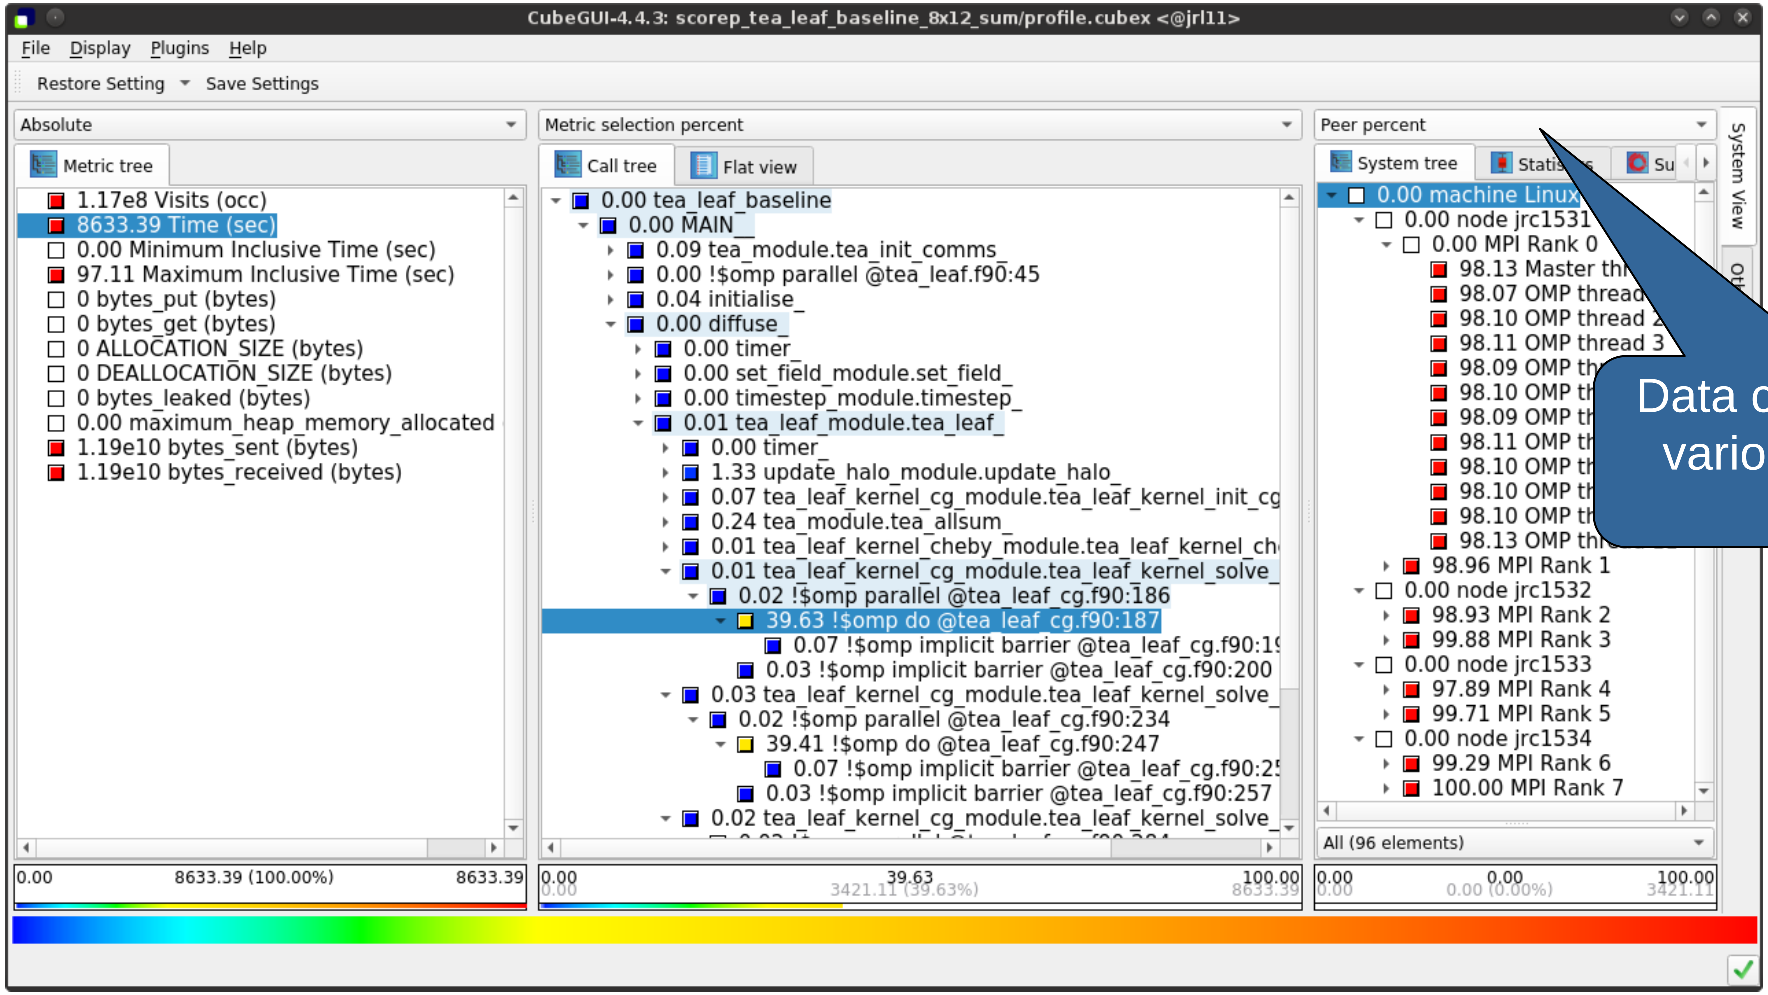

CubeGUI-4.4.3: scorep_tea_leaf_baseline_8x12_sum/profile.cubex <@jrl11>

Absolute | Metric selection percent | Peer percent

```
1.17e8 Visits (occ)
8633.39 Time (sec)
0.00 Minimum Inclusive Time (sec)
97.11 Maximum Inclusive Time (sec)
0 bytes_put (bytes)
0 bytes_get (bytes)
0 ALLOCATION_SIZE (bytes)
0 DEALLOCATION_SIZE (bytes)
0 bytes_leaked (bytes)
0.00 maximum_heap_memory_allocated
1.19e10 bytes_sent (bytes)
1.19e10 bytes_received (bytes)
```

```
0.00 tea_leaf_baseline
  0.00 MAIN__
    0.09 tea_module.tea_init_comms_
    0.00 !$omp parallel @tea_leaf.f90:45
    0.04 initialise_
    0.00 diffuse_
      0.00 timer_
      0.00 set_field_module.set_field_
      0.00 timestep_module.timestep_
      0.01 tea_leaf_module.tea_leaf_
        0.00 timer_
        1.33 update_halo_module.update_halo_
        0.07 tea_leaf_kernel_cg_module.tea_leaf_kernel_init_cg
        0.24 tea_module.tea_allsum_
        0.01 tea_leaf_kernel_cheby_module.tea_leaf_kernel_ch
        0.01 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_
          0.02 !$omp parallel @tea_leaf_cg.f90:186
            39.63 !$omp do @tea_leaf_cg.f90:187
              0.07 !$omp implicit barrier @tea_leaf_cg.f90:19
            0.03 !$omp implicit barrier @tea_leaf_cg.f90:200
        0.03 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_
          0.02 !$omp parallel @tea_leaf_cg.f90:234
            39.41 !$omp do @tea_leaf_cg.f90:247
              0.07 !$omp implicit barrier @tea_leaf_cg.f90:25
            0.03 !$omp implicit barrier @tea_leaf_cg.f90:257
        0.02 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_
```

```
0.00 machine Linux
  0.00 node jrc1531
    0.00 MPI Rank 0
      98.13 Master thr
      98.07 OMP thread
      98.10 OMP thread 2
      98.11 OMP thread 3
      98.09 OMP th
      98.10 OMP th
      98.09 OMP th
      98.11 OMP th
      98.10 OMP th
      98.10 OMP th
      98.10 OMP th
      98.13 OMP thr
    98.96 MPI Rank 1
  0.00 node jrc1532
    98.93 MPI Rank 2
    99.88 MPI Rank 3
  0.00 node jrc1533
    97.89 MPI Rank 4
    99.71 MPI Rank 5
  0.00 node jrc1534
    99.29 MPI Rank 6
    100.00 MPI Rank 7
All (96 elements)
```

```
0.00   8633.39 (100.00%)   8633.39
0.00   39.63   100.00
0.00   3421.11 (39.63%)   8633.39
0.00   0.00   100.00
0.00   0.00 (0.00%)   3421.11
```

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

Source-code view via context menu

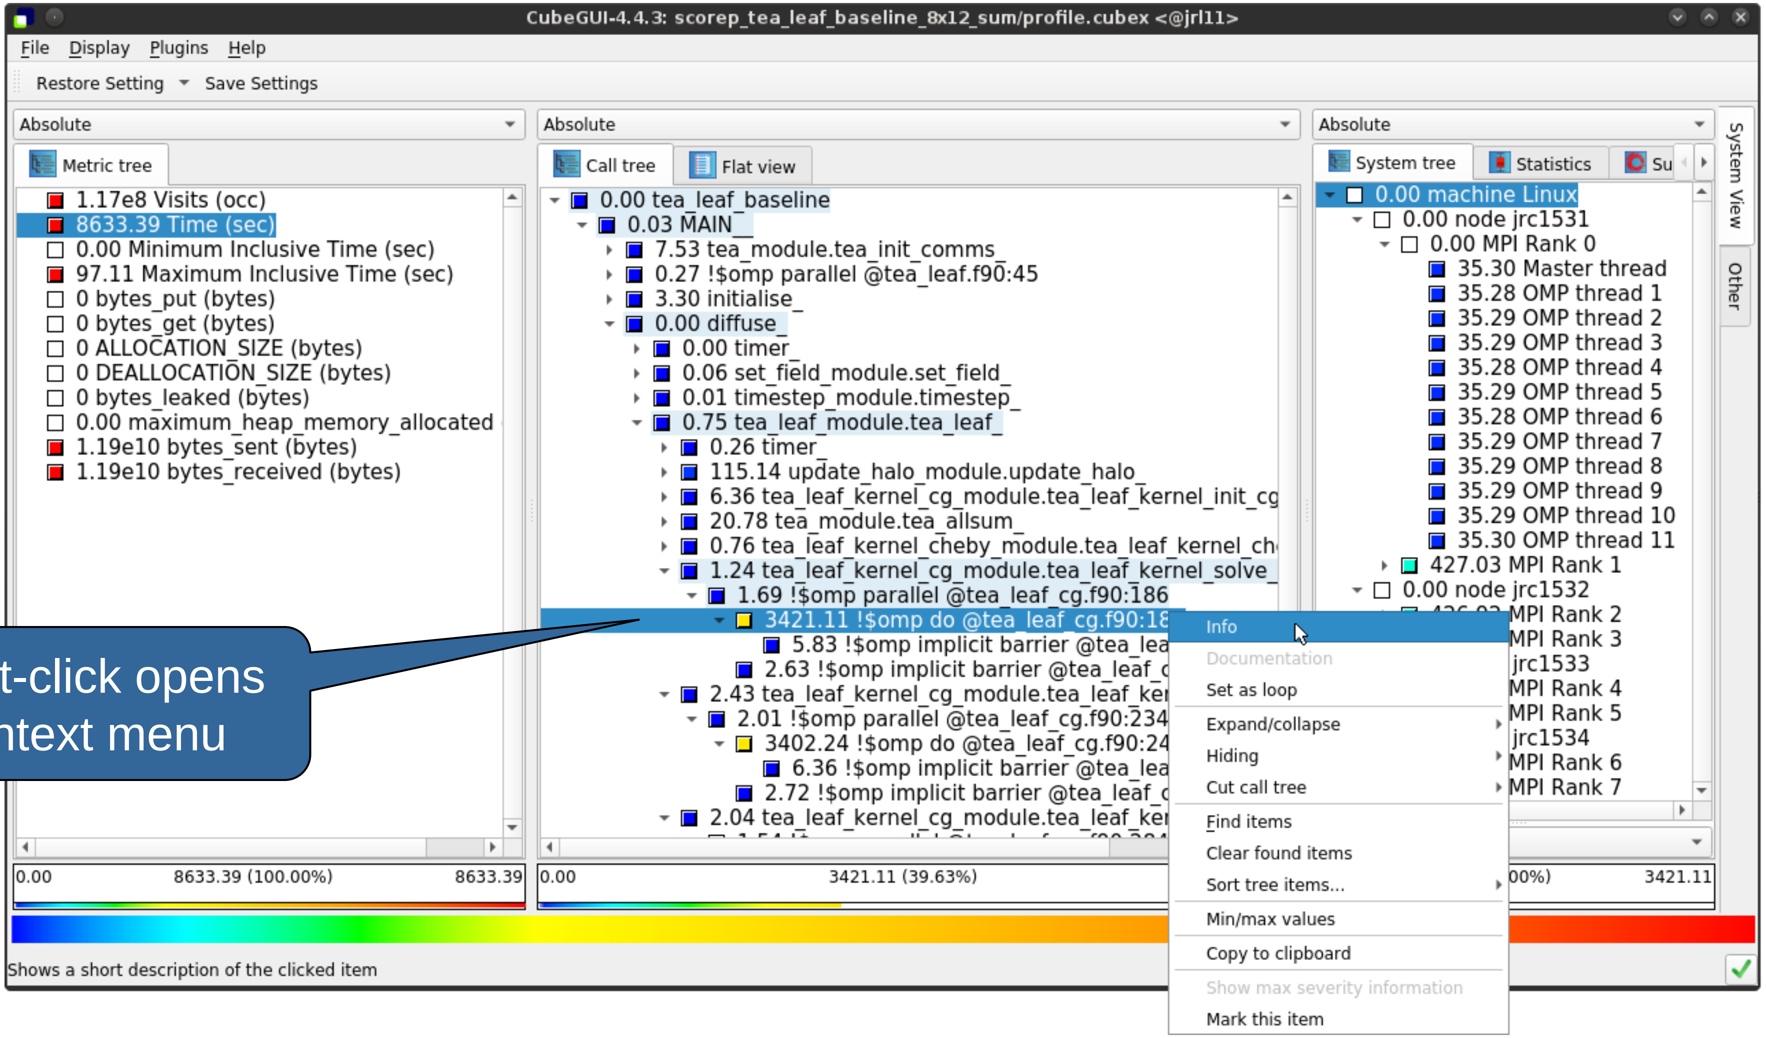

CubeGUI-4.4.3: scorep_tea_leaf_baseline_8x12_sum/profile.cubex <@jrl11>

```
1.17e8 Visits (occ)
8633.39 Time (sec)
0.00 Minimum Inclusive Time (sec)
97.11 Maximum Inclusive Time (sec)
0 bytes_put (bytes)
0 bytes_get (bytes)
0 ALLOCATION_SIZE (bytes)
0 DEALLOCATION_SIZE (bytes)
0 bytes_leaked (bytes)
0.00 maximum_heap_memory_allocated
1.19e10 bytes_sent (bytes)
1.19e10 bytes_received (bytes)
```

```
0.00 tea_leaf_baseline
  0.03 MAIN__
    7.53 tea_module.tea_init_comms_
    0.27 !$omp parallel @tea_leaf.f90:45
    3.30 initialise_
    0.00 diffuse_
      0.00 timer_
      0.06 set_field_module.set_field_
      0.01 timestep_module.timestep_
      0.75 tea_leaf_module.tea_leaf_
        0.26 timer_
        115.14 update_halo_module.update_halo_
        6.36 tea_leaf_kernel_cg_module.tea_leaf_kernel_init_cg
        20.78 tea_module.tea_allsum_
        0.76 tea_leaf_kernel_cheby_module.tea_leaf_kernel_ch
        1.24 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_
          1.69 !$omp parallel @tea_leaf_cg.f90:186
            3421.11 !$omp do @tea_leaf_cg.f90:18
              5.83 !$omp implicit barrier @tea_lea
            2.63 !$omp implicit barrier @tea_leaf_c
        2.43 tea_leaf_kernel_cg_module.tea_leaf_ker
          2.01 !$omp parallel @tea_leaf_cg.f90:234
            3402.24 !$omp do @tea_leaf_cg.f90:24
              6.36 !$omp implicit barrier @tea_lea
            2.72 !$omp implicit barrier @tea_leaf_c
        2.04 tea_leaf_kernel_cg_module.tea_leaf_ker
```

```
0.00 machine Linux
  0.00 node jrc1531
    0.00 MPI Rank 0
      35.30 Master thread
      35.28 OMP thread 1
      35.29 OMP thread 2
      35.29 OMP thread 3
      35.28 OMP thread 4
      35.29 OMP thread 5
      35.28 OMP thread 6
      35.29 OMP thread 7
      35.29 OMP thread 8
      35.29 OMP thread 9
      35.29 OMP thread 10
      35.30 OMP thread 11
    427.03 MPI Rank 1
  0.00 node jrc1532
    MPI Rank 2
    MPI Rank 3
  jrc1533
    MPI Rank 4
    MPI Rank 5
  jrc1534
    MPI Rank 6
    MPI Rank 7
```

0.00 8633.39 (100.00%) 8633.39

0.00 3421.11 (39.63%)

00%) 3421.11

Shows a short description of the clicked item

Right-click opens context menu

Source-code view

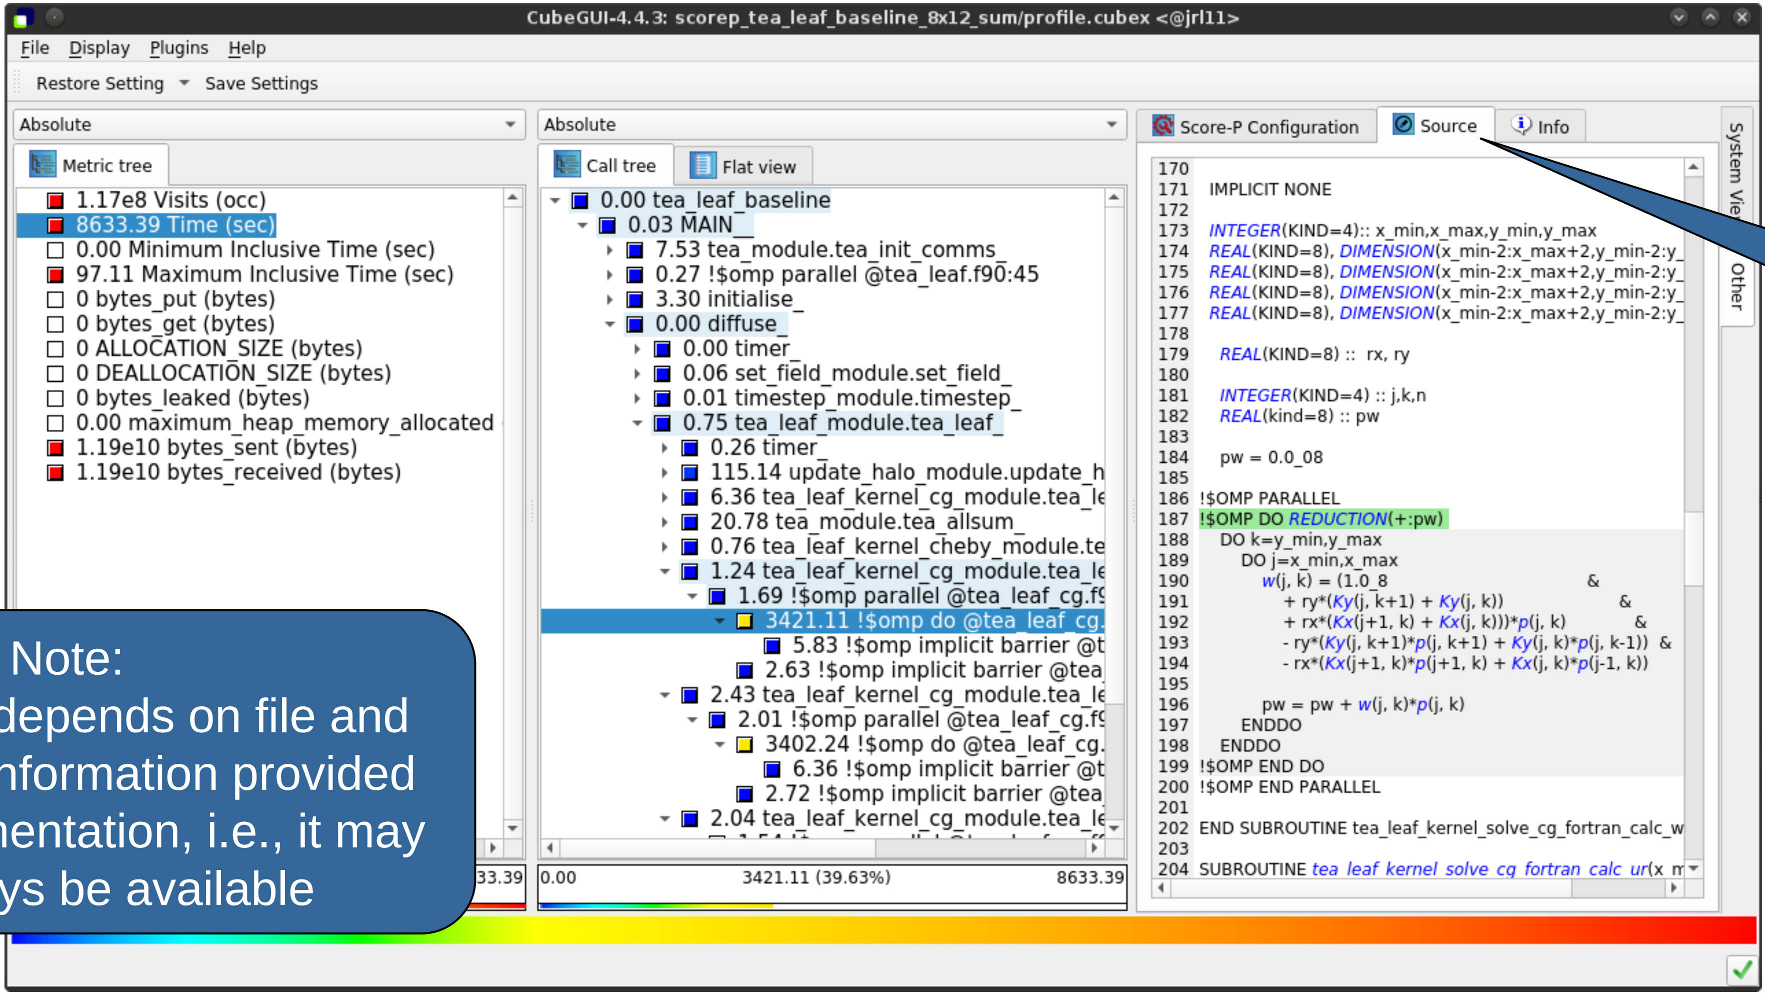

```
1.17e8 Visits (occ)
8633.39 Time (sec)
0.00 Minimum Inclusive Time (sec)
97.11 Maximum Inclusive Time (sec)
0 bytes_put (bytes)
0 bytes_get (bytes)
0 ALLOCATION_SIZE (bytes)
0 DEALLOCATION_SIZE (bytes)
0 bytes_leaked (bytes)
0.00 maximum_heap_memory_allocated
1.19e10 bytes_sent (bytes)
1.19e10 bytes_received (bytes)
```

```
0.00 tea_leaf_baseline
  0.03 MAIN_
    7.53 tea_module.tea_init_comms_
    0.27 !$omp parallel @tea_leaf.f90:45
    3.30 initialise_
    0.00 diffuse_
      0.00 timer_
      0.06 set_field_module.set_field_
      0.01 timestep_module.timestep_
      0.75 tea_leaf_module.tea_leaf_
        0.26 timer_
        115.14 update_halo_module.update_h
        6.36 tea_leaf_kernel_cg_module.tea_le
        20.78 tea_module.tea_allsum_
        0.76 tea_leaf_kernel_cheby_module.te
        1.24 tea_leaf_kernel_cg_module.tea_le
          1.69 !$omp parallel @tea_leaf_cg.f9
            3421.11 !$omp do @tea_leaf_cg.
              5.83 !$omp implicit barrier @t
            2.63 !$omp implicit barrier @tea
        2.43 tea_leaf_kernel_cg_module.tea_le
          2.01 !$omp parallel @tea_leaf_cg.f9
            3402.24 !$omp do @tea_leaf_cg.
              6.36 !$omp implicit barrier @t
            2.72 !$omp implicit barrier @tea
        2.04 tea_leaf_kernel_cg_module.tea_le
```

```
33.39
0.00        3421.11 (39.63%)        8633.39
```

```
  IMPLICIT NONE

  INTEGER(KIND=4):: x_min,x_max,y_min,y_max
  REAL(KIND=8), DIMENSION(x_min-2:x_max+2,y_min-2:y_
  REAL(KIND=8), DIMENSION(x_min-2:x_max+2,y_min-2:y_
  REAL(KIND=8), DIMENSION(x_min-2:x_max+2,y_min-2:y_
  REAL(KIND=8), DIMENSION(x_min-2:x_max+2,y_min-2:y_

   REAL(KIND=8) ::  rx, ry

   INTEGER(KIND=4) :: j,k,n
   REAL(kind=8) :: pw

   pw = 0.0_08

!$OMP PARALLEL
!$OMP DO REDUCTION(+:pw)
   DO k=y_min,y_max
     DO j=x_min,x_max
       w(j, k) = (1.0_8                             &
           + ry*(Ky(j, k+1) + Ky(j, k))                    &
           + rx*(Kx(j+1, k) + Kx(j, k)))*p(j, k)          &
           - ry*(Ky(j, k+1)*p(j, k+1) + Ky(j, k)*p(j, k-1))  &
           - rx*(Kx(j+1, k)*p(j+1, k) + Kx(j, k)*p(j-1, k))

       pw = pw + w(j, k)*p(j, k)
     ENDDO
   ENDDO
!$OMP END DO
!$OMP END PARALLEL

END SUBROUTINE tea_leaf_kernel_solve_cg_fortran_calc_w

SUBROUTINE tea_leaf_kernel_solve_cg_fortran_calc_ur(x_m
```

Select “Source” tab

Note:
This feature depends on file and line number information provided by the instrumentation, i.e., it may not always be available

Context-sensitive help

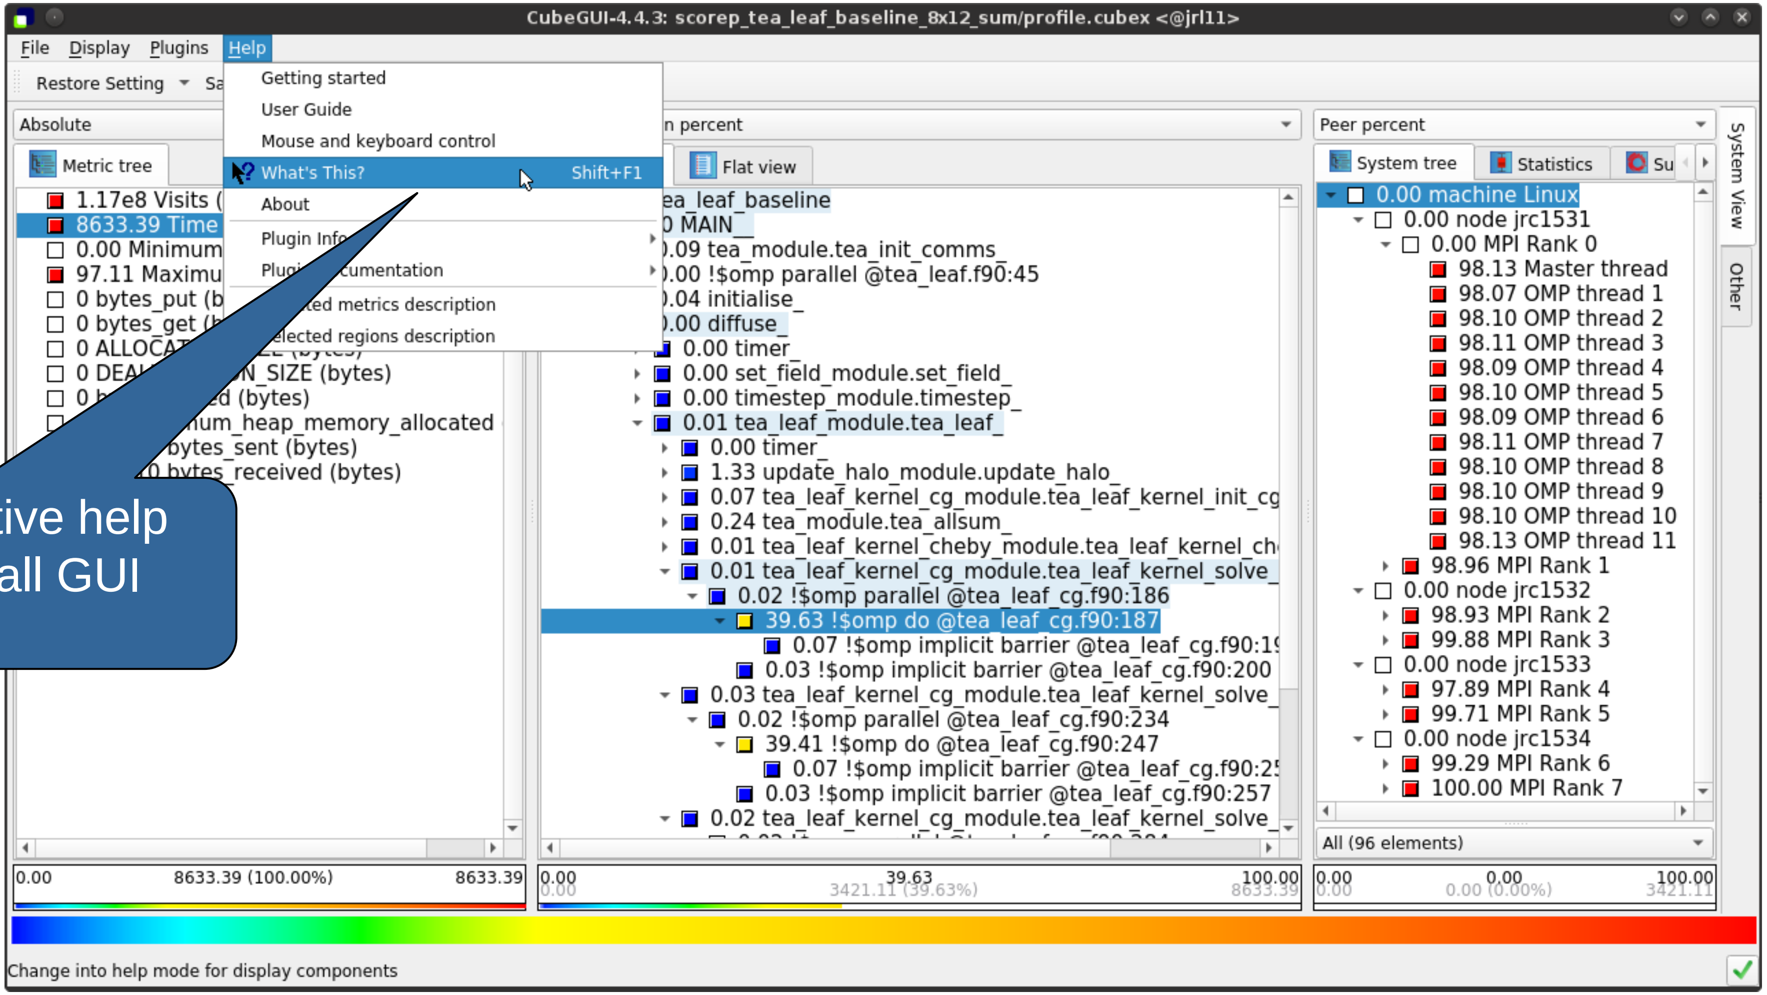

```
1.17e8 Visits (
8633.39 Time
0.00 Minimum
97.11 Maximu
0 bytes_put (b
0 bytes_get (
0 ALLOCA
0 DEAL                N_SIZE (bytes)
0                   d (bytes)
                 mum_heap_memory_allocated
                 bytes_sent (bytes)
               0 bytes_received (bytes)
```

```
ea_leaf_baseline
0 MAIN_
.09 tea_module.tea_init_comms_
.00 !$omp parallel @tea_leaf.f90:45
.04 initialise_
.00 diffuse_
  0.00 timer_
  0.00 set_field_module.set_field_
  0.00 timestep_module.timestep_
  0.01 tea_leaf_module.tea_leaf_
    0.00 timer_
    1.33 update_halo_module.update_halo_
    0.07 tea_leaf_kernel_cg_module.tea_leaf_kernel_init_cg
    0.24 tea_module.tea_allsum_
    0.01 tea_leaf_kernel_cheby_module.tea_leaf_kernel_ch
    0.01 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_
      0.02 !$omp parallel @tea_leaf_cg.f90:186
        39.63 !$omp do @tea_leaf_cg.f90:187
          0.07 !$omp implicit barrier @tea_leaf_cg.f90:19
        0.03 !$omp implicit barrier @tea_leaf_cg.f90:200
    0.03 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_
      0.02 !$omp parallel @tea_leaf_cg.f90:234
        39.41 !$omp do @tea_leaf_cg.f90:247
          0.07 !$omp implicit barrier @tea_leaf_cg.f90:25
        0.03 !$omp implicit barrier @tea_leaf_cg.f90:257
    0.02 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_
```

```
0.00 machine Linux
  0.00 node jrc1531
    0.00 MPI Rank 0
      98.13 Master thread
      98.07 OMP thread 1
      98.10 OMP thread 2
      98.11 OMP thread 3
      98.09 OMP thread 4
      98.10 OMP thread 5
      98.09 OMP thread 6
      98.11 OMP thread 7
      98.10 OMP thread 8
      98.10 OMP thread 9
      98.10 OMP thread 10
      98.13 OMP thread 11
    98.96 MPI Rank 1
  0.00 node jrc1532
    98.93 MPI Rank 2
    99.88 MPI Rank 3
  0.00 node jrc1533
    97.89 MPI Rank 4
    99.71 MPI Rank 5
  0.00 node jrc1534
    99.29 MPI Rank 6
    100.00 MPI Rank 7
```

```
0.00        8633.39 (100.00%)        8633.39
0.00        39.63                     100.00
0.00        3421.11 (39.63%)          8633.39
0.00        0.00                      100.00
0.00        0.00 (0.00%)              3421.11
```

Change into help mode for display components

Context-sensitive help available for all GUI items

Scalasca report post-processing

- Scalasca’s report post-processing derives additional metrics and generates a structured metric hierarchy

- Automatically run (if needed) when using the **square** convenience command:

```
% square scorep_tea_leaf_baseline_8x12_sum
INFO: Post-processing runtime summarization report (profile.cubex)...
INFO: Displaying ./scorep_tea_leaf_baseline_8x12_sum/summary.cubex...

                    [GUI showing post-processed summary analysis report]
```

Post-processed summary analysis report

Split base metrics into more specific metrics

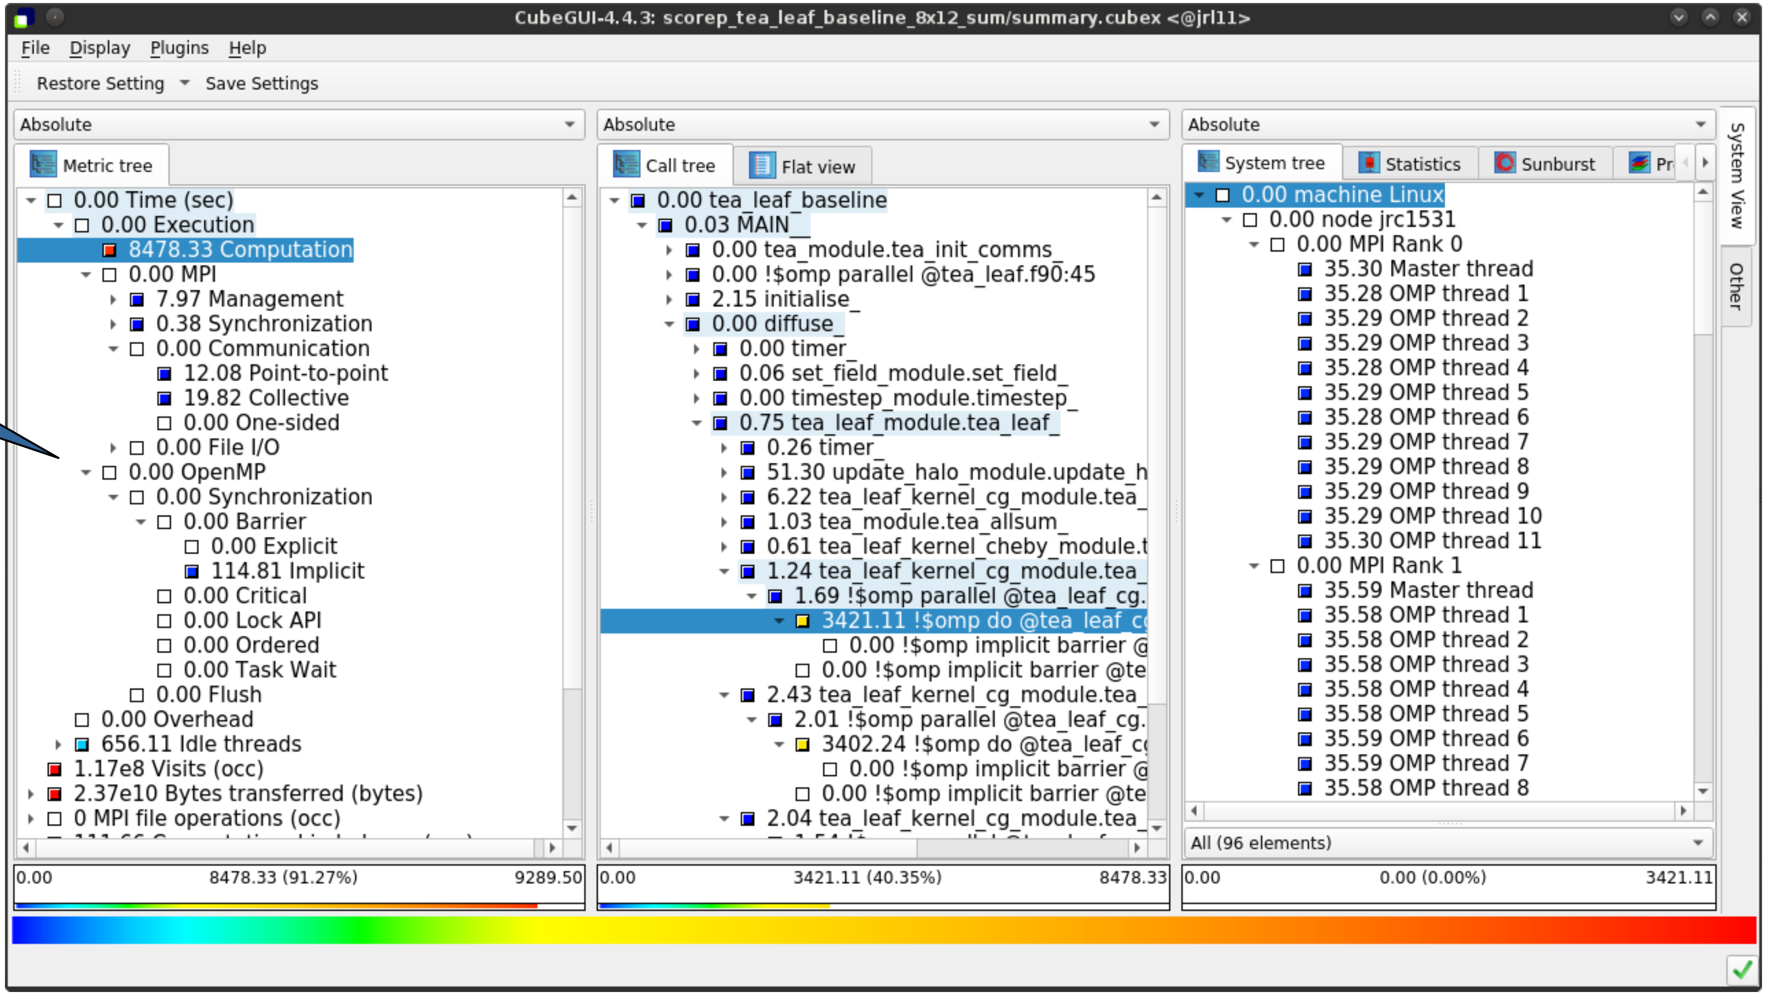

CubeGUI-4.4.3: scorep_tea_leaf_baseline_8x12_sum/summary.cubex <@jrl11>

```
□ 0.00 Time (sec)
  □ 0.00 Execution
    ■ 8478.33 Computation
    □ 0.00 MPI
      ■ 7.97 Management
      ■ 0.38 Synchronization
      □ 0.00 Communication
        ■ 12.08 Point-to-point
        ■ 19.82 Collective
        □ 0.00 One-sided
      □ 0.00 File I/O
    □ 0.00 OpenMP
      □ 0.00 Synchronization
        □ 0.00 Barrier
          □ 0.00 Explicit
          ■ 114.81 Implicit
        □ 0.00 Critical
        □ 0.00 Lock API
        □ 0.00 Ordered
        □ 0.00 Task Wait
      □ 0.00 Flush
  □ 0.00 Overhead
  ■ 656.11 Idle threads
■ 1.17e8 Visits (occ)
■ 2.37e10 Bytes transferred (bytes)
□ 0 MPI file operations (occ)
```

```
0.00        8478.33 (91.27%)        9289.50
```

```
■ 0.00 tea_leaf_baseline
  ■ 0.03 MAIN__
    ■ 0.00 tea_module.tea_init_comms_
    ■ 0.00 !$omp parallel @tea_leaf.f90:45
    ■ 2.15 initialise_
    ■ 0.00 diffuse_
      ■ 0.00 timer_
      ■ 0.06 set_field_module.set_field_
      ■ 0.00 timestep_module.timestep_
      ■ 0.75 tea_leaf_module.tea_leaf_
        ■ 0.26 timer_
        ■ 51.30 update_halo_module.update_h
        ■ 6.22 tea_leaf_kernel_cg_module.tea_
        ■ 1.03 tea_module.tea_allsum_
        ■ 0.61 tea_leaf_kernel_cheby_module.t
        ■ 1.24 tea_leaf_kernel_cg_module.tea_
          ■ 1.69 !$omp parallel @tea_leaf_cg.
            ■ 3421.11 !$omp do @tea_leaf_c
              □ 0.00 !$omp implicit barrier @
            □ 0.00 !$omp implicit barrier @te
        ■ 2.43 tea_leaf_kernel_cg_module.tea_
          ■ 2.01 !$omp parallel @tea_leaf_cg.
            ■ 3402.24 !$omp do @tea_leaf_c
              □ 0.00 !$omp implicit barrier @
            □ 0.00 !$omp implicit barrier @te
        ■ 2.04 tea_leaf_kernel_cg_module.tea_
```

```
0.00        3421.11 (40.35%)        8478.33
```

```
□ 0.00 machine Linux
  □ 0.00 node jrc1531
    □ 0.00 MPI Rank 0
      ■ 35.30 Master thread
      ■ 35.28 OMP thread 1
      ■ 35.29 OMP thread 2
      ■ 35.29 OMP thread 3
      ■ 35.28 OMP thread 4
      ■ 35.29 OMP thread 5
      ■ 35.28 OMP thread 6
      ■ 35.29 OMP thread 7
      ■ 35.29 OMP thread 8
      ■ 35.29 OMP thread 9
      ■ 35.29 OMP thread 10
      ■ 35.30 OMP thread 11
    □ 0.00 MPI Rank 1
      ■ 35.59 Master thread
      ■ 35.58 OMP thread 1
      ■ 35.58 OMP thread 2
      ■ 35.58 OMP thread 3
      ■ 35.58 OMP thread 4
      ■ 35.58 OMP thread 5
      ■ 35.59 OMP thread 6
      ■ 35.59 OMP thread 7
      ■ 35.58 OMP thread 8
```

```
0.00        0.00 (0.00%)        3421.11
```

TeaLeaf summary report analysis (I)

91% of the execution time is computation…

…almost entirely spent in 3 OpenMP do loops…

…with a slight imbalance across ranks & threads

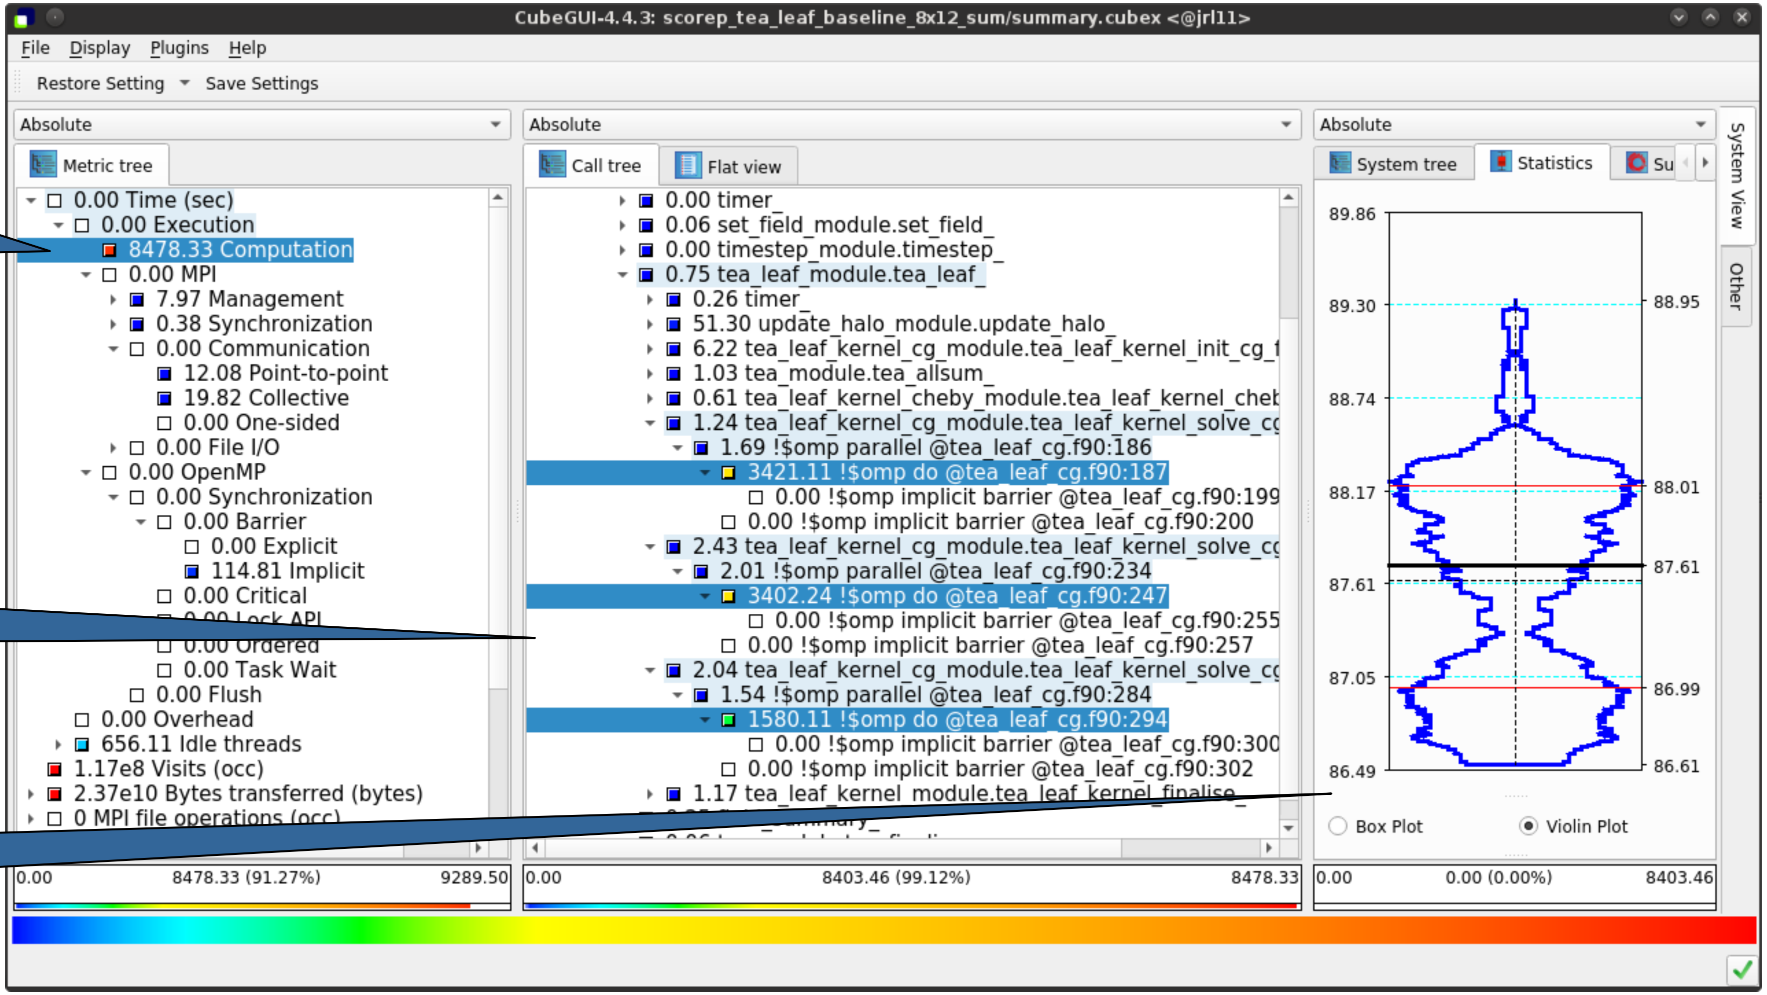

CubeGUI-4.4.3: scorep_tea_leaf_baseline_8x12_sum/summary.cubex <@jrl11>

```
▾ □ 0.00 Time (sec)
  ▾ □ 0.00 Execution
      ■ 8478.33 Computation
    ▾ □ 0.00 MPI
      ▸ ■ 7.97 Management
      ▸ ■ 0.38 Synchronization
      ▾ □ 0.00 Communication
          ■ 12.08 Point-to-point
          ■ 19.82 Collective
          □ 0.00 One-sided
      ▸ □ 0.00 File I/O
    ▾ □ 0.00 OpenMP
      ▾ □ 0.00 Synchronization
        ▾ □ 0.00 Barrier
            □ 0.00 Explicit
            ■ 114.81 Implicit
          □ 0.00 Critical
          □ 0.00 Lock API
          □ 0.00 Ordered
          □ 0.00 Task Wait
        □ 0.00 Flush
    □ 0.00 Overhead
  ▸ ■ 656.11 Idle threads
■ 1.17e8 Visits (occ)
▸ ■ 2.37e10 Bytes transferred (bytes)
▸ □ 0 MPI file operations (occ)
```

```
▸ ■ 0.00 timer_
▸ ■ 0.06 set_field_module.set_field_
▸ ■ 0.00 timestep_module.timestep_
▾ ■ 0.75 tea_leaf_module.tea_leaf_
  ▸ ■ 0.26 timer_
  ▸ ■ 51.30 update_halo_module.update_halo_
  ▸ ■ 6.22 tea_leaf_kernel_cg_module.tea_leaf_kernel_init_cg_f
  ▸ ■ 1.03 tea_module.tea_allsum_
  ▸ ■ 0.61 tea_leaf_kernel_cheby_module.tea_leaf_kernel_cheb
  ▾ ■ 1.24 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_cg
    ▾ ■ 1.69 !$omp parallel @tea_leaf_cg.f90:186
      ▾ ■ 3421.11 !$omp do @tea_leaf_cg.f90:187
          □ 0.00 !$omp implicit barrier @tea_leaf_cg.f90:199
        □ 0.00 !$omp implicit barrier @tea_leaf_cg.f90:200
  ▾ ■ 2.43 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_cg
    ▾ ■ 2.01 !$omp parallel @tea_leaf_cg.f90:234
      ▾ ■ 3402.24 !$omp do @tea_leaf_cg.f90:247
          □ 0.00 !$omp implicit barrier @tea_leaf_cg.f90:255
        □ 0.00 !$omp implicit barrier @tea_leaf_cg.f90:257
  ▾ ■ 2.04 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_cg
    ▾ ■ 1.54 !$omp parallel @tea_leaf_cg.f90:284
      ▾ ■ 1580.11 !$omp do @tea_leaf_cg.f90:294
          □ 0.00 !$omp implicit barrier @tea_leaf_cg.f90:300
        □ 0.00 !$omp implicit barrier @tea_leaf_cg.f90:302
  ▸ ■ 1.17 tea_leaf_kernel_module.tea_leaf_kernel_finalise_
```

0.00	8478.33 (91.27%)	9289.50
0.00	8403.46 (99.12%)	8478.33
0.00	0.00 (0.00%)	8403.46

TeaLeaf summary report analysis (II)

7% of the execution time are lost due to idle threads…

…in non-OpenMP parallelized code regions

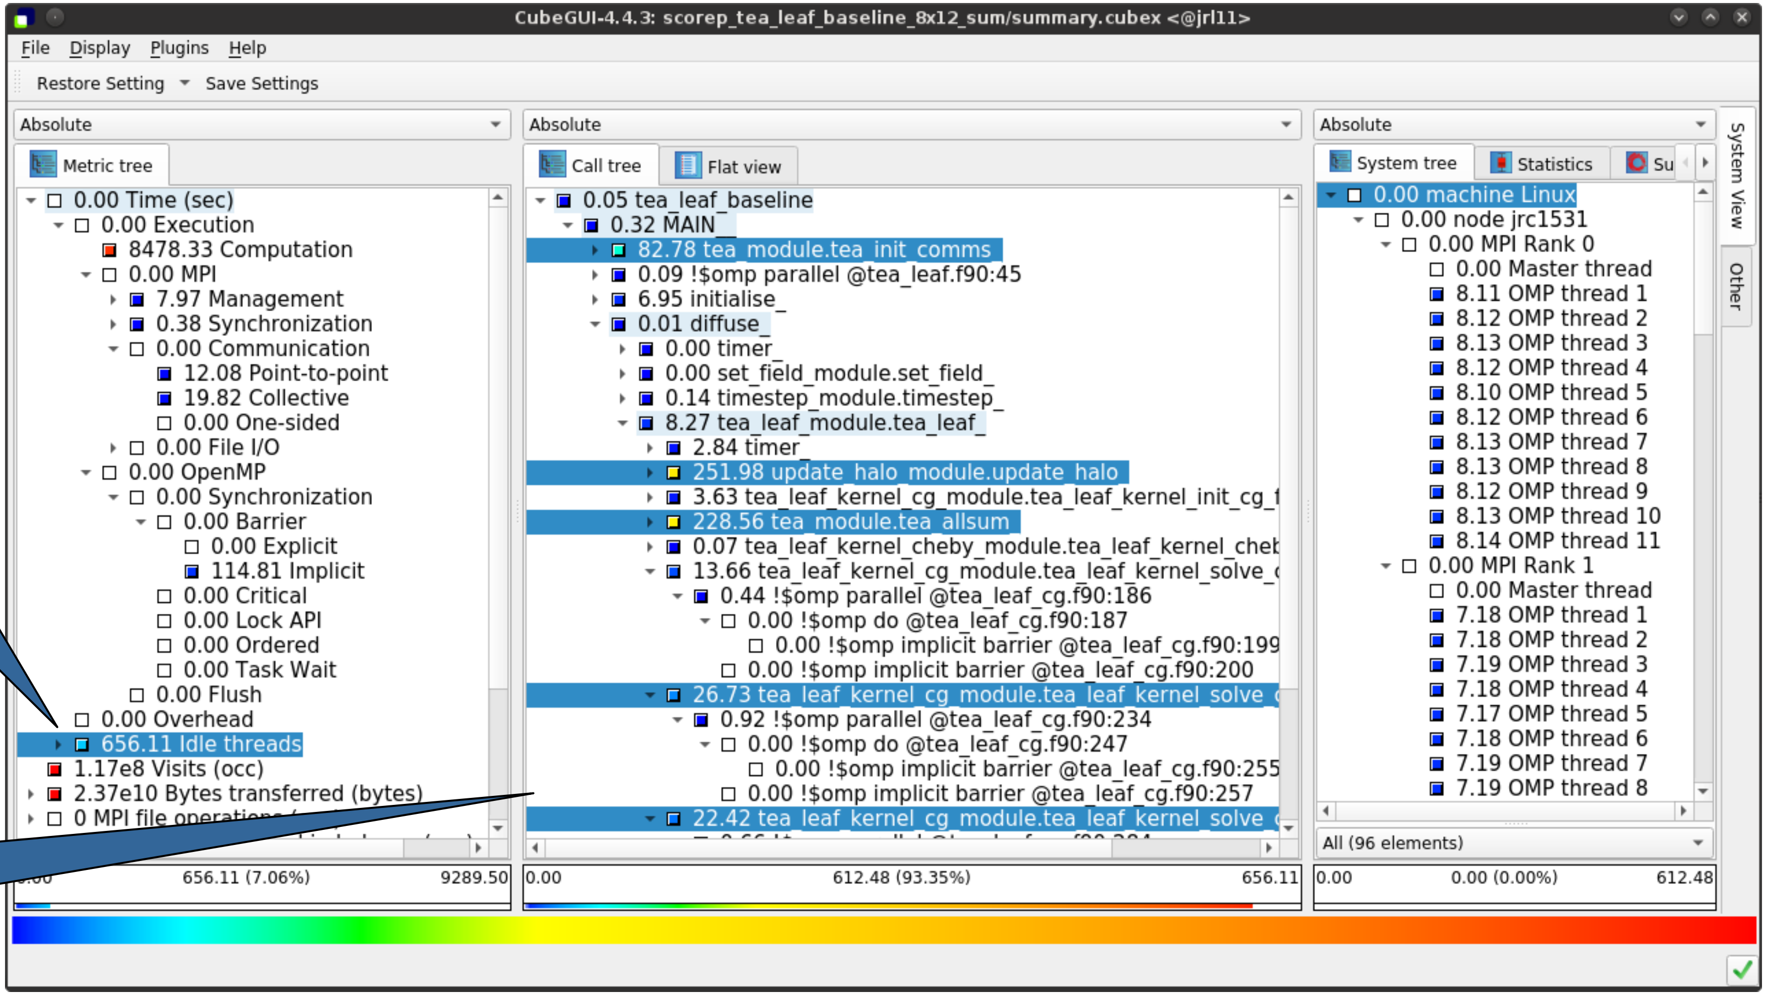
CubeGUI-4.4.3: scorep_tea_leaf_baseline_8x12_sum/summary.cubex <@jrl11>
File Display Plugins Help
Restore Setting Save Settings
Absolute
Metric tree
Absolute
Call tree
Flat view
Absolute
System tree
Statistics
Su
System View
Other
All (96 elements)

```
Metric tree
□ 0.00 Time (sec)
  □ 0.00 Execution
    ■ 8478.33 Computation
    □ 0.00 MPI
      ■ 7.97 Management
      ■ 0.38 Synchronization
      □ 0.00 Communication
        ■ 12.08 Point-to-point
        ■ 19.82 Collective
        □ 0.00 One-sided
      □ 0.00 File I/O
    □ 0.00 OpenMP
      □ 0.00 Synchronization
        □ 0.00 Barrier
          □ 0.00 Explicit
          ■ 114.81 Implicit
        □ 0.00 Critical
        □ 0.00 Lock API
        □ 0.00 Ordered
        □ 0.00 Task Wait
      □ 0.00 Flush
  □ 0.00 Overhead
  ■ 656.11 Idle threads
■ 1.17e8 Visits (occ)
■ 2.37e10 Bytes transferred (bytes)
□ 0 MPI file operations

0.00        656.11 (7.06%)        9289.50
```

```
Call tree
■ 0.05 tea_leaf_baseline
  ■ 0.32 MAIN_
    ■ 82.78 tea_module.tea_init_comms_
    ■ 0.09 !$omp parallel @tea_leaf.f90:45
    ■ 6.95 initialise_
    ■ 0.01 diffuse_
      ■ 0.00 timer_
      ■ 0.00 set_field_module.set_field_
      ■ 0.14 timestep_module.timestep_
      ■ 8.27 tea_leaf_module.tea_leaf_
        ■ 2.84 timer_
        ■ 251.98 update_halo_module.update_halo_
        ■ 3.63 tea_leaf_kernel_cg_module.tea_leaf_kernel_init_cg_f
        ■ 228.56 tea_module.tea_allsum_
        ■ 0.07 tea_leaf_kernel_cheby_module.tea_leaf_kernel_cheb
        ■ 13.66 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_c
          ■ 0.44 !$omp parallel @tea_leaf_cg.f90:186
            □ 0.00 !$omp do @tea_leaf_cg.f90:187
              □ 0.00 !$omp implicit barrier @tea_leaf_cg.f90:199
            □ 0.00 !$omp implicit barrier @tea_leaf_cg.f90:200
        ■ 26.73 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_c
          ■ 0.92 !$omp parallel @tea_leaf_cg.f90:234
            □ 0.00 !$omp do @tea_leaf_cg.f90:247
              □ 0.00 !$omp implicit barrier @tea_leaf_cg.f90:255
            □ 0.00 !$omp implicit barrier @tea_leaf_cg.f90:257
        ■ 22.42 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_c

0.00        612.48 (93.35%)        656.11
```

```
System tree
□ 0.00 machine Linux
  □ 0.00 node jrc1531
    □ 0.00 MPI Rank 0
      □ 0.00 Master thread
      ■ 8.11 OMP thread 1
      ■ 8.12 OMP thread 2
      ■ 8.13 OMP thread 3
      ■ 8.12 OMP thread 4
      ■ 8.10 OMP thread 5
      ■ 8.12 OMP thread 6
      ■ 8.13 OMP thread 7
      ■ 8.13 OMP thread 8
      ■ 8.12 OMP thread 9
      ■ 8.13 OMP thread 10
      ■ 8.14 OMP thread 11
    □ 0.00 MPI Rank 1
      □ 0.00 Master thread
      ■ 7.18 OMP thread 1
      ■ 7.18 OMP thread 2
      ■ 7.19 OMP thread 3
      ■ 7.18 OMP thread 4
      ■ 7.17 OMP thread 5
      ■ 7.18 OMP thread 6
      ■ 7.19 OMP thread 7
      ■ 7.19 OMP thread 8

0.00        0.00 (0.00%)        612.48
```

TeaLeaf summary report analysis (III)

MPI communication time is negligible (0.34%); communication is only on the master threads (MPI_THREAD_FUNNELED)

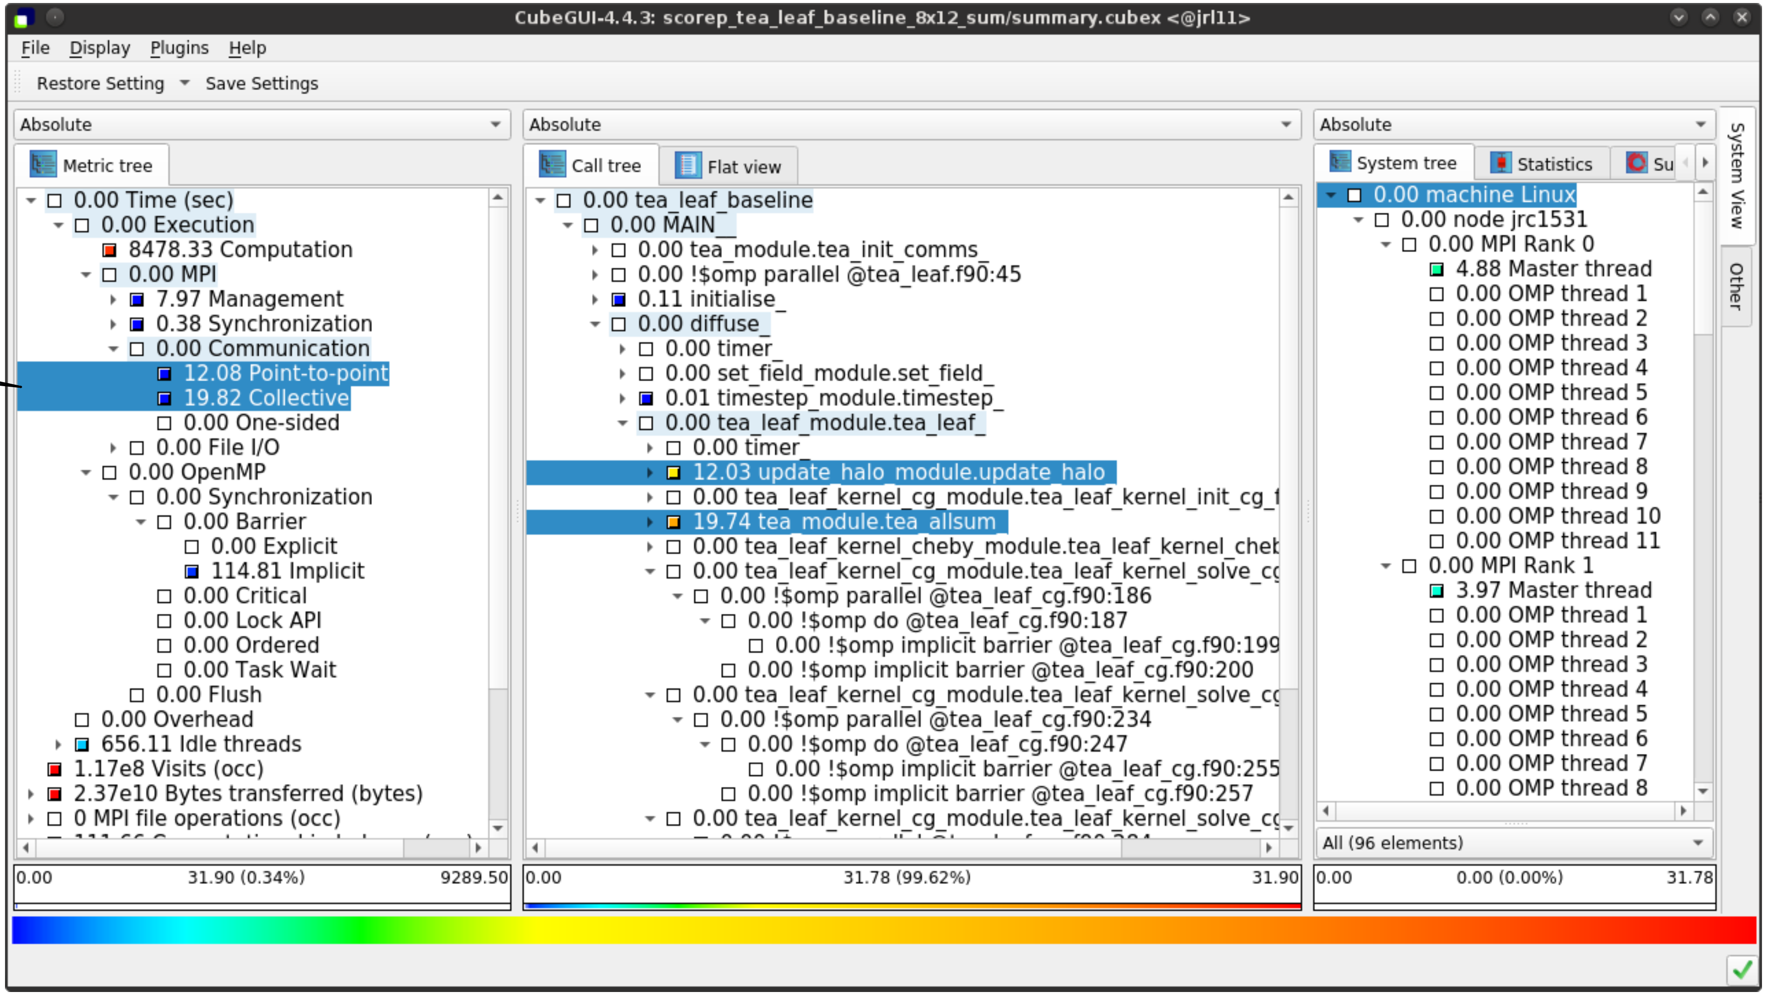

CubeGUI-4.4.3: scorep_tea_leaf_baseline_8x12_sum/summary.cubex <@jrl11>

```
▾ □ 0.00 Time (sec)
  ▾ □ 0.00 Execution
      ■ 8478.33 Computation
    ▾ □ 0.00 MPI
      ▸ ■ 7.97 Management
      ▸ ■ 0.38 Synchronization
      ▾ □ 0.00 Communication
          ■ 12.08 Point-to-point
          ■ 19.82 Collective
          □ 0.00 One-sided
      ▸ □ 0.00 File I/O
    ▾ □ 0.00 OpenMP
      ▾ □ 0.00 Synchronization
        ▾ □ 0.00 Barrier
            □ 0.00 Explicit
            ■ 114.81 Implicit
          □ 0.00 Critical
          □ 0.00 Lock API
          □ 0.00 Ordered
          □ 0.00 Task Wait
        □ 0.00 Flush
    □ 0.00 Overhead
  ▸ ■ 656.11 Idle threads
■ 1.17e8 Visits (occ)
▸ ■ 2.37e10 Bytes transferred (bytes)
▸ □ 0 MPI file operations (occ)

0.00    31.90 (0.34%)    9289.50
```

```
▾ □ 0.00 tea_leaf_baseline
  ▾ □ 0.00 MAIN__
    ▸ □ 0.00 tea_module.tea_init_comms_
    ▸ □ 0.00 !$omp parallel @tea_leaf.f90:45
    ▸ ■ 0.11 initialise_
    ▾ □ 0.00 diffuse_
      ▸ □ 0.00 timer_
      ▸ □ 0.00 set_field_module.set_field_
      ▸ ■ 0.01 timestep_module.timestep_
      ▾ □ 0.00 tea_leaf_module.tea_leaf_
        ▸ □ 0.00 timer_
        ▸ ■ 12.03 update_halo_module.update_halo
        ▸ □ 0.00 tea_leaf_kernel_cg_module.tea_leaf_kernel_init_cg_f
        ▸ ■ 19.74 tea_module.tea_allsum
        ▸ □ 0.00 tea_leaf_kernel_cheby_module.tea_leaf_kernel_chet
        ▾ □ 0.00 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_cg
          ▾ □ 0.00 !$omp parallel @tea_leaf_cg.f90:186
            ▾ □ 0.00 !$omp do @tea_leaf_cg.f90:187
                □ 0.00 !$omp implicit barrier @tea_leaf_cg.f90:199
              □ 0.00 !$omp implicit barrier @tea_leaf_cg.f90:200
        ▾ □ 0.00 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_cg
          ▾ □ 0.00 !$omp parallel @tea_leaf_cg.f90:234
            ▾ □ 0.00 !$omp do @tea_leaf_cg.f90:247
                □ 0.00 !$omp implicit barrier @tea_leaf_cg.f90:255
              □ 0.00 !$omp implicit barrier @tea_leaf_cg.f90:257
        ▾ □ 0.00 tea_leaf_kernel_cg_module.tea_leaf_kernel_solve_cg

0.00    31.78 (99.62%)    31.90
```

```
▾ □ 0.00 machine Linux
  ▾ □ 0.00 node jrc1531
    ▾ □ 0.00 MPI Rank 0
        ■ 4.88 Master thread
        □ 0.00 OMP thread 1
        □ 0.00 OMP thread 2
        □ 0.00 OMP thread 3
        □ 0.00 OMP thread 4
        □ 0.00 OMP thread 5
        □ 0.00 OMP thread 6
        □ 0.00 OMP thread 7
        □ 0.00 OMP thread 8
        □ 0.00 OMP thread 9
        □ 0.00 OMP thread 10
        □ 0.00 OMP thread 11
    ▾ □ 0.00 MPI Rank 1
        ■ 3.97 Master thread
        □ 0.00 OMP thread 1
        □ 0.00 OMP thread 2
        □ 0.00 OMP thread 3
        □ 0.00 OMP thread 4
        □ 0.00 OMP thread 5
        □ 0.00 OMP thread 6
        □ 0.00 OMP thread 7
        □ 0.00 OMP thread 8

All (96 elements)

0.00    0.00 (0.00%)    31.78
```

Cube: Further information

- Parallel program analysis report exploration tools
  - Libraries for Cube report reading & writing
  - Algebra utilities for report processing
  - GUI for interactive analysis exploration
- Available under 3-clause BSD open-source license
- Documentation & sources:
  - http://www.scalasca.org
- User guide also part of installation:
  - <prefix>/share/doc/cubegui/CubeUserGuide.pdf
- Contact:
  - mailto: scalasca@fz-juelich.de

Reference material

Derived metrics

- Derived metrics are defined using CubePL expressions, e.g.:

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

- Values of derived metrics are not stored, but calculated on-the-fly
- Types of derived metrics:
  - Prederived: evaluation of the CubePL expression is performed before aggregation
  - Postderived: evaluation of the CubePL expression is performed after aggregation

- Examples:
  - “Average execution time”: Postderived metric with expression

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

  - “Number of FLOP per second”: Postderived metric with expression

    metric::FLOP()/metric::time()

Derived metrics in Cube GUI

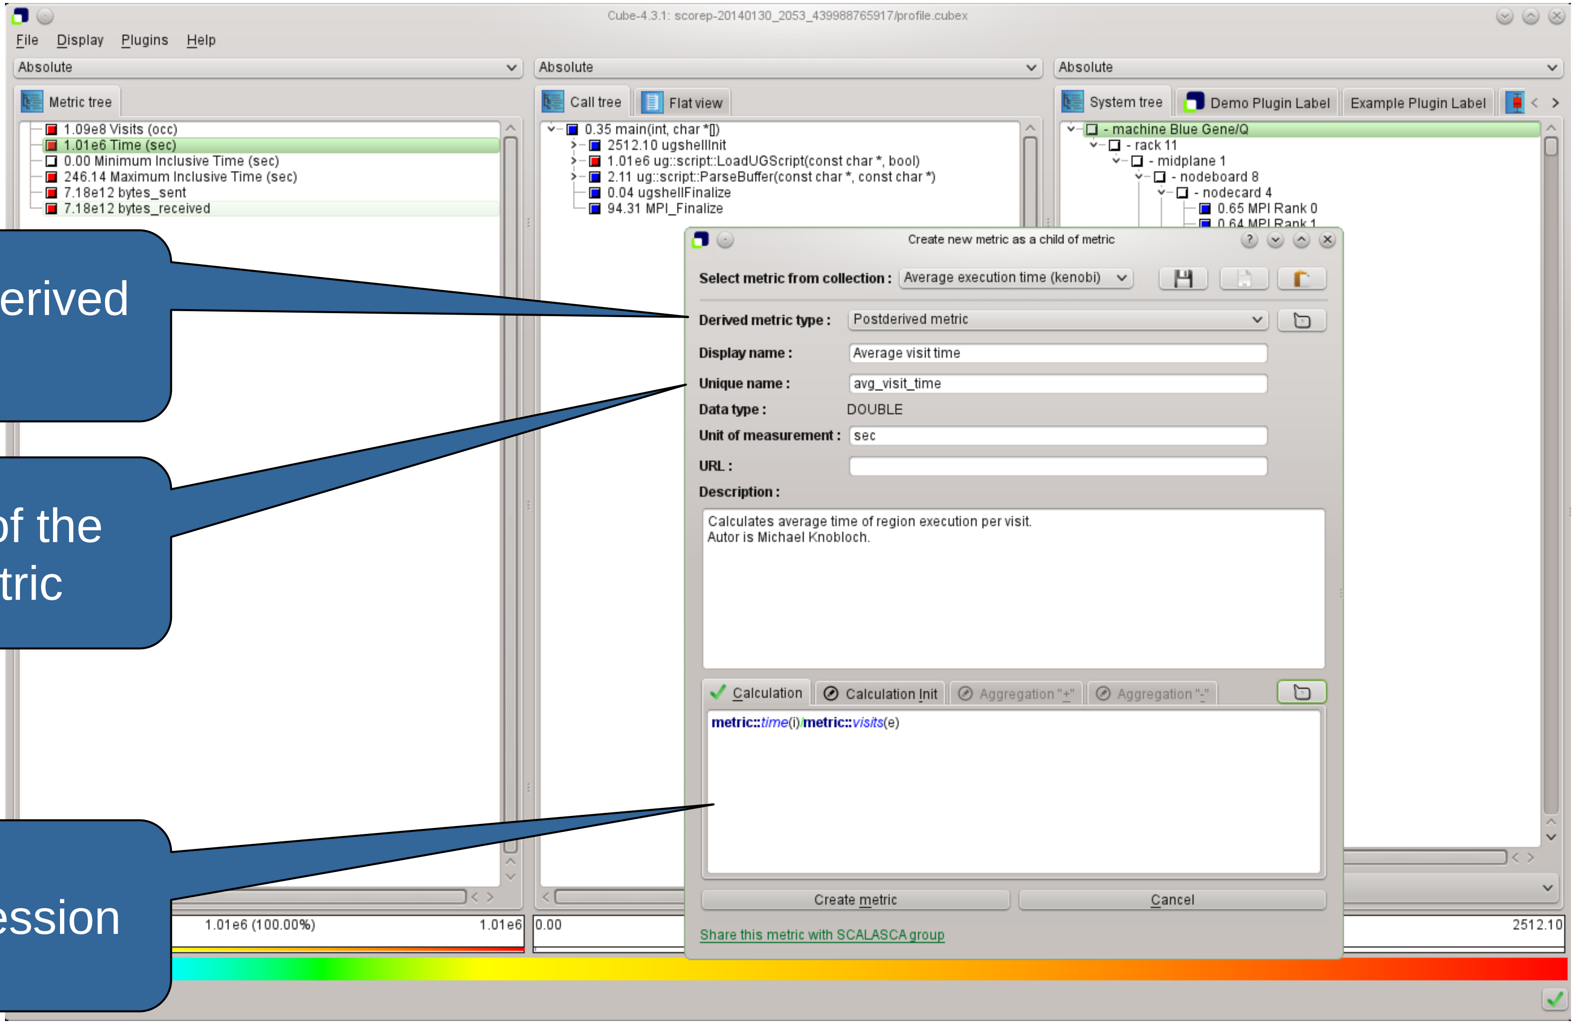

Metric tree:

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

Call tree:

```
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
```

System tree:

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

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: FLOPS based on PAPI_FP_OPS and time

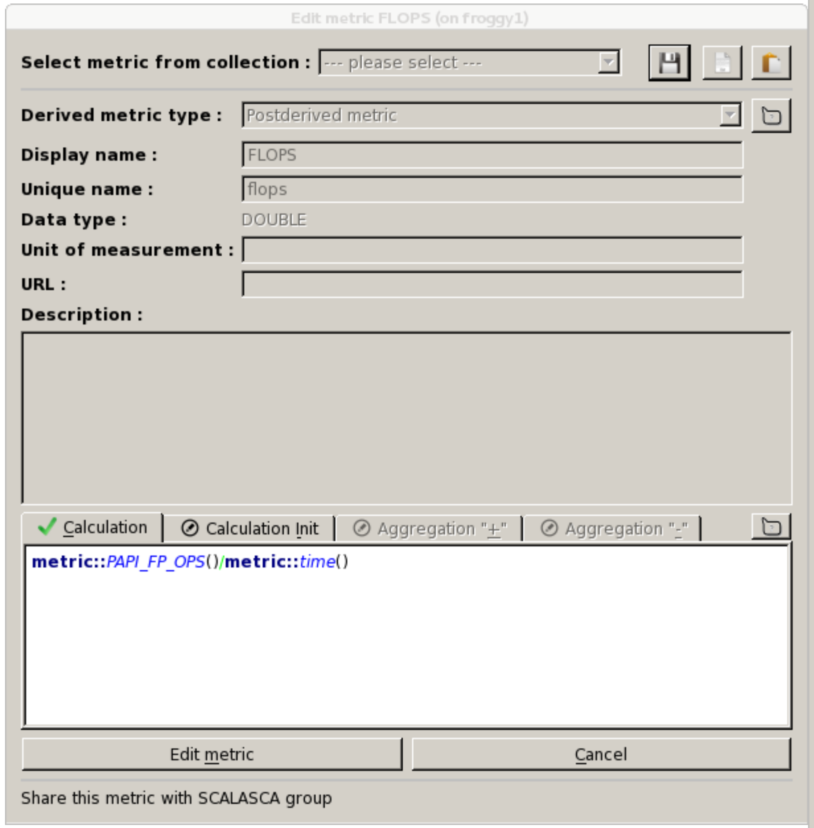

Edit metric FLOPS (on froggy1)

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 :

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

Share this metric with SCALASCA group

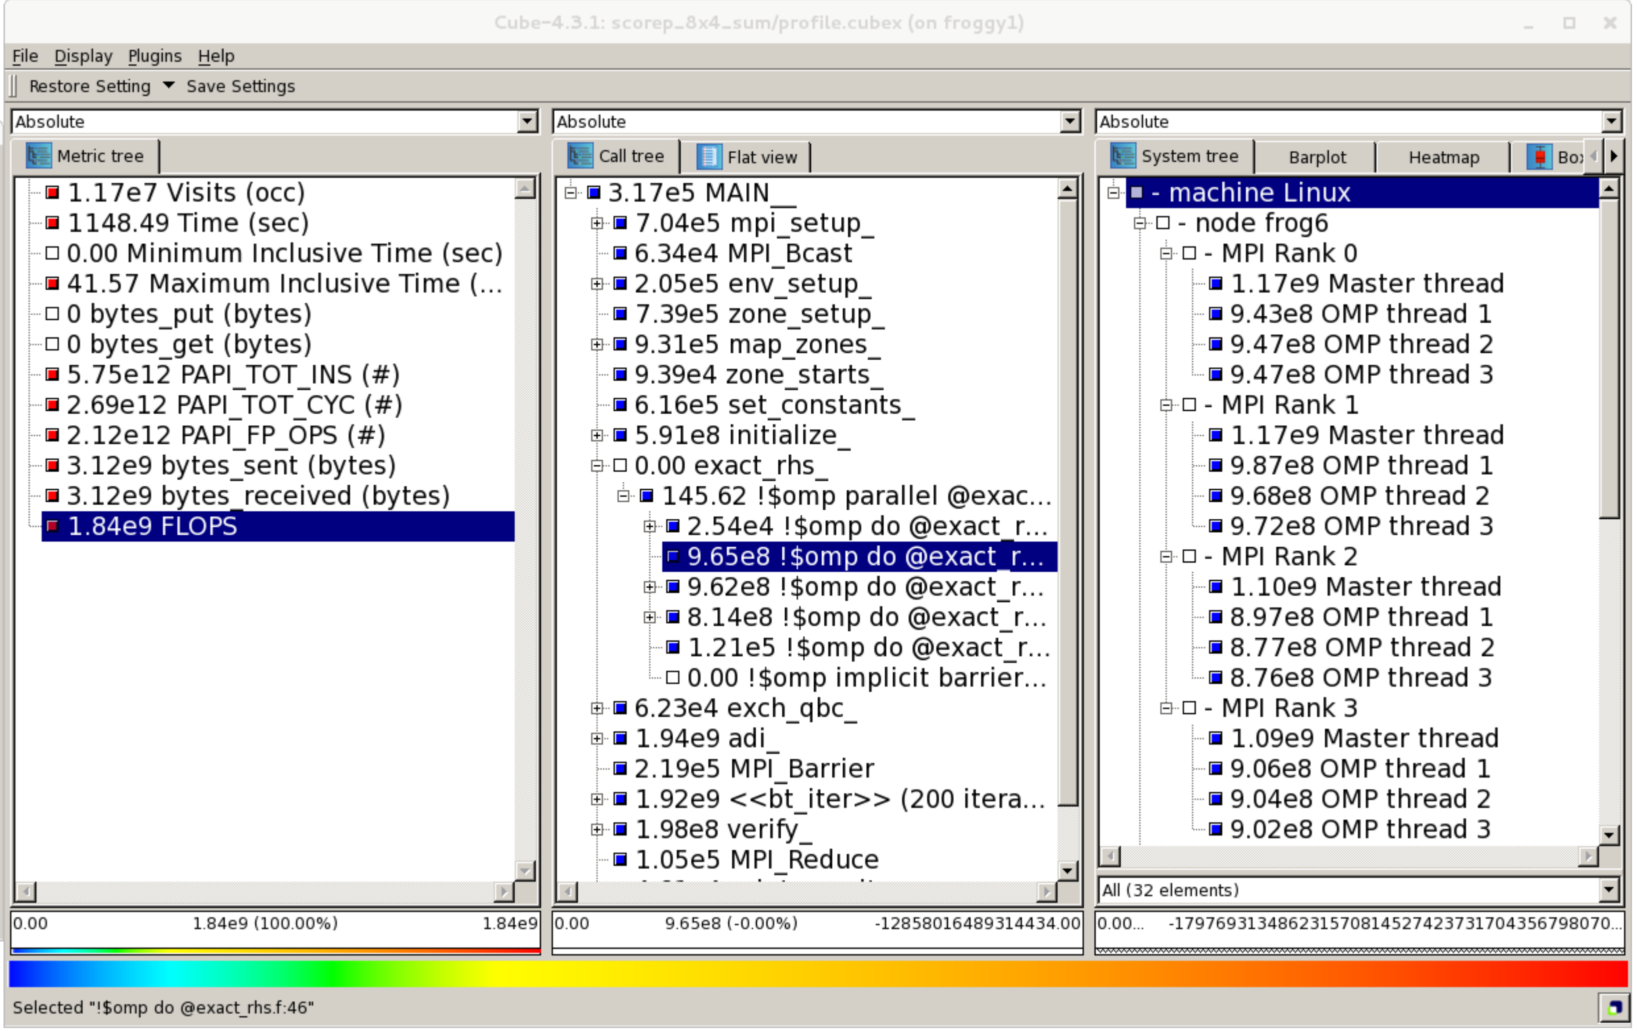

Cube-4.3.1: scorep_8x4_sum/profile.cubex (on froggy1)

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:

```
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:

```
- 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...    -1797693134862315708145274237317043567980 70...
```

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

CUBE algebra utilities

- Extracting solver sub-tree from analysis report

```
% cube_cut  -r '<<ITERATION>>'  scorep_bt-mz_C_32x4_sum/profile.cubex
Writing cut.cubex... done.
```

- Calculating difference of two reports

```
% cube_diff  scorep_bt-mz_C_32x4_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

Iteration profiling

- Show time dependent behavior by “unrolling” iterations

- Preparations:
  - Mark loop body by using Score-P instrumentation API in your source code

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

- 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

Iteration profiling: Barplot

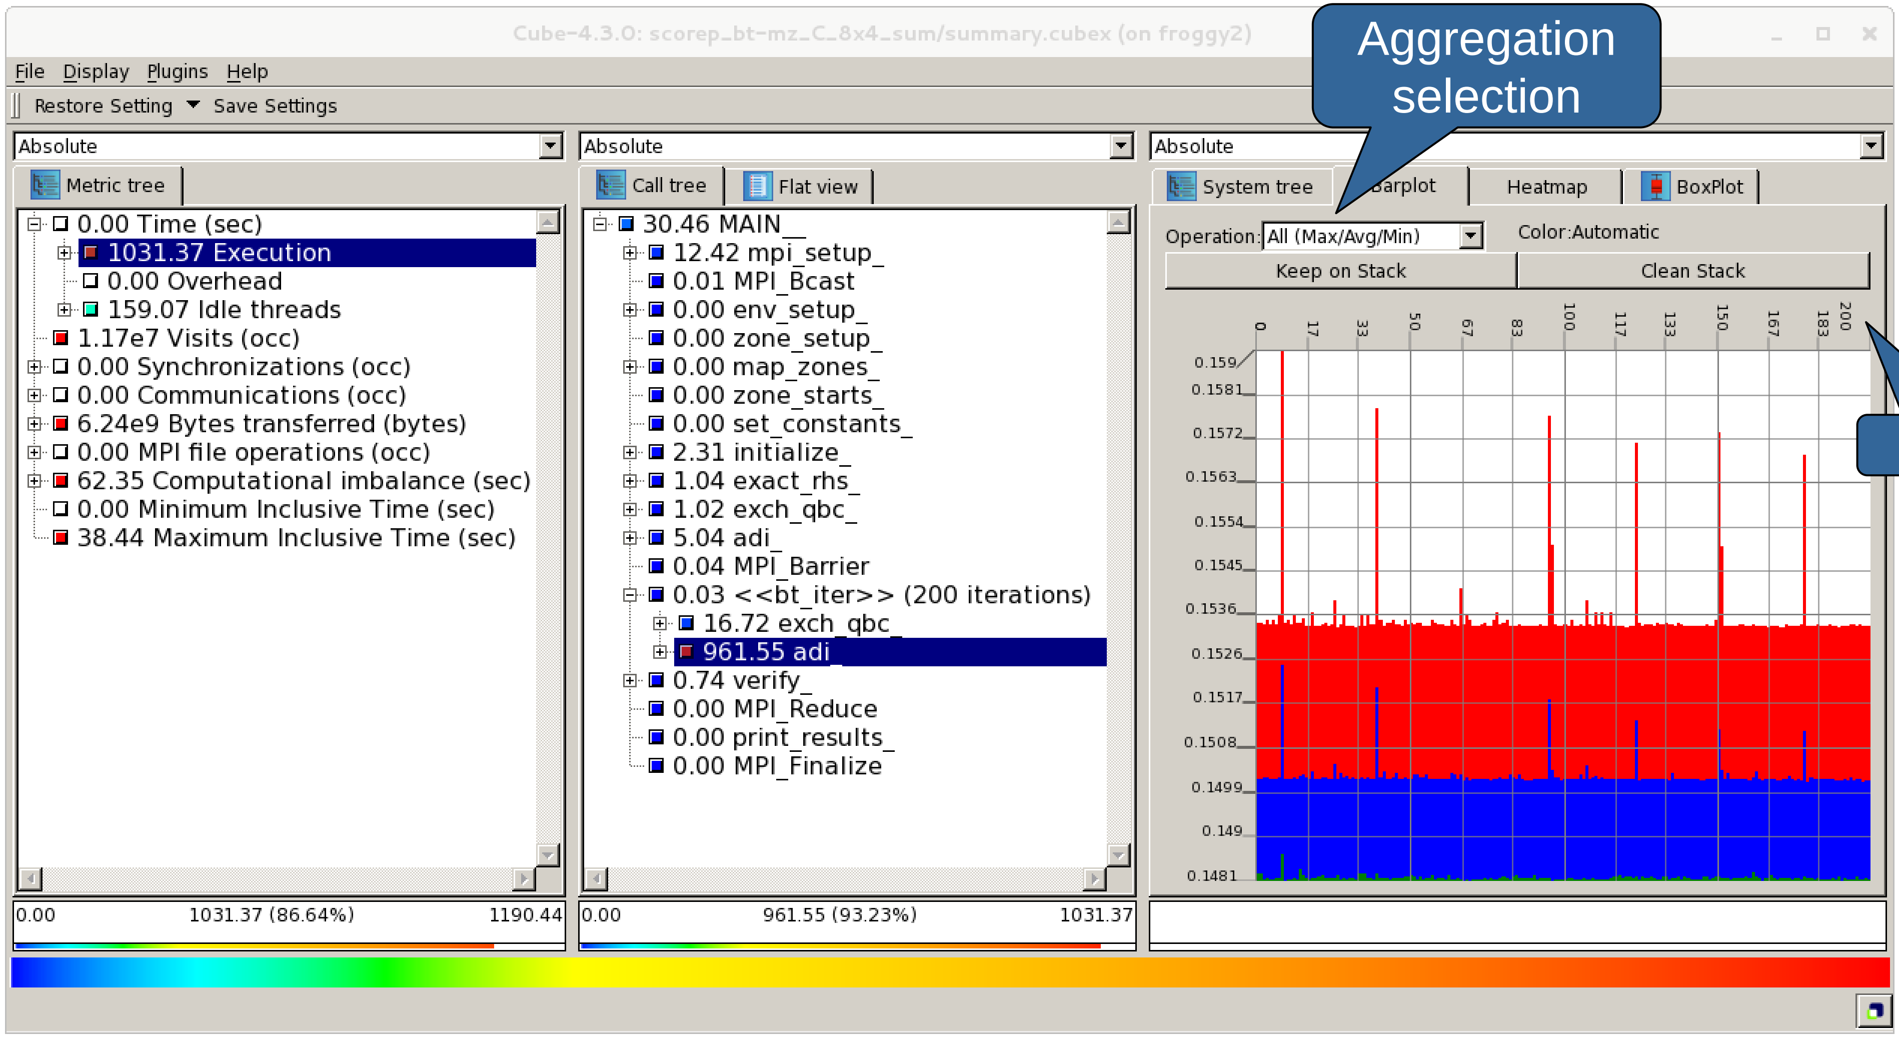

- 0.00 Time (sec)
  - 1031.37 Execution
  - 0.00 Overhead
  - 159.07 Idle threads
- 1.17e7 Visits (occ)
- 0.00 Synchronizations (occ)
- 0.00 Communications (occ)
- 6.24e9 Bytes transferred (bytes)
- 0.00 MPI file operations (occ)
- 62.35 Computational imbalance (sec)
- 0.00 Minimum Inclusive Time (sec)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

Aggregation selection

Iterations

Iteration profiling: Heatmap

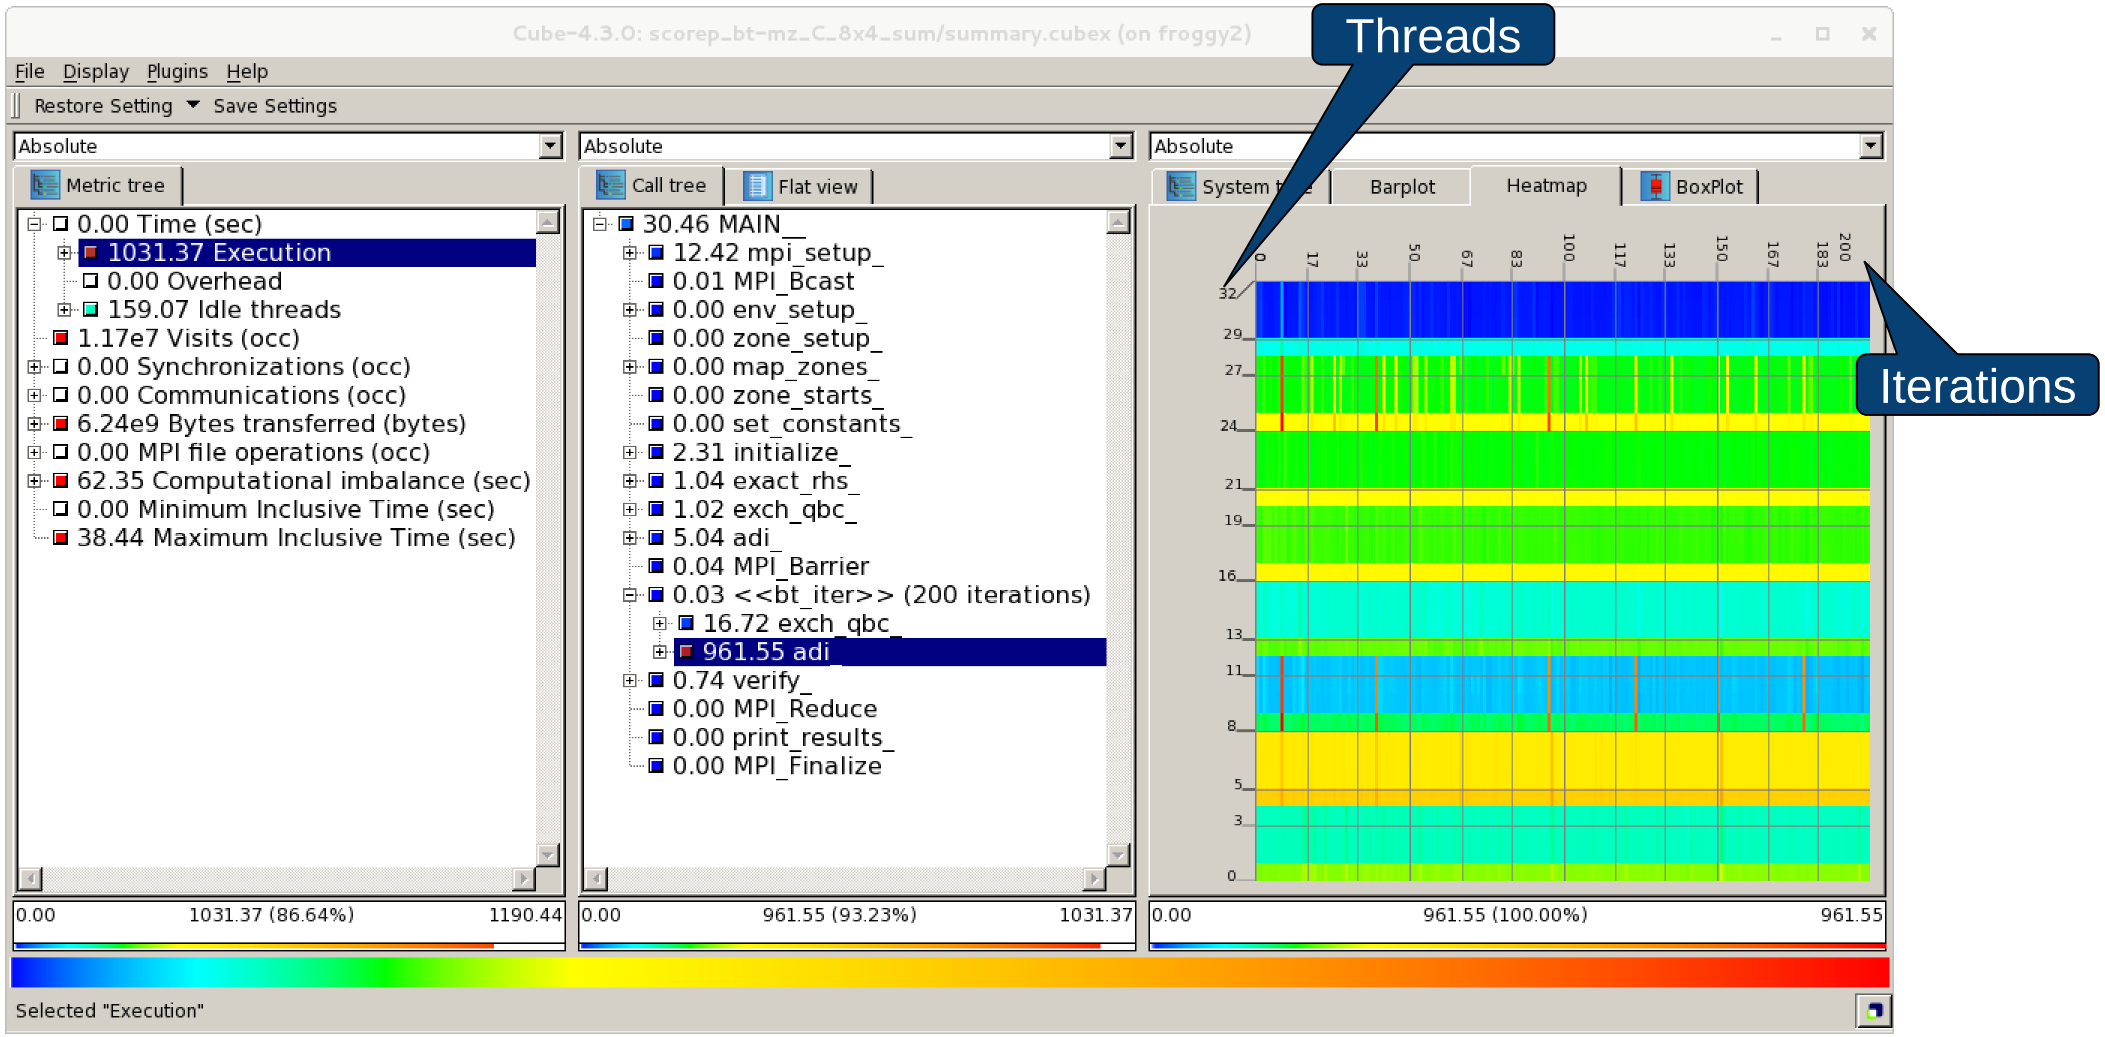

Cube-4.3.0: scorep_bt-mz_C_8x4_sum/summary.cubex (on froggy2)

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 | Flat view

- 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

System tree | Barplot | Heatmap | BoxPlot

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"

Threads

Iterations