

# Automatic trace analysis with the Scalasca Trace Tools

---

Markus Geimer  
Jülich Supercomputing Centre

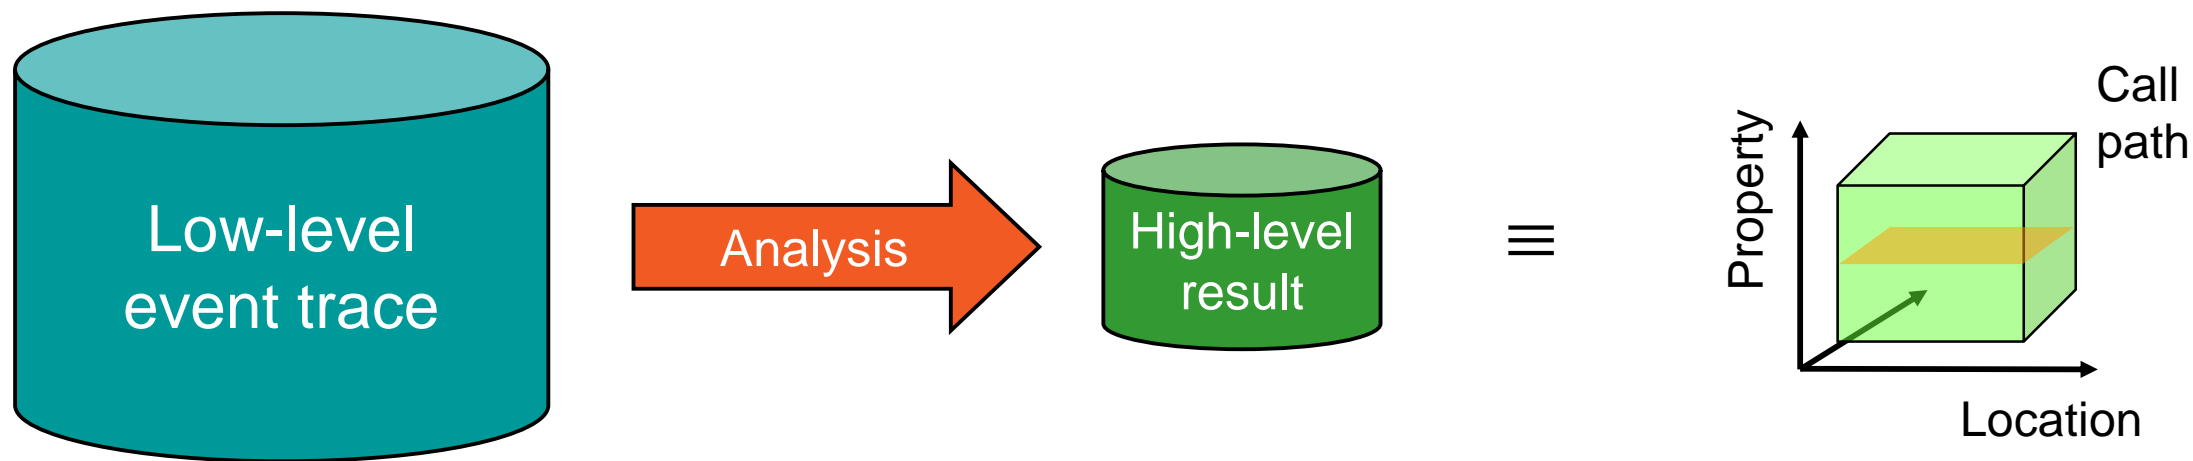




# Automatic trace analysis

## ▪ Idea

- Automatic search for patterns of inefficient behaviour
- Classification of behaviour & quantification of significance
- Identification of delays as root causes of inefficiencies



- Guaranteed to cover the entire event trace
- Quicker than manual/visual trace analysis
- Parallel replay analysis exploits available memory & processors to deliver scalability



# Scalasca Trace Tools: Objective

---

- Development of a **scalable trace-based** performance analysis toolset for the most popular parallel programming paradigms
  - Current focus: MPI, OpenMP, and POSIX threads
- Specifically targeting large-scale parallel applications
  - Such as those running on IBM Blue Gene or Cray systems with one million or more processes/threads
- Latest release:
  - Scalasca v2.5 coordinated with Score-P v5.0 (March 2019), also works with Score-P v6.0



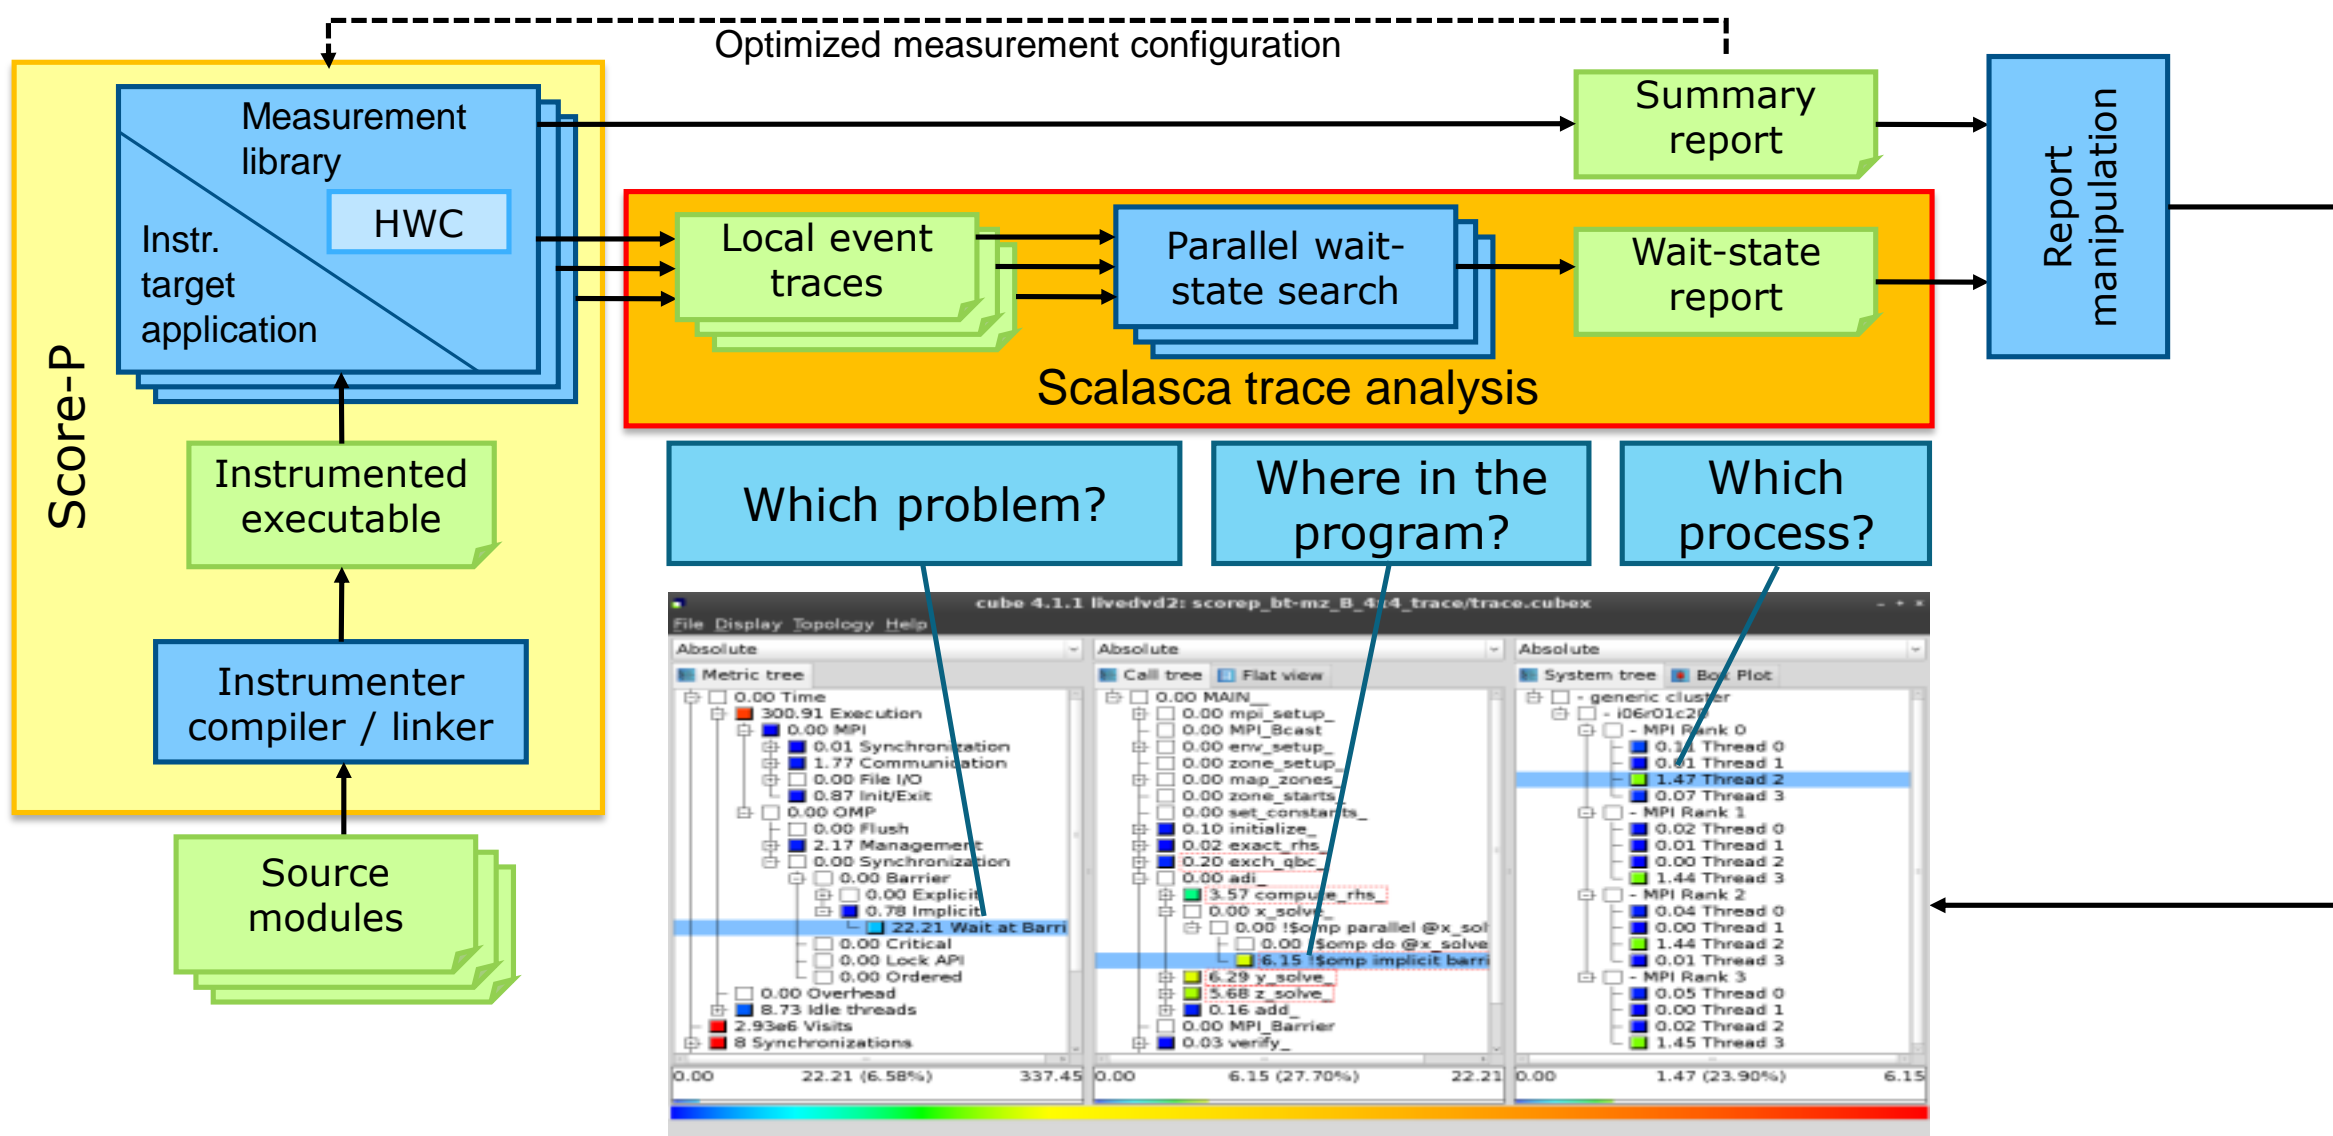
# Scalasca Trace Tools: Features

---

- Open source, 3-clause BSD license
- Fairly portable
  - IBM Blue Gene, Cray XT/XE/XK/XC, SGI Altix, Fujitsu FX systems, Linux clusters (x86, Power, ARM), Intel Xeon Phi, ...
- Uses Score-P instrumenter & measurement libraries
  - Scalasca v2 core package focuses on trace-based analyses
  - Supports common data formats
    - Reads event traces in OTF2 format
    - Writes analysis reports in CUBE4 format
- Current limitations:
  - Unable to handle traces
    - With MPI thread level exceeding MPI\_THREAD\_FUNNELED
    - Containing Memory events, CUDA/OpenCL device events (kernel, memcpy), SHMEM, or OpenMP nested parallelism
  - PAPI/rusage metrics for trace events are ignored

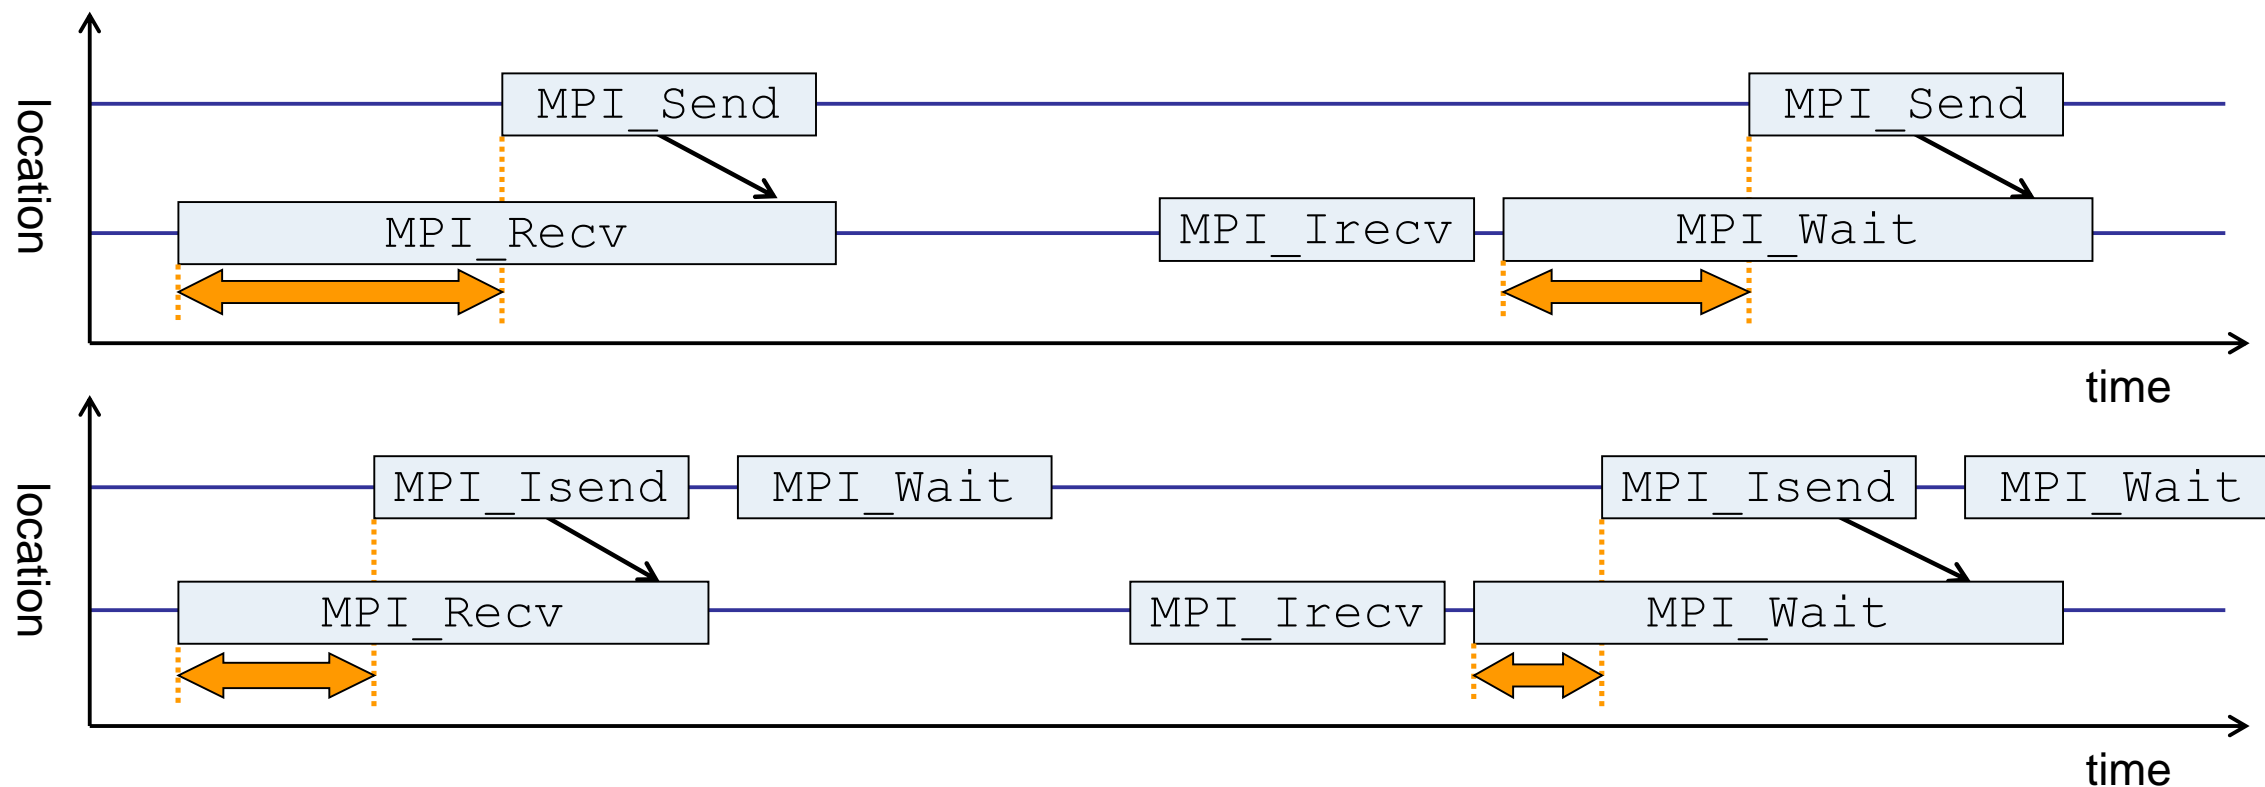


# Scalasca workflow





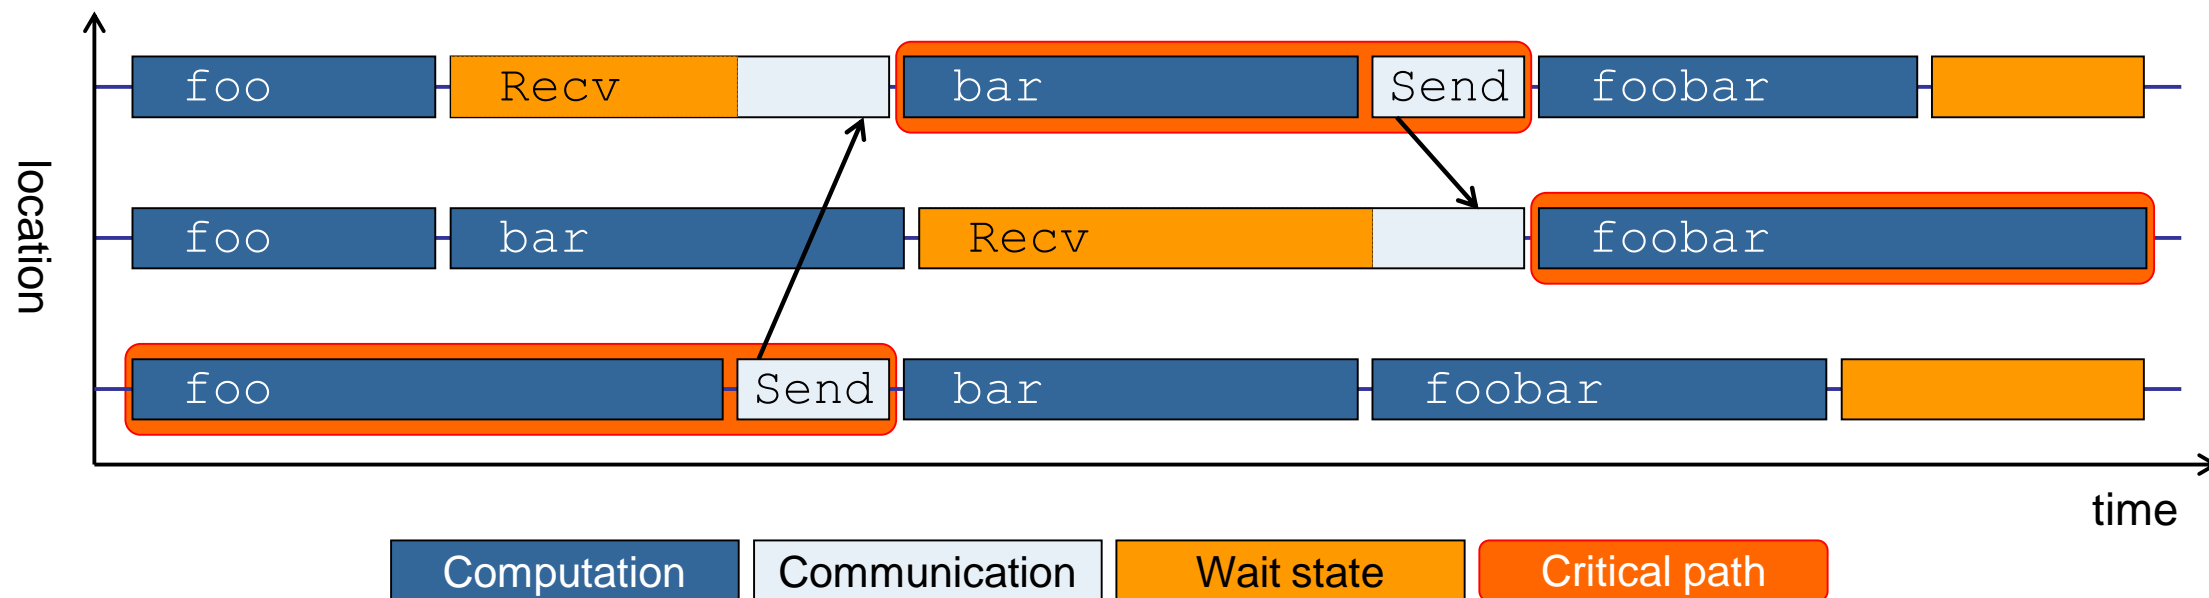
## Example: “*Late Sender*” wait state



- Waiting time caused by a blocking receive operation posted earlier than the corresponding send
- Applies to blocking as well as non-blocking communication




## Example: Critical path



- Shows call paths and processes/threads that are responsible for the program's wall-clock runtime
- Identifies good optimization candidates and parallelization bottlenecks



## Example: Root-cause analysis



- Classifies wait states into direct and indirect (i.e., caused by other wait states)
- Identifies *delays* (excess computation/communication) as root causes of wait states
- Attributes wait states as *delay costs*



# Demo: TeaLeaf case study

---





## 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: <https://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, Score-P 5.0, and Scalasca 2.5
  - 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



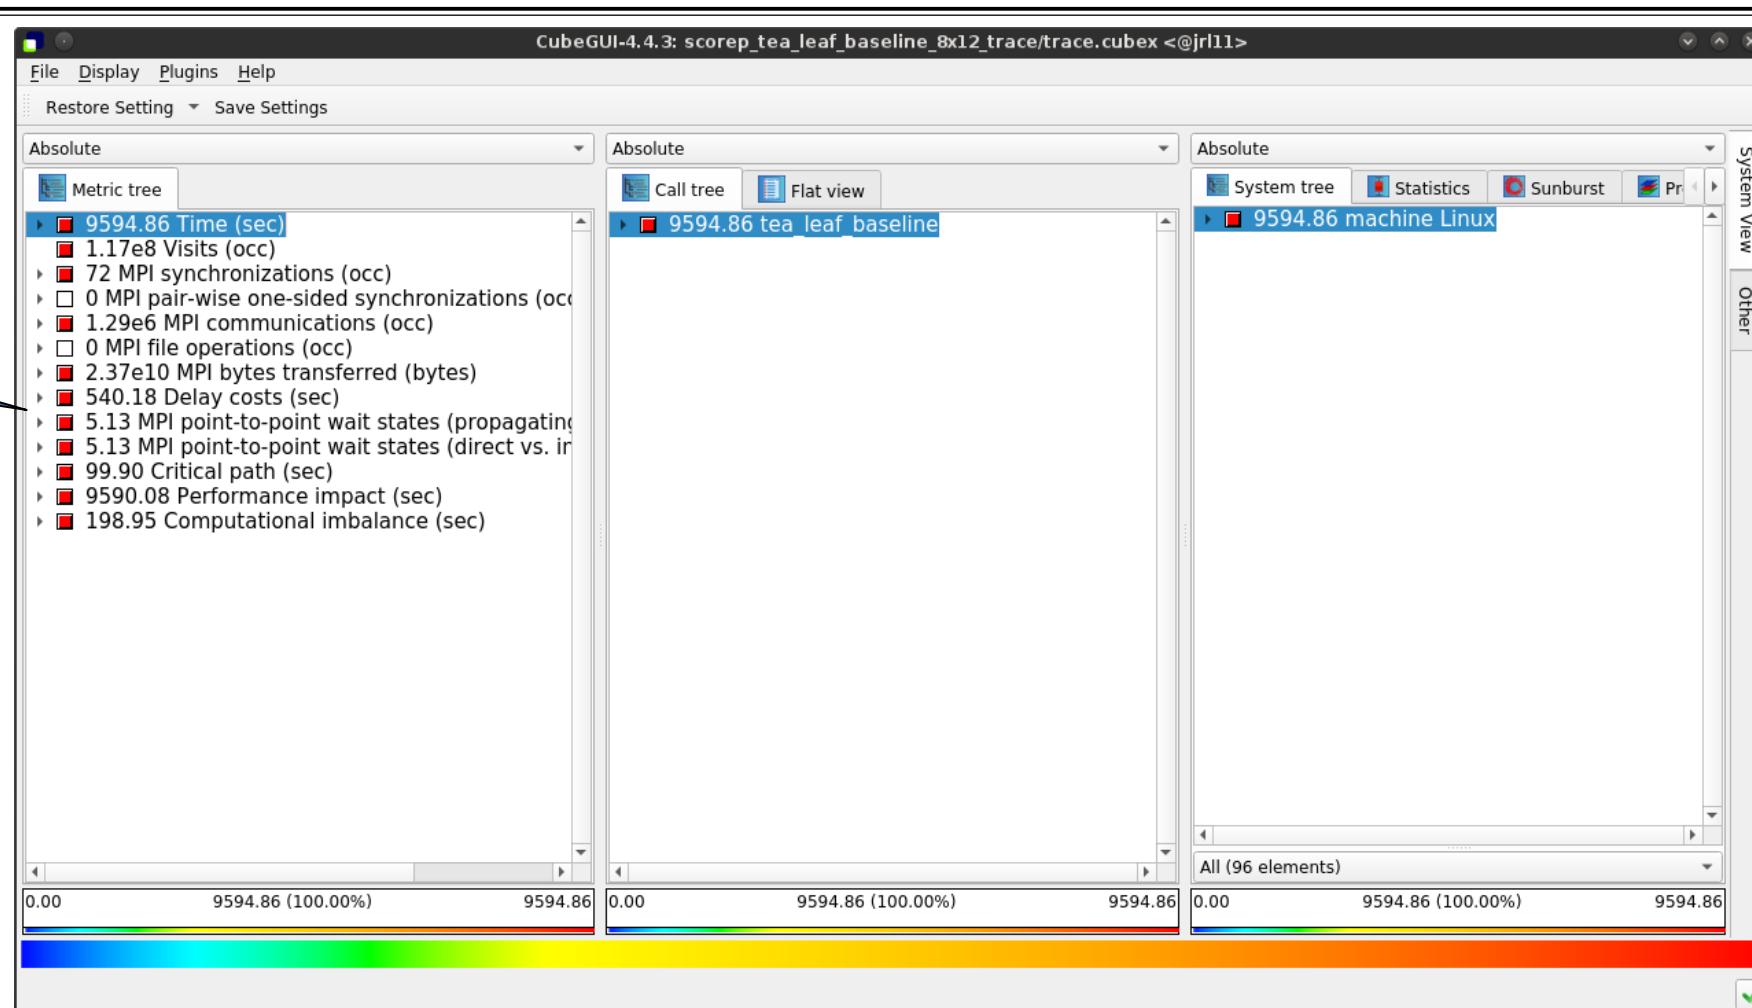
```
% scp -r marconi:/m100_work/tra20_TW36/experiments/
                                     jureca/scorep_tea_leaf_baseline_8x12_trace .
% cube scorep_tea_leaf_baseline_8x12_trace/trace.cubex
                                     [GUI showing post-processed trace analysis report]
```



# Scalasca analysis report exploration (opening view)



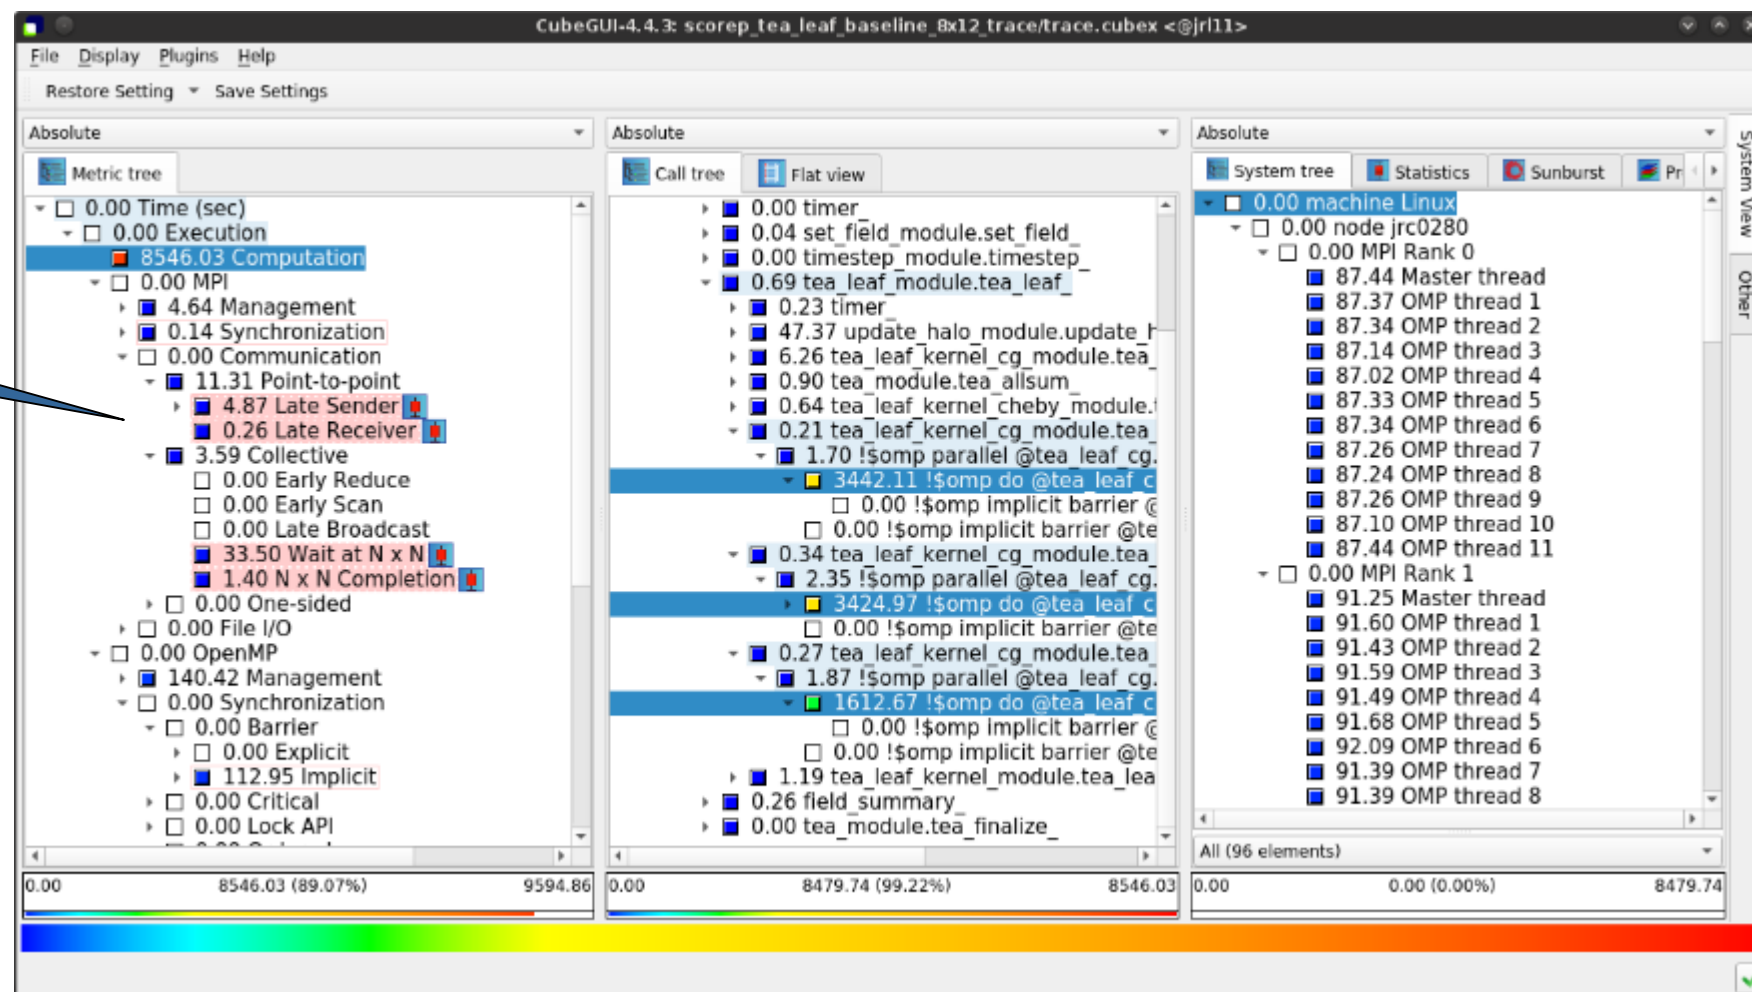
Additional top-level metrics produced by the trace analysis...





# Scalasca wait-state metrics

...plus additional wait-state metrics as part of the “Time” hierarchy

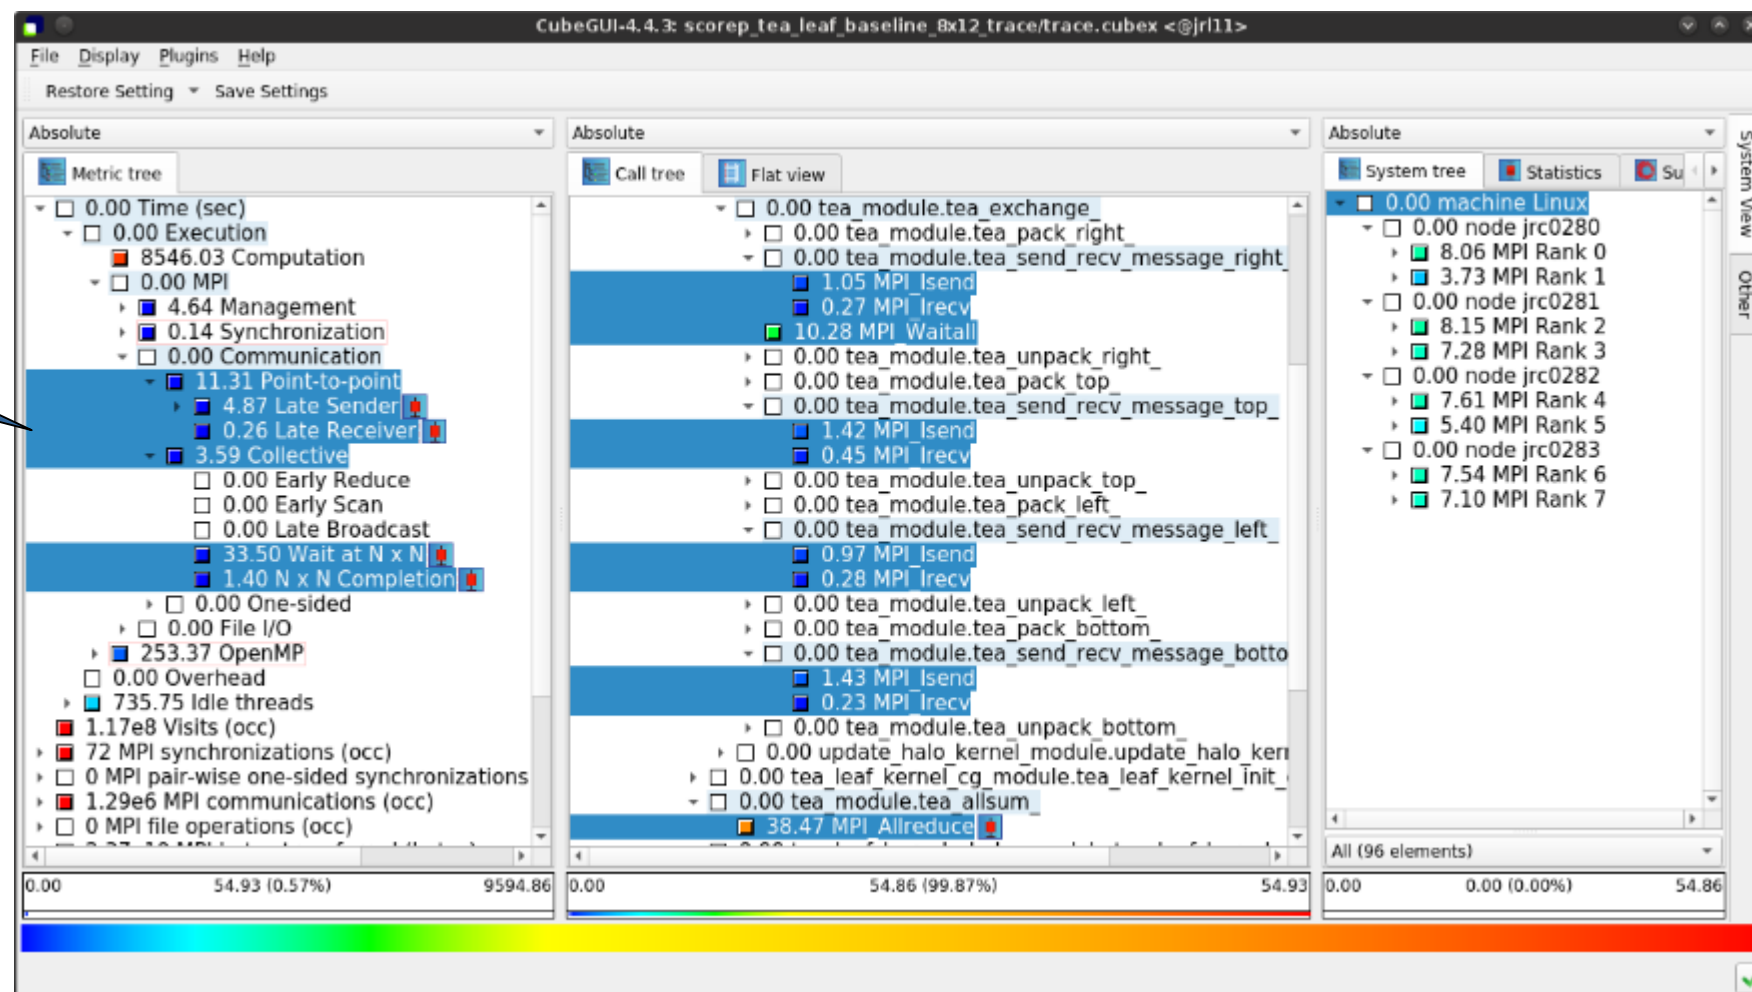




# TeaLeaf Scalasca report analysis (I)



While MPI communication time and wait states are small (~0.6% of the total execution time)...

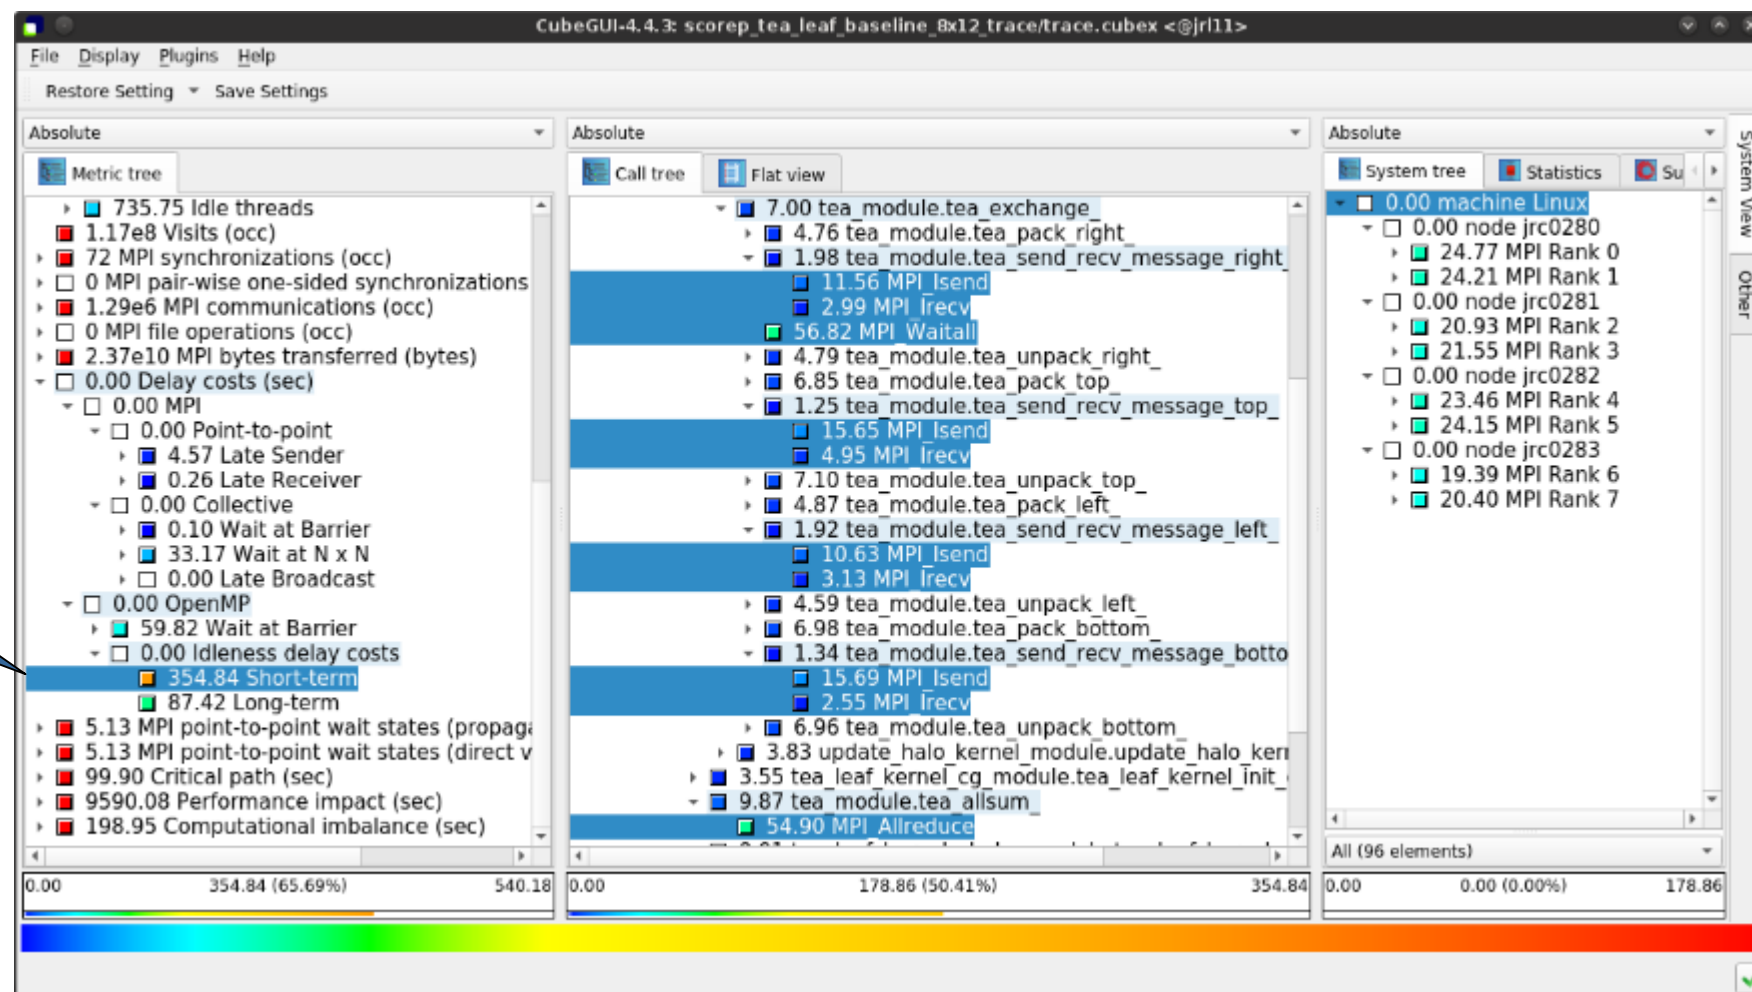




## TeaLeaf Scalasca report analysis (II)



...they directly cause a significant amount of the OpenMP thread idleness

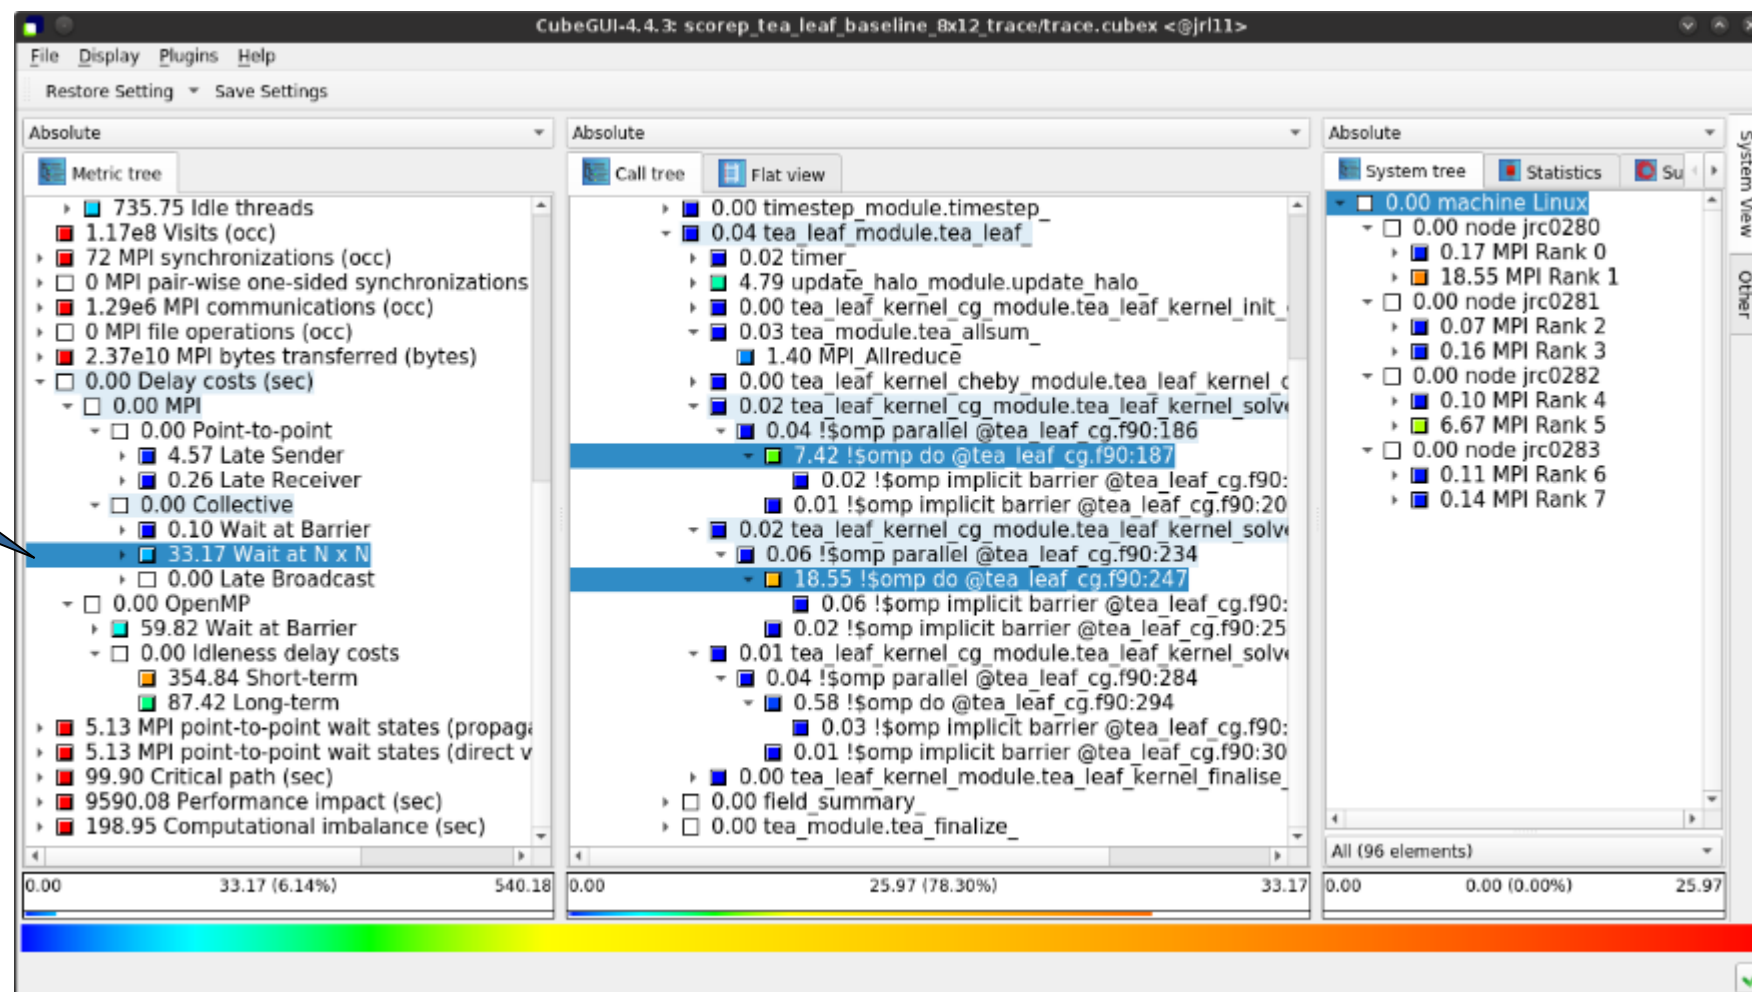




# TeaLeaf Scalasca report analysis (III)



The “Wait at NxN” collective wait states are mostly caused by the first 2 OpenMP `do` loops of the solver (on ranks 5 & 1, resp.)...

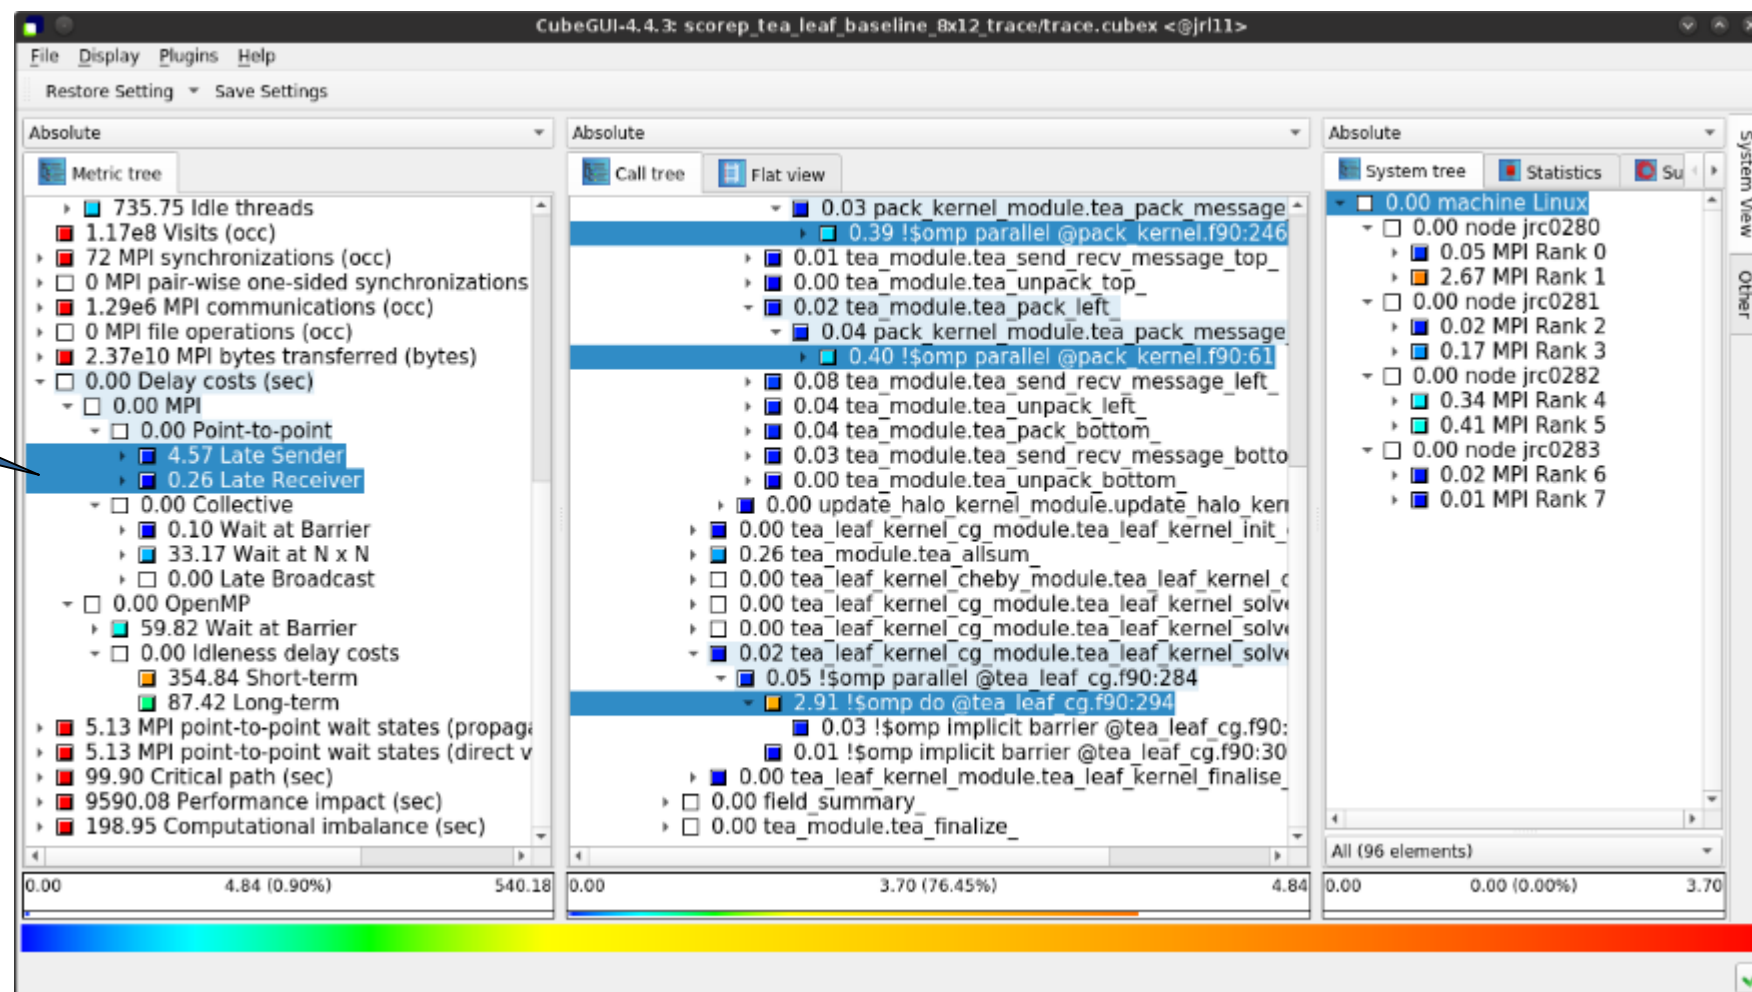




## TeaLeaf Scalasca report analysis (IV)



...while the MPI point-to-point wait states are caused by the 3<sup>rd</sup> solver do loop (on rank 1) and two loops in the halo exchange

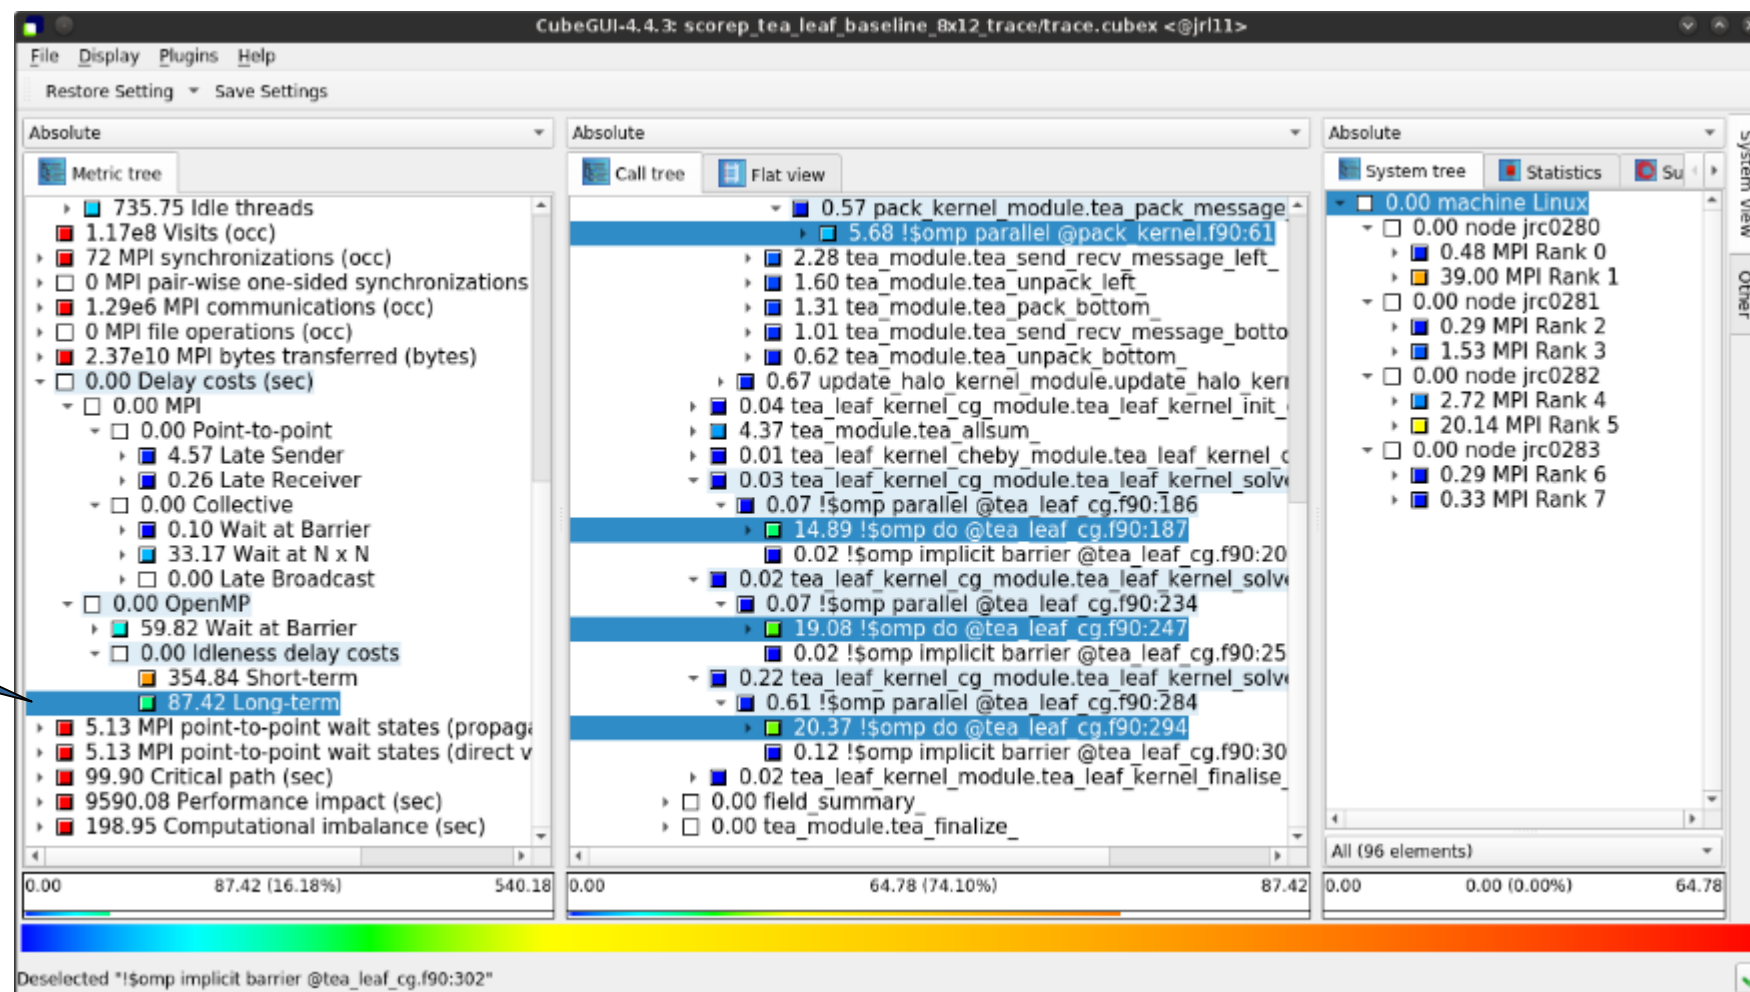




# TeaLeaf Scalasca report analysis (V)



Various OpenMP do loops (incl. the solver loops) also cause OpenMP thread idleness on other ranks via propagation






# TeaLeaf Scalasca report analysis (VI)



The Critical Path also highlights the three solver loops...

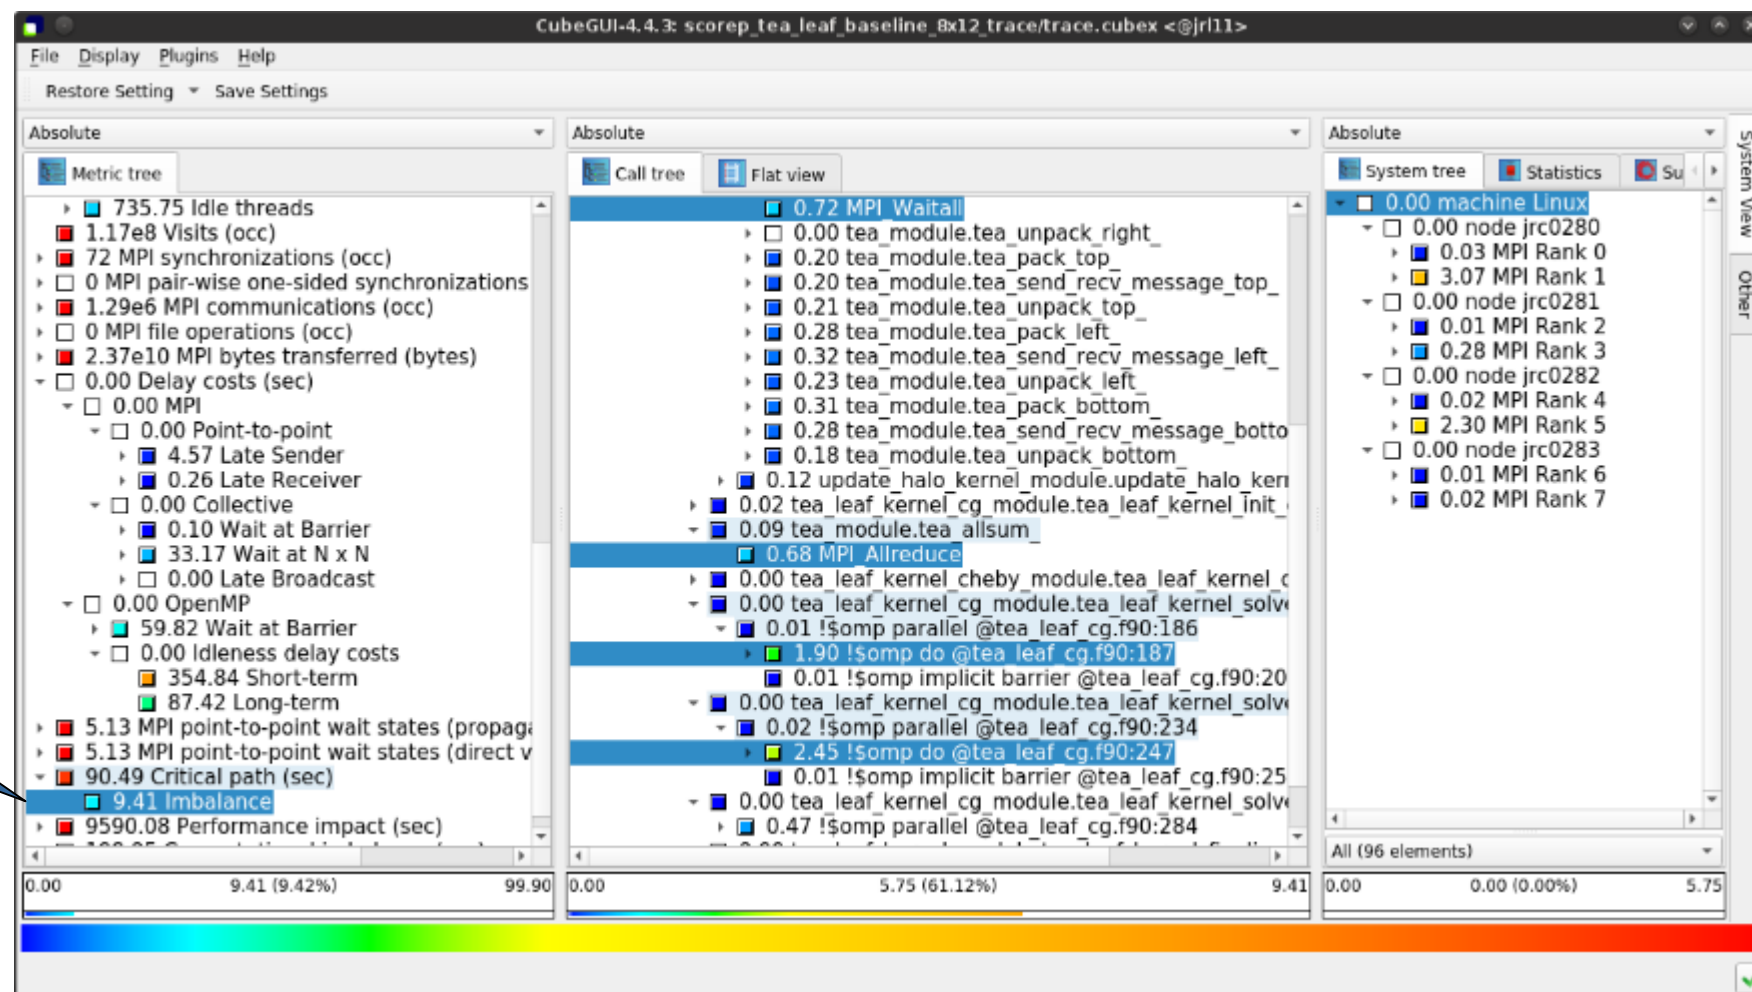




# TeaLeaf Scalasca report analysis (VII)



...with imbalance (time on critical path above average) mostly in the first two loops and MPI communication

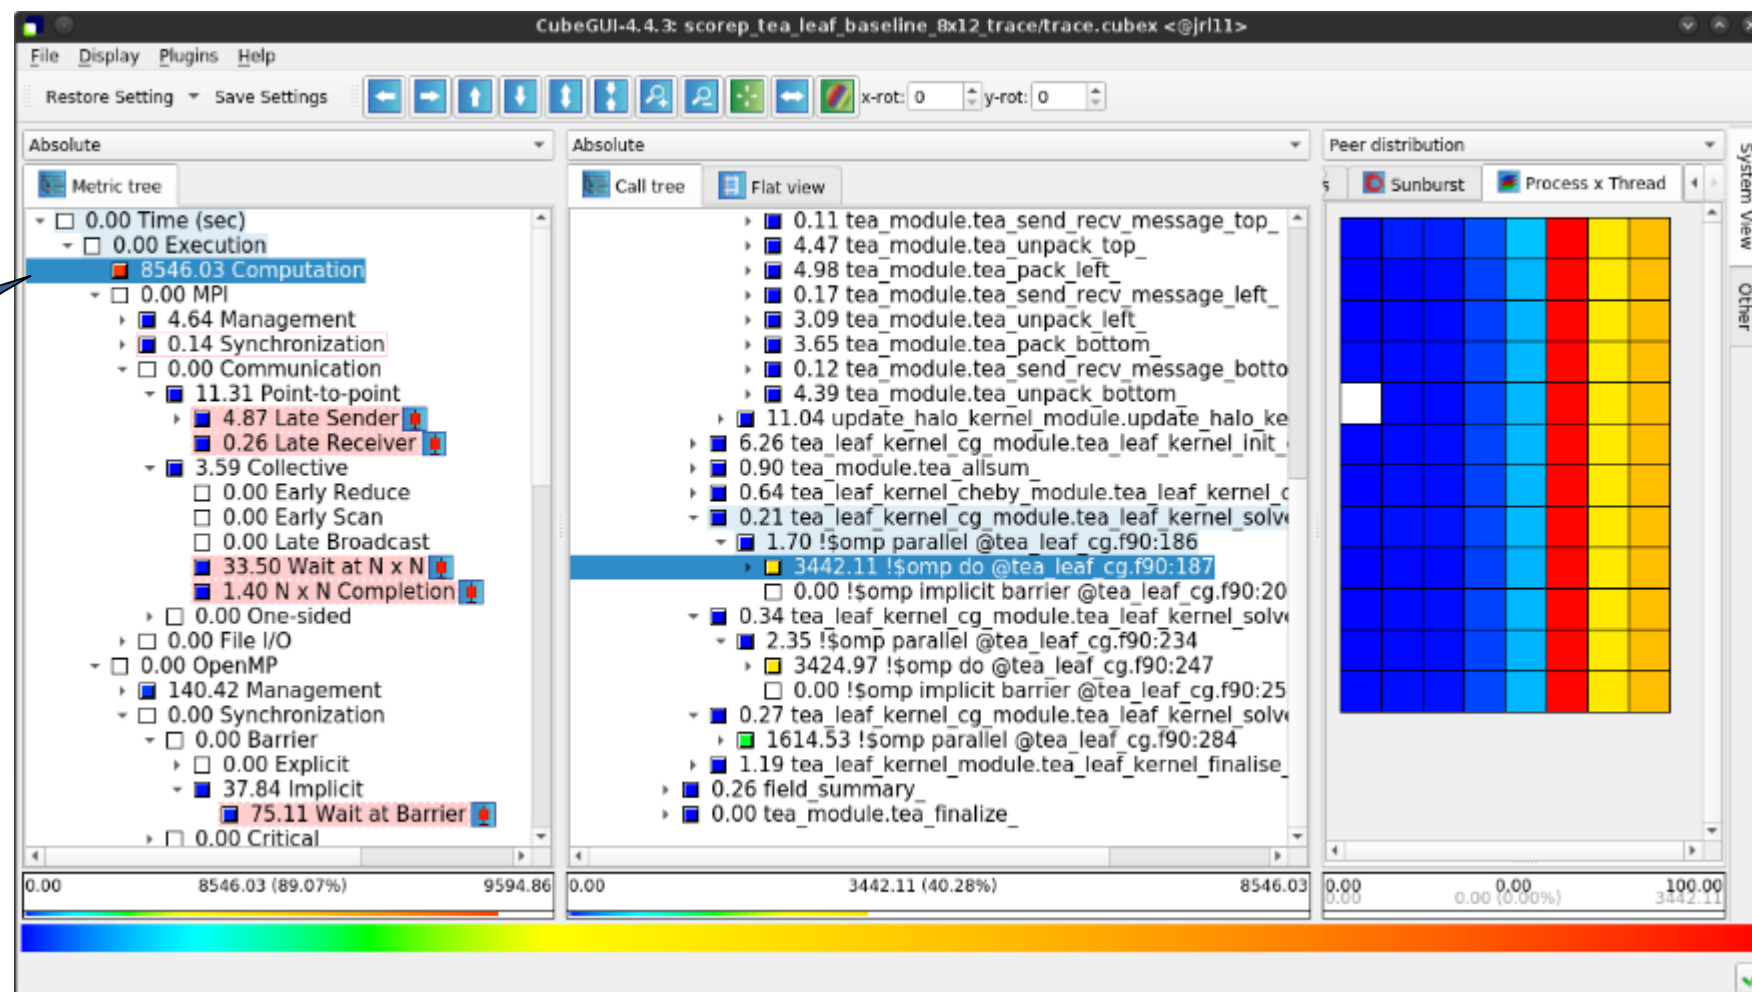




# TeaLeaf Scalasca report analysis (VIII)



Computation time of  
1<sup>st</sup>...

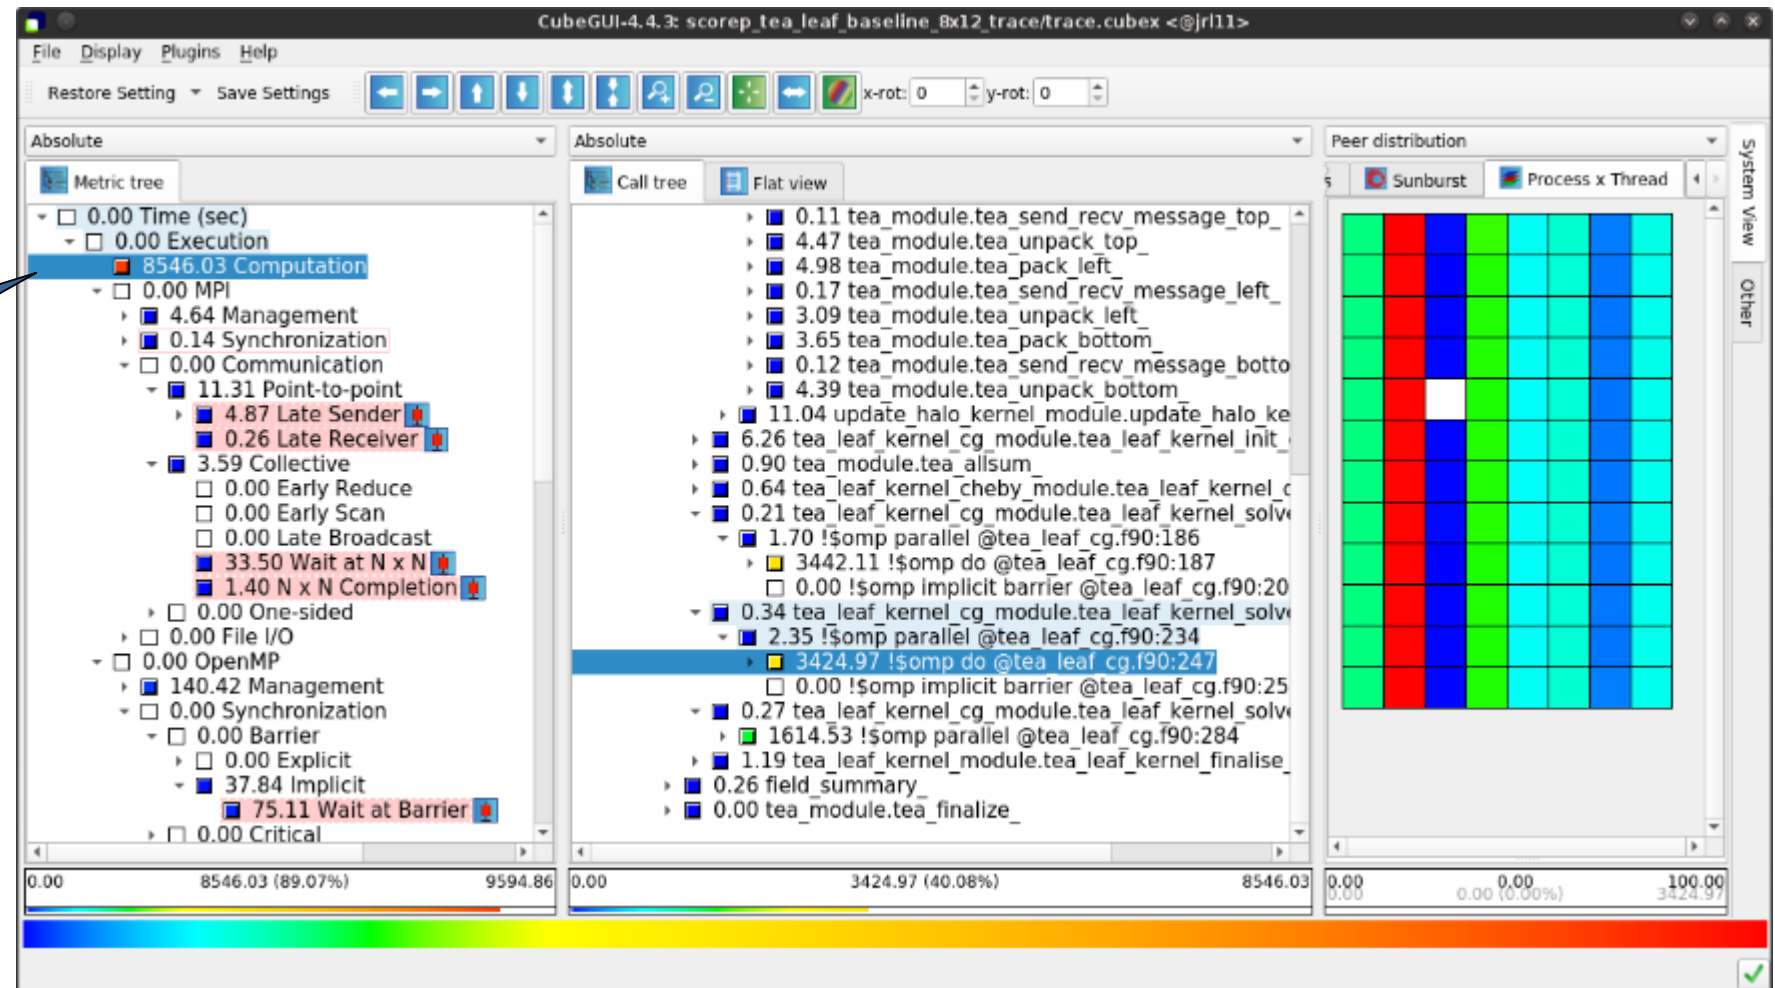




# TeaLeaf Scalasca report analysis (IX)



...and 2<sup>nd</sup> do loop  
mostly balanced within  
each rank, but vary  
considerably across  
ranks...

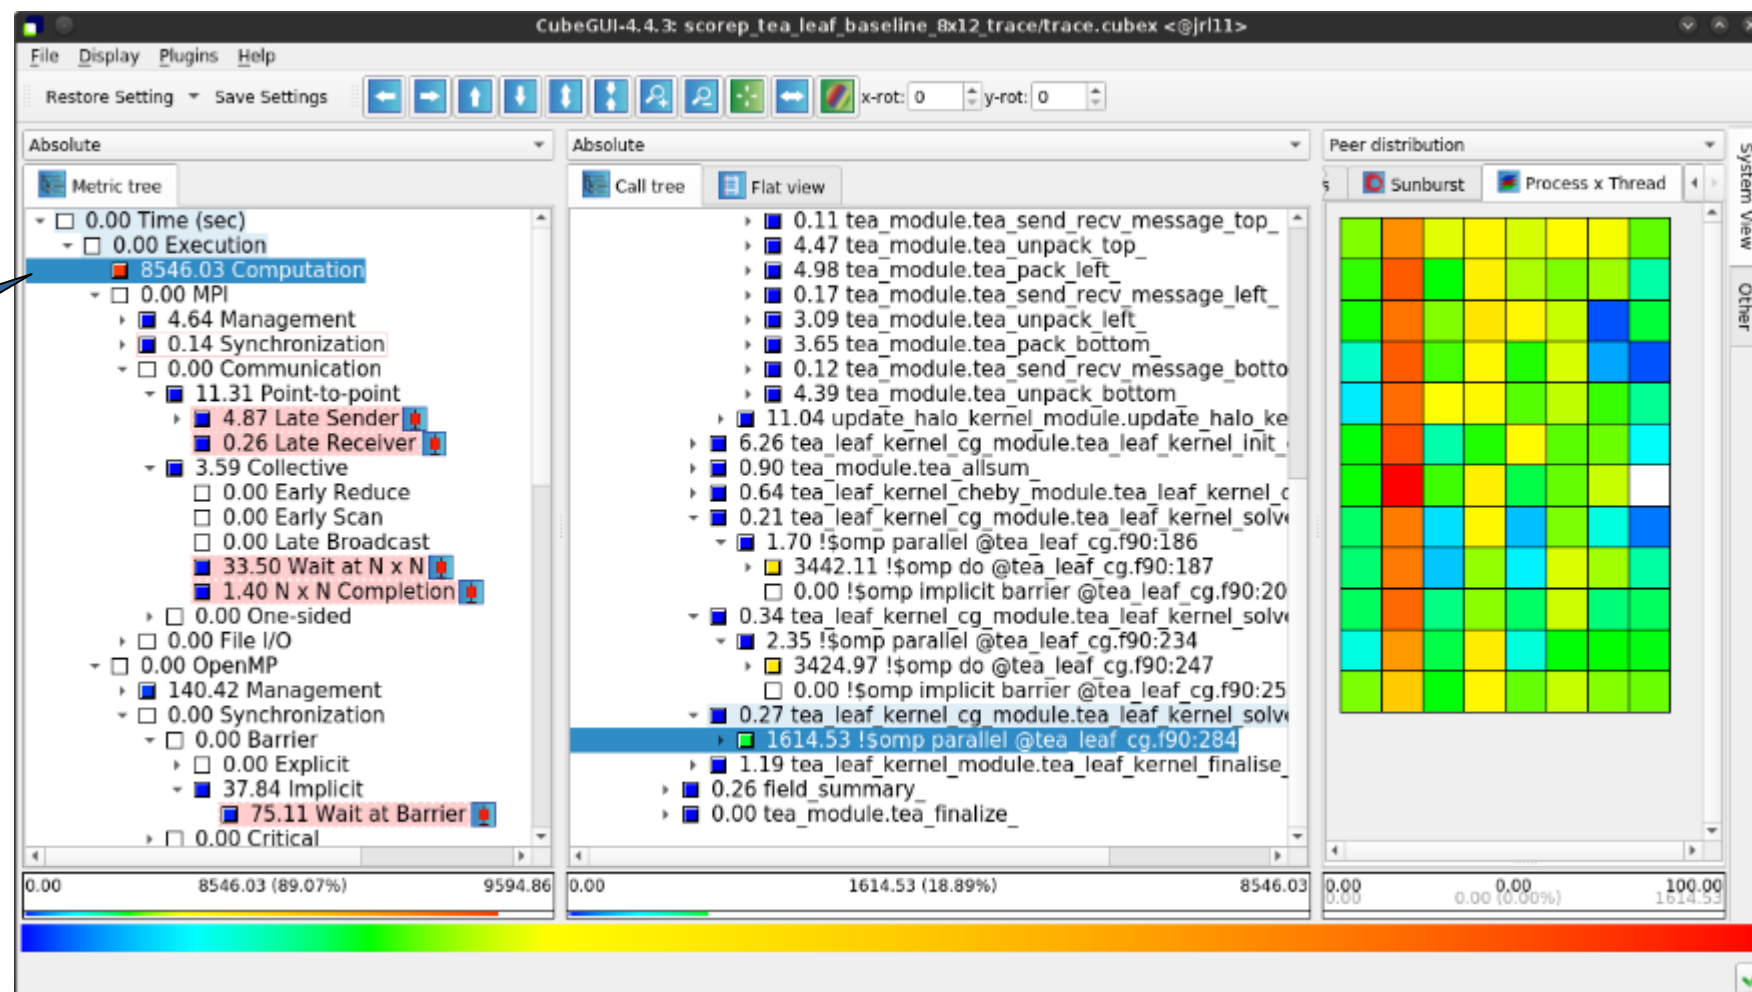




# TeaLeaf Scalasca report analysis (X)



...while the 3<sup>rd</sup> do loop  
also shows imbalance  
within each rank





## TeaLeaf analysis summary

---

- The first two OpenMP do loops of the solver are well balanced within a rank, but are imbalanced across ranks
  - Requires a global load balancing strategy
- The third OpenMP do loop, however, is imbalanced within ranks,
  - causing direct “Wait at OpenMP Barrier” wait states,
  - which cause indirect MPI point-to-point wait states,
  - which in turn cause OpenMP thread idleness
  - Low-hanging fruit
- Adding a `SCHEDULE(guided)` clause reduced
  - the MPI point-to-point wait states by ~66%
  - the MPI collective wait states by ~50%
  - the OpenMP “Wait at Barrier” wait states by ~55%
  - the OpenMP thread idleness by ~11%
  - **Overall runtime (wall-clock) reduction by ~5%**



# Hands-on: TeaLeaf MPI+CUDA

---





## Scalasca command – One command for (almost) everything

```
% scalasca
Scalasca 2.5
Toolset for scalable performance analysis of large-scale parallel applications
usage: scalasca [OPTION]... ACTION <argument>...
  1. prepare application objects and executable for measurement:
    scalasca -instrument <compile-or-link-command> # skin (using scorep)
  2. run application under control of measurement system:
    scalasca -analyze <application-launch-command> # scan
  3. interactively explore measurement analysis report:
    scalasca -examine <experiment-archive|report> # square

Options:
  -c, --show-config      show configuration summary and exit
  -h, --help             show this help and exit
  -n, --dry-run          show actions without taking them
                        --quickref      show quick reference guide and exit
                        --remap-specfile show path to remapper specification file and exit
  -v, --verbose          enable verbose commentary
  -V, --version          show version information and exit
```

- The '`scalasca -instrument`' command is deprecated and only provided for backwards compatibility with Scalasca 1.x., recommended: use Score-P instrumenter directly



## Scalasca convenience command: scan / scalasca -analyze

```
% scan
Scalasca 2.5: measurement collection & analysis nexus
usage: scan {options} [launchcmd [launchargs]] target [targetargs]
      where {options} may include:
  -h      Help          : show this brief usage message and exit.
  -v      Verbose       : increase verbosity.
  -n      Preview       : show command(s) to be launched but don't execute.
  -q      Quiescent     : execution with neither summarization nor tracing.
  -s      Summary       : enable runtime summarization. [Default]
  -t      Tracing       : enable trace collection and analysis.
  -a      Analyze       : skip measurement to (re-)analyze an existing trace.
  -e exptdir           : Experiment archive to generate and/or analyze.
                       (overrides default experiment archive title)
  -f filtfile         : File specifying measurement filter.
  -l lockfile         : File that blocks start of measurement.
  -R #runs            : Specify the number of measurement runs per config.
  -M cfgfile          : Specify a config file for a multi-run measurement.
```

- Scalasca measurement collection & analysis nexus



## Scalasca convenience command: square / scalasca -examine

```
% square
Scalasca 2.5: analysis report explorer
usage: square [OPTIONS] <experiment archive | cube file>
  -c <none | quick | full> : Level of sanity checks for newly created reports
  -F                        : Force remapping of already existing reports
  -f filtfiler              : Use specified filter file when doing scoring (-s)
  -s                        : Skip display and output textual score report
  -v                        : Enable verbose mode
  -n                        : Do not include idle thread metric
  -S <mean | merge>         : Aggregation method for summarization results of
                             each configuration (default: merge)
  -T <mean | merge>         : Aggregation method for trace analysis results of
                             each configuration (default: merge)
  -A                        : Post-process every step of a multi-run experiment
```

### ▪ Scalasca analysis report explorer (Cube)



# Automatic measurement configuration

---

- scan configures Score-P measurement by automatically setting some environment variables and exporting them
  - E.g., experiment title, profiling/tracing mode, filter file, ...
  - Precedence order:
    - Command-line arguments
    - Environment variables already set
    - Automatically determined values
- Also, scan includes consistency checks and prevents corrupting existing experiment directories
- For tracing experiments, after trace collection completes then automatic parallel trace analysis is initiated
  - Uses identical launch configuration to that used for measurement (i.e., the same allocated compute resources)



## Recap: Local installation (Marconi100)

---

- VI-HPS tools not yet installed system-wide
  - Select appropriate environment
  - Required for each shell session

```
% . /m100_work/tra20_TW36/tools/sourceme.scorep-scalasca.gnu-spectrummpi
```

- Change to directory containing TeaLeaf\_CUDA sources
  - Existing instrumented executable in bin.scorep/ directory can be reused

```
% cd $CINECA_SCRATCH/TeaLeaf_CUDA
```



## TeaLeaf\_CUDA summary measurement collection...

```
% cd bin.scorep
% cp ../jobscript/marconi100/scalasca.sbatch .
% cat scalasca.sbatch

# Score-P measurement configuration
export SCOREP_CUDA_ENABLE=runtime,driver
export SCOREP_CUDA_BUFFER=48M
export SCOREP_FILTERING_FILE=scorep.filt
#export SCOREP_TOTAL_MEMORY=72M

# Scalasca configuration
#export SCAN_ANALYZE_OPTS="--time-correct"

# Run the application
export OMP_NUM_THREADS=${SLURM_CPUS_PER_TASK}
scan -s mpirun ./tea_leaf
```

```
% sbatch scalasca.sbatch
```

- Change to directory with the Score-P instrumented executable and edit the job script

### Hint:

```
scan = scalasca -analyze
-s = profile/summary (def)
```

- Submit the job



# TeaLeaf\_CUDA summary measurement

---

```
S=C=A=N: Scalasca 2.5 runtime summarization
S=C=A=N: ./scorep_tea_leaf_8x3_sum experiment archive
S=C=A=N: Mon Sep 28 12:07:38 2020: Collect start
mpirun ./tea_leaf

    Tea version      1.400

[... More application output ...]

S=C=A=N: Mon Sep 28 12:07:59 2020: Collect done (status=0) 21s
S=C=A=N: ./scorep_tea_leaf_8x3_sum complete.
```

- Run the application using the Scalasca measurement collection & analysis nexus prefixed to launch command
- Creates experiment directory:  
scorep\_tea\_leaf\_8x3\_sum



# TeaLeaf\_CUDA summary analysis report examination

---

- Score summary analysis report

```
% square -s scorep_tea_leaf_8x3_sum  
INFO: Post-processing runtime summarization result (profile.cubex)...  
INFO: Score report written to ./scorep_tea_leaf_8x3_sum/scorep.score
```

- Post-processing and interactive exploration with Cube

```
% square scorep_tea_leaf_8x3_sum  
INFO: Displaying ./scorep_tea_leaf_8x3_sum/summary.cubex...
```

```
[GUI showing summary analysis report]
```

**Hint:**

Copy 'profile.cubex' to local system (laptop) using 'scp' to improve responsiveness of GUI

- The post-processing derives additional metrics and generates a structured metric hierarchy



## TeaLeaf\_CUDA trace measurement collection...

```
% cd bin.scorep
% cp ../jobscript/marconi100/scalasca.sbatch .
% vim scalasca.sbatch

# Score-P measurement configuration
export SCOREP_CUDA_ENABLE=runtime,driver
export SCOREP_CUDA_BUFFER=48M
export SCOREP_FILTERING_FILE=scorep.filt
export SCOREP_TOTAL_MEMORY=72M

# Scalasca configuration
export SCAN_ANALYZE_OPTS="--time-correct"

# Run the application
export OMP_NUM_THREADS=${SLURM_CPUS_PER_TASK}
scan -t mpirun ./tea_leaf

% sbatch scalasca.sbatch
```

- Change to directory with the Score-P instrumented executable and edit the job script
- Add "-t" to the scan command
- Submit the job



## TeaLeaf\_CUDA trace measurement ... collection

---

```
S=C=A=N: Scalasca 2.5 trace collection and analysis  
S=C=A=N: Mon Sep 28 12:15:02 2020: Collect start  
mpirun ./tea_leaf
```

```
Tea version      1.400
```

```
[... More application output ...]
```

```
S=C=A=N: Mon Sep 28 12:15:24 2020: Collect done (status=0) 22s
```

- Starts measurement with collection of trace files ...



# TeaLeaf\_CUDA trace measurement ... analysis

```
...
S=C=A=N: Mon Sep 28 12:15:24 2020: Analyze start
mpirun scout.hyb --time-correct ./scorep_tea_leaf_8x3_trace/traces.otf2

SCOUT    (Scalasca 2.5)

Analyzing experiment archive ./scorep_tea_leaf_8x3_trace/traces.otf2

Opening experiment archive ... done (0.154s).
Reading definition data    ... done (0.040s).
Reading event trace data  ... done (1.707s).
Preprocessing              ... done (1.286s).
Timestamp correction       ... done (0.872s).
Analyzing trace data       ... done (18.027s).
Writing analysis report    ... done (0.201s).

Max. memory usage          : 360.625MB

      # passes              : 1
      # violated            : 0

Total processing time       : 22.495s
S=C=A=N: Mon Sep 28 12:15:48 2020: Analyze done (status=0) 24s
```

- Continues with automatic (parallel) analysis of trace files



## TeaLeaf CUDA trace analysis report exploration

---

- Produces trace analysis report in the experiment directory containing trace-based wait-state metrics

```
% square scorep_tea_leaf_8x3_trace
INFO: Post-processing runtime summarization report (profile.cubex)...
INFO: Post-processing trace analysis report (scout.cubex)...
INFO: Displaying ./scorep_tea_leaf_8x3_trace/trace.cubex...

[GUI showing trace analysis report]
```

**Hint:**

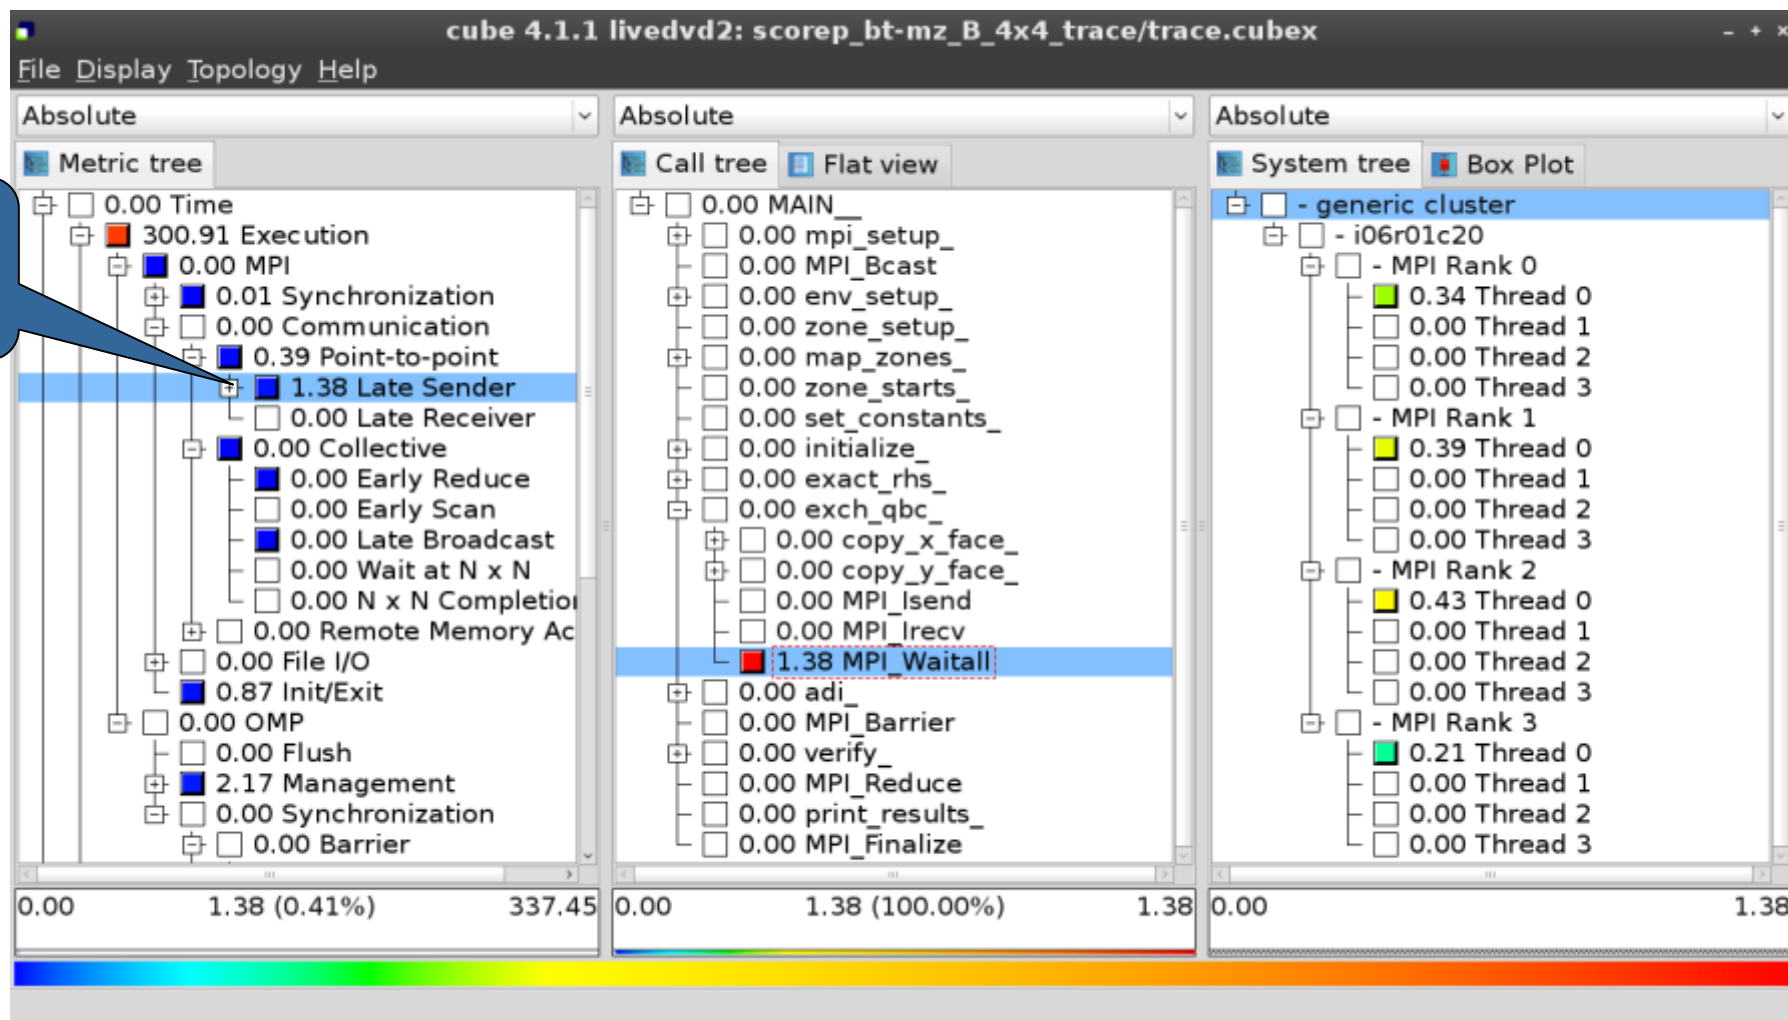
Run 'square -s' first and then copy 'trace.cubex' to local system (laptop) using 'scp' to improve responsiveness of GUI



# Post-processed trace analysis report



Additional trace-based metrics in metric hierarchy

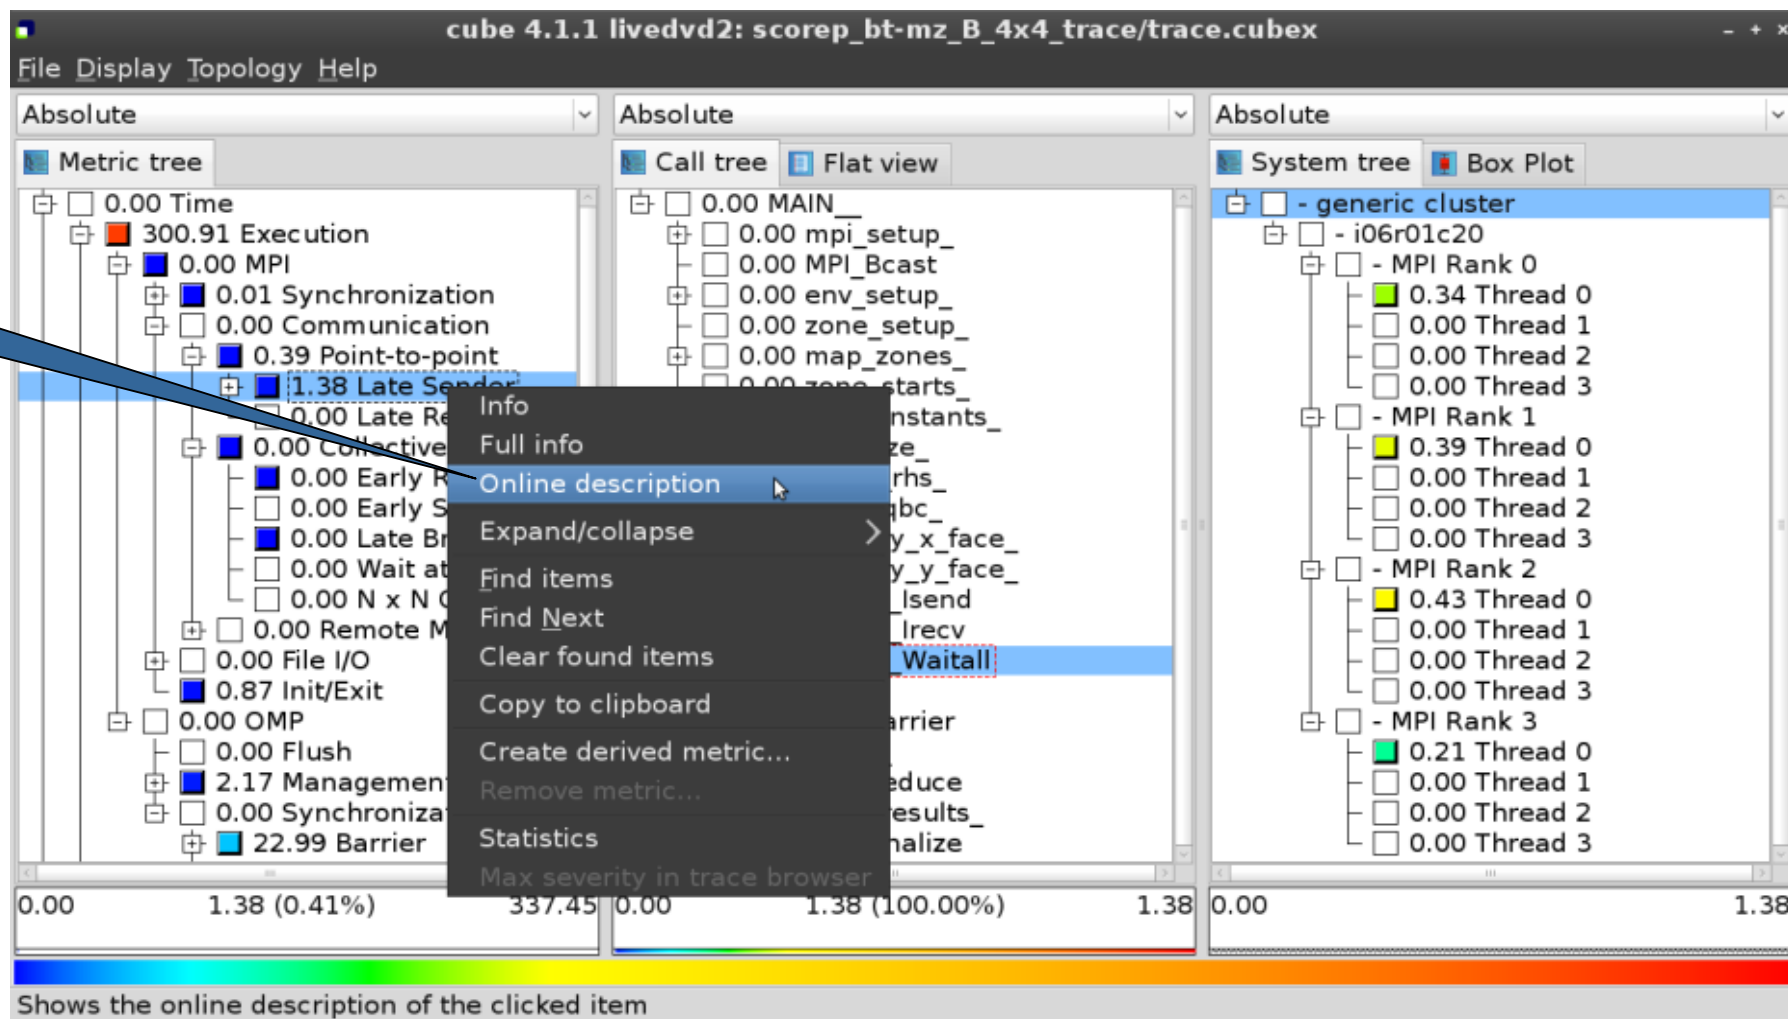




# Online metric description



Access online metric description via context menu





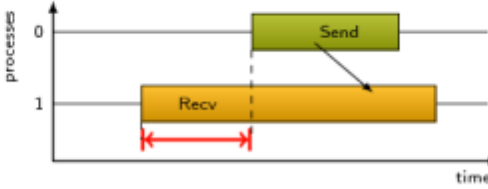
# Online metric description



**Performance properties**

**Late Sender Time**

**Description:**  
Refers to the time lost waiting caused by a blocking receive operation (e.g., `MPI_Recv` or `MPI_Wait`) that is posted earlier than the corresponding send operation.



If the receiving process is waiting for multiple messages to arrive (e.g., in an call to `MPI_Waitall`), the maximum waiting time is accounted, i.e., the waiting time due to the latest sender.

**Unit:**  
Seconds

**Diagnosis:**  
Try to replace `MPI_Recv` with a non-blocking receive `MPI_Irecv` that can be posted earlier, proceed concurrently with computation, and complete with a wait operation after the message is expected to have been sent. Try to post sends earlier, such that they are available when receivers need them. Note that outstanding messages (i.e., sent before the receiver is ready) will occupy internal message buffers, and that large numbers of posted receive buffers will also introduce message management overhead, therefore moderation is advisable.

**Parent:**  
[MPI Point-to-point Communication Time](#)

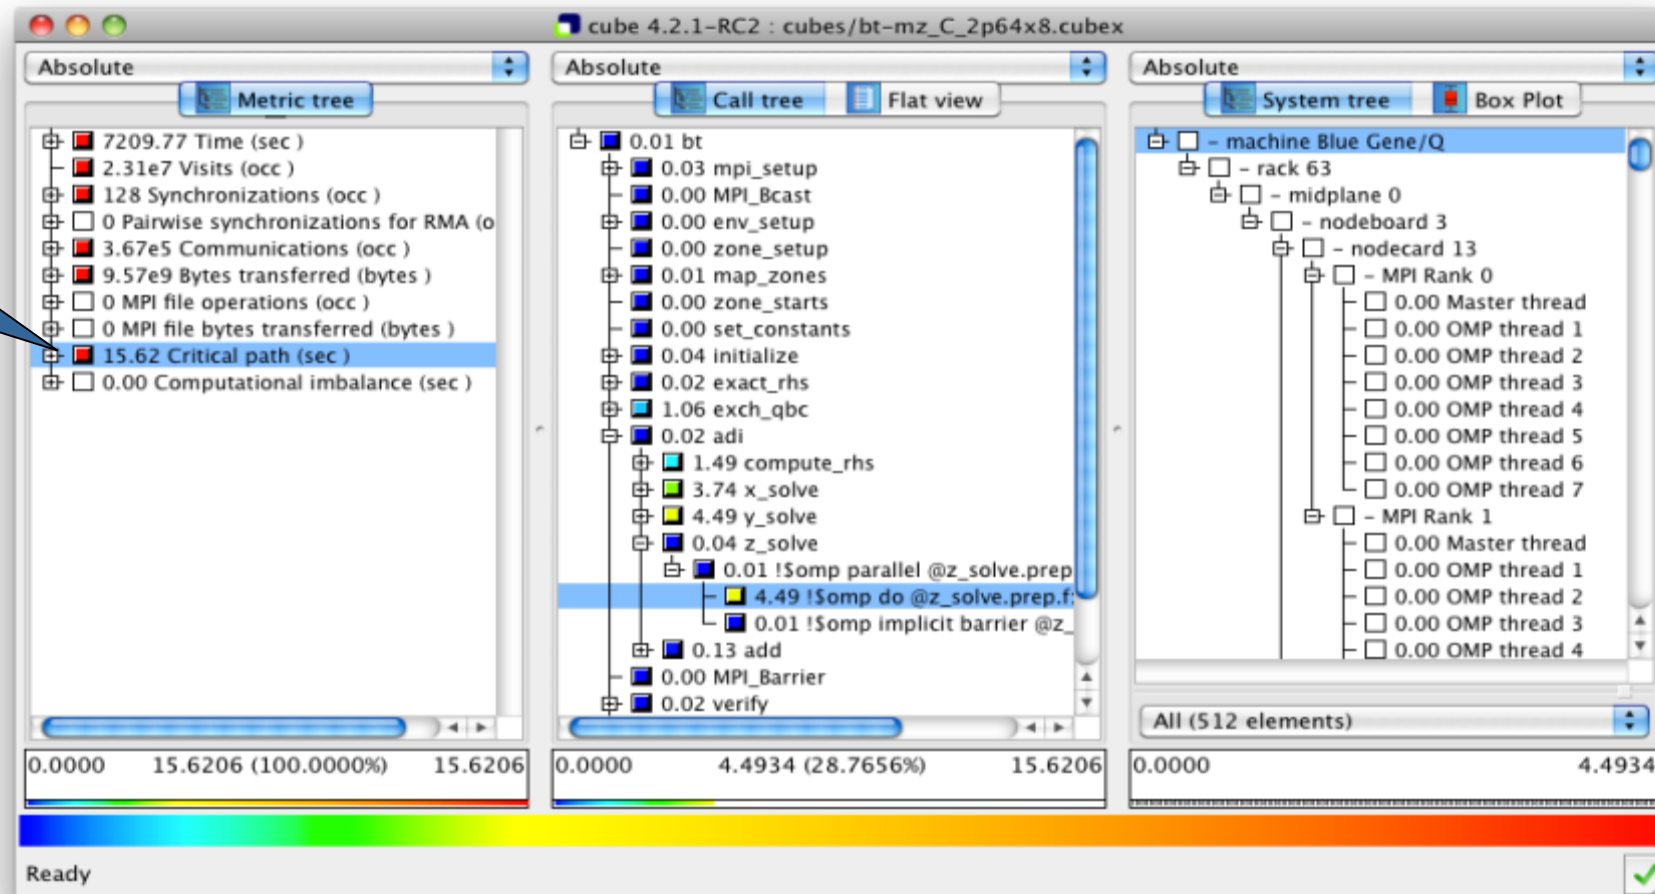
**Children:**

Close



# Critical-path analysis

Critical-path profile shows wall-clock time impact

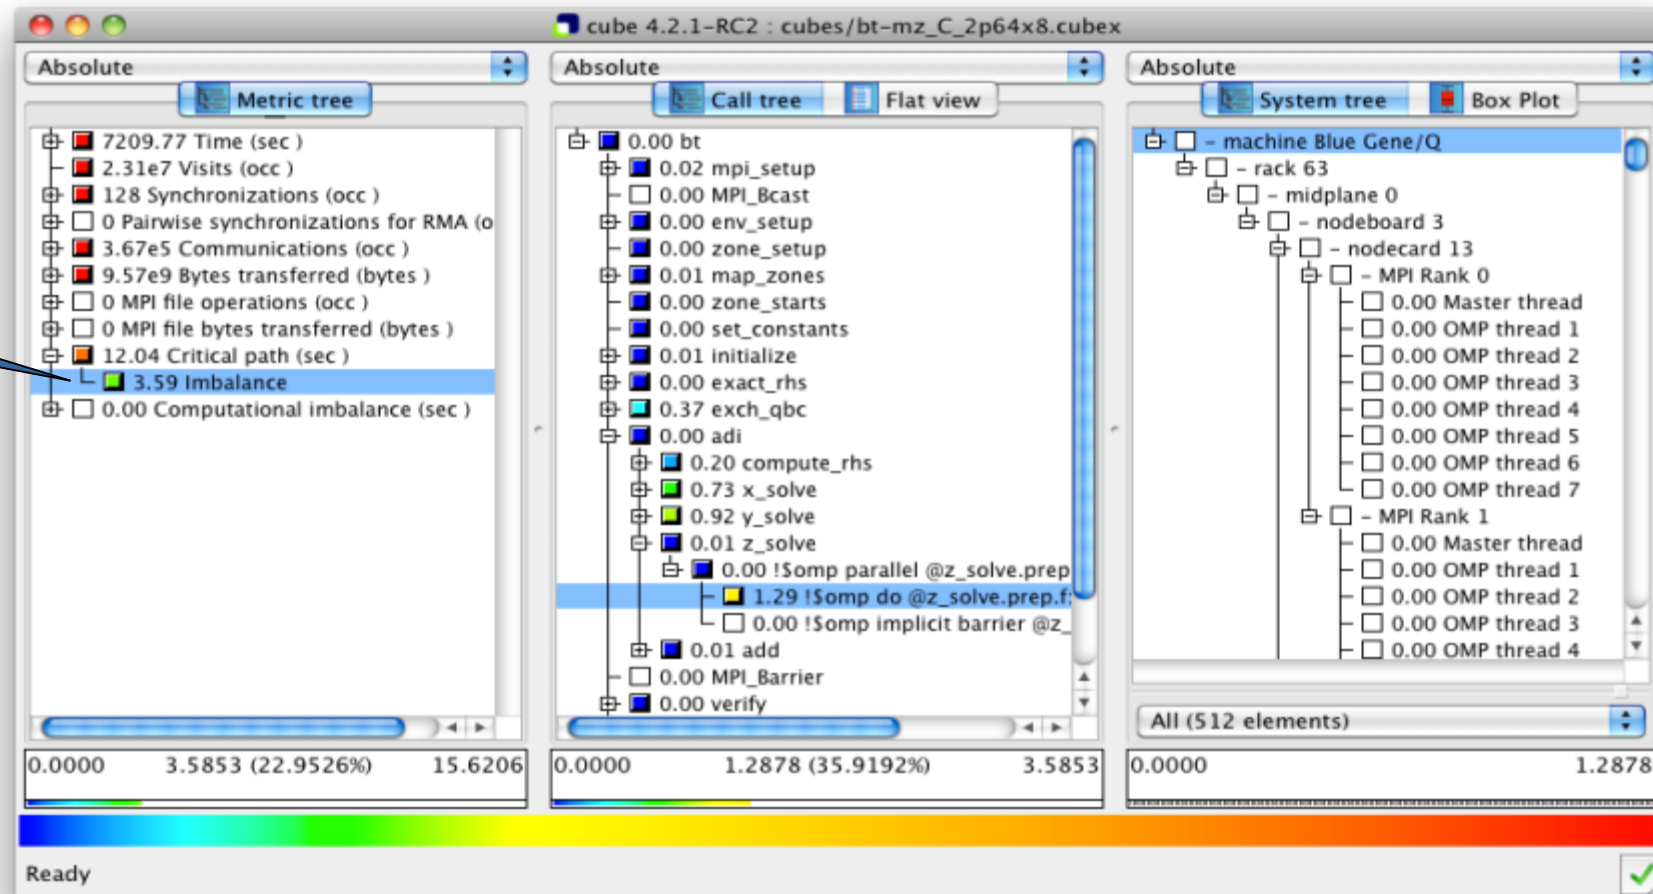




# Critical-path analysis

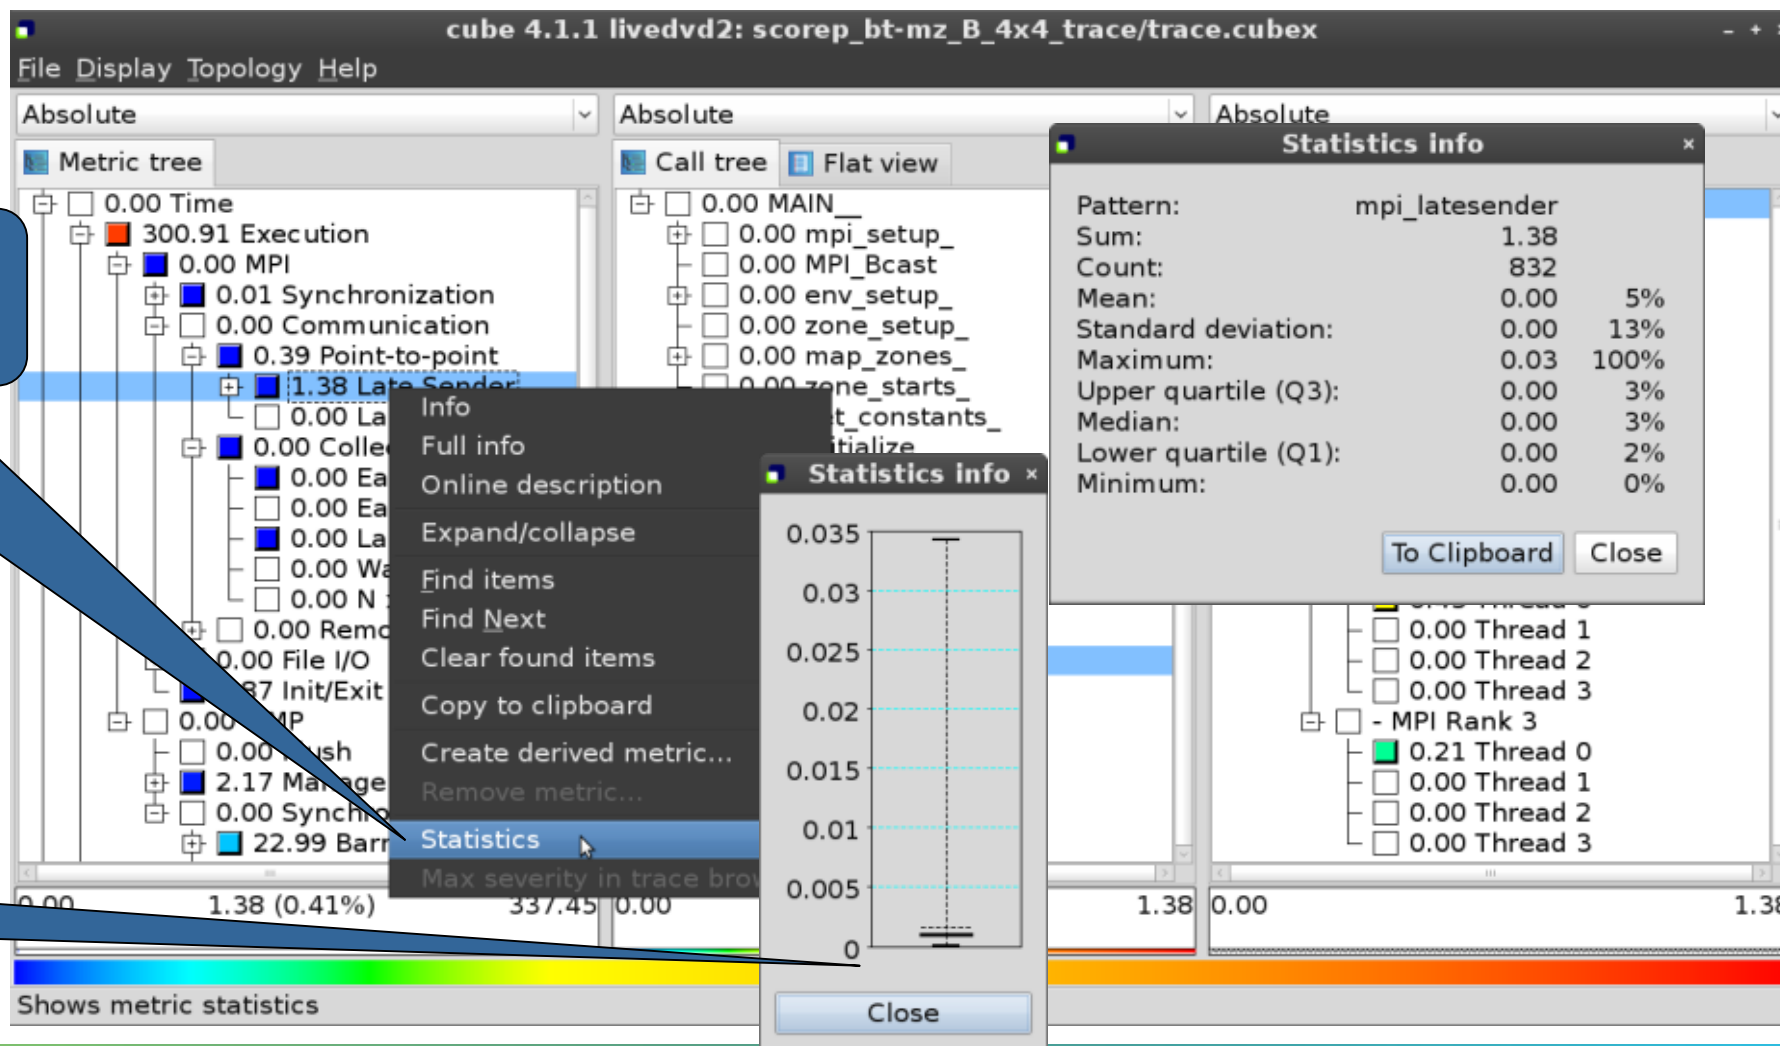


Critical-path imbalance highlights inefficient parallelism





# Pattern instance statistics





## Scalasca Trace Tools: Further information

---

- Collection of trace-based performance tools
  - Specifically designed for large-scale systems
  - Features an automatic trace analyzer providing wait-state, critical-path, and delay analysis
  - Supports MPI, OpenMP, POSIX threads, and hybrid MPI+OpenMP/Pthreads
- Available under 3-clause BSD open-source license
- Documentation & sources:
  - <http://www.scalasca.org>
- Contact:
  - mailto: [scalasca@fz-juelich.de](mailto:scalasca@fz-juelich.de)

