

Ferramentas de Avaliação de Desempenho

Roteiro

1 Scalasca, Score-p, Cube

2 hpctoolkit

Scalasca

<https://www.scalasca.org>



Scalasca

Exemplo: NAS Parallel Benchmarks (NPB)

```
scalasca/  
  NPB3.3.1-MZ/  
    NPB3.3-MZ-MPI/  
    NPB3.3-MZ-OMP/  
    NPB3.3-MZ-SER/  
  Changes.log  
  env_scalasca  
  README
```

Scalasca

Exemplo: NAS Parallel Benchmarks (NPB)

```
scalasca/  
  NPB3.3.1-MZ/  
    NPB3.3-MZ-MPI/  
    NPB3.3-MZ-OMP/  
    NPB3.3-MZ-SER/  
  Changes.log  
  env_scalasca  
  README
```

Scalasca

Preparando o ambiente

```
$ cat env_scalasca
```

```
module load openmpi/gnu/2.0.4.2  
module load scalasca/2.4_openmpi_gnu  
module load papi/5.5.1.0  
module load papi-devel/5.5.1.0
```

Preparando o ambiente

```
$ source env_scalasca
scalasca 2.4 for GNU OpenMPI loaded
Compiled with openMPI 2.0.4.2 and GNU compilers Red Hat 4.8.5-36

$ scalasca
Scalasca 2.4
Toolset for scalable performance analysis of large-scale parallel applications
usage: scalasca [-v][-n] action
  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
```

Scalasca

```
scalasca -instrument (skin/scorep)
```

```
scalasca/  
  NPB3.3.1-MZ/  
    NPB3.3-MZ-MPI/  
    NPB3.3-MZ-OMP/  
    NPB3.3-MZ-SER/  
  Changes.log  
  env_scalasca  
  README
```

Scalasca

```
scalasca -instrument (skin/scorep)
```

```
scalasca/  
  NPB3.3.1-MZ/  
    NPB3.3-MZ-MPI/  
      bin/  
      BT-MZ/  
      common/  
      config/  
      LU-MZ/  
      SP-MZ/  
      sys/  
      Makefile  
      README  
      README.install  
    NPB3.3-MZ-OMP/  
    NPB3.3-MZ-SER/  
  Changes.log  
  env_scalasca  
  README
```

Scalasca

```
scalasca -instrument (skin/scorep)
```

```
scalasca/  
  NPB3.3.1-MZ/  
    NPB3.3-MZ-MPI/  
      bin/  
      BT-MZ/  
      common/  
      config/  
      LU-MZ/  
      SP-MZ/  
      sys/  
      Makefile  
      README  
      README.install  
    NPB3.3-MZ-OMP/  
    NPB3.3-MZ-SER/  
  Changes.log  
  env_scalasca  
  README
```

Scalasca

```
scalasca -instrument (skin/scorep)
```

```
config/  
  NAS.samples  
  make.def -> make_scalasca.def  
  make.def.template  
  make_scalasca.def  
  suite.def  
  suite.def.template
```

Scalasca

```
scalasca -instrument (skin/scorep)
```

```
config/  
  NAS.samples  
  make.def -> make_scalasca.def  
  make.def.template  
  make_scalasca.def  
  suite.def  
  suite.def.template
```

Scalasca

```
scalasca -instrument (skin/scorep)
```

```
$ cat make_scalasca.def
```

```
#-----  
# This is the fortran compiler used for fortran programs  
#-----  
#F77 = mpif77  
F77 = scalasca -instrument mpif77  
#F77 = scorep mpif77  
  
#-----  
# This is the C compiler used for C programs  
#-----  
#CC = mpicc  
CC = scalasca -instrument mpicc  
#CC = scorep mpicc
```

Scalasca

NPB: benchmark, classe e número de processos MPI

```
config/  
  NAS.samples  
  make.def -> make_scalasca.def  
  make.def.template  
  make_scalasca.def  
  suite.def  
  suite.def.template
```

Scalasca

NPB: benchmark, classe e número de processos MPI

```
config/  
  NAS.samples  
  make.def -> make_scalasca.def  
  make.def.template  
  make_scalasca.def  
  suite.def  
  suite.def.template
```

Estudo de caso

NPB: benchmark, classe e número de processos MPI

```
$ cat suite.def
```

```
# config/suite.def
# This file is used to build several benchmarks with a single command.
# Typing "make suite" in the main directory will build all the benchmarks
# specified in this file.
# Each line of this file contains a benchmark name, class, and number
# of nodes. The name is one of "sp-mz", "bt-mz", and "lu-mz".
# The class is one of "S", "W", and "A" through "F".
# No blank lines.
# The following example builds serial sample sizes of all benchmarks.
#sp-mz S 1
#lu-mz S 1
#bt-mz S 2
bt-mz S 1
bt-mz S 2
bt-mz S 4
bt-mz W 1
bt-mz W 2
bt-mz W 4
bt-mz W 8
bt-mz W 16
```

Estudo de caso

NPB: compilação

```
$ cd ..  
$ make suite %compila o NPB  
$ cd bin
```

Estudo de caso

NPB: compilação

```
$ ls -Al
bt-mz.S.1
bt-mz.S.2
bt-mz.S.4
bt-mz.W.1
bt-mz.W.2
bt-mz.W.4
BULL_srun_scan_prof.sh
BULL_srun_scan_trace.sh
BULL_srun_scan_trace_filt.sh
```

Scalasca

scalasca -analyze: coleta de dados da execução

```
$ ls -Al
bt-mz.S.1
bt-mz.S.2
bt-mz.S.4
bt-mz.W.1
bt-mz.W.2
bt-mz.W.4
BULL_srun_scan_prof.sh
BULL_srun_scan_trace.sh
BULL_srun_scan_trace_filt.sh
```

Scalasca

```
BULL_srun_scalasca.sh
```

Estudo de caso

NPB: submetendo job

```
$ sbatch BULL_srun_scan_prof.sh bt-mz S
```

```
Submitted batch job 437607
```

```
$ squeue -u $USER
```

JOBID	PARTITION	NAME	USER	ST	TIME	NODES	NODELIST (REASON)
437607	treinamen	NPB_BT-M	professo	R	0:02	1	sdumont3000

Estudo de caso

NPB: perfil de desempenho

```
$ ls -Al
bt-mz.S.1
bt-mz.S.2
bt-mz.S.4
bt-mz.W.1
bt-mz.W.2
bt-mz.W.4
BULL_srun_scan_prof.sh
BULL_srun_scan_trace.sh
BULL_srun_scan_trace_filt.sh
scorep_bt-mz_S_sum_MPI-1_OMP-1_JOBID-437607/
```

Estudo de caso

NPB: perfil de desempenho

```
scorep_bt-mz_S_sum_MPI-1_OMP-1_JOBID-437607/
```

```
profile.cubex
```

```
summary.cubex
```

```
scorep.score
```

```
scorep.cfg
```

```
scorep.log
```

```
slurm-437607.out
```

Estudo de caso

NPB: perfil de desempenho

`scorep_bt-mz_S_sum_MPI-1_OMP-1_JOBID-437607/`

<code>profile.cubex</code>	--> análise básica, a partir de dados coletados durante execução
<code>summary.cubex</code>	--> análise mais detalhada
<code>scorep.score</code>	--> relatório formato texto com a análise
<code>scorep.cfg</code>	--> configuração da coleta de dados
<code>scorep.log</code>	--> output da aplicação
<code>slurm-437607.out</code>	--> output do SLURM

Estudo de caso

NPB: perfil de desempenho

```
scorep_bt-mz_S_sum_MPI-1_OMP-1_JOBID-437607/
```

```
profile.cubex
```

```
summary.cubex
```

```
scorep.score
```

```
scorep.cfg
```

```
scorep.log
```

```
slurm-437607.out
```

Estudo de caso

```
$ cat slurm-437607.out
```

```
Cluster configuration:
```

```
===
```

```
Partition: treinamento
```

```
Number of nodes: 1
```

```
Number of MPI processes: 1 ( 1 nodes)
```

```
Number of MPI processes per node: 1
```

```
Number of threads per MPI process: 1
```

```
NPB Benchmark: bt-mz
```

```
Bechmark class problem: S
```

```
scalasca 2.4 for GNU OpenMPI loaded
```

```
Compiled with openMPI 2.0.4.2 and GNU compilers Red Hat 4.8.5-36
```

```
S=C=A=N: Scalasca 2.4 runtime summarization
```

```
S=C=A=N: ./scorep_bt-mz_lx1_sum experiment archive
```

```
S=C=A=N: Tue Jan 28 14:17:02 2020: Collect start
```

```
/usr/bin/srun --resv-ports -n 1 /scratch/treinamento/professor/MC1-I/tools/scalasca/NPB3.3.1-MZ/NP
```

```
[1580231823.240921] [sdumont5000:73441:0] mxm.c:196 MXM WARN The 'ulimit -s' on the sys
```

```
[1580231823.242634] [sdumont5000:73441:0] mxm.c:196 MXM WARN The 'ulimit -s' on the sys
```

```
NAS Parallel Benchmarks (NPB3.3-MZ-MPI) - BT-MZ MPI+OpenMP Benchmark
```

```
Number of zones: 2 x 2
```

```
Iterations: 60 dt: 0.010000
```

```
Number of active processes: 1
```

Estudo de caso (cont.)

Use the default load factors with threads

Total number of threads: 1 (1.0 threads/process)

Calculated speedup = 1.00

Time step 1

Time step 20

Time step 40

Time step 60

Verification being performed for class S

accuracy setting for epsilon = 0.10000000000000E-07

Comparison of RMS-norms of residual

1	0.1047687395830E+04	0.1047687395830E+04	0.1751386499571E-12
2	0.9419911314792E+02	0.9419911314792E+02	0.1478425555772E-13
3	0.2124737403068E+03	0.2124737403068E+03	0.9002435039286E-13
4	0.1422173591794E+03	0.1422173591794E+03	0.3089634277625E-12
5	0.1135441572375E+04	0.1135441572375E+04	0.3103895484466E-13

Comparison of RMS-norms of solution error

1	0.1775416062982E+03	0.1775416062982E+03	0.1922618237923E-12
2	0.1875540250835E+02	0.1875540250835E+02	0.1558955269742E-12
3	0.3863334844506E+02	0.3863334844506E+02	0.1105356386074E-12
4	0.2634713890362E+02	0.2634713890362E+02	0.3991337551951E-13
5	0.1965566269675E+03	0.1965566269675E+03	0.2336704854379E-12

Verification Successful

BT-MZ Benchmark Completed.

Class	=			S
Size	=	24x	24x	6
Iterations	=			60

Estudo de caso (cont.)

```
Time in seconds =                0.35
Total processes =                  1
Total threads   =                  1
Mop/s total     =                1093.41
Mop/s/thread    =                1093.41
Operation type  =                floating point
Verification    =                SUCCESSFUL
Version         =                  3.3.1
Compile date    =                21 Jan 2020
```

Compile options:

```
F77           = scalasca -instrument mpif77
FLINK         = $(F77)
F_LIB         = (none)
F_INC         = (none)
FFLAGS       = -O3 -fopenmp
FLINKFLAGS    = $(FFLAGS)
RAND         = (none)
```

Please send all errors/feedbacks to:

NPB Development Team
npb@nas.nasa.gov

```
S=C=A=N: Tue Jan 28 14:17:03 2020: Collect done (status=0) 1s
```

```
S=C=A=N: ./scorep_bt-mz_lx1_sum complete.
```

```
INFO: Post-processing runtime summarization report...
```

```
/opt/bullxde/utils/scalasca/openmpi-gnu/scorep/bin/scorep-score -r ./scorep_bt-mz_S_sum_MPI-1_OMP-
```

Estudo de caso (cont.)

INFO: Score report written to `./scorep_bt-mz_S_sum_MPI-1_OMP-1_JOBID-437607/scorep.score`

Estudo de caso

NPB: submetendo job

```
$ sbatch BULL_srun_scan_prof.sh bt-mz W
```

```
Submitted batch job 437632
```

```
$ squeue -u $USER
```

JOBID	PARTITION	NAME	USER	ST	TIME	NODES	NODELIST (REASON)
% 437632	treinamen	NPB_BT-M	professo	R	0:02	1	sdumont3000

Estudo de caso

NPB: perfil de desempenho

```
scorep_bt-mz_W_sum_MPI-1_OMP-1_JOBID-437632/  
  profile.cubex  
  summary.cubex  
  scorep.score  
  scorep.cfg  
  scorep.log  
  slurm-437632.out
```

Estudo de caso

```
sbatch --nodes=1 --ntasks=1 BULL_srun_scan_prof.sh bt-mz W
```

```
Number of zones:   4 x   4
Iterations: 200   dt: 0.000800
Number of active processes:   1
```

Use the default load factors with threads

```
Total number of threads:   1 ( 1.0 threads/process)
```

```
Calculated speedup =   1.00
```

BT-MZ Benchmark Completed.

```
Class           =                               W
Size            =          64x   64x   8
Iterations      =                               200
Time in seconds =                               12.59
Total processes =                               1
Total threads   =                               1
Mop/s total    =          1140.40
Mop/s/thread   =          1140.40
Operation type  =          floating point
Verification   =          SUCCESSFUL
Version        =                               3.3.1
```

```
S=C=A=N: Tue Jan 28 14:57:17 2020: Collect done (status=0) 14s
```

Estudo de caso

```
sbatch --nodes=1 --ntasks=2 BULL_srun_scan_prof.sh bt-mz W
```

```
Number of zones:   4 x 4  
Iterations: 200   dt: 0.000800  
Number of active processes: 2
```

Use the default load factors with threads

```
Total number of threads: 2 ( 1.0 threads/process)
```

```
Calculated speedup = 1.98
```

BT-MZ Benchmark Completed.

```
Class      = W  
Size       = 64x 64x 8  
Iterations = 200  
Time in seconds = 6.46  
Total processes = 2  
Total threads = 2  
Mop/s total = 2219.05  
Mop/s/thread = 1109.53  
Operation type = floating point  
Verification = SUCCESSFUL  
Version      = 3.3.1
```

```
S=C=A=N: Tue Jan 28 15:01:09 2020: Collect done (status=0) 8s
```

Estudo de caso

```
sbatch --nodes=1 --ntasks=4 BULL_srun_scan_prof.sh bt-mz W
```

```
Number of zones:   4 x 4  
Iterations: 200   dt: 0.000800  
Number of active processes: 4
```

Use the default load factors with threads

```
Total number of threads: 4 ( 1.0 threads/process)
```

```
Calculated speedup = 3.95
```

BT-MZ Benchmark Completed.

```
Class           = W  
Size            = 64x 64x 8  
Iterations      = 200  
Time in seconds = 3.40  
Total processes = 4  
Total threads   = 4  
Mop/s total     = 4206.28  
Mop/s/thread    = 1051.57  
Operation type  = floating point  
Verification    = SUCCESSFUL  
Version         = 3.3.1
```

```
S=C=A=N: Tue Jan 28 15:01:34 2020: Collect done (status=0) 5s
```

Estudo de caso

```
sbatch --nodes=1 --ntasks=8 BULL_srun_scan_prof.sh bt-mz W
```

```
Number of zones:   4 x   4  
Iterations: 200   dt: 0.000800  
Number of active processes:      8
```

```
Use the default load factors with threads  
Total number of threads:      8 ( 1.0 threads/process)
```

```
Calculated speedup =      4.87
```

```
BT-MZ Benchmark Completed.
```

```
Class      =                               W  
Size       =          64x   64x   8  
Iterations =                               200  
Time in seconds =          2.82  
Total processes =                               8  
Total threads =                               8  
Mop/s total =          5086.41  
Mop/s/thread =          635.80  
Operation type =          floating point  
Verification =          SUCCESSFUL  
Version     =                               3.3.1
```

```
S=C=A=N: Wed Jan 29 10:51:43 2020: Collect done (status=0) 4s
```

Estudo de caso

```
sbatch --nodes=1 --ntasks=16 BULL_srun_scan_prof.sh bt-mz W
```

```
Number of zones:   4 x   4  
Iterations: 200   dt: 0.000800  
Number of active processes:   16
```

```
Use the default load factors with threads  
Total number of threads:   16 ( 1.0 threads/process)
```

```
Calculated speedup =   4.87
```

```
BT-MZ Benchmark Completed.
```

```
Class           =                               W  
Size            =          64x   64x   8  
Iterations      =                               200  
Time in seconds =                               2.84  
Total processes =                               16  
Total threads   =                               16  
Mop/s total     =          5047.62  
Mop/s/thread    =             315.48  
Operation type  =          floating point  
Verification    =          SUCCESSFUL  
Version         =             3.3.1
```

```
S=C=A=N: Wed Jan 29 10:52:03 2020: Collect done (status=0) 6s
```

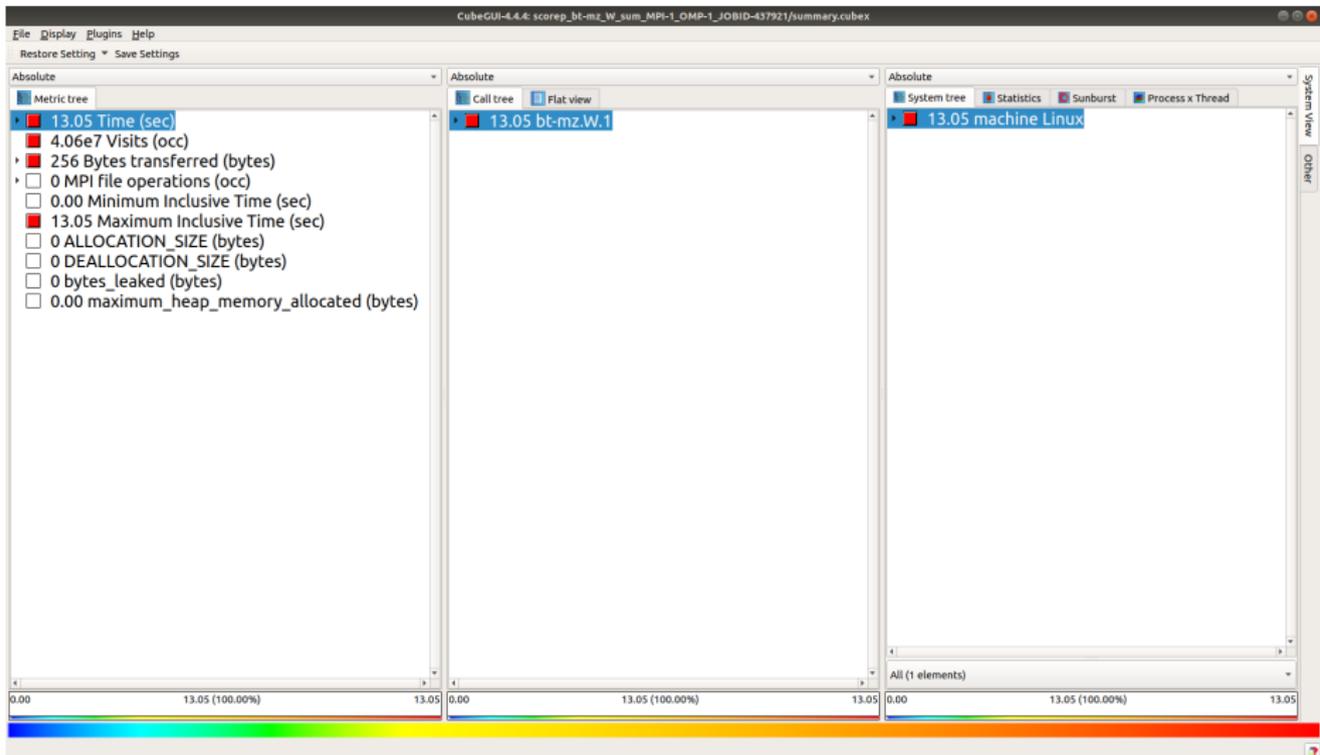
Visualizando no Cube

NPB: estudo de caso

```
$ cd profiling/NUMNODES-1/scorep_bt-mz_W_sum_MPI-1_OMP-1_JOBID-437921  
$ cube summary.cubex
```

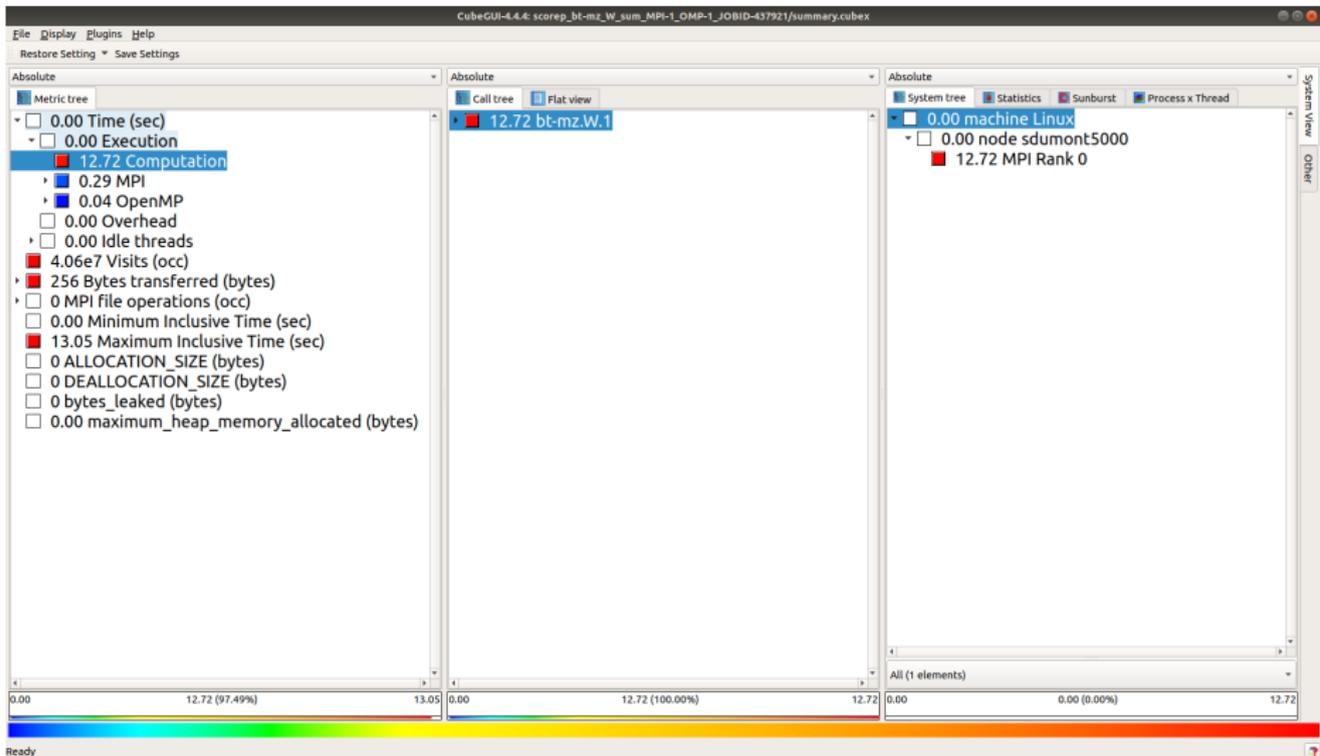
Visualizando no Cube

`-nodes=1 -ntasks=1`



Visualizando no Cube

`-nodes=1 -ntasks=1`



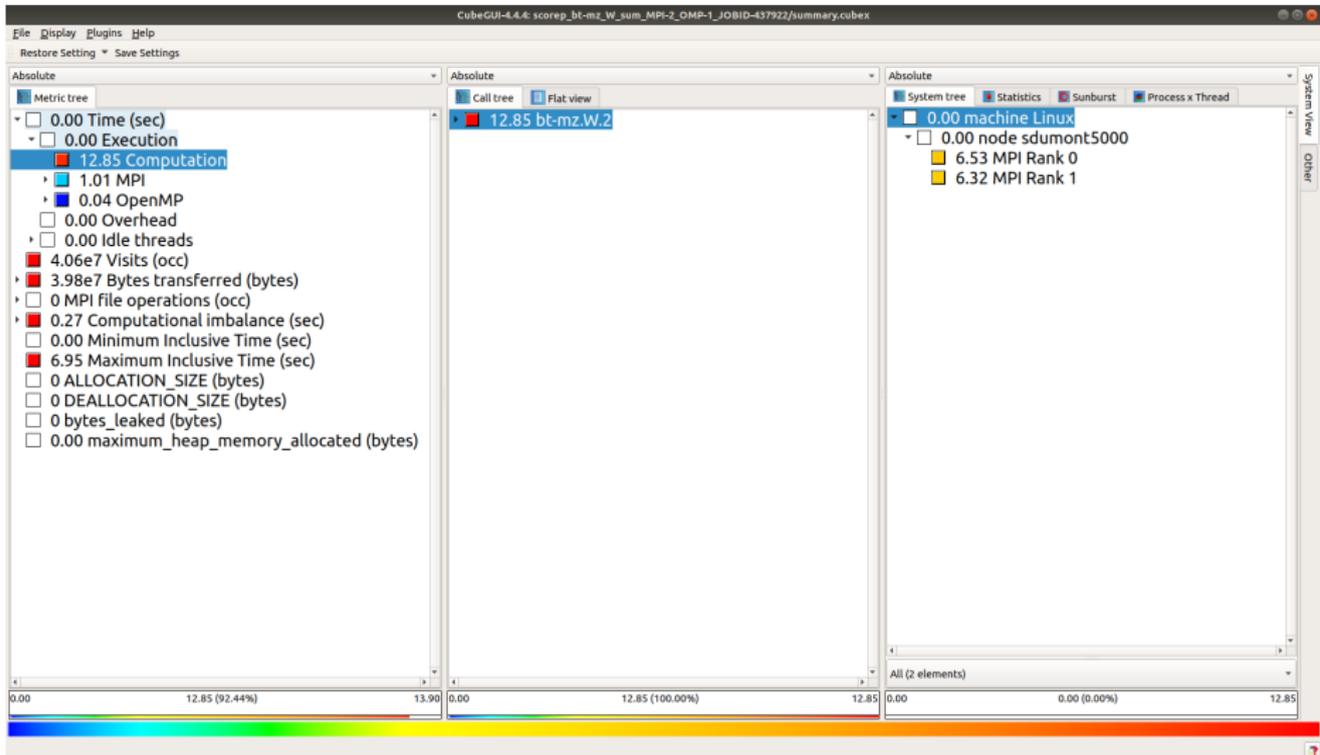
Visualizando no Cube

divisão do tempo de processamento

- Tempo de computação: 12.72s
- Tempo de MPI: 0.29s
- Tempo de OpenMP: 0.04s

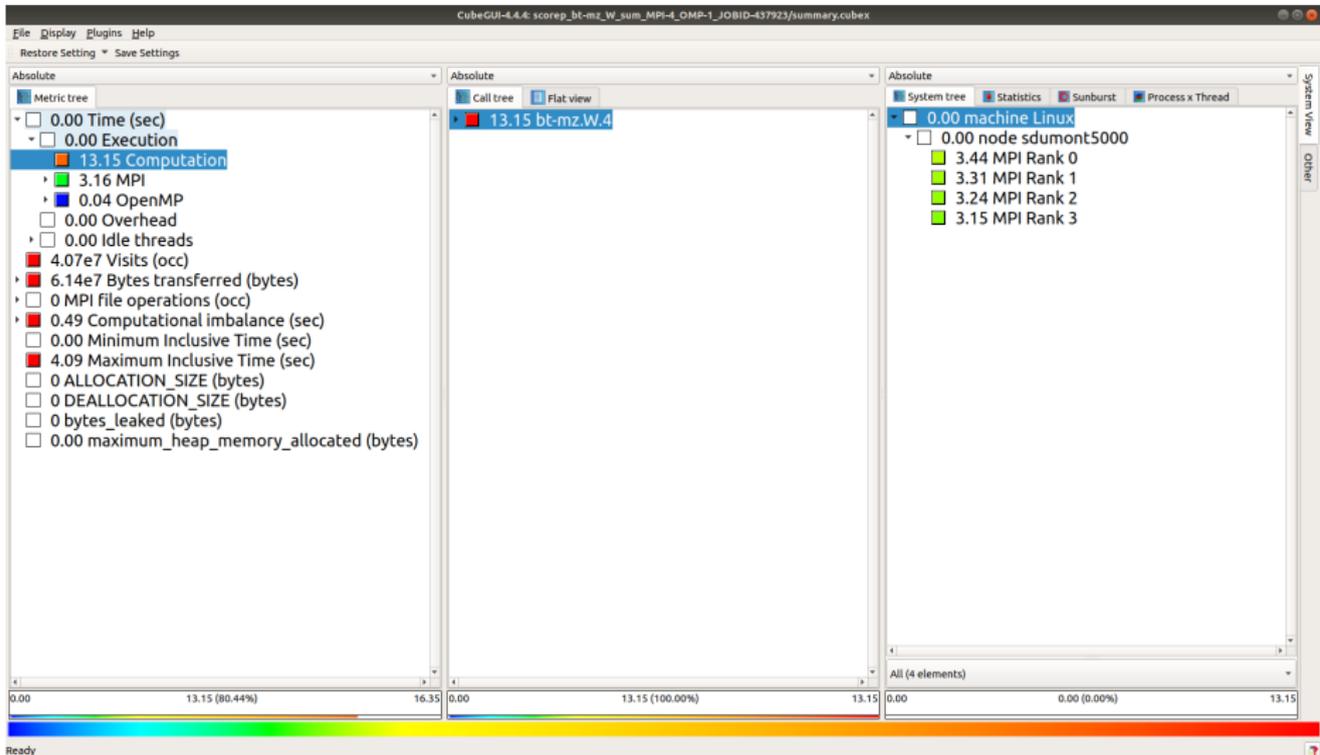
Visualizando no Cube

`-nodes=1 -ntasks=2`



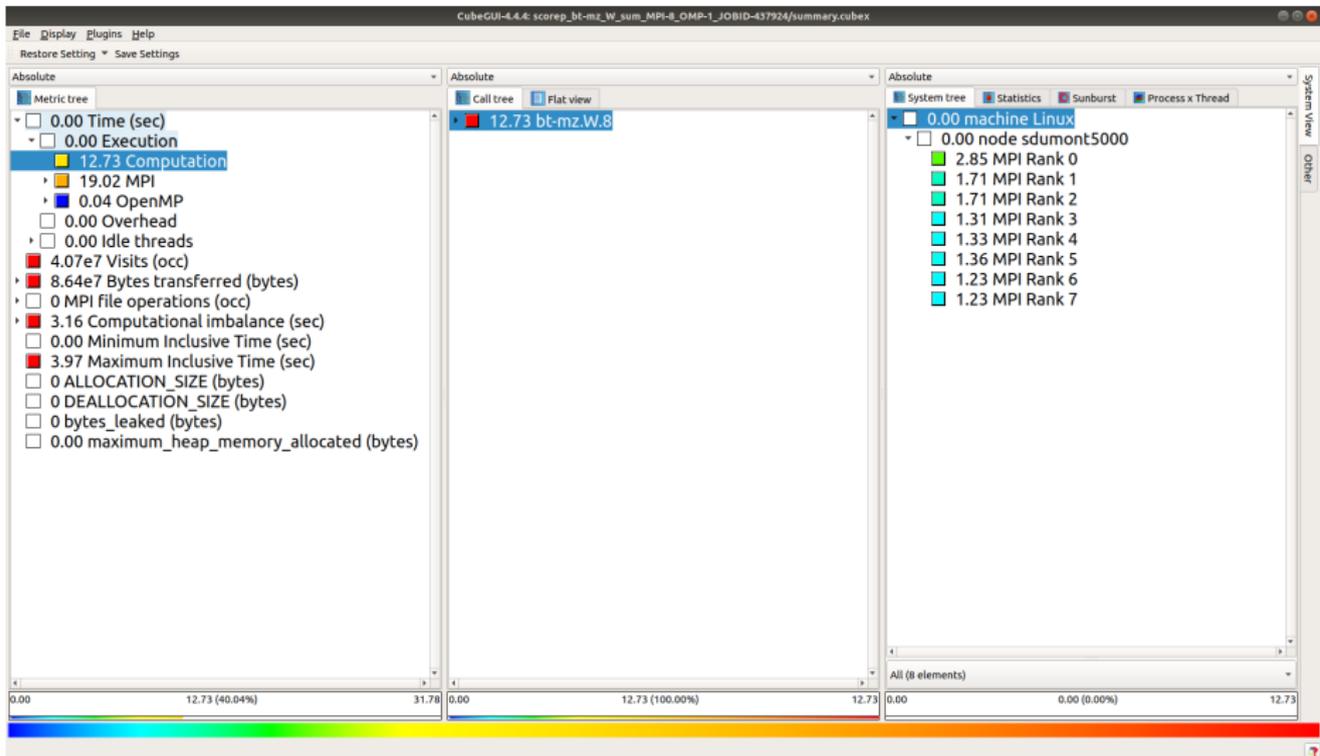
Visualizando no Cube

`-nodes=1 -ntasks=4`



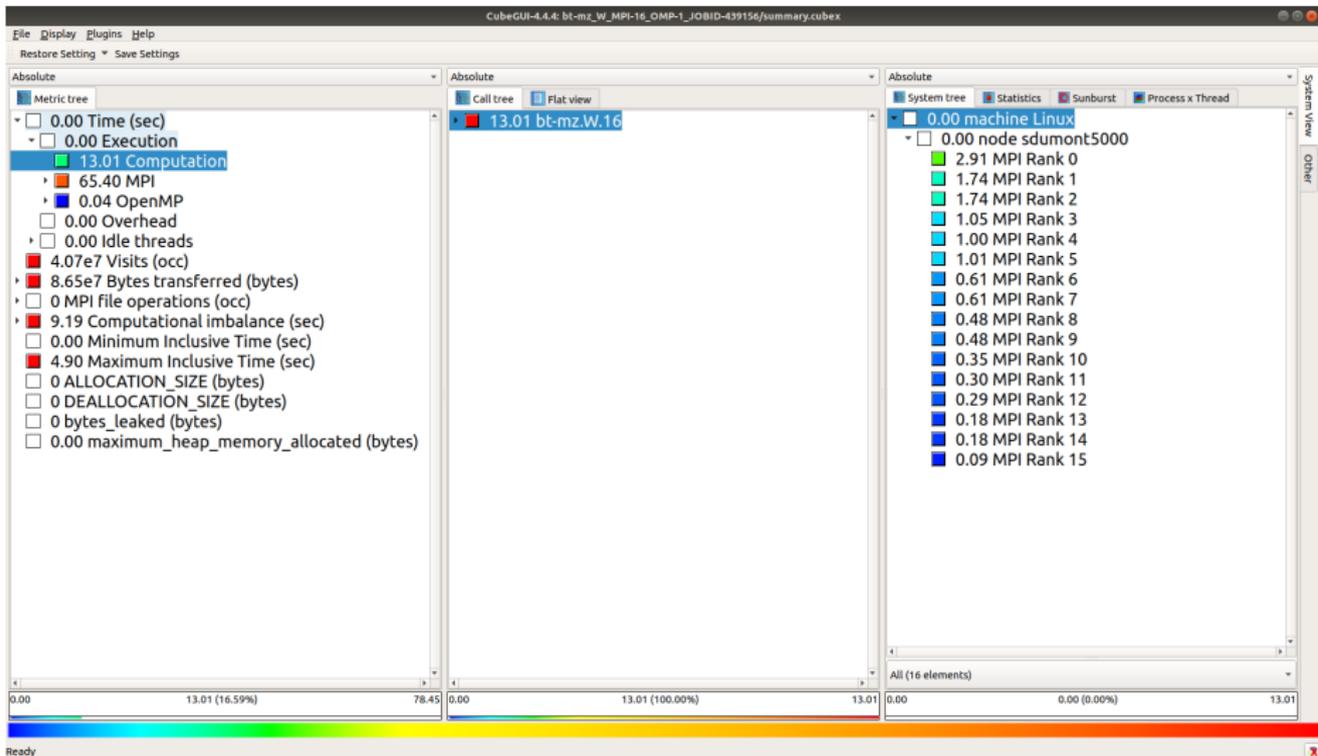
Visualizando no Cube

`-nodes=1 -ntasks=8`



Visualizando no Cube

`-nodes=1 -ntasks=16`



BT-MZ *benchmark*: divisão de domínio

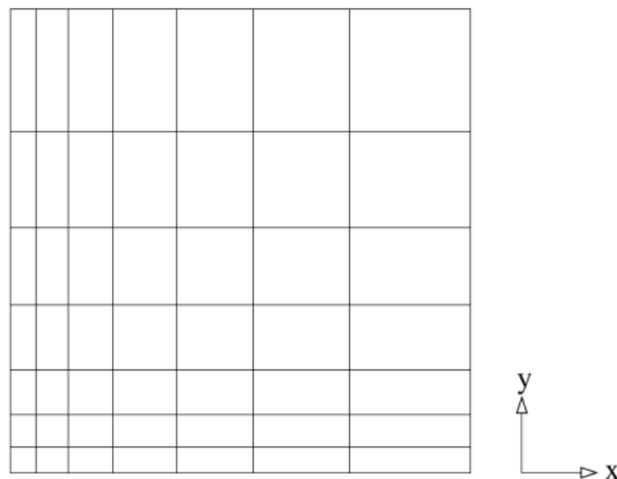


Figure 3: Example of uneven mesh tiling (horizontal cut through mesh system) for the BT-MZ benchmark.

Definição

Balanco de carga de computação (LB):

$$LB = \frac{avg(tcomp)}{max(tcomp)}$$

FONTE: <https://pop-coe.eu>

Visualizando no Cube

`-nodes=1 -ntasks=16`

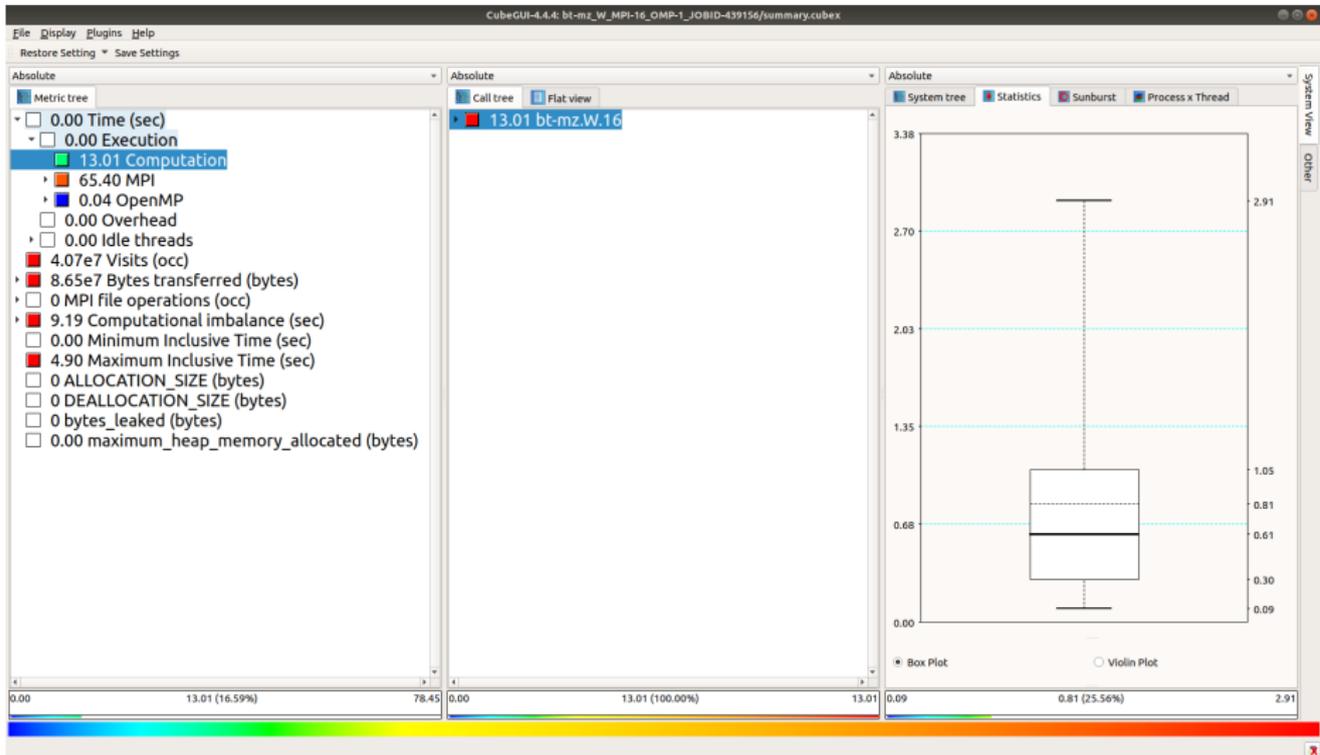
The screenshot displays the Cube GUI interface for a job titled "CubeGUI-4.4.4: bt-ml_W_MPI-16_OMP-1_JOBID-439156/summary.cubex". The interface is divided into three main panels, each showing a different view of the performance data:

- Left Panel (Metric tree):** Shows a hierarchical tree of metrics. The total time is 13.01 seconds, with 13.01 seconds spent on computation. Other metrics include 65.40 MPI, 0.04 OpenMP, 4.07e7 visits, and 8.65e7 bytes transferred.
- Middle Panel (Call tree):** Shows a call tree view with a single node: "13.01 bt-mz.W.16".
- Right Panel (System tree):** Shows a system tree view for the "0.00 machine Linux" node, listing 16 MPI ranks with their respective times: MPI Rank 0 (2.91s), MPI Rank 1 (1.74s), MPI Rank 2 (1.74s), MPI Rank 3 (1.05s), MPI Rank 4 (1.00s), MPI Rank 5 (1.01s), MPI Rank 6 (0.61s), MPI Rank 7 (0.61s), MPI Rank 8 (0.48s), MPI Rank 9 (0.48s), MPI Rank 10 (0.35s), MPI Rank 11 (0.30s), MPI Rank 12 (0.29s), MPI Rank 13 (0.18s), MPI Rank 14 (0.18s), and MPI Rank 15 (0.09s).

At the bottom of the GUI, there is a color-coded progress bar and a status indicator that reads "Ready".

Visualizando no Cube

`-nodes=1 -ntasks=16`



Cálculo

Balanço de carga de computação (LB)

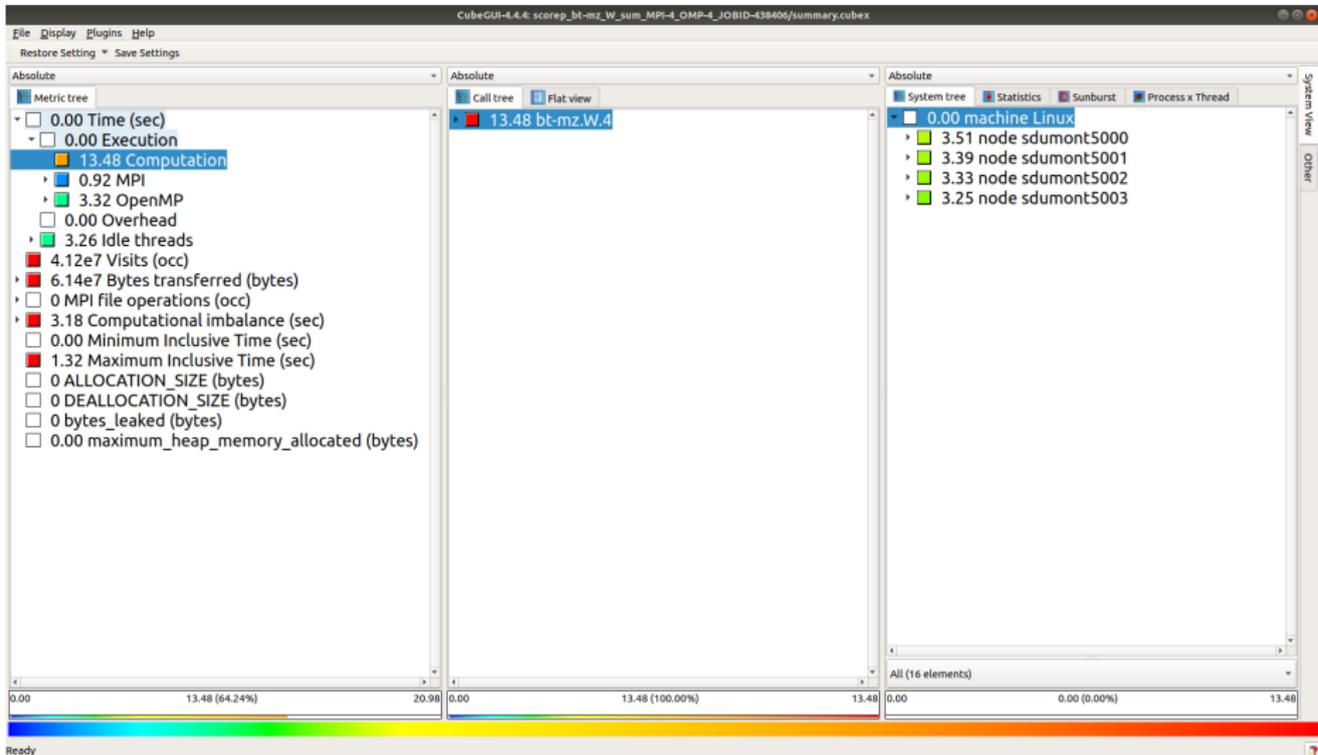
$$LB = \frac{avg(tcomp)}{max(tcomp)}$$

$$LB = \frac{0.81}{2.91}$$

$$LB = 0.28$$

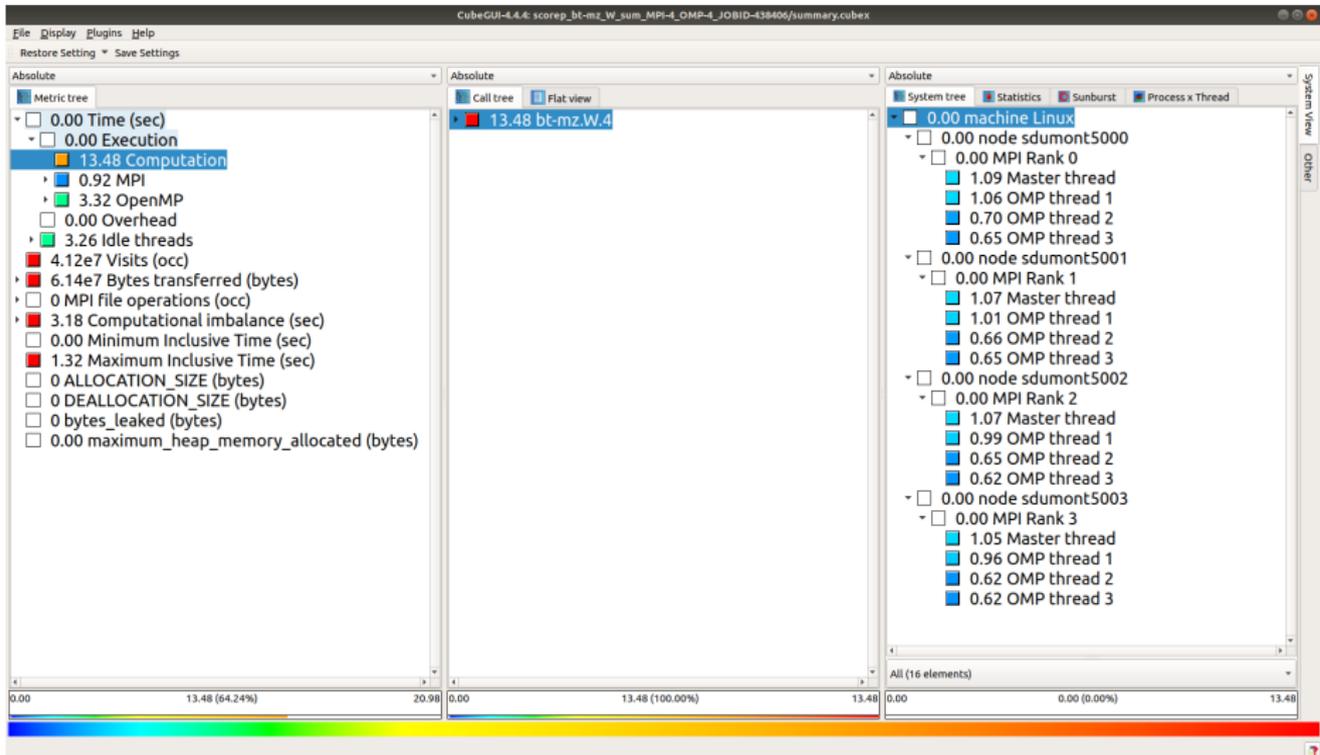
Visualizando no Cube

`-nodes=4 -ntasks=4 -cpus-per-task=4`



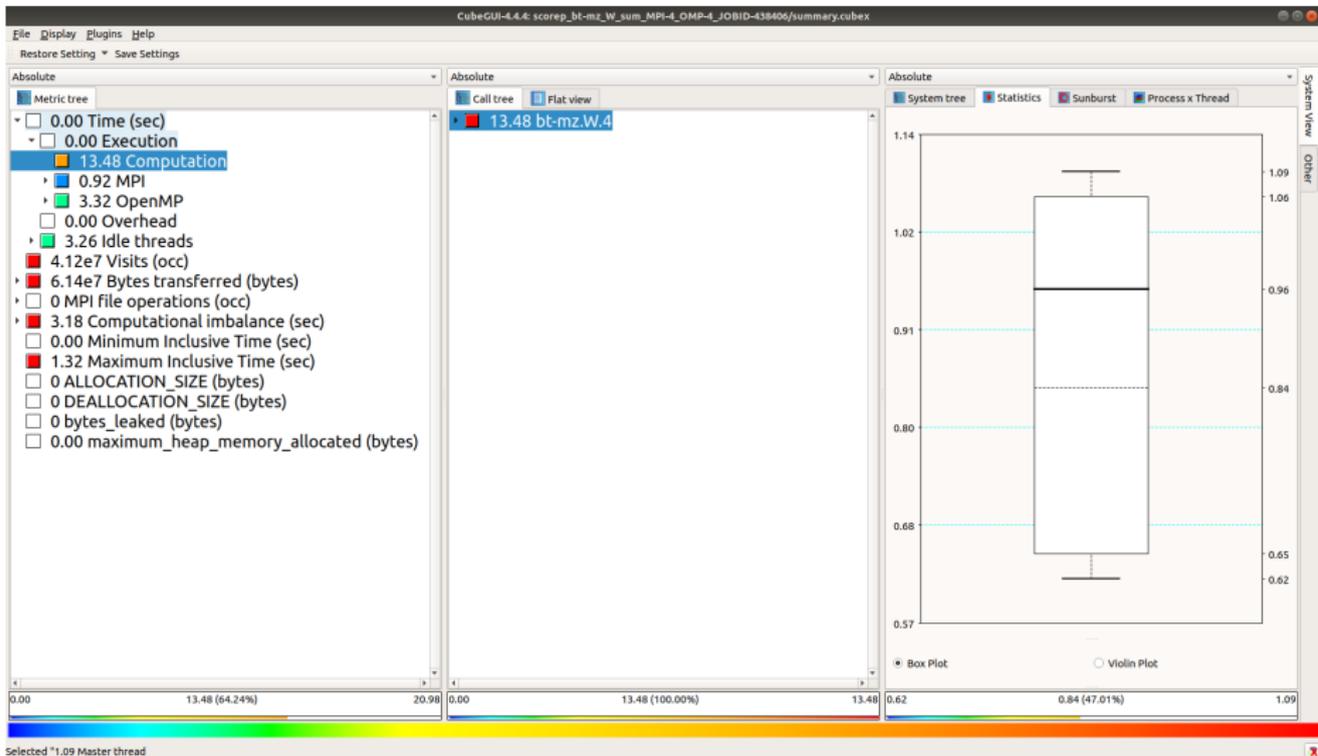
Visualizando no Cube

-nodes=4 -ntasks=4 -cpus-per-task=4



Visualizando no Cube

`-nodes=4 -ntasks=4 -cpus-per-task=4`



Cálculo

Balanço de carga de computação (LB)

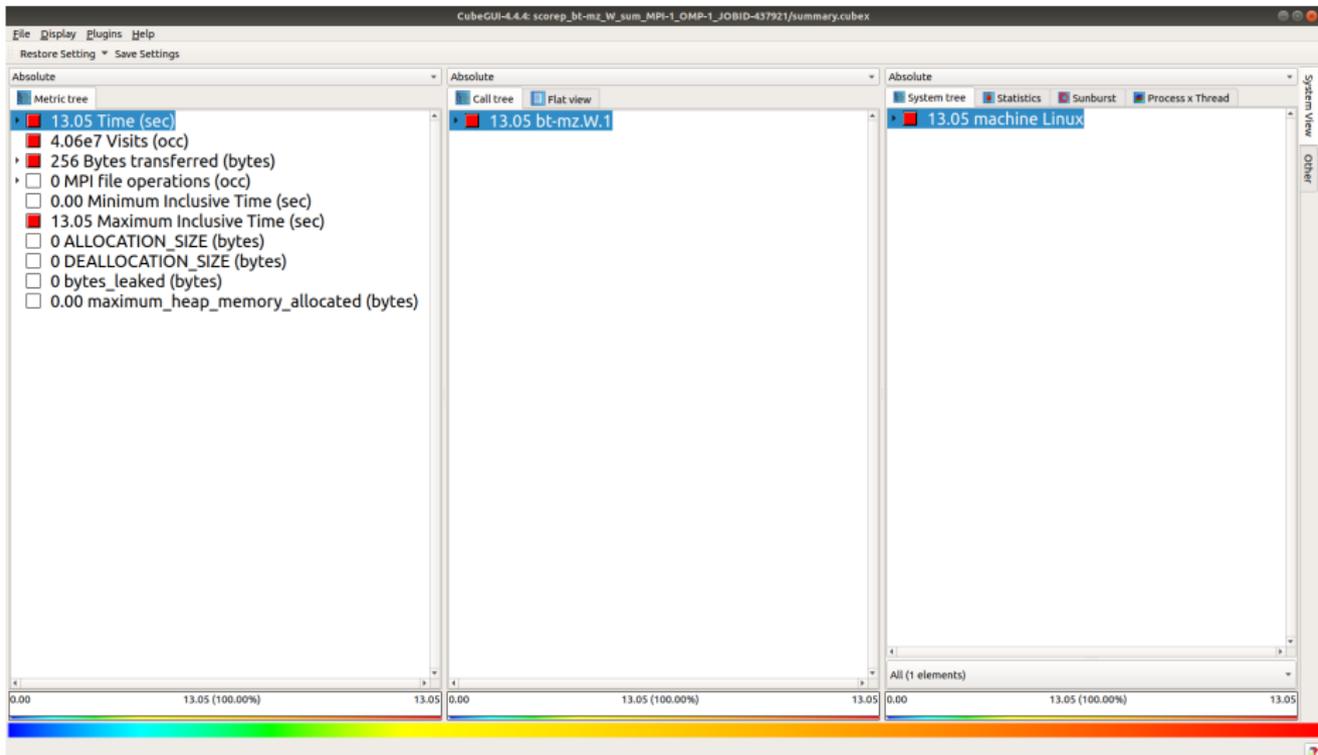
$$LB = \frac{avg(tcomp)}{max(tcomp)}$$

$$LB = \frac{0.84}{1.09}$$

$$LB = 0.77$$

Visualizando no Cube

`-nodes=1 -ntasks=1 / Absolute`



Visualizando no Cube

`-nodes=1 -ntasks=1 / Absolute`

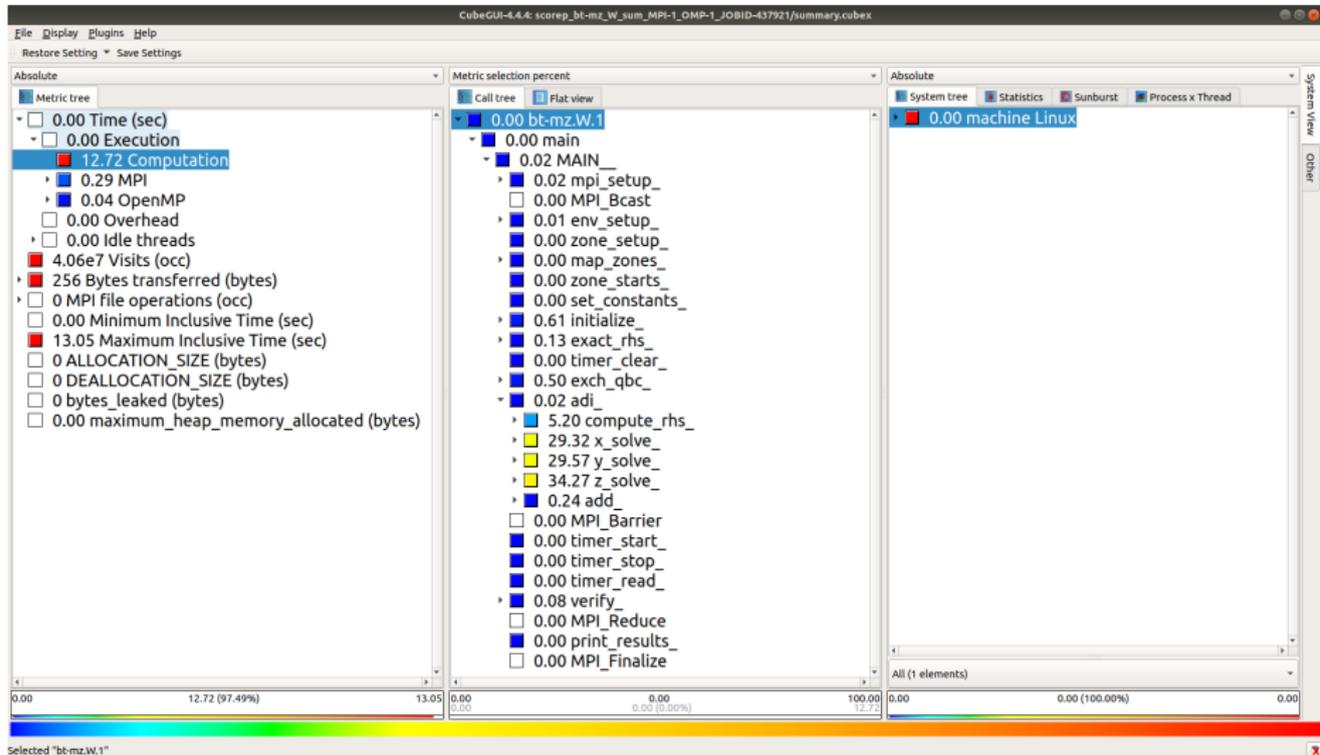
The screenshot displays the Cube GUI interface for a performance analysis. The window title is "CubeGUI-4.4.4: scorep_bt-mz_w_sum_MPI-1_OMP-1_JOBID-437921/summary.cubex". The interface is divided into three main panels:

- Left Panel (Metric tree):** Shows a hierarchical view of performance metrics. The "Computation" metric is highlighted in red, indicating it is the most significant. Other metrics include MPI (0.29), OpenMP (0.04), Overhead (0.00), Idle threads (0.00), Visits (4.06e7), Bytes transferred (256), MPI file operations (0.00), Minimum Inclusive Time (0.00), Maximum Inclusive Time (13.05), ALLOCATION_SIZE (0), DEALLOCATION_SIZE (0), bytes_leaked (0), and maximum_heap_memory_allocated (0.00).
- Middle Panel (Call tree):** Shows a detailed view of the "0.00 bt-mz.W.1" task. It is expanded to show sub-tasks such as "0.00 MAIN_" (0.00), "0.00 mpi_setup_" (0.00), "0.00 env_setup_" (0.00), "0.00 zone_setup_" (0.00), "0.00 zone_starts_" (0.00), "0.00 set_constants_" (0.00), "0.08 initialize_" (0.08), "0.02 exact_rhs_" (0.02), "0.00 timer_clear_" (0.00), "0.06 exch_qbc_" (0.06), "0.00 adj_" (0.00), "0.66 compute_rhs_" (0.66), "3.73 x_solve_" (3.73), "3.76 y_solve_" (3.76), "4.36 z_solve_" (4.36), "0.03 add_" (0.03), "0.00 MPI_Barrier" (0.00), "0.00 timer_start_" (0.00), "0.00 timer_stop_" (0.00), "0.00 timer_read_" (0.00), "0.01 verify_" (0.01), "0.00 MPI_Reduce" (0.00), "0.00 print_results_" (0.00), and "0.00 MPI_Finalize" (0.00).
- Right Panel (System tree):** Shows the system tree with "0.00 machine Linux" selected.

At the bottom of the interface, there is a "Ready" status bar and a set of navigation icons.

Visualizando no Cube

`-nodes=1 -ntasks=1 / Metric Own percent`



Visualizando no Cube

3 hotspots de computação

- **x_solve**: 3.73s (29.32%)
- **y_solve**: 3.76s (29.57%)
- **z_solve**: 4.36s (34.27%)

Medindo contadores de hardware (PAPI)

É possível obter acessar contadores de hardware (hwc) por meio da biblioteca PAPI (<https://icl.utk.edu/papi/>)

Para saber os contadores disponíveis, dar o comando `papi_avail`:

```
$ module load papi/5.5.1.0
$ papi_avail
```

```
Available PAPI preset and user defined events plus hardware information.
```

```
-----
PAPI Version           : 5.5.1.0
Vendor string and code : GenuineIntel (1)
Model string and code  : Intel(R) Xeon(R) CPU E5-2695 v2 @ 2.40GHz (62)
CPU Revision           : 4.000000
CPUID Info             : Family: 6 Model: 62 Stepping: 4
CPU Max Megahertz     : 2401
CPU Min Megahertz     : 1200
Hdw Threads per core  : 1
Cores per Socket      : 12
Sockets               : 2
NUMA Nodes            : 2
CPUs per Node         : 12
Total CPUs            : 24
Running in a VM       : no
Number Hardware Counters : 11
Max Multiplex Counters : 32
-----
```

Medindo contadores de hardware (PAPI) (cont.)

```
=====
PAPI Preset Events
=====
```

Name	Code	Avail	Deriv	Description (Note)
PAPI_L1_DCM	0x80000000	Yes	No	Level 1 data cache misses
PAPI_L1_ICM	0x80000001	Yes	No	Level 1 instruction cache misses
PAPI_L2_DCM	0x80000002	Yes	Yes	Level 2 data cache misses
PAPI_L2_ICM	0x80000003	Yes	No	Level 2 instruction cache misses
PAPI_L3_DCM	0x80000004	No	No	Level 3 data cache misses
PAPI_L3_ICM	0x80000005	No	No	Level 3 instruction cache misses
PAPI_L1_TCM	0x80000006	Yes	Yes	Level 1 cache misses
PAPI_L2_TCM	0x80000007	Yes	No	Level 2 cache misses
PAPI_L3_TCM	0x80000008	Yes	No	Level 3 cache misses
PAPI_CA_SNP	0x80000009	No	No	Requests for a snoop
PAPI_CA_SHR	0x8000000a	No	No	Requests for exclusive access to shared cache line
PAPI_CA_CLN	0x8000000b	No	No	Requests for exclusive access to clean cache line
PAPI_CA_INV	0x8000000c	No	No	Requests for cache line invalidation
PAPI_CA_ITV	0x8000000d	No	No	Requests for cache line intervention
PAPI_L3_LDM	0x8000000e	No	No	Level 3 load misses
PAPI_L3_STM	0x8000000f	No	No	Level 3 store misses
PAPI_BRU_IDL	0x80000010	No	No	Cycles branch units are idle
PAPI_FXU_IDL	0x80000011	No	No	Cycles integer units are idle
PAPI_FPU_IDL	0x80000012	No	No	Cycles floating point units are idle
PAPI_LSU_IDL	0x80000013	No	No	Cycles load/store units are idle
PAPI_TLB_DM	0x80000014	Yes	Yes	Data translation lookaside buffer misses
PAPI_TLB_IM	0x80000015	Yes	No	Instruction translation lookaside buffer misses
PAPI_TLB_TL	0x80000016	No	No	Total translation lookaside buffer misses
PAPI_L1_LDM	0x80000017	Yes	No	Level 1 load misses
PAPI_L1_STM	0x80000018	Yes	No	Level 1 store misses
PAPI_L2_LDM	0x80000019	No	No	Level 2 load misses

Medindo contadores de hardware (PAPI) (cont.)

PAPI_L2_STM	0x8000001a	Yes	No	Level 2 store misses
PAPI_BTAC_M	0x8000001b	No	No	Branch target address cache misses
PAPI_PRF_DM	0x8000001c	No	No	Data prefetch cache misses
PAPI_L3_DCH	0x8000001d	No	No	Level 3 data cache hits
PAPI_TLB_SD	0x8000001e	No	No	Translation lookaside buffer shutdowns
PAPI_CSR_FAL	0x8000001f	No	No	Failed store conditional instructions
PAPI_CSR_SUC	0x80000020	No	No	Successful store conditional instructions
PAPI_CSR_TOT	0x80000021	No	No	Total store conditional instructions
PAPI_MEM_SCY	0x80000022	No	No	Cycles Stalled Waiting for memory accesses
PAPI_MEM_RCY	0x80000023	No	No	Cycles Stalled Waiting for memory Reads
PAPI_MEM_WCY	0x80000024	No	No	Cycles Stalled Waiting for memory writes
PAPI_STL_ICY	0x80000025	Yes	No	Cycles with no instruction issue
PAPI_FUL_ICY	0x80000026	No	No	Cycles with maximum instruction issue
PAPI_STL_CCY	0x80000027	No	No	Cycles with no instructions completed
PAPI_FUL_CCY	0x80000028	No	No	Cycles with maximum instructions completed
PAPI_HW_INT	0x80000029	No	No	Hardware interrupts
PAPI_BR_UCN	0x8000002a	Yes	Yes	Unconditional branch instructions
PAPI_BR_CN	0x8000002b	Yes	No	Conditional branch instructions
PAPI_BR_TKN	0x8000002c	Yes	Yes	Conditional branch instructions taken
PAPI_BR_NTK	0x8000002d	Yes	No	Conditional branch instructions not taken
PAPI_BR_MSP	0x8000002e	Yes	No	Conditional branch instructions mispredicted
PAPI_BR_PRC	0x8000002f	Yes	Yes	Conditional branch instructions correctly predicted
PAPI_FMA_INS	0x80000030	No	No	FMA instructions completed
PAPI_TOT_IIS	0x80000031	No	No	Instructions issued
PAPI_TOT_INS	0x80000032	Yes	No	Instructions completed
PAPI_INT_INS	0x80000033	No	No	Integer instructions
PAPI_FP_INS	0x80000034	Yes	Yes	Floating point instructions
PAPI_LD_INS	0x80000035	Yes	No	Load instructions
PAPI_SR_INS	0x80000036	Yes	No	Store instructions
PAPI_BR_INS	0x80000037	Yes	No	Branch instructions

Medindo contadores de hardware (PAPI) (cont.)

PAPI_VEC_INS	0x80000038	No	No	Vector/SIMD instructions (could include integer)
PAPI_RES_STL	0x80000039	No	No	Cycles stalled on any resource
PAPI_FP_STAL	0x8000003a	No	No	Cycles the FP unit(s) are stalled
PAPI_TOT_CYC	0x8000003b	Yes	No	Total cycles
PAPI_LST_INS	0x8000003c	No	No	Load/store instructions completed
PAPI_SYC_INS	0x8000003d	No	No	Synchronization instructions completed
PAPI_L1_DCH	0x8000003e	No	No	Level 1 data cache hits
PAPI_L2_DCH	0x8000003f	Yes	Yes	Level 2 data cache hits
PAPI_L1_DCA	0x80000040	No	No	Level 1 data cache accesses
PAPI_L2_DCA	0x80000041	Yes	No	Level 2 data cache accesses
PAPI_L3_DCA	0x80000042	Yes	Yes	Level 3 data cache accesses
PAPI_L1_DCR	0x80000043	No	No	Level 1 data cache reads
PAPI_L2_DCR	0x80000044	Yes	No	Level 2 data cache reads
PAPI_L3_DCR	0x80000045	Yes	No	Level 3 data cache reads
PAPI_L1_DCW	0x80000046	No	No	Level 1 data cache writes
PAPI_L2_DCW	0x80000047	Yes	No	Level 2 data cache writes
PAPI_L3_DCW	0x80000048	Yes	No	Level 3 data cache writes
PAPI_L1_ICH	0x80000049	No	No	Level 1 instruction cache hits
PAPI_L2_ICH	0x8000004a	Yes	No	Level 2 instruction cache hits
PAPI_L3_ICH	0x8000004b	No	No	Level 3 instruction cache hits
PAPI_L1_ICA	0x8000004c	No	No	Level 1 instruction cache accesses
PAPI_L2_ICA	0x8000004d	Yes	No	Level 2 instruction cache accesses
PAPI_L3_ICA	0x8000004e	Yes	No	Level 3 instruction cache accesses
PAPI_L1_ICR	0x8000004f	No	No	Level 1 instruction cache reads
PAPI_L2_ICR	0x80000050	Yes	No	Level 2 instruction cache reads
PAPI_L3_ICR	0x80000051	Yes	No	Level 3 instruction cache reads
PAPI_L1_ICW	0x80000052	No	No	Level 1 instruction cache writes
PAPI_L2_ICW	0x80000053	No	No	Level 2 instruction cache writes
PAPI_L3_ICW	0x80000054	No	No	Level 3 instruction cache writes
PAPI_L1_TCH	0x80000055	No	No	Level 1 total cache hits

Medindo contadores de hardware (PAPI) (cont.)

PAPI_L2_TCH	0x80000056	No	No	Level 2 total cache hits
PAPI_L3_TCH	0x80000057	No	No	Level 3 total cache hits
PAPI_L1_TCA	0x80000058	No	No	Level 1 total cache accesses
PAPI_L2_TCA	0x80000059	Yes	Yes	Level 2 total cache accesses
PAPI_L3_TCA	0x8000005a	Yes	No	Level 3 total cache accesses
PAPI_L1_TCR	0x8000005b	No	No	Level 1 total cache reads
PAPI_L2_TCR	0x8000005c	Yes	Yes	Level 2 total cache reads
PAPI_L3_TCR	0x8000005d	Yes	Yes	Level 3 total cache reads
PAPI_L1_TCW	0x8000005e	No	No	Level 1 total cache writes
PAPI_L2_TCW	0x8000005f	Yes	No	Level 2 total cache writes
PAPI_L3_TCW	0x80000060	Yes	No	Level 3 total cache writes
PAPI_FML_INS	0x80000061	No	No	Floating point multiply instructions
PAPI_FAD_INS	0x80000062	No	No	Floating point add instructions
PAPI_FDV_INS	0x80000063	Yes	No	Floating point divide instructions
PAPI_FSQ_INS	0x80000064	No	No	Floating point square root instructions
PAPI_FNV_INS	0x80000065	No	No	Floating point inverse instructions
PAPI_FP_OPS	0x80000066	Yes	Yes	Floating point operations
PAPI_SP_OPS	0x80000067	Yes	Yes	Floating point operations; optimized to count scaled single pr
PAPI_DP_OPS	0x80000068	Yes	Yes	Floating point operations; optimized to count scaled double pr
PAPI_VEC_SP	0x80000069	Yes	Yes	Single precision vector/SIMD instructions
PAPI_VEC_DP	0x8000006a	Yes	Yes	Double precision vector/SIMD instructions
PAPI_REF_CYC	0x8000006b	Yes	No	Reference clock cycles

Of 108 possible events, 50 are available, of which 17 are derived.

avail.c

PASSED

Medindo contadores de hardware (PAPI)

Contadores a serem utilizados

PAPI_TOT_INS	0x80000032	Yes	No	Instructions completed
PAPI_TOT_CYC	0x8000003b	Yes	No	Total cycles
PAPI_L2_DCM	0x80000002	Yes	Yes	Level 2 data cache misses
PAPI_L2_DCA	0x80000041	Yes	No	Level 2 data cache accesses

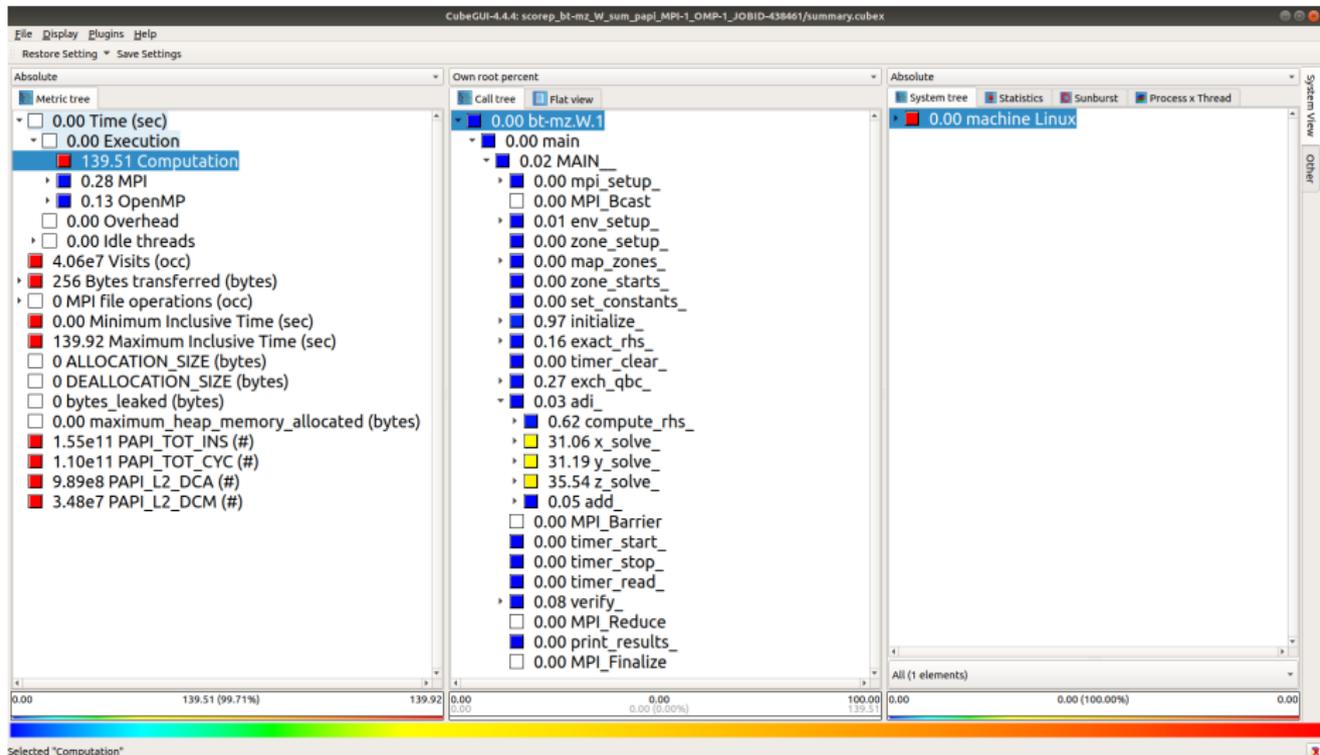
Medindo contadores de hardware (PAPI)

Definindo a variável de ambiente do Scalasca **SCOREP_METRIC_PAPI**.

```
BULL_srun_scalasca.sh
```

Medindo contadores de hardware (PAPI)

`-nodes=1 -ntasks=1 / BT-MZ.W`



Medindo contadores de hardware (PAPI)

```
-nodes=1 -ntasks=1 / BT-MZ.W
```

- Tempo de processamento aproximadamente 10X maior!

Medindo contadores de hardware (PAPI)

```
-nodes=1 -ntasks=1 / BT-MZ.W
```

- Tempo de processamento aproximadamente 10X maior!
- *Overhead* (custo adicional) devido a PAPI.

Medindo contadores de hardware (PAPI)

```
-nodes=1 -ntasks=1 / BT-MZ.W
```

- Tempo de processamento aproximadamente 10X maior!
- *Overhead* (custo adicional) devido a PAPI.
- Rodar em uma instância menor do problema: **BT-MZ.S**

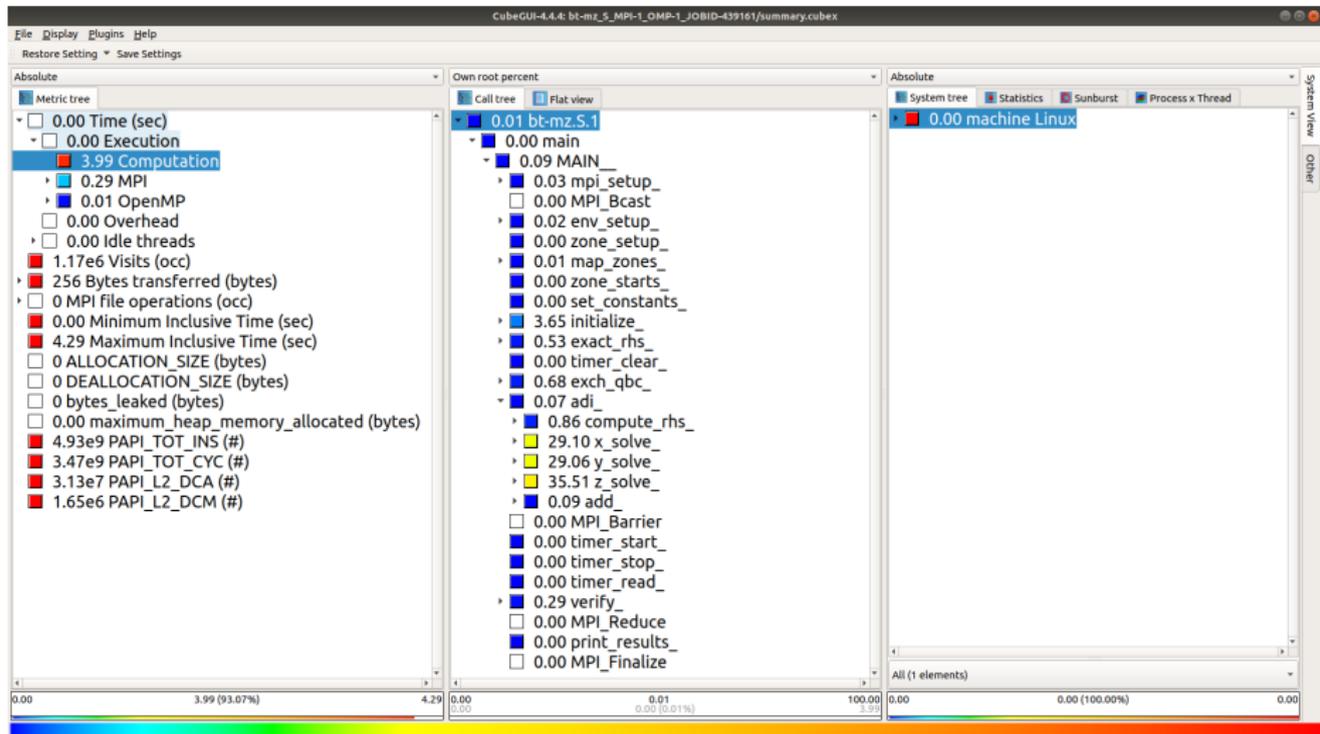
Medindo contadores de hardware (PAPI)

```
-nodes=1 -ntasks=1 / BT-MZ.W
```

- Tempo de processamento aproximadamente 10X maior!
- *Overhead* (custo adicional) devido a PAPI.
- Rodar em uma instância menor do problema: **BT-MZ.S**
- O perfil de desempenho muitas vezes é similar.

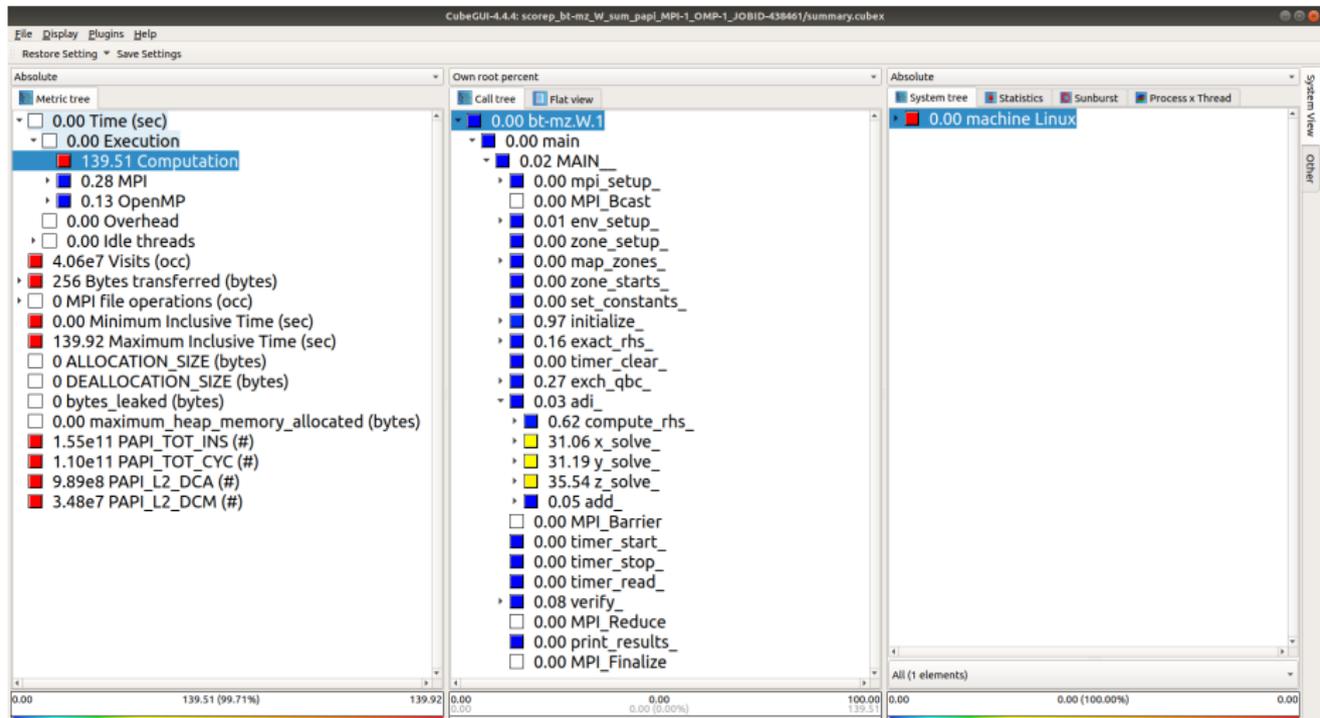
Medindo contadores de hardware (PAPI)

`-nodes=1 -ntasks=1 / BT-MZ.S`



Medindo contadores de hardware (PAPI)

`-nodes=1 -ntasks=1 / BT-MZ.W`



Selected "Computation"

Medindo contadores de hardware (PAPI)

Cálculo de novas métricas

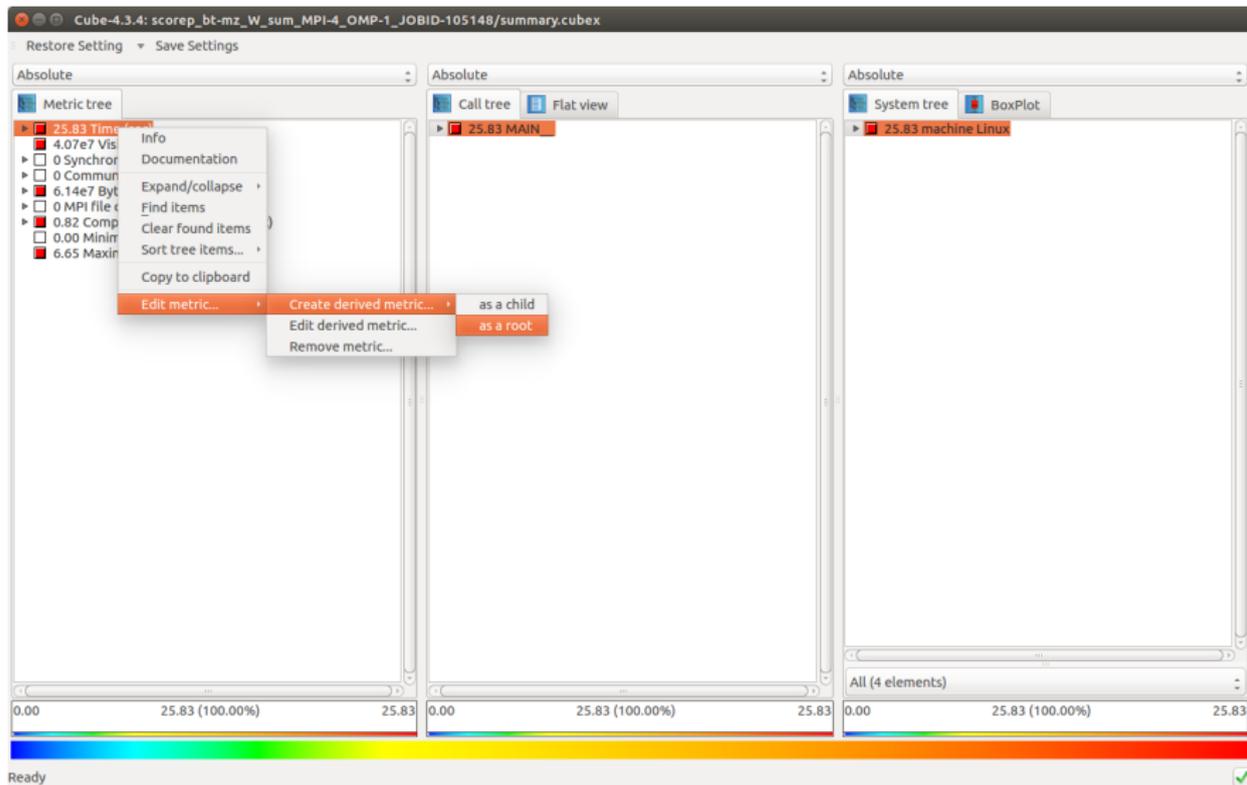
Razão de cache miss L2:

$$\text{L2 cache miss ratio} = \text{PAPI_L2_DCM} / \text{PAPI_L2_DCA}$$

Instruções por ciclo (IPC):

$$\text{IPC} = \text{PAPI_TOT_INS} / \text{PAPI_TOT_CYC}$$

Criando nova métrica no Cube



Criando nova métrica no Cube

`metric::PAPI_L2_DCM(i)/metric::PAPI_L2_DCA(i)`

Create new metric as a child of metric

Select metric from collection: -- please select --

Derived metric type: Postderived metric

Display name: L2 cache miss ratio

Unique name: cache_miss_ratio_l2

Data type: DOUBLE

Unit of measurement:

URL:

Description:

5: Calculation Calculation init Aggregation "+" Aggregation "-"

`metric::PAPI_L2_DCM(i)/metric::PAPI_L2_DCA(i)`

Create metric Cancel

Share this metric with SCALASCA group
Ok

Criando nova métrica no Cube

L2 cache miss ratio / BT-MZ.S

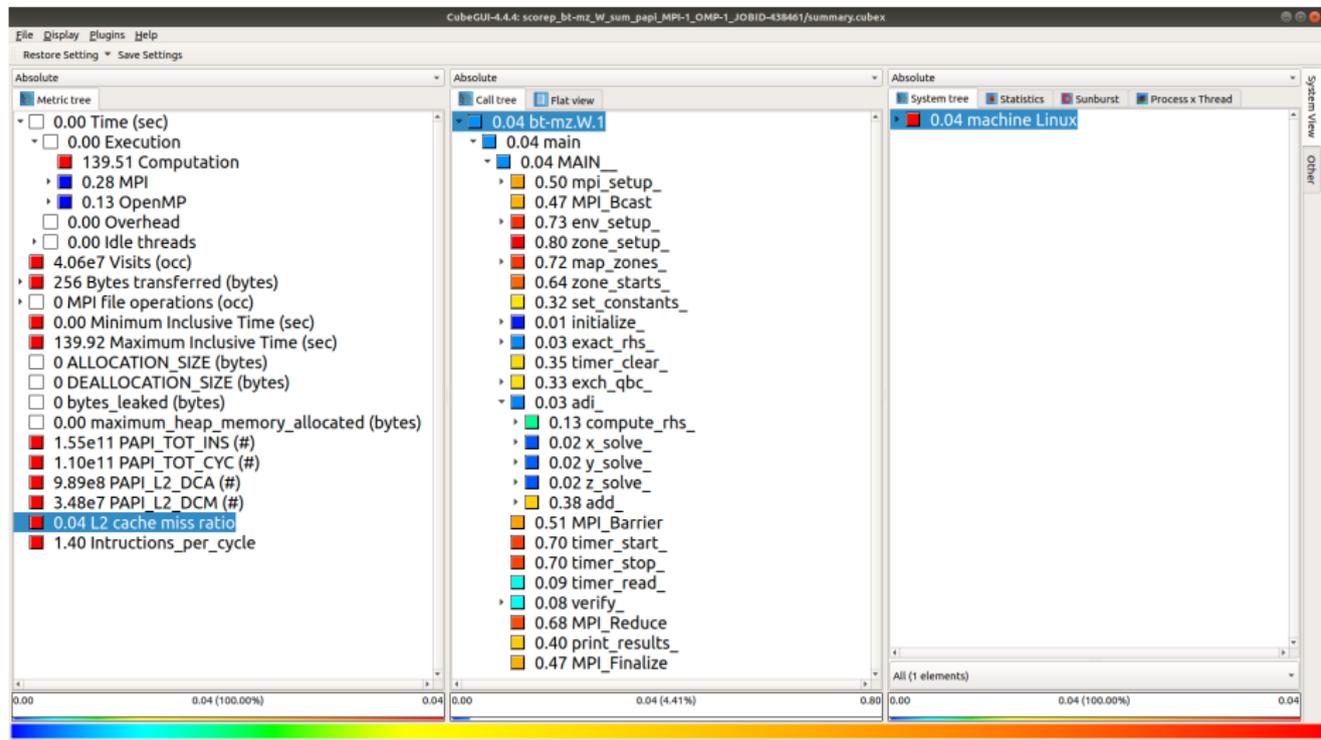
The screenshot displays the CubeGUI-4.4.4 interface for analyzing performance data from a job named 'bt-mz_S_MPI-1_OMP-1_JOBID-639161'. The main window is divided into three panels:

- Metric tree (left):** Shows a hierarchical list of metrics. The '0.05 L2 cache miss ratio' metric is highlighted in blue. Other visible metrics include '0.00 Time (sec)', '0.00 Execution', '3.99 Computation', '0.29 MPI', '0.01 OpenMP', '0.00 Overhead', '0.00 Idle threads', '1.17e6 Visits (occ)', '256 Bytes transferred (bytes)', '0 MPI file operations (occ)', '0.00 Minimum Inclusive Time (sec)', '4.29 Maximum Inclusive Time (sec)', '0 ALLOCATION_SIZE (bytes)', '0 DEALLOCATION_SIZE (bytes)', '0 bytes_leaked (bytes)', '0.00 maximum_heap_memory_allocated (bytes)', '4.93e9 PAPI_TOT_INS (#)', '3.47e9 PAPI_TOT_CYC (#)', '3.13e7 PAPI_L2_DCA (#)', '1.65e6 PAPI_L2_DCM (#)', and '1.42 Instructions per cycle'.
- Call tree (middle):** Shows a detailed view of the selected metric. The root is '0.05 bt-mz.S.1', which branches into '0.05 main' and '0.05 MAIN_'. The 'main' node contains a list of sub-tasks with their respective values: '0.49 mpi_setup_', '0.66 MPI_Bcast', '0.67 env_setup_', '0.92 zone_setup_', '0.72 map_zones_', '0.58 zone_starts_', '0.75 set_constants_', '0.01 initialize_', '0.03 exact_rhs_', '0.30 timer_clear_', '0.19 exch_qbc_', '0.03 adi_', '0.12 compute_rhs_', '0.02 x_solve_', '0.03 y_solve_', '0.02 z_solve_', '0.23 add_', '0.58 MPI_Barrier', '0.82 timer_start_', '0.66 timer_stop_', '0.13 timer_read_', '0.08 verify_', '0.64 MPI_Reduce', '0.37 print_results_', and '0.59 MPI_Finalize_'.
- System tree (right):** Shows the system context for the selected metric, currently displaying '0.05 machine Linux'.

At the bottom of the interface, there is a color-coded progress bar and a status bar showing 'All (1 elements)'.

Criando nova métrica no Cube

L2 cache miss ratio / BT-MZ.W



Criando nova métrica no Cube

Instructions per cycle / BT-MZ.S

Criando nova métrica no Cube

Instructions per cycle / BT-MZ.W

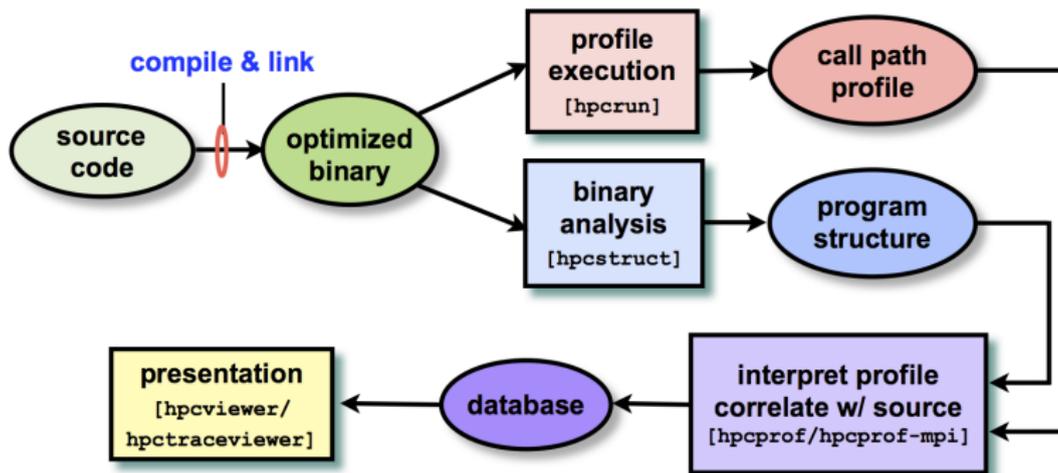
Roteiro

1 Scalasca, Score-p, Cube

2 hpctoolkit

HPCToolkit

<http://hpctoolkit.org>



HPCToolkit

Exemplo: NAS Parallel Benchmarks (NPB)

```
hpctoolkit/  
  NPB3.3.1-MZ/  
    NPB3.3-MZ-MPI/  
    NPB3.3-MZ-OMP/  
    NPB3.3-MZ-SER/  
    Changes.log  
    env_hpctoolkit  
    README
```

HPCToolkit

Exemplo: NAS Parallel Benchmarks (NPB)

```
hpctoolkit/  
  NPB3.3.1-MZ/  
    NPB3.3-MZ-MPI/  
    NPB3.3-MZ-OMP/  
    NPB3.3-MZ-SER/  
  Changes.log  
  env_hpctoolkit  
  README
```

HPCToolkit

Preparando o ambiente

```
$ cat env_hpctoolkit
```

```
module load openmpi/gnu/2.0.4.2  
module load hpctoolkit/5.3.2_4712  
module load papi/5.5.1.0  
module load papi-devel/5.5.1.0
```

HPCToolkit

```
hpctoolkit/  
  NPB3.3.1-MZ/  
    NPB3.3-MZ-MPI/  
    NPB3.3-MZ-OMP/  
    NPB3.3-MZ-SER/  
  Changes.log  
  env_hpctoolkit  
  README
```

HPCToolkit

```
hpctoolkit/  
  NPB3.3.1-MZ/  
    NPB3.3-MZ-MPI/  
      bin/  
      BT-MZ/  
      common/  
      config/  
      LU-MZ/  
      SP-MZ/  
      sys/  
      Makefile  
      README  
      README.install  
  NPB3.3-MZ-OMP/  
  NPB3.3-MZ-SER/  
  Changes.log  
  env_hpctoolkit  
  README
```

HPCToolkit

```
hpctoolkit/  
  NPB3.3.1-MZ/  
    NPB3.3-MZ-MPI/  
      bin/  
      BT-MZ/  
      common/  
      config/  
      LU-MZ/  
      SP-MZ/  
      sys/  
      Makefile  
      README  
      README.install  
  NPB3.3-MZ-OMP/  
  NPB3.3-MZ-SER/  
  Changes.log  
  env_hpctoolkit  
  README
```

HPCToolkit

```
config/  
  NAS.samples  
  make.def -> make_hpctoolkit.def  
  make.def.template  
  make_hpctoolkit.def  
  suite.def  
  suite.def.template
```

HPCToolkit

```
config/  
  NAS.samples  
  make.def -> make_hpctoolkit.def  
  make.def.template  
  make_hpctoolkit.def  
  suite.def  
  suite.def.template
```

HPCToolkit

```
$ cat make_hpctoolkit.def
```

```
#-----  
# This is the fortran compiler used for fortran programs  
#-----  
F77 = mpif77  
#F77 = scalasca -instrument mpif77  
#F77 = scorep mpif77  
  
#-----  
# This is the C compiler used for C programs  
#-----  
CC = mpicc  
#CC = scalasca -instrument mpicc  
#CC = scorep mpicc
```

HPCToolkit

NPB: benchmark, classe e número de processos MPI

```
config/  
  NAS.samples  
  make.def -> make_hpctoolkit.def  
  make.def.template  
  make_hpctoolkit.def  
  suite.def  
  suite.def.template
```

HPCToolkit

NPB: benchmark, classe e número de processos MPI

```
config/  
  NAS.samples  
  make.def -> make_hpctoolkit.def  
  make.def.template  
  make_hpctoolkit.def  
  suite.def  
  suite.def.template
```

Estudo de caso

NPB: benchmark, classe e número de processos MPI

```
$ cat suite.def

# config/suite.def
# This file is used to build several benchmarks with a single command.
# Typing "make suite" in the main directory will build all the benchmarks
# specified in this file.
# Each line of this file contains a benchmark name, class, and number
# of nodes. The name is one of "sp-mz", "bt-mz", and "lu-mz".
# The class is one of "S", "W", and "A" through "F".
# No blank lines.
# The following example builds serial sample sizes of all benchmarks.
#sp-mz S 1
#lu-mz S 1
#bt-mz S 2
bt-mz S 1
bt-mz S 2
bt-mz S 4
bt-mz W 1
bt-mz W 2
bt-mz W 4
bt-mz W 8
bt-mz W 16
```

Estudo de caso

NPB: compilação

```
$ cd ..  
$ make suite %compila o NPB  
$ cd bin
```

Estudo de caso

NPB: compilação

```
$ ls -Al
bt-mz.S.1
bt-mz.S.2
bt-mz.S.4
bt-mz.W.1
bt-mz.W.2
bt-mz.W.4
bt-mz.W.8
bt-mz.W.16
BULL_srun_hpctoolkit.sh
```

HPCToolkit

BULL_srun_hpctoolkit.sh

Estudo de caso

NPB: submetendo job

```
$ sbatch BULL_srun_hpctoolkit.sh bt-mz W
```

```
Submitted batch job 438988
```

```
$ squeue -u $USER
```

JOBID	PARTITION	NAME	USER	ST	TIME	NODES	NODELIST (REASON)
438988	treinamen	NPB_BT-M	professo	R	0:02	1	sdumont3000

Estudo de caso

NPB: perfil de desempenho

```
$ ls -Al
bt-mz.S.1
bt-mz.S.2
bt-mz.S.4
bt-mz.W.1
bt-mz.W.2
bt-mz.W.4
bt-mz.W.8
bt-mz.W.16
BULL_srun_hpctoolkit.sh
hpctoolkit/
```

Estudo de caso

NPB: perfil de desempenho

```
hpctoolkit/  
NUMNODES-1  
  bt-mz_W_MPI-1_OMP-1_JOBID-438988  
    bt-mz.W.1.hpcstruct  
    hpctoolkit-bt-mz.W.1-database-438988  
    hpctoolkit-bt-mz.W.1-measurements-438988  
    slurm-438988.out
```

Estudo de caso

NPB: perfil de desempenho

```
hpctoolkit/  
NUMNODES-1  
  bt-mz_W_MPI-1_OMP-1_JOBID-438988  
    bt-mz.W.1.hpcstruct  
    hpctoolkit-bt-mz.W.1-database-438988  
    hpctoolkit-bt-mz.W.1-measurements-438988  
    slurm-438988.out
```

Estudo de caso

NPB: perfil de desempenho

```
hpctoolkit/  
NUMNODES-1  
  bt-mz_W_MPI-1_OMP-1_JOBID-438988  
    bt-mz.W.1.hpctruct  
    hpctoolkit-bt-mz.W.1-database-438988  
    hpctoolkit-bt-mz.W.1-measurements-438988  
    slurm-438988.out
```

Estudo de caso

```
$ cat slurm-438988.out
```

```
Cluster configuration:
```

```
===
```

```
Partition: treinamento
```

```
Number of nodes: 1
```

```
Number of MPI processes: 1 ( 1 nodes)
```

```
Number of MPI processes per node:
```

```
Number of threads per MPI process: 1
```

```
NPB Benchmark: bt-mz
```

```
Benchmark class problem: W
```

```
[1580508594.340437] [sdumont5000:81015:0]
```

```
mxm.c:196 MXM WARN The 'ulimit -s' on the sys
```

```
[1580508594.342077] [sdumont5000:81015:0]
```

```
mxm.c:196 MXM WARN The 'ulimit -s' on the sys
```

```
NAS Parallel Benchmarks (NPB3.3-MZ-MPI) - BT-MZ MPI+OpenMP Benchmark
```

```
Number of zones: 4 x 4
```

```
Iterations: 200 dt: 0.000800
```

```
Number of active processes: 1
```

```
Use the default load factors with threads
```

```
Total number of threads: 1 ( 1.0 threads/process)
```

```
Calculated speedup = 1.00
```

```
Time step 1
```

```
Time step 20
```

Estudo de caso (cont.)

Time step 40
Time step 60
Time step 80
Time step 100
Time step 120
Time step 140
Time step 160
Time step 180
Time step 200

Verification being performed for class W
accuracy setting for epsilon = 0.100000000000000E-07

Comparison of RMS-norms of residual

1	0.5562611195402E+05	0.5562611195402E+05	0.2289019558898E-13
2	0.5151404119932E+04	0.5151404119932E+04	0.3195605260010E-13
3	0.1080453907954E+05	0.1080453907954E+05	0.4314917838667E-12
4	0.6576058591929E+04	0.6576058591929E+04	0.2033067669511E-13
5	0.4528609293561E+05	0.4528609293561E+05	0.3100863263992E-13

Comparison of RMS-norms of solution error

1	0.7185154786403E+04	0.7185154786403E+04	0.4961924046085E-13
2	0.7040472738068E+03	0.7040472738068E+03	0.3326408529931E-13
3	0.1437035074443E+04	0.1437035074443E+04	0.1887614294376E-12
4	0.8570666307849E+03	0.8570666307849E+03	0.3143720636440E-13
5	0.5991235147368E+04	0.5991235147368E+04	0.6770467641700E-13

Verification Successful

BT-MZ Benchmark Completed.

Class	=			W
Size	=	64x	64x	8
Iterations	=			200

Estudo de caso (cont.)

```
Time in seconds =          5.58
Total processes =           1
Total threads   =           1
Mop/s total    =          2572.99
Mop/s/thread   =          2572.99
Operation type = floating point
Verification   = SUCCESSFUL
Version       =           3.3.1
Compile date  =          31 Jan 2020
```

Compile options:

```
F77           = mpif77
FLINK         = $(F77)
F_LIB        = (none)
FFLAGS       = -O3 -fopenmp -g
FLINKFLAGS   = $(FFLAGS)
RAND         = (none)
```

Please send all errors/feedbacks to:

NPB Development Team
npb@nas.nasa.gov

```
msg: STRUCTURE: /scratch/treinamento/professor/modulo1/MC1-I/tools/hpctoolkit/NPB3.3.1-MZ/NPB3.3-M
msg: Line map : /opt/bullxde/profilers/hpctoolkit/5.3.2_4712/lib/hpctoolkit/ext-libs/libmonitor.so
msg: Line map : /opt/mpi/openmpi-gnu/2.0.4.2/lib/libmpi_mpifh.so.20.2.1
msg: Line map : /opt/mpi/openmpi-gnu/2.0.4.2/lib/libmpi.so.20.0.4
msg: Line map : /usr/lib64/libgomp.so.1.0.0
```

Estudo de caso (cont.)

```
msg: Line map : /usr/lib64/libc-2.17.so
msg: Line map : /opt/mpi/openmpi-gnu/2.0.4.2/lib/libopen-rte.so.20.1.2
msg: Line map : /opt/mpi/openmpi-gnu/2.0.4.2/lib/libopen-pal.so.20.2.2
msg: Line map : /opt/mpi/openmpi-gnu/2.0.4.2/lib/openmpi/mca_ess_pmi.so
msg: Line map : /opt/mpi/openmpi-gnu/2.0.4.2/lib/openmpi/mca_pml_cm.so
msg: Line map : /opt/mpi/openmpi-gnu/2.0.4.2/lib/openmpi/mca_pml_yalla.so
msg: Line map : /opt/mellanox/mxm/lib/libmxm.so.2.0.32
msg: Line map : /opt/mpi/openmpi-gnu/2.0.4.2/lib/openmpi/mca_mtl_mxm.so
msg: Populating Experiment database: /scratch/treinamento/professor/modulo1/MC1-I/tools/hpctoolkit
```

Estudo de caso

```
sbatch --nodes=1 --ntasks=1 BULL_srun_hpctoolkit.sh bt-mz W
```

```
Number of zones:   4 x   4  
Iterations: 200   dt: 0.000800  
Number of active processes:   1
```

Use the default load factors with threads

```
Total number of threads:   1 ( 1.0 threads/process)
```

```
Calculated speedup =   1.00
```

BT-MZ Benchmark Completed.

```
Class           =                               W  
Size            =          64x   64x   8  
Iterations      =                               200  
Time in seconds =                               5.58  
Total processes =                               1  
Total threads   =                               1  
Mop/s total     =          2572.99  
Mop/s/thread    =          2572.99  
Operation type  =          floating point  
Verification    =          SUCCESSFUL  
Version         =                               3.3.1
```

Estudo de caso

```
sbatch --nodes=1 --ntasks=2 BULL_srun_hpctoolkit.sh bt-mz W
```

```
Number of zones:   4 x   4  
Iterations: 200   dt: 0.000800  
Number of active processes:   2
```

Use the default load factors with threads

```
Total number of threads:   2 ( 1.0 threads/process)
```

```
Calculated speedup =   1.98
```

BT-MZ Benchmark Completed.

```
Class           =                               W  
Size            =          64x   64x   8  
Iterations      =                               200  
Time in seconds =                               2.84  
Total processes =                               2  
Total threads   =                               2  
Mop/s total     =          5056.23  
Mop/s/thread    =          2528.12  
Operation type  =          floating point  
Verification    =          SUCCESSFUL  
Version         =                               3.3.1
```

Estudo de caso

```
sbatch --nodes=1 --ntasks=4 BULL_srun_hpctoolkit.sh bt-mz W
```

```
Number of zones:   4 x   4  
Iterations: 200   dt: 0.000800  
Number of active processes:   4
```

Use the default load factors with threads

```
Total number of threads:   4 ( 1.0 threads/process)
```

```
Calculated speedup =   3.95
```

BT-MZ Benchmark Completed.

Class	=		W
Size	=	64x	64x 8
Iterations	=		200
Time in seconds	=		1.49
Total processes	=		4
Total threads	=		4
Mop/s total	=	9610.43	
Mop/s/thread	=	2402.61	
Operation type	=	floating point	
Verification	=	SUCCESSFUL	
Version	=	3.3.1	

Estudo de caso

```
sbatch --nodes=1 --ntasks=8 BULL_srun_hpctoolkit.sh bt-mz W
```

```
Number of zones:   4 x   4  
Iterations: 200   dt: 0.000800  
Number of active processes:      8
```

Use the default load factors with threads

```
Total number of threads:      8 ( 1.0 threads/process)
```

```
Calculated speedup =      4.87
```

BT-MZ Benchmark Completed.

Class	=		W
Size	=	64x	64x 8
Iterations	=		200
Time in seconds	=		1.29
Total processes	=		8
Total threads	=		8
Mop/s total	=	11125.59	
Mop/s/thread	=	1390.70	
Operation type	=	floating point	
Verification	=	SUCCESSFUL	
Version	=	3.3.1	

Estudo de caso

```
sbatch --nodes=1 --ntasks=16 BULL_srun_hpctoolkit.sh bt-mz W
```

```
Number of zones:   4 x   4  
Iterations: 200   dt: 0.000800  
Number of active processes:   16
```

Use the default load factors with threads

```
Total number of threads:   16 ( 1.0 threads/process)
```

```
Calculated speedup =   4.87
```

BT-MZ Benchmark Completed.

Class	=		W
Size	=	64x	64x 8
Iterations	=		200
Time in seconds	=		1.28
Total processes	=		16
Total threads	=		16
Mop/s total	=	11181.49	
Mop/s/thread	=	698.84	
Operation type	=	floating point	
Verification	=	SUCCESSFUL	
Version	=	3.3.1	

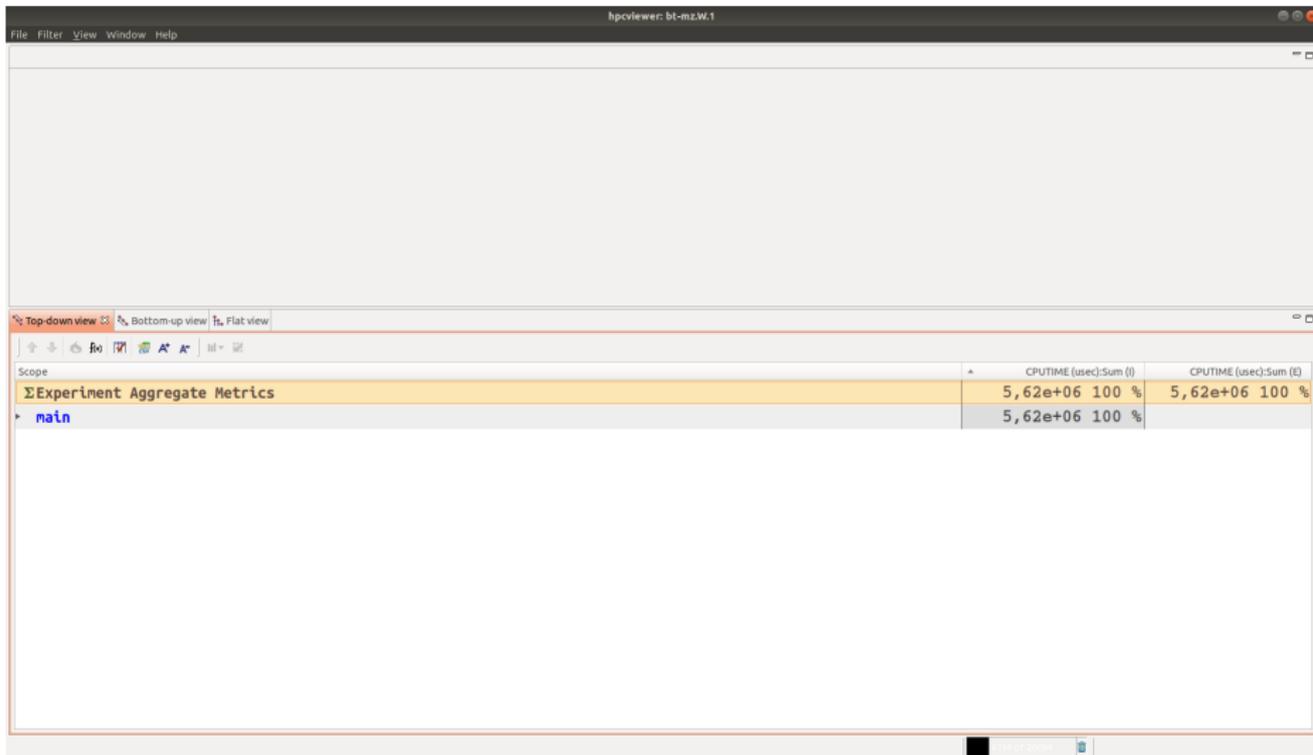
Visualizando no hpcviewer

NPB: estudo de caso

```
$ cd profiling/hpctoolkit/NUMNODES-1/bt-mz_W_MPI-1_OMP-1_JOBID-438988  
$ hpcviewer hpctoolkit-bt-mz.W.1-database-438988/
```

Visualizando no hpcviewer

`-nodes=1 -ntasks=1 / Top-down view`



The screenshot shows the hpcviewer application window titled "hpcviewer: bt-m2w.1". The interface includes a menu bar (File, Filter, View, Window, Help) and a toolbar with navigation icons. The main display area is currently empty. Below the main area, there are view mode tabs: "Top-down view" (selected), "Bottom-up view", and "Flat view". A table at the bottom of the window displays performance metrics for the "main" process.

Scope	CPUTIME (usec):Sum (l)	CPUTIME (usec):Sum (R)
ΣExperiment Aggregate Metrics	5,62e+06 100 %	5,62e+06 100 %
main	5,62e+06 100 %	

Visualizando no hpcviewer

`-nodes=1 -ntasks=1 / Top-down view`

The screenshot shows the hpcviewer application window titled "hpcviewer: bt-m2.w.1". The main area is a code editor displaying Fortran code. The code includes MPI-related operations like `mpi_send`, `mpi_recv`, and `mpi_barrier`. The line numbers range from 300 to 317. The code is as follows:

```
300      end do
301      ip = ip + 1
302      if (ip .lt. num_procs) then
303          call mpi_send(myid, 1, MPI_INTEGER, ip, 1000,
304                      comm_setup, ierror)
305      > call mpi_recv(trecs, t_last, dp_type, ip, 1001,
306                  comm_setup, statuses, ierror)
307      > write(*,*)
308          goto 910
309      endif
310
311 999 continue
312 call mpi_barrier(MPI_COMM_WORLD, ierror)
313 call mpi_finalize(ierror)
314
315 end
316
317
```

Below the code editor, there is a navigation bar with tabs for "Top-down view", "Bottom-up view", and "Flat view". The "Top-down view" tab is selected. Below the navigation bar is a toolbar with various icons. The main area below the toolbar is a performance metrics table:

Scope	CPUTIME (usec):Sum (l)	CPUTIME (usec):Sum (R)
ΣExperiment Aggregate Metrics	5,62e+06 100 %	5,62e+06 100 %
main	5,62e+06 100 %	

Visualizando no hpcviewer

`-nodes=1 -ntasks=1 / Top-down view`

The screenshot shows the hpcviewer interface. The top pane displays the source code of a program named 'bt'. The code includes headers, defines variables, and contains a loop. The bottom pane shows the performance metrics for the 'main' function, indicating a CPU time of 5,62e+06 seconds (100%) for the 'bt' function.

```
43 c      H. Jin
44 c
45 c .....
46
47 c .....
48 c      program bt
49 c .....
50
51 c      include 'header.h'
52 c      include 'mpi_stuff.h'
53
54 c      integer num_zones
55 c      parameter (num_zones=x_zones*y_zones)
56
57 c      integer nx(num_zones), nymax(num_zones), ny(num_zones),
58 c              nz(num_zones)
59
60 c .....
61 c      define all_flatd_array as non-dimensional array to be reduced
```

Scope	CPU TIME (usec):Sum (I)	CPU TIME (usec):Sum (O)
Σ Experiment Aggregate Metrics	5,62e+06 100 %	5,62e+06 100 %
main	5,62e+06 100 %	
▶ 315: bt	5,62e+06 100 %	

Visualizando no hpcviewer

-nodes=1 -ntasks=1 / Top-down view

The screenshot displays the hpcviewer interface. The top pane shows Fortran code for a subroutine named `adi(rho_1, us, vs, ws, qs, square, rhs)`. The code includes variable declarations for `nx, nmax, ny, nz`, `rho_1`, `us`, `vs`, `ws`, `qs`, `square`, `rhs`, and `u`. The bottom pane shows a top-down performance view with a tree structure of execution events and a table of metrics.

Scope	CPUTIME (usec):Sum (%)	CPUTIME (usec):Sum (%)
ΣExperiment Aggregate Metrics	5,62e+06 100 %	5,62e+06 100 %
main	5,62e+06 100 %	
bt	5,62e+06 100 %	
loop at bt.f: 207	5,55e+06 98,7%	
loop at bt.f: 224	5,52e+06 98,2%	
adi	5,49e+06 97,6%	
z_solve	1,69e+06 30,1%	
y_solve	1,59e+06 28,3%	
x_solve	1,54e+06 27,4%	5,99e+03 0,1%
compute_rhs	6,64e+05 11,8%	
add	3,59e+04 0,6%	
exch_qbc	2,99e+04 0,5%	
loop at bt.f: 187	2,85e+04 0,5%	
~unknown-proc~	1,81e+04 0,3%	

Visualizando no hpcviewer

`-nodes=1 -ntasks=1 / Top-down view`

The screenshot displays the hpcviewer interface. The top pane shows the source code for a function named `z_solve`. The code includes comments and a loop structure for computing indices and solving a system. The bottom pane shows a top-down view of the performance metrics.

```
49 c Compute the indices for storing the block-diagonal matrix;
50 c determine c (labeled f) and s jacobians
51 c -----
52 OMP DO
53 do j = 1, ny-2
54 do i = 1, nx-2
55 do k = 0, ksize
56
57 tmp1 = 1.d0 / u(1,1,j,k)
58 tmp2 = tmp1 * tmp1
59 tmp3 = tmp1 * tmp2
60
61 fjac(1,1,k) = 0.d0
62 fjac(1,2,k) = 0.d0
63 fjac(1,3,k) = 0.d0
64 fjac(1,4,k) = 1.d0
65 fjac(1,5,k) = 0.d0
66
67 fjac(2,1,k) = - fjac(1,1,k)u(1,1,k)u(1,1,k)
68
69
```

The bottom pane shows the performance metrics for the `z_solve` function. The table below summarizes the data shown in the screenshot.

Scope	CPUTIME (usec):Sum (l)	CPUTIME (usec):Sum (R)
ΣExperiment Aggregate Metrics	5,62e+06 100 %	5,62e+06 100 %
main	5,62e+06 100 %	
bt	5,62e+06 100 %	
loop at bt.f: 207	5,55e+06 98,7%	
loop at bt.f: 224	5,52e+06 98,2%	
adi	5,49e+06 97,6%	
z_solve	1,69e+06 30,1%	
[I] z_solve_	1,69e+06 30,1%	
z_solve_._omp_fn.0	1,69e+06 30,1%	7,72e+05 13,7%
[I] z_solve_._omp_fn.0	1,69e+06 30,1%	
loop at z_solve.f: 313	1,69e+06 30,1%	
loop at z_solve.f: 54	1,69e+06 30,1%	
loop at z_solve.f: 351	7,84e+05 13,9%	1,80e+04 0,3%
loop at z_solve.f: 146	4,91e+05 8,7%	4,91e+05 8,7%

Visualizando no hpcviewer

`-nodes=1 -ntasks=1 / Bottom-up view`

The screenshot shows the hpcviewer interface with the 'Bottom-up view' selected. The top pane displays the source code of a program, with the 'binvrhs' subroutine highlighted in orange. The bottom pane shows a table of execution metrics for various subroutines.

Scope	CPUTIME (usec):Sum (%)	CPUTIME (usec):Sum (%)
ΣExperiment Aggregate Metrics	5,62e+06 100 %	5,62e+06 100 %
binvrhs	1,46e+06 26,0%	1,46e+06 26,0%
matmul_sub	8,00e+05 14,2%	8,00e+05 14,2%
z_solve_-_omp_fn.0	1,71e+06 30,4%	7,77e+05 13,8%
x_solve_-_omp_fn.0	1,54e+06 27,5%	7,60e+05 13,5%
y_solve_-_omp_fn.0	1,58e+06 28,2%	7,42e+05 13,2%
compute_rhs_-_omp_fn.0	6,64e+05 11,8%	6,64e+05 11,8%
matvec_sub	1,85e+05 3,3%	1,85e+05 3,3%
binvrhs	5,99e+04 1,1%	5,99e+04 1,1%
lhsinit	5,39e+04 1,0%	5,39e+04 1,0%
add_-_omp_fn.0	3,59e+04 0,6%	3,59e+04 0,6%
~unknown-proc~	3,59e+04 0,6%	3,59e+04 0,6%
exact_solution	1,19e+04 0,2%	1,19e+04 0,2%
opal_hash_table_remove_all	1,03e+04 0,2%	1,03e+04 0,2%

Visualizando no hpcviewer

`-nodes=1 -ntasks=1 / Bottom-up view`

The screenshot shows the hpcviewer interface. The top pane displays Fortran code with line numbers 201 to 218. Line 206 is highlighted in orange and contains the subroutine definition: `subroutine binvrchs(lhs,c,r)`. The bottom pane shows the 'Bottom-up view' of the execution tree. The table below summarizes the performance metrics for the execution tree.

Scope	CPUTIME (usec):Sum (%)	CPUTIME (usec):Sum (%)
ΣExperiment Aggregate Metrics	5,62e+06 100 %	5,62e+06 100 %
binvrchs	1,46e+06 26,0%	1,46e+06 26,0%
375: [I] z_solve__omp_fn.0	5,32e+05 9,5%	5,32e+05 9,5%
313: z_solve__omp_fn.0	5,32e+05 9,5%	5,32e+05 9,5%
45: [I] z_solve__	5,32e+05 9,5%	5,32e+05 9,5%
45: z_solve	5,32e+05 9,5%	5,32e+05 9,5%
28: adi	5,32e+05 9,5%	5,32e+05 9,5%
187: bt	5,32e+05 9,5%	5,32e+05 9,5%
315: main	5,32e+05 9,5%	5,32e+05 9,5%
336: x_solve__omp_fn.0	4,67e+05 8,3%	4,67e+05 8,3%
331: y_solve__omp_fn.0	4,61e+05 8,2%	4,61e+05 8,2%
matmul_sub	8,00e+05 14,2%	8,00e+05 14,2%
z_solve__omp_fn.0	1,71e+06 30,4%	7,77e+05 13,8%
x_solve__omp_fn.0	1,54e+06 27,5%	7,60e+05 13,5%

Visualizando no hpcviewer

`-nodes=1 -ntasks=1 / Flat view`

The screenshot shows the hpcviewer interface. The top window displays the source code for a program named 'solve_subs.f'. The code includes headers, defines variables for zones, and sets MPI-related parameters. The bottom window shows the 'Flat view' of the performance metrics, which is a table with columns for Scope, CPUTIME (usec):Sum (l), and CPUTIME (usec):Sum (r).

Scope	CPUTIME (usec):Sum (l)	CPUTIME (usec):Sum (r)
Experiment Aggregate Metrics	5,62e+06 100 %	5,62e+06 100 %
<code>/scratch/treInamento/professor/modulo1/MC1-I/tools/hpctoolkit/NPB3.3.1-MZ/NPB3.3-MZ-MPI/bin/bt-mz.W.1</code>	<code>5,62e+06 100 %</code>	<code>5,57e+06 99,0%</code>
<code>/usr/lib64/libgomp.so.1.0.0</code>	<code>3,59e+04 0,6%</code>	<code>3,59e+04 0,6%</code>
<code>/opt/mpi/openmpi-gnu/2.0.4.2/lib/libopen-pal.so.20.2.2</code>	<code>1,81e+04 0,3%</code>	<code>1,03e+04 0,2%</code>
<code>/usr/lib64/libc-2.17.so</code>	<code>7,79e+03 0,1%</code>	<code>7,79e+03 0,1%</code>
<code>/opt/bullxde/profilers/hpctoolkit/5.3.2_4712/lib/hpctoolkit/ext-libs/libmonitor.so.0.0.0</code>	<code>1,81e+04 0,3%</code>	
<code>/opt/mellanox/mxm/lib/libmxm.so.2.0.32</code>	<code>7,79e+03 0,1%</code>	
<code>/opt/mpi/openmpi-gnu/2.0.4.2/lib/libmpi.so.20.0.4</code>	<code>1,81e+04 0,3%</code>	
<code>/opt/mpi/openmpi-gnu/2.0.4.2/lib/libmpi_pifh.so.20.2.1</code>	<code>1,81e+04 0,3%</code>	
<code>/opt/mpi/openmpi-gnu/2.0.4.2/lib/openmpi/mca_pml_yalla.so</code>	<code>7,79e+03 0,1%</code>	

Visualizando no hpcviewer

`-nodes=1 -ntasks=1 / Flat view`

The screenshot shows the hpcviewer interface. The top pane displays source code for a file named 'bt.f'. The code includes comments and a subroutine definition: 'subroutine btvrhs(lhs,c,r)'. The bottom pane shows a 'Flat view' of the performance metrics for the experiment.

Scope	CPUTIME (usec):Sum (I)	CPUTIME (usec):Sum (D)
ΣExperiment Aggregate Metrics	5,62e+06 100 %	5,62e+06 100 %
/scratch/treInamento/professor/modulo1/MC1-I/tools/hpctoolkit/NPB3.3.1-MZ/NPB3.3-MZ-MPI/bin/bt-mz.W.1	5,62e+06 100 %	5,57e+06 99,0 %
▶ solve_subs.f	2,51e+06 44,6%	2,51e+06 44,6%
▶ z_solve.f	1,71e+06 30,4%	7,77e+05 13,8%
▶ x_solve.f	1,55e+06 27,6%	7,66e+05 13,6%
▶ y_solve.f	1,59e+06 28,3%	7,42e+05 13,2%
▶ rhs.f	6,70e+05 11,9%	6,64e+05 11,8%
▶ initialize.f	7,19e+04 1,3%	5,99e+04 1,1%
▶ add.f	3,59e+04 0,6%	3,59e+04 0,6%
▶ exact_solution.f	1,19e+04 0,2%	1,19e+04 0,2%
▶ exch_qbc.f	2,99e+04 0,5%	5,99e+03 0,1%
▶ adi.f	5,51e+06 98,1%	
▶ bt.f	5,62e+06 100 %	
▶ verify.f	5,99e+03 0,1%	

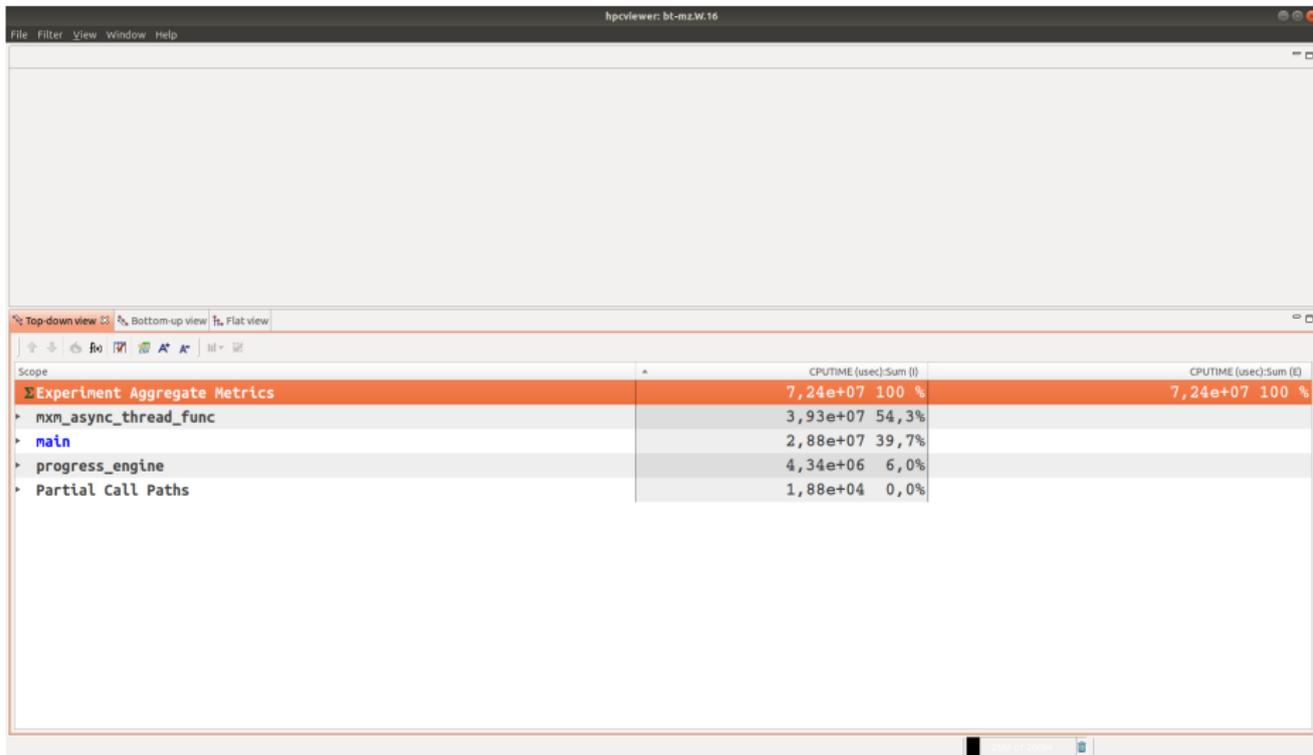
Visualizando no `hpcviewer`

NPB: estudo de caso

```
$ cd profiling/hpctoolkit/NUMNODES-1/bt-mz_W_MPI-16_OMP-1_JOBID-439032  
$ hpcviewer hpctoolkit-bt-mz.W.16-database-439032
```

Visualizando no hpcviewer

`-nodes=1 -ntasks=16 / Top-down view`



The screenshot shows the hpcviewer application window titled "hpcviewer: bt-mz.W.16". The interface includes a menu bar (File, Filter, View, Window, Help) and a toolbar with navigation icons. The main display area is currently empty. Below the toolbar, there are tabs for "Top-down view", "Bottom-up view", and "Flat view", with "Top-down view" selected. A table of performance metrics is displayed below the tabs.

Scope	CPU TIME (usec):Sum (I)	CPU TIME (usec):Sum (O)
Experiment Aggregate Metrics	7,24e+07 100 %	7,24e+07 100 %
mxm_async_thread_func	3,93e+07 54,3%	
main	2,88e+07 39,7%	
progress_engine	4,34e+06 6,0%	
Partial Call Paths	1,88e+04 0,0%	

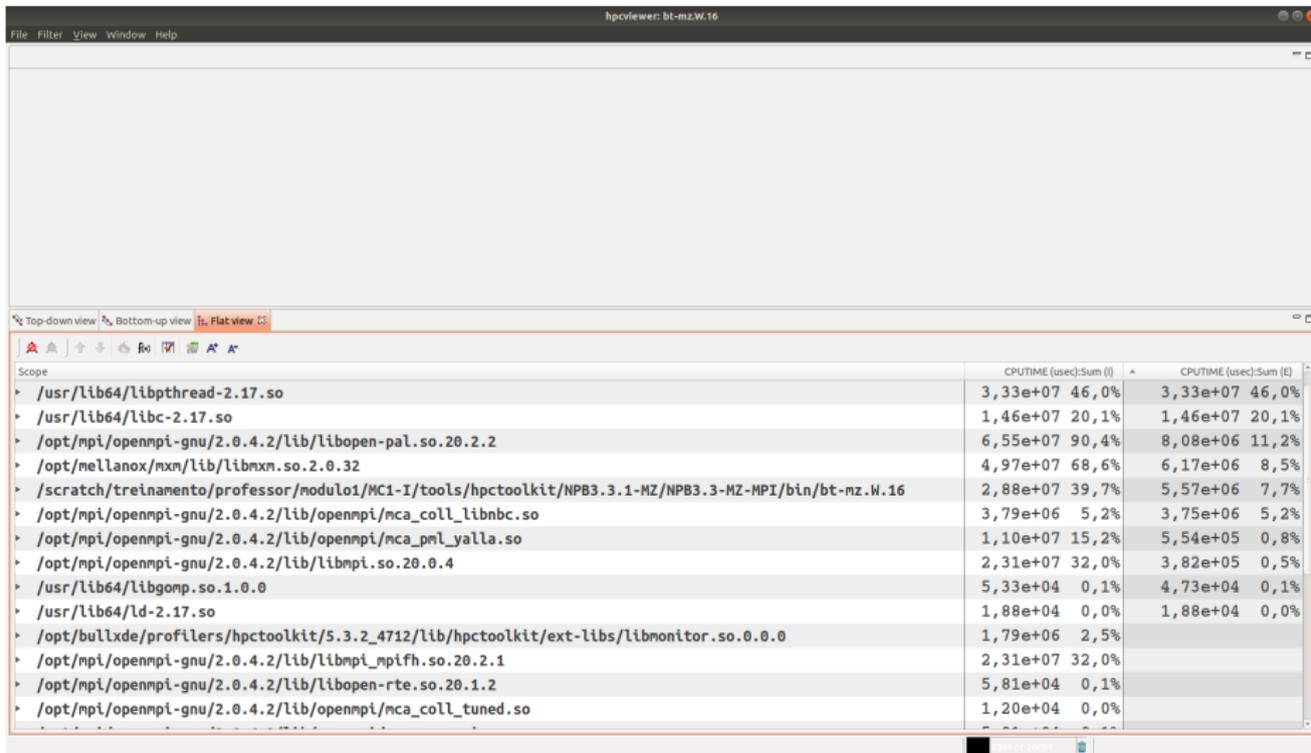
Visualizando no hpcviewer

`-nodes=1 -ntasks=16 / Bottom-up view`

Scope	CPUTIME (usec):Sum (%)	CPUTIME (usec):Sum (%)
__read_nocancel	2,98e+07 41,2%	2,98e+07 41,2%
__IO_vfsconf	9,83e+06 13,6%	6,57e+06 9,1%
opal_timer_linux_get_cycles_sys_timer	5,99e+06 8,3%	5,99e+06 8,3%
__xstat64	4,26e+06 5,9%	4,26e+06 5,9%
__GI___strtoul_l_internal	3,25e+06 4,5%	3,25e+06 4,5%
pthread_spin_lock	3,14e+06 4,3%	3,14e+06 4,3%
opal_atomic_cmpset_32	3,10e+06 4,3%	3,10e+06 4,3%
opal_progress	2,10e+07 28,9%	1,68e+06 2,3%
binvcrhs	1,49e+06 2,1%	1,49e+06 2,1%
mxm_shm_ep_poll_rcv_queue	9,08e+05 1,3%	9,08e+05 1,3%
mxm_shm_progress_pending_sends	8,48e+05 1,2%	8,48e+05 1,2%
mxm_shm_ep_progress	1,68e+06 2,3%	8,21e+05 1,1%
matmul_sub	8,21e+05 1,1%	8,21e+05 1,1%
mxm_progress	8,16e+05 1,1%	8,16e+05 1,1%

Visualizando no hpcviewer

`-nodes=1 -ntasks=16 / Flat view`



hpcviewer: bt-mz.W.16

File Filter View Window Help

Top-down view Bottom-up view **Flat view**

Scope	CPUTIME (usec):Sum (I)	CPUTIME (usec):Sum (D)
/usr/lib64/libpthread-2.17.so	3,33e+07 46,0%	3,33e+07 46,0%
/usr/lib64/libc-2.17.so	1,46e+07 20,1%	1,46e+07 20,1%
/opt/mpl/openmpi-gnu/2.0.4.2/lib/libopen-pal.so.20.2.2	6,55e+07 90,4%	8,08e+06 11,2%
/opt/mellanox/nxm/lib/libnmxn.so.2.0.32	4,97e+07 68,6%	6,17e+06 8,5%
/scratch/treinamento/professor/modulo1/MC1-I/tools/hpctoolkit/NPB3.3.1-MZ/NPB3.3-MZ-MPI/bin/bt-mz.W.16	2,88e+07 39,7%	5,57e+06 7,7%
/opt/mpl/openmpi-gnu/2.0.4.2/lib/openmpi/mca_coll_libnbc.so	3,79e+06 5,2%	3,75e+06 5,2%
/opt/mpl/openmpi-gnu/2.0.4.2/lib/openmpi/mca_pml_yalla.so	1,10e+07 15,2%	5,54e+05 0,8%
/opt/mpl/openmpi-gnu/2.0.4.2/lib/libmpi.so.20.0.4	2,31e+07 32,0%	3,82e+05 0,5%
/usr/lib64/libgomp.so.1.0.0	5,33e+04 0,1%	4,73e+04 0,1%
/usr/lib64/ld-2.17.so	1,88e+04 0,0%	1,88e+04 0,0%
/opt/bullxde/profilers/hpctoolkit/5.3.2_4712/lib/hpctoolkit/ext-libs/libmonitor.so.0.0.0	1,79e+06 2,5%	
/opt/mpl/openmpi-gnu/2.0.4.2/lib/libmpi_mpfh.so.20.2.1	2,31e+07 32,0%	
/opt/mpl/openmpi-gnu/2.0.4.2/lib/libopen-rte.so.20.1.2	5,81e+04 0,1%	
/opt/mpl/openmpi-gnu/2.0.4.2/lib/openmpi/mca_coll_tuned.so	1,20e+04 0,0%	

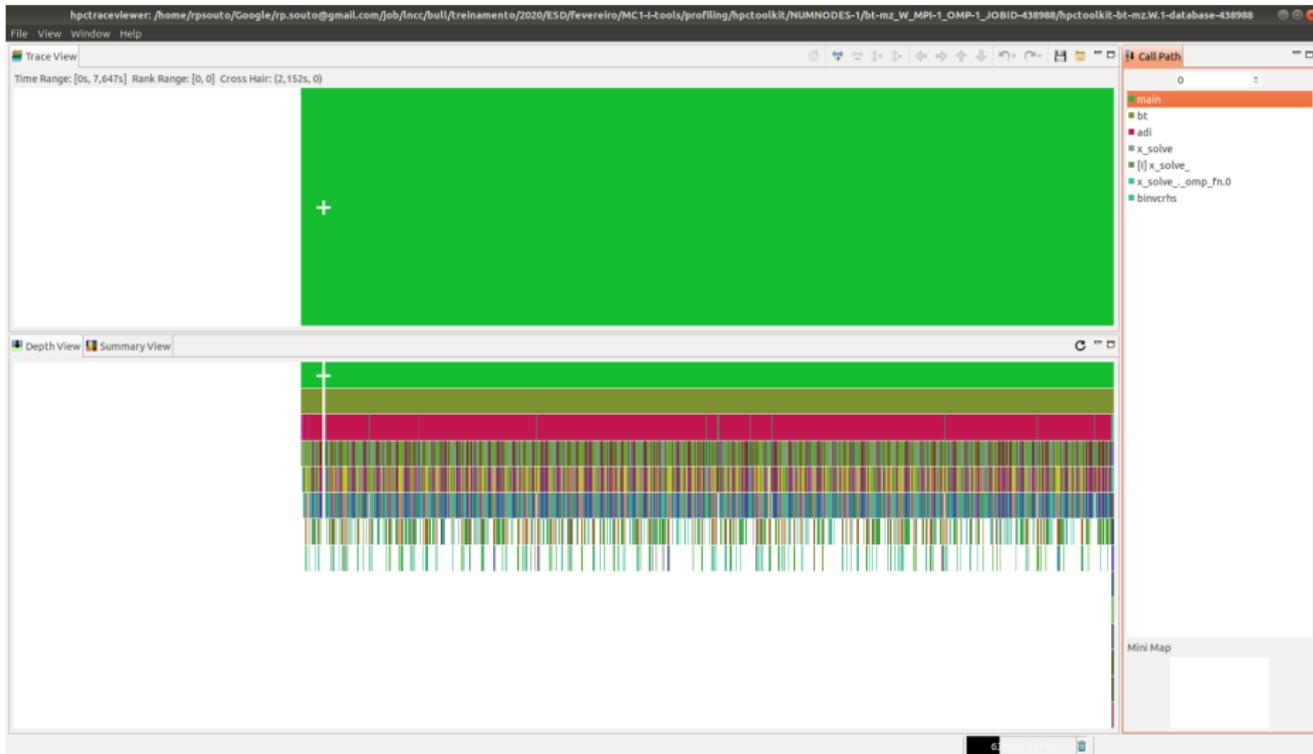
Visualizando no hpctraceview

NPB: estudo de caso

```
$ cd profiling/hpctoolkit/NUMNODES-1/bt-mz_W_MPI-1_OMP-1_JOBID-438988  
$ hpctraceview hpctoolkit-bt-mz.W.1-database-438988/
```

Visualizando no hpctraceview

`-nodes=1 -ntasks=1`



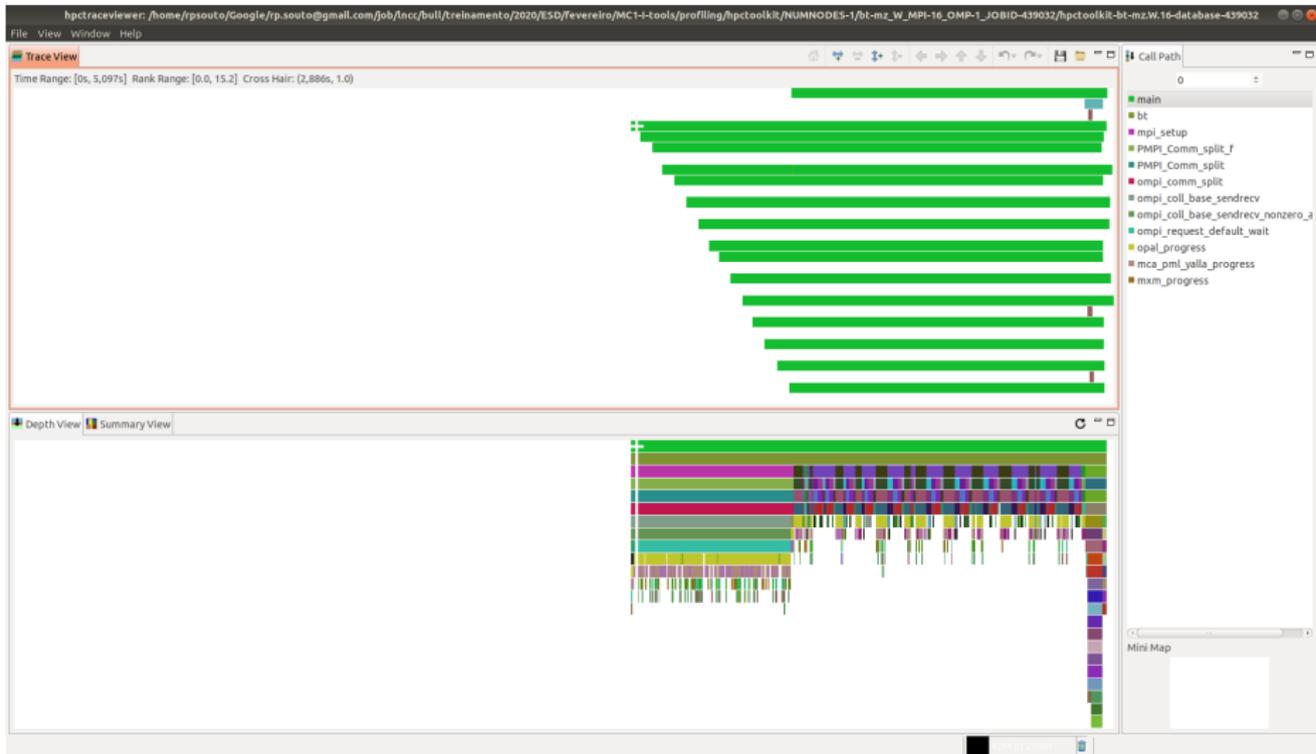
Visualizando no `hpctraceview`

NPB: estudo de caso

```
$ cd profiling/hpctoolkit/NUMNODES-1/bt-mz_W_MPI-16_OMP-1_JOBID-439032  
$ hpctraceview hpctoolkit-bt-mz.W.16-database-439032
```

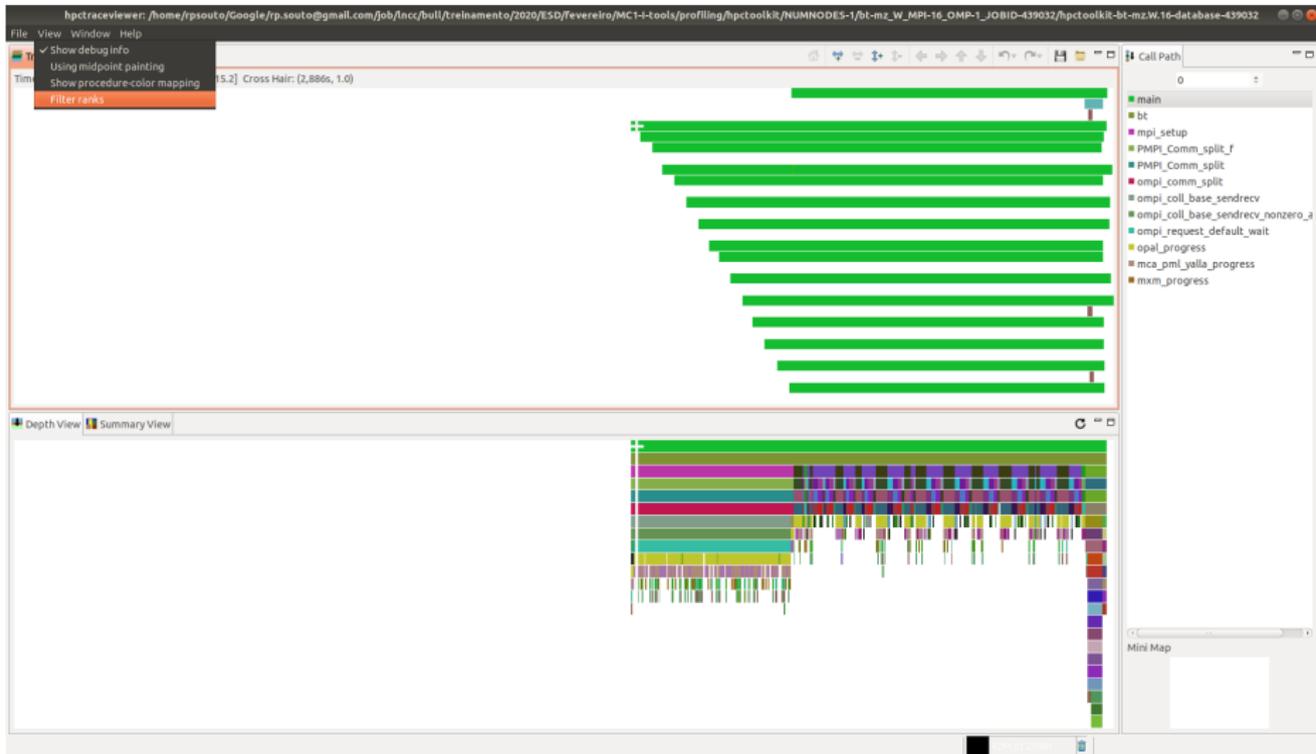
Visualizando no hpctraceview

`-nodes=1 -ntasks=16`



Visualizando no hpctraceview

`-nodes=1 -ntasks=16`



Visualizando no hpctraceview

`-nodes=1 -ntasks=16`

The screenshot displays the hpctraceview application interface. The main window shows a process trace with a 'Trace View' tab selected. The top status bar indicates 'Time Range: [0s, 5,097s]', 'Rank Range: [0,0, 15,2]', and 'Cross Hair: (2,886s, 1,0)'. A 'Call Path' panel on the right lists various process components such as 'main', 'bt', 'mpi_setup', and 'PMPI_Comm_split_f'. A 'Filter patterns' dialog box is open in the center, allowing users to add or remove glob patterns to filter displayed processes. The dialog includes a 'Mode of filter' section with 'To show' (selected) and 'To hide' options, and a 'Filter' input field with 'Add', 'Remove', and 'Remove all' buttons. The bottom of the application shows 'Depth View' and 'Summary View' tabs.

Visualizando no hpctraceview

`-nodes=1 -ntasks=16`

The screenshot displays the hpctraceview application interface. At the top, the title bar shows the file path: `hpctraceviewer: /home/rpsouto/Google/rp_souto@gmail.com/job/incc/dull/treinamento/2020/ESD/revereiro/MCT-I-tools/profiling/hpctrackIR/NUMNODES-1/bl-mz_W_MPI-16_OMP-1_JOBID-439032/hpctrackIR-bl-mz.W.16-database-439032`. The main window is titled "Trace View" and shows a Gantt chart with green bars representing process execution. The "Time Range" is [0s, 5,097s] and the "Rank Range" is [0.0, 15.2]. A "Cross Hair" is positioned at (2,886s, 1.0). On the right, a "Call Path" panel lists the following components: main, bt, mpi_setup, PMPI_Comm_split_f, PMPI_Comm_split, ompi_coll_base_sendrecv, ompi_coll_base_sendrecv_nonzero_a, ompi_request_default_wait, opal_progress, mca_pml_yalla_progress, and mxm_progress. A "Filter patterns" dialog box is open in the center, containing the following text: "Please type a pattern in the format minimum:maximum:stride. Any omitted or invalid sections will match as many processes or threads as possible." Below this, it provides examples: "For instance, 3:7:2 in the process box with the thread box empty will match all threads of processes 3, 5, and 7. 1 in the thread box with the process box empty will match thread 1 of all processes. 1:2 in the process box and 2:4:2 in the thread box will match 1,2, 1,4, 3,2, 3,4, 5,2 ...". The dialog has input fields for "Process" (containing "0:15:1") and "Thread" (containing "0"). Buttons for "Add", "Remove", "Remove all", "Cancel", and "OK" are visible.

Visualizando no hpctraceview

`-nodes=1 -ntasks=16`

The screenshot displays the hpctraceview application interface. At the top, the title bar shows the file path: `hpctraceviewer: /home/rpsouto/Google/rp_souto@gmail.com/job/incc/dull/treinamento/2020/ESD/fevereiro/MCT-I-tools/profiling/hpctrackit/NUMNODES-1/ht-mz_W_MPI-16_OMP-1_JOBID-439032/hpctoolkit-ht-mz.W.16-database-439032`. The main window is titled "Trace View" and shows a horizontal bar chart representing process execution over time. The top part of the chart shows a few long green bars, while the bottom part shows a dense, multi-colored forest of bars. A "Filter patterns" dialog box is open in the center, with the following content:

Filter patterns
Add/remove glob patterns to filter displayed processes

Mode of filter
 To show To hide

Selecting the "To show" radio button will show matching processes, while selecting the "To hide" button will hide them.

Filter
Add 0:15:1,0

Remove

Remove all

Cancel OK

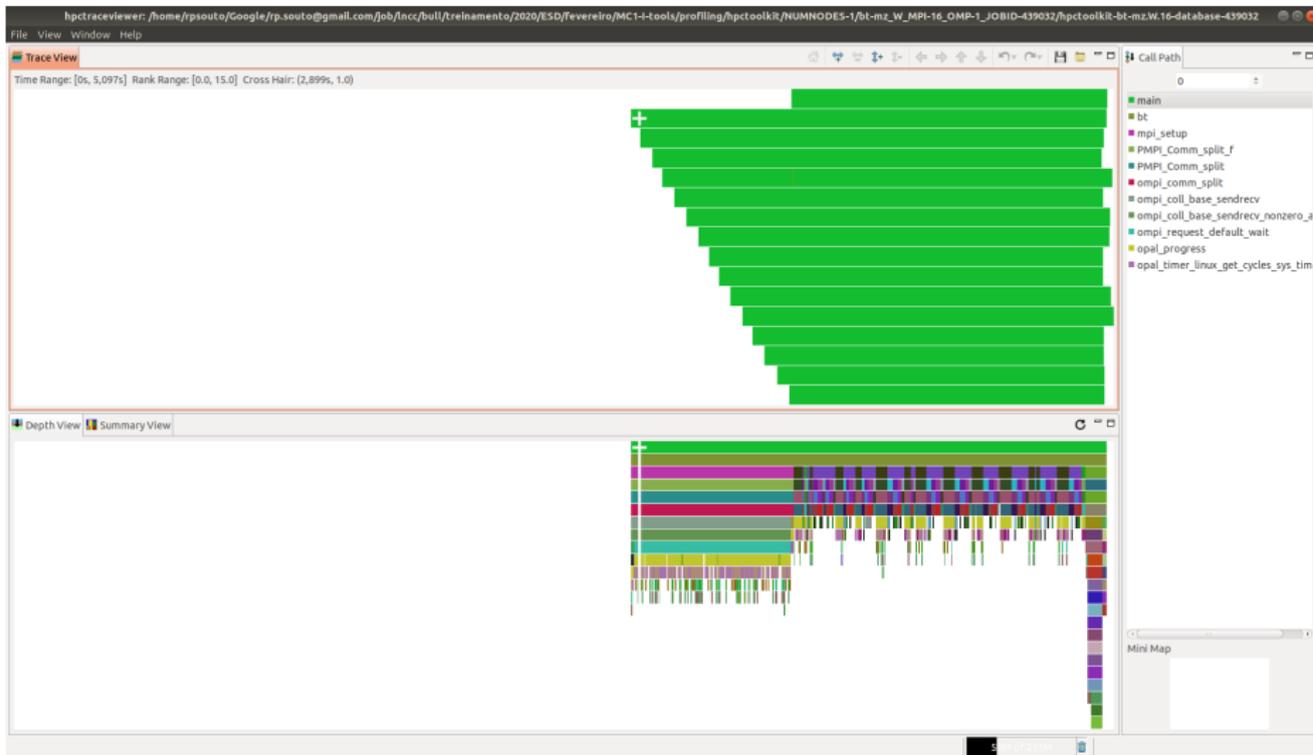
On the right side, a "Call Path" panel lists the following components:

- main
- bt
- mpi_setup
- PMPI_Comm_split_f
- PMPI_Comm_split
- ompi_comm_split
- ompi_coll_base_sendrecv
- ompi_coll_base_sendrecv_nonzero_a
- ompi_request_default_wait
- opal_progress
- mca_pml_yalla_progress
- mxm_progress

At the bottom left, there are tabs for "Depth View" and "Summary View". At the bottom right, there is a "Mini Map" section.

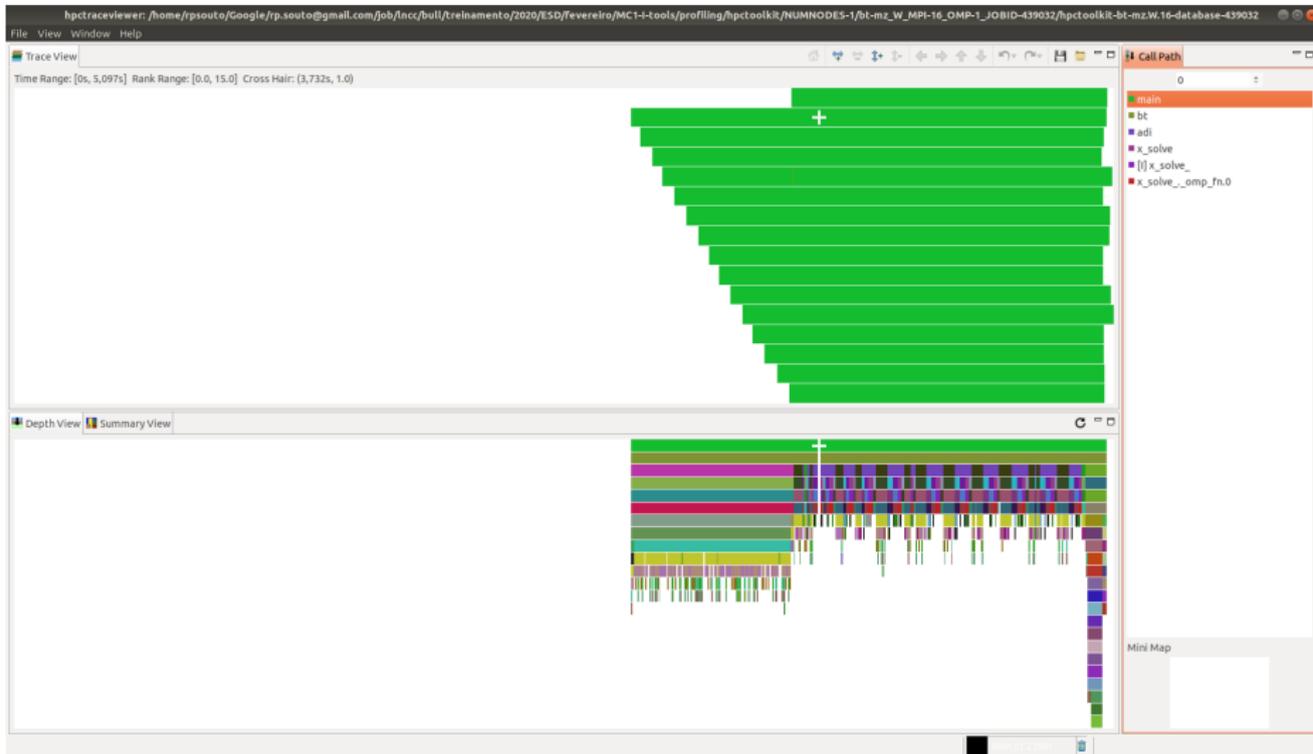
Visualizando no hpctraceview

`-nodes=1 -ntasks=16`



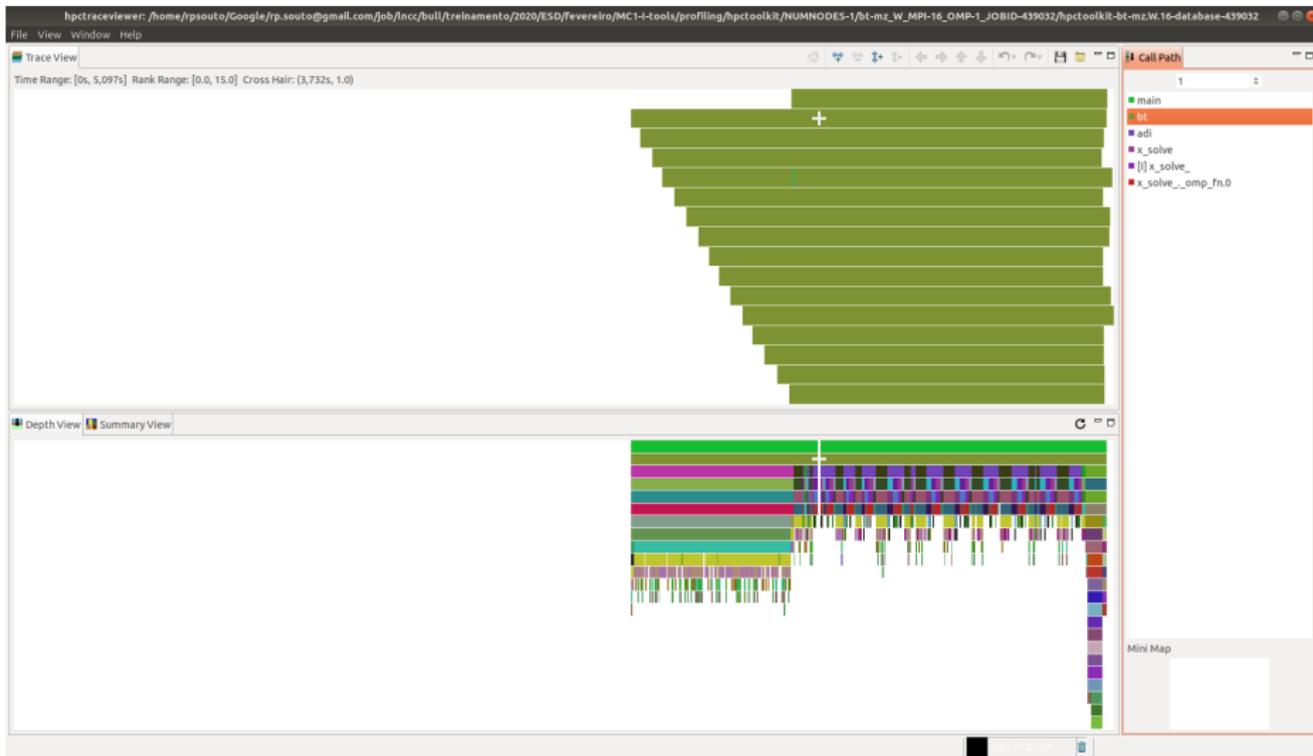
Visualizando no hpctraceview

`-nodes=1 -ntasks=16 - função main`



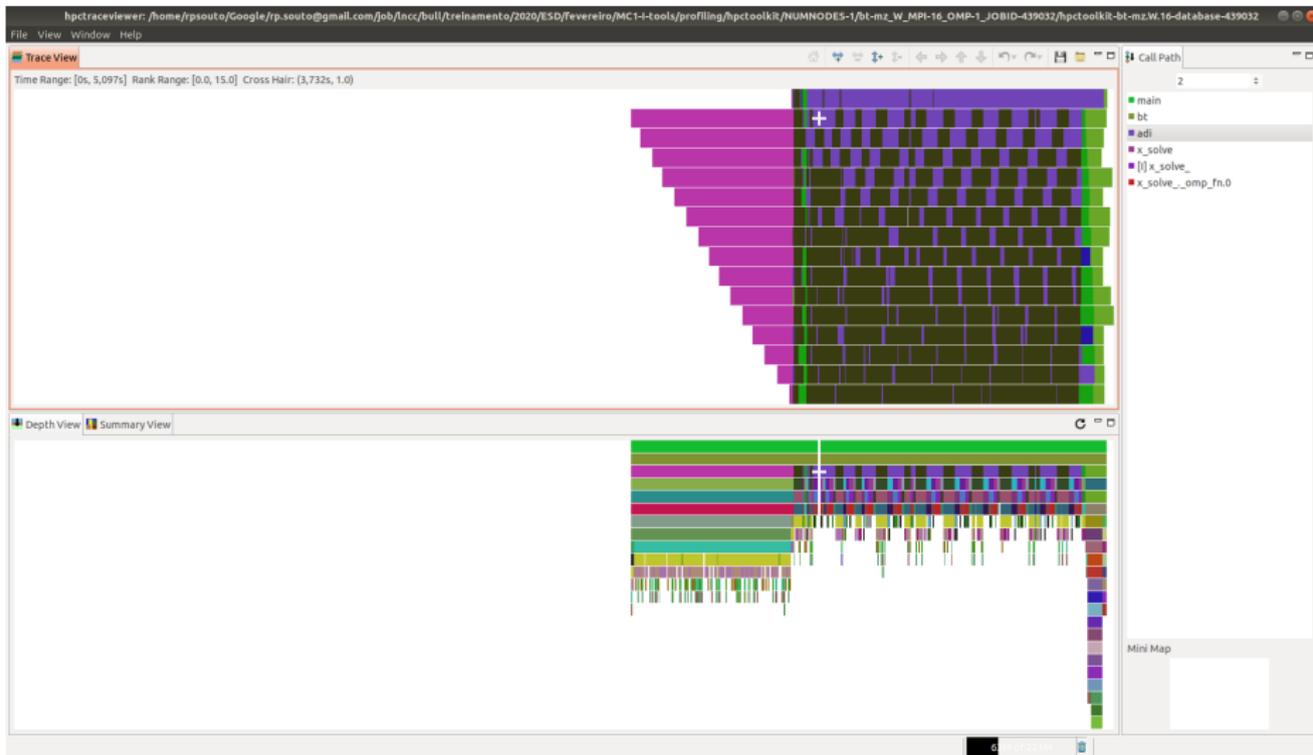
Visualizando no hpctraceview

`-nodes=1 -ntasks=16 - função bt`



Visualizando no hpctraceview

`-nodes=1 -ntasks=16` - função adi



Visualizando no hpctraceview

`-nodes=1 -ntasks=16` - função `xsolve`

