Performance Analysis with Vampir

Bert Wesarg
Technische Universität Dresden

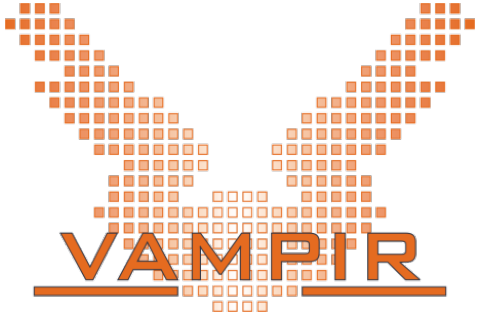

Outline

- **Part I: Welcome to the Vampir Tool Suite**
  - Mission
  - Event trace visualization
  - Vampir & VampirServer
  - The Vampir displays
- **Part II: Vampir hands-on**
  - Visualizing and analyzing NPB-MZ-MPI / BT

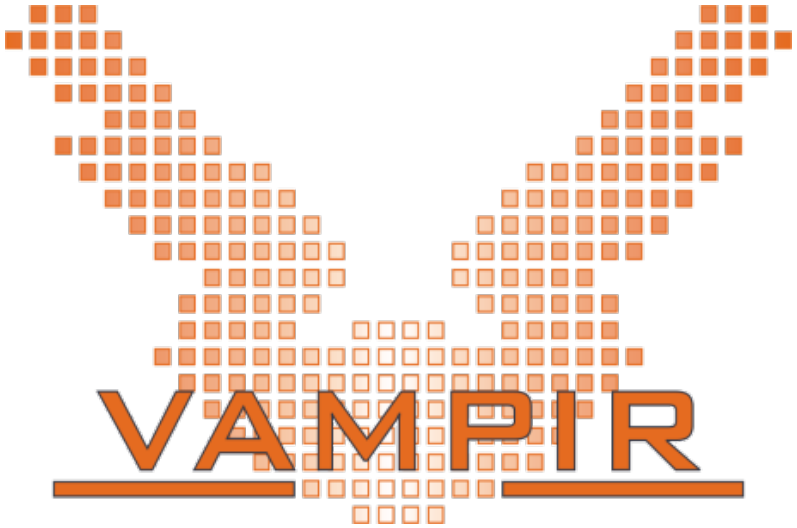

Event trace visualization with Vampir

- Alternative and supplement to automatic analysis
- Show dynamic run-time behavior graphically at any level of detail
- Provide statistics and performance metrics

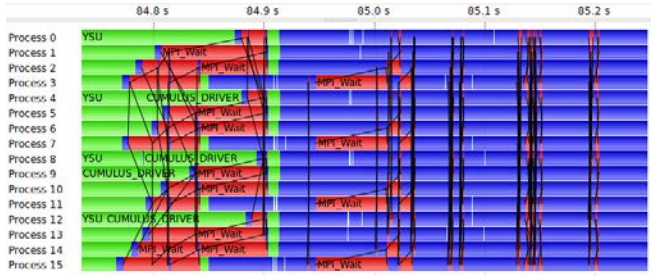

- **Timeline charts**
  - Show application activities and communication along a time axis

- **Summary charts**
  - Provide quantitative results for the currently selected time interval

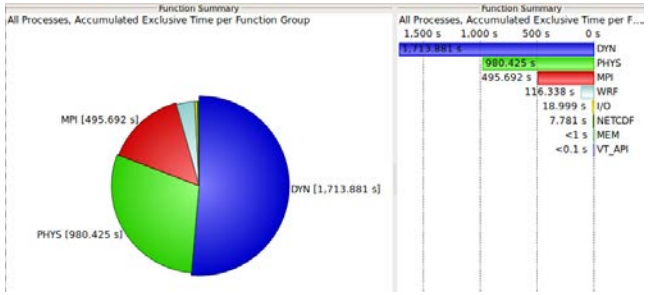

Visualization modes (1)
Directly on front end or local machine

```
% vampir
```

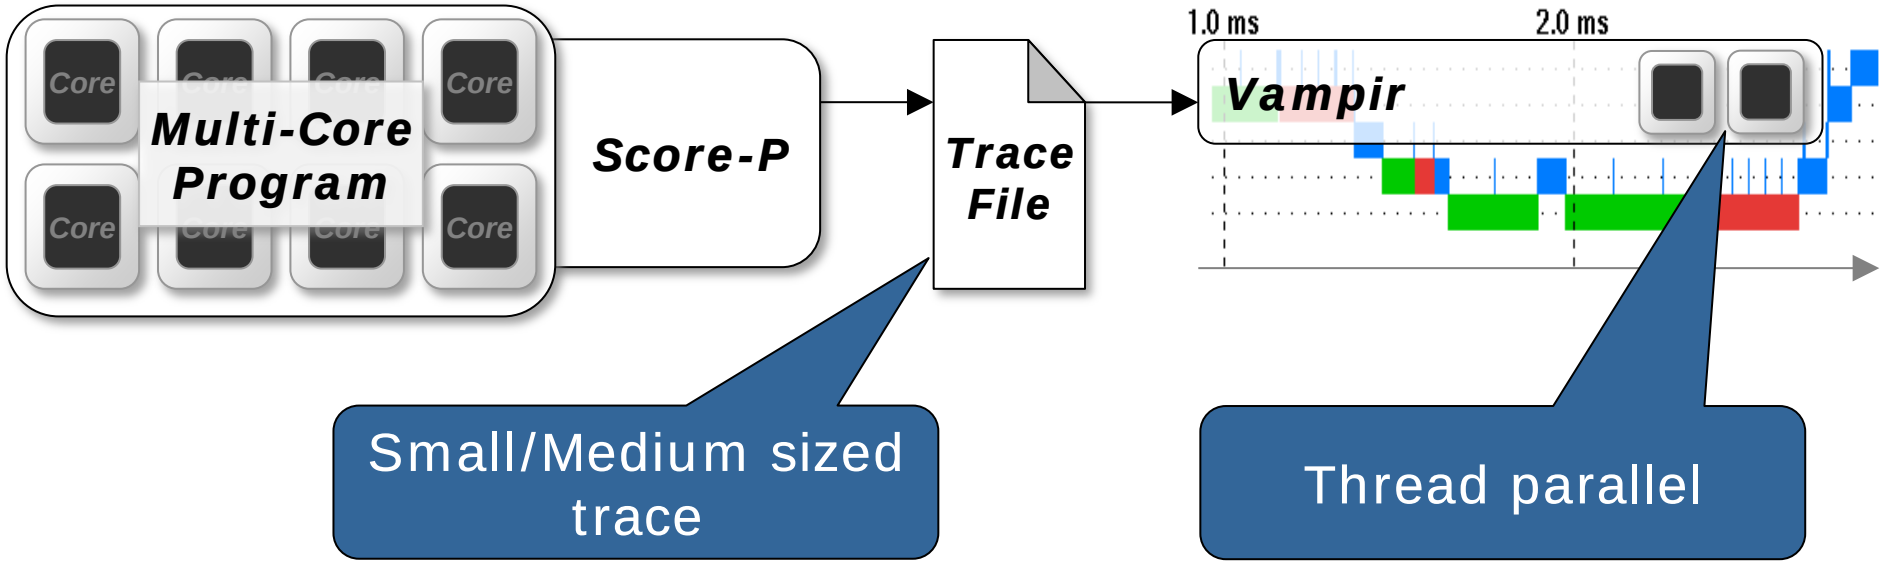

Visualization Modes (2)
On local machine with remote VampirServer

```
% vampirserver start
```

```
% vampir
```

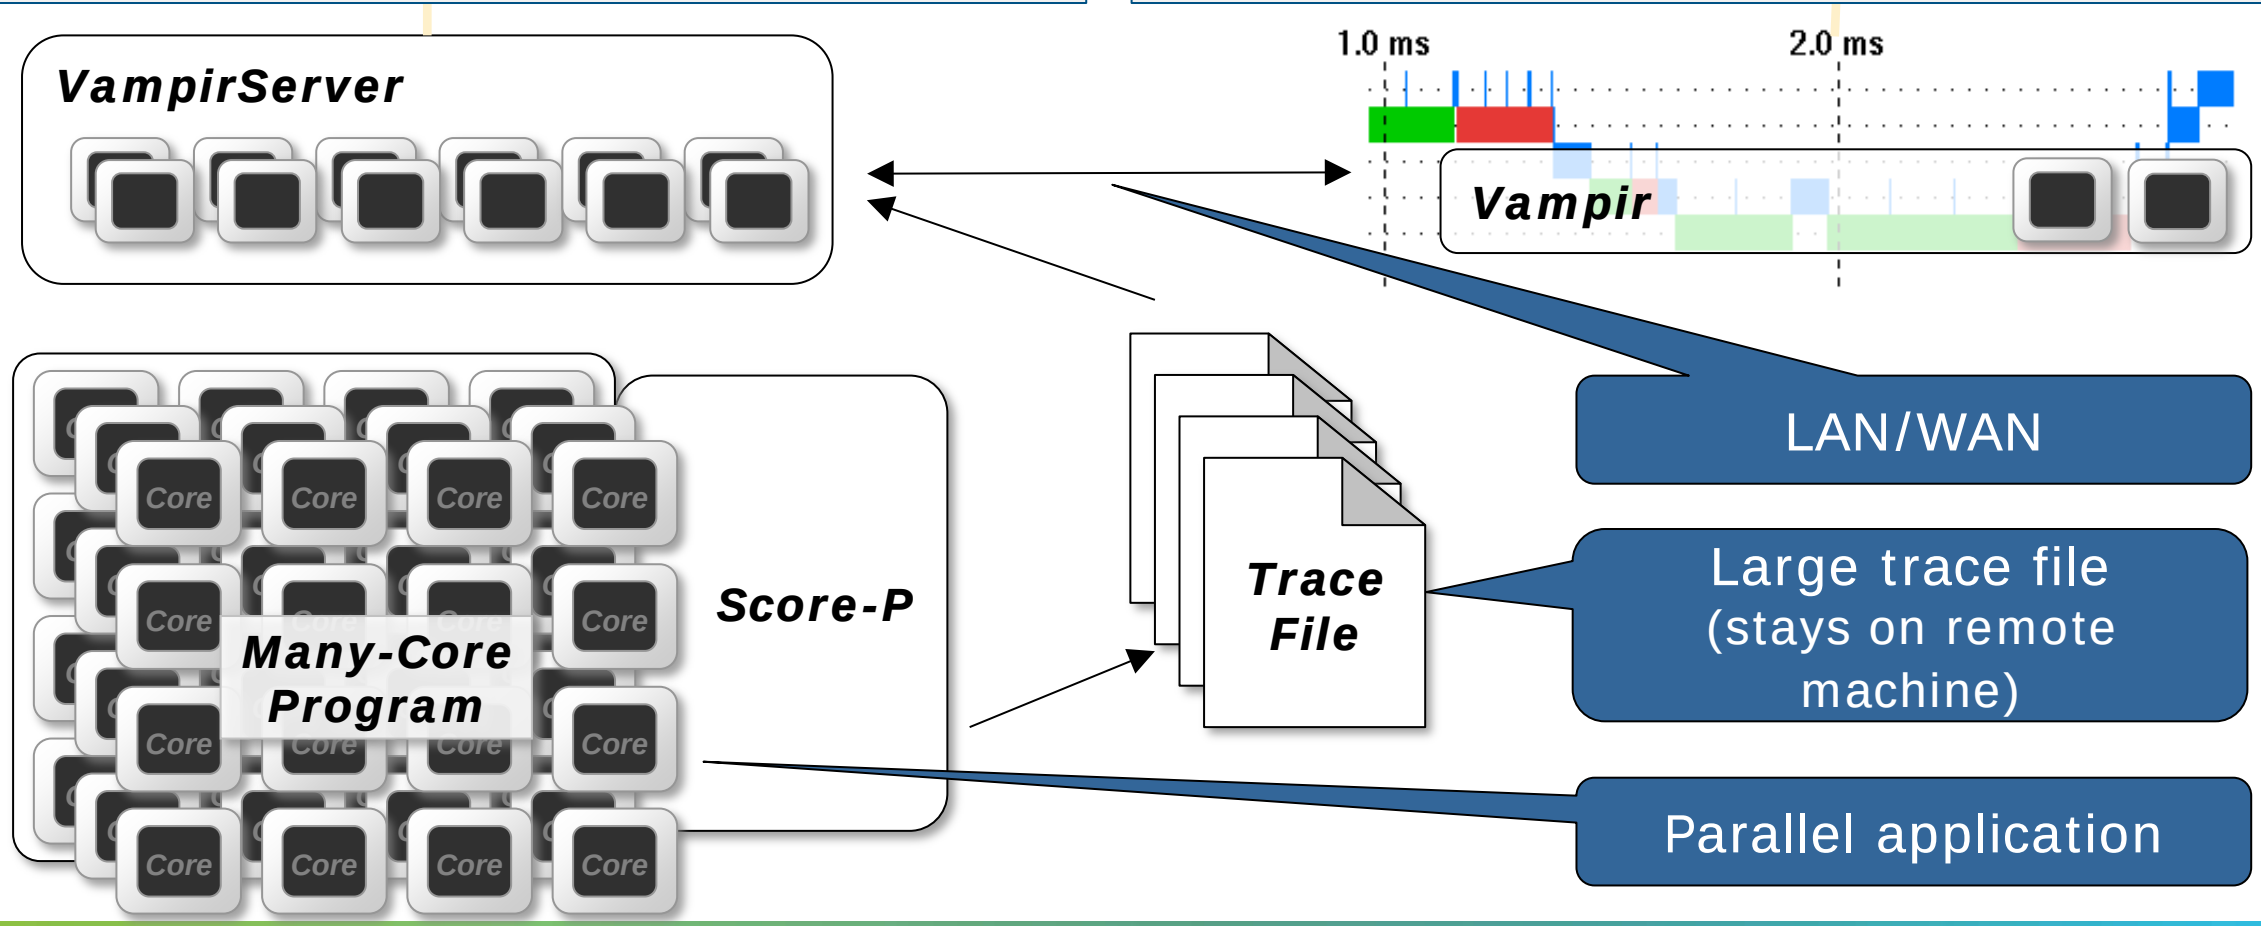

The main displays of Vampir

- Timeline Charts:
  -  Master Timeline
  -  Process Timeline
  -  Counter Data Timeline
  -  Performance Radar
- Summary Charts:
  -  Function Summary
  -  Message Summary
  -  Process Summary
  -  Communication Matrix View

Hands-on:
Visualizing and analyzing NPB-MZ-MPI / BT

Start Vampir

```
% vampirserver start
Launching VampirServer...
Submitting slurm 45 minutes job (this might take a while)...
salloc: Granted job allocation 26559
VampirServer 9.0.0  (r9950)
Licensed to VI-HPS Tools Workshop 04/2016
Running 4 analysis processes... (abort with vampirserver stop \
> 14457)
VampirServer <14457> listens on: uv2:30009
```

Remember host and port

Remember host and port

```
% ssh -YAC lxlogin1.lrz.de –l <login>
% source ~lu23bim/LRZ-VIHPSTW21/tools/source-me.scorep-2.0.1.mpt.sh
% vampir
```

Start the Vampir GUI on the login node or on your local machine

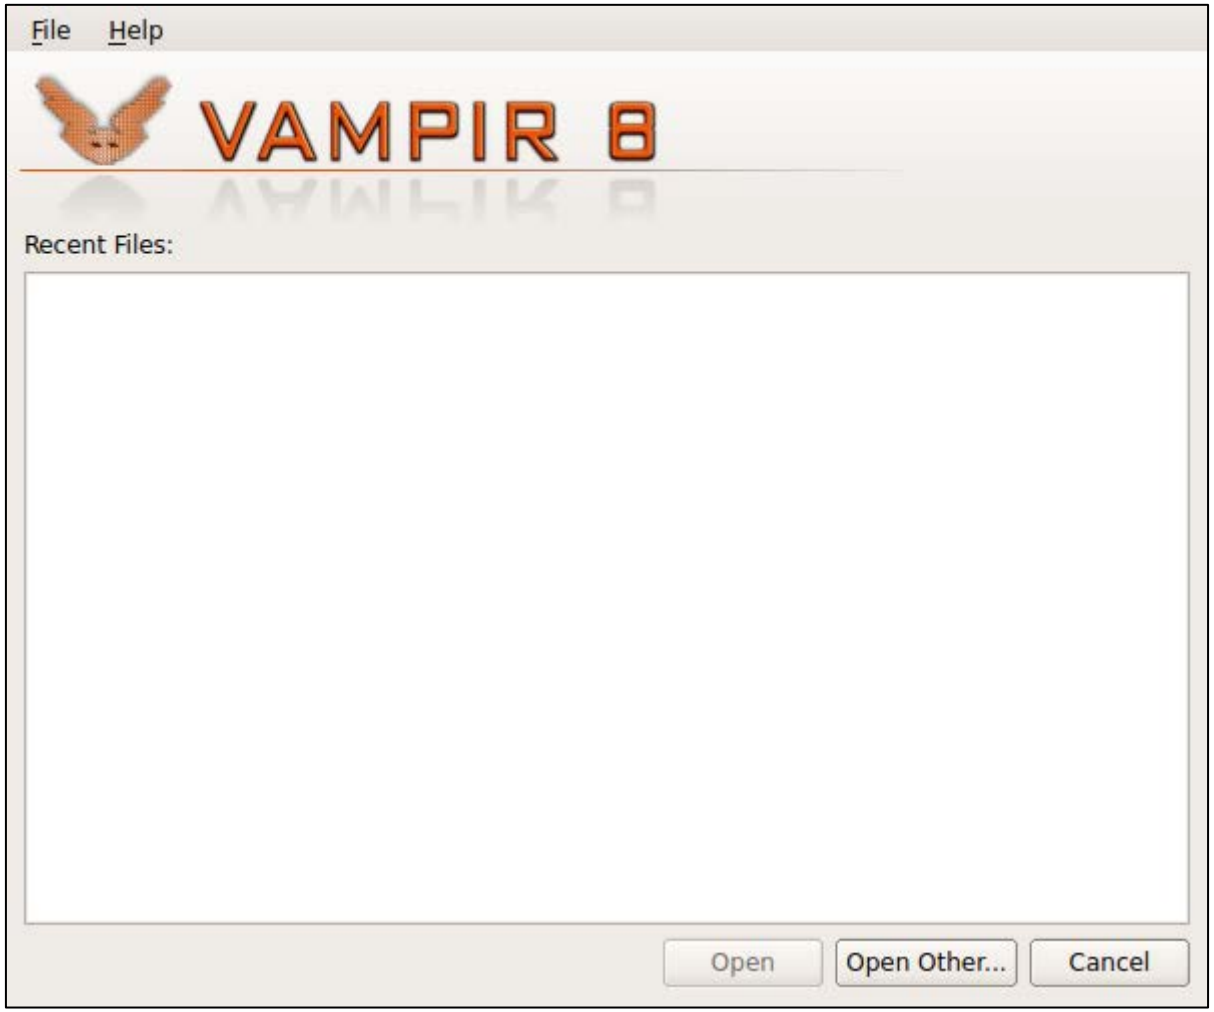
File
Help
VAMPIR 8
Recent Files:
Open
Open Other...
Cancel

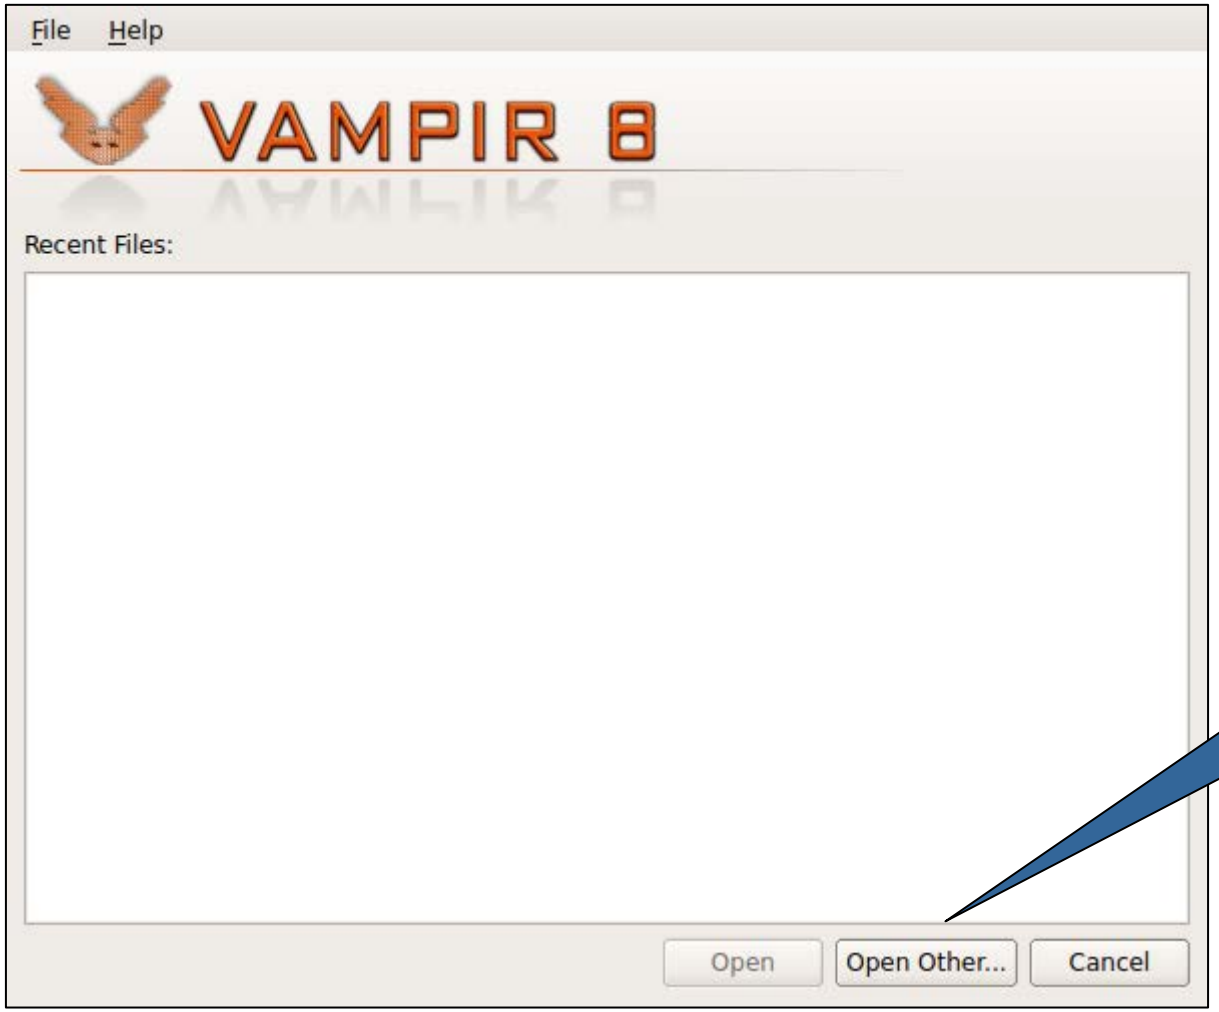

Use the “Open Other” option

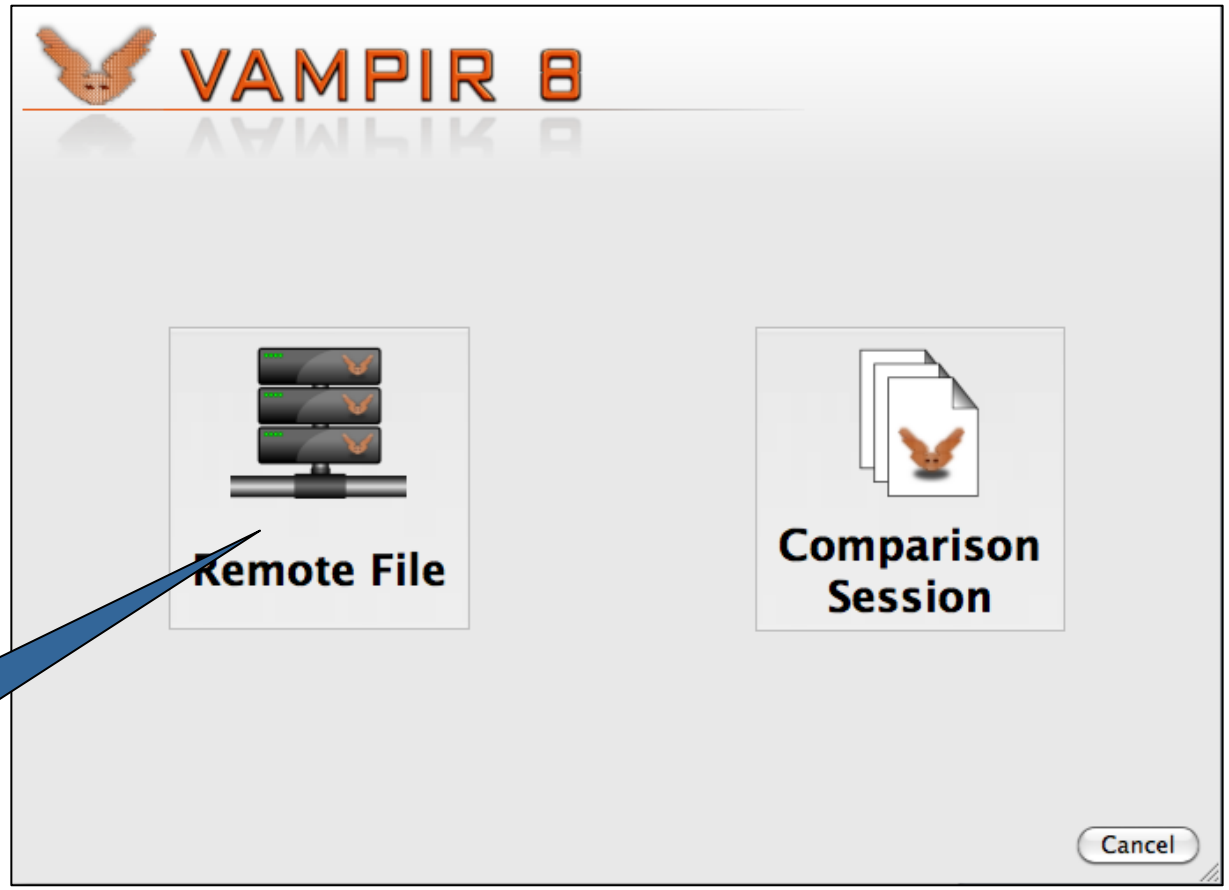

Select “Remote File”

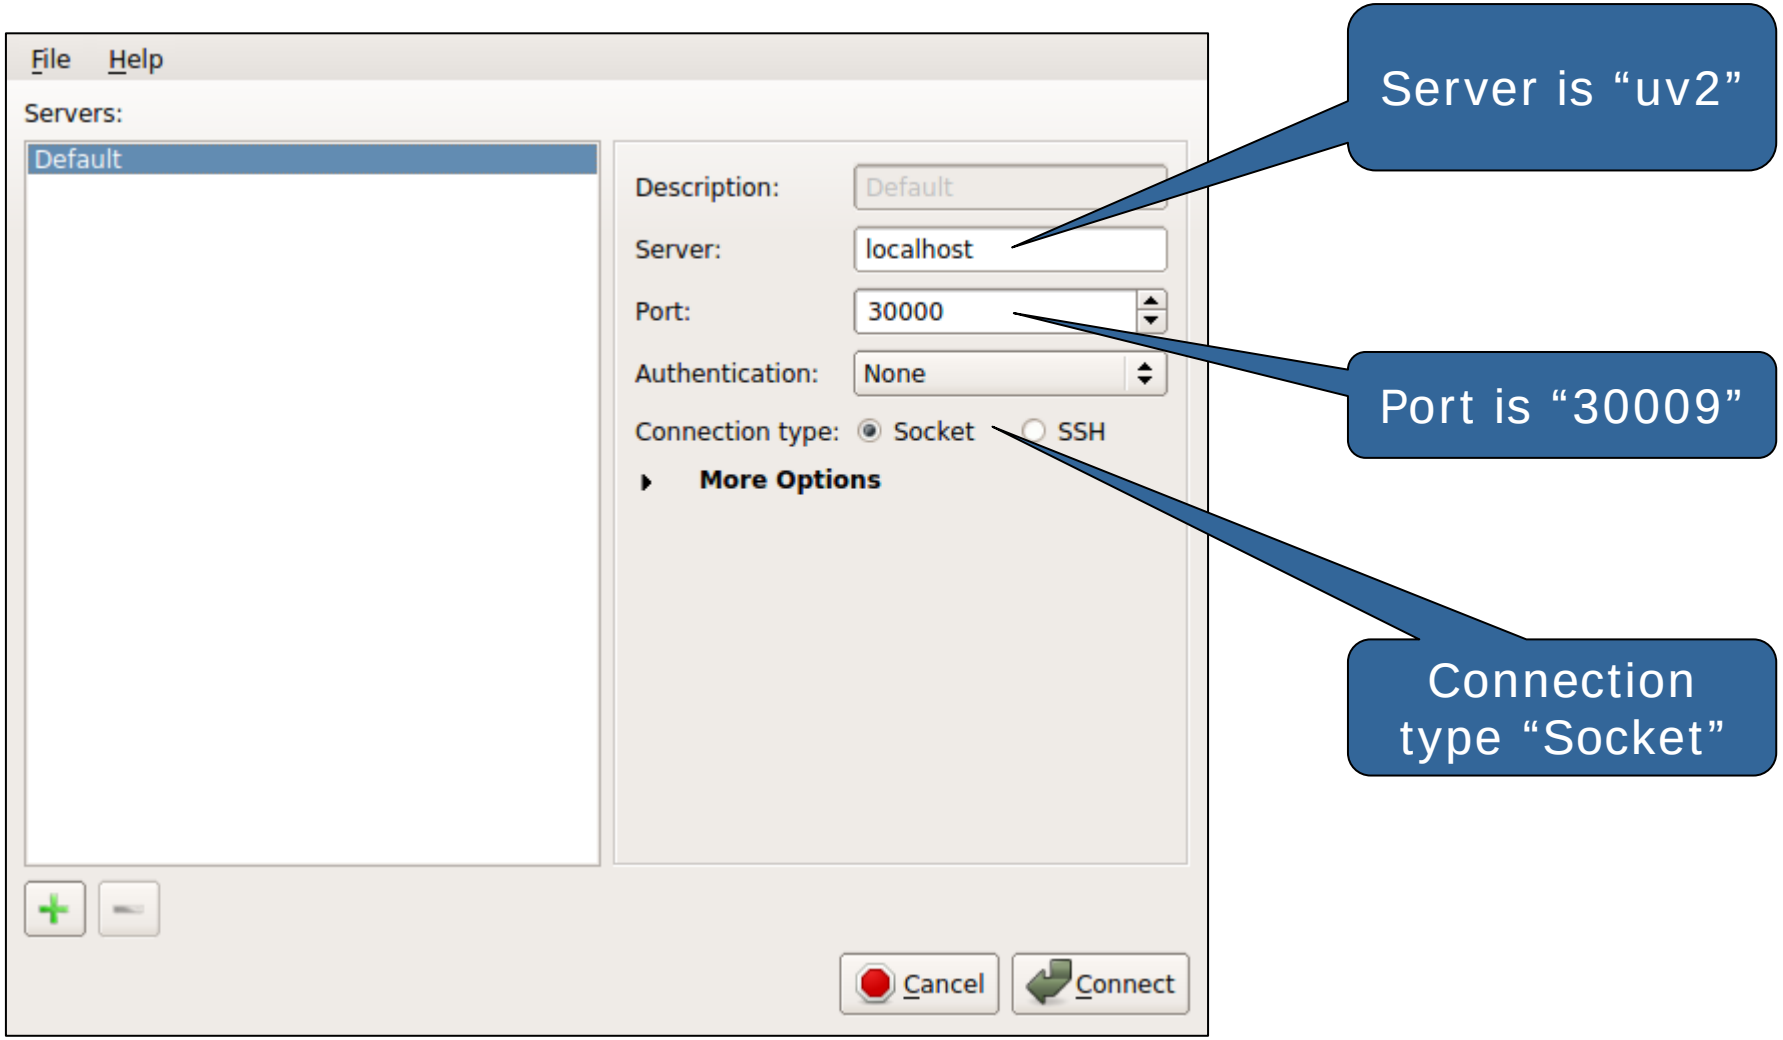
File
Help
Servers:
Default
Description:
Default
Server:
localhost
Port:
30000
Authentication:
None
Connection type:
Socket
SSH
More Options
Cancel
Connect
Server is “uv2”
Port is “30009”
Connection type “Socket”

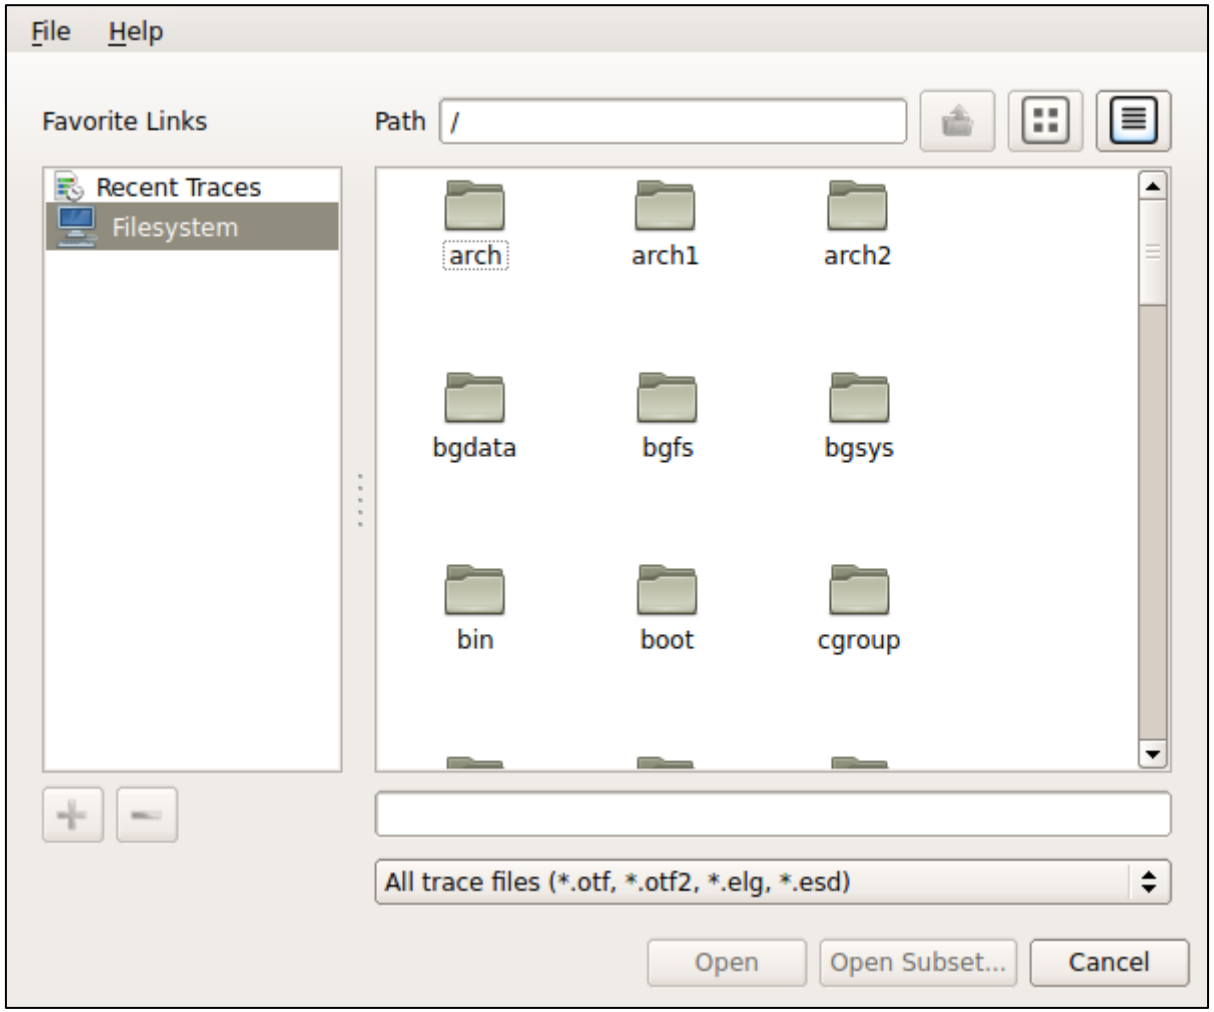
File
Help
Favorite Links
Path
/
Recent Traces
Filesystem
arch
arch1
arch2
bgdata
bgfs
bgsys
bin
boot
cgroup
All trace files (*.otf, *.otf2, *.elg, *.esd)
Open
Open Subset...
Cancel

Serial exercises

- Extract exercise traces:

```
% cd ~
% tar xf ~lu23bud/LRZ-VIHPSTW21/vampir/scorep_bt_C_traces.tar.xz
```

- Open `~/scorep_bt_C_serial_trace/traces.otf2` in Vampir

1. What are the functions with the longest running time in total?
   1. Color the top 3 with different colors
2. Which function called these functions?
3. How often is this function called?
4. What is the average call time of this function?

OpenMP exercises

- Open `~/scorep_bt_C_omp_4_trace/traces.otf2` in Vampir

1. Total time spent in all OpenMP barriers?
   1. How much in the barriers from the {x,y,z}_solve parallel regions
2. Which thread spends the most time in these barriers?

MPI exercises

- Open `~/scorep_bt_C_mpi_16_trace/traces.otf2` in Vampir

1. Find `MPI_Init`
2. How many nodes did the job used?
3. Total wait time in MPI functions?
4. Communication pattern and cliques?
5. Any MPI collectives used?

BT-MZ-MPI Trace

- Open `scorep_bt-mz_4x4_trace/traces.otf2` in Vampir

Visualization of the NPB-MZ-MPI / BT trace

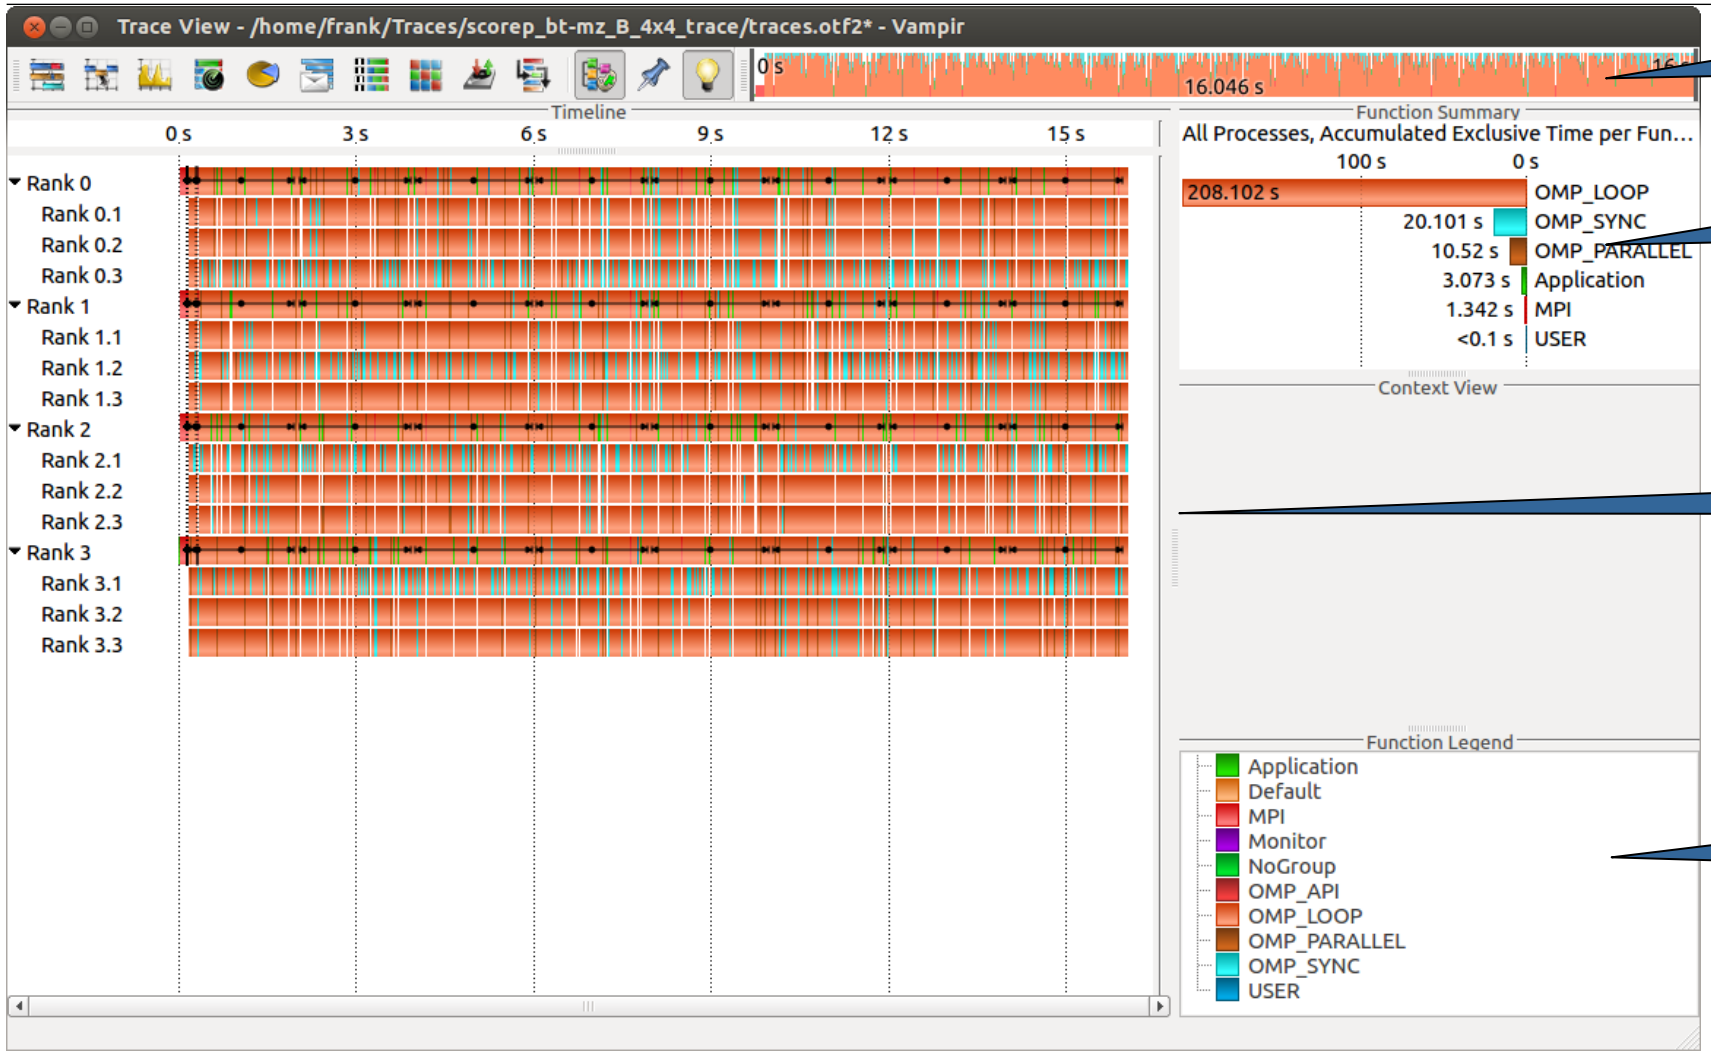

Navigation Toolbar

Function Summary

Master Timeline

Function Legend

Visualization of the NPB-MZ-MPI / BT trace
Master Timeline

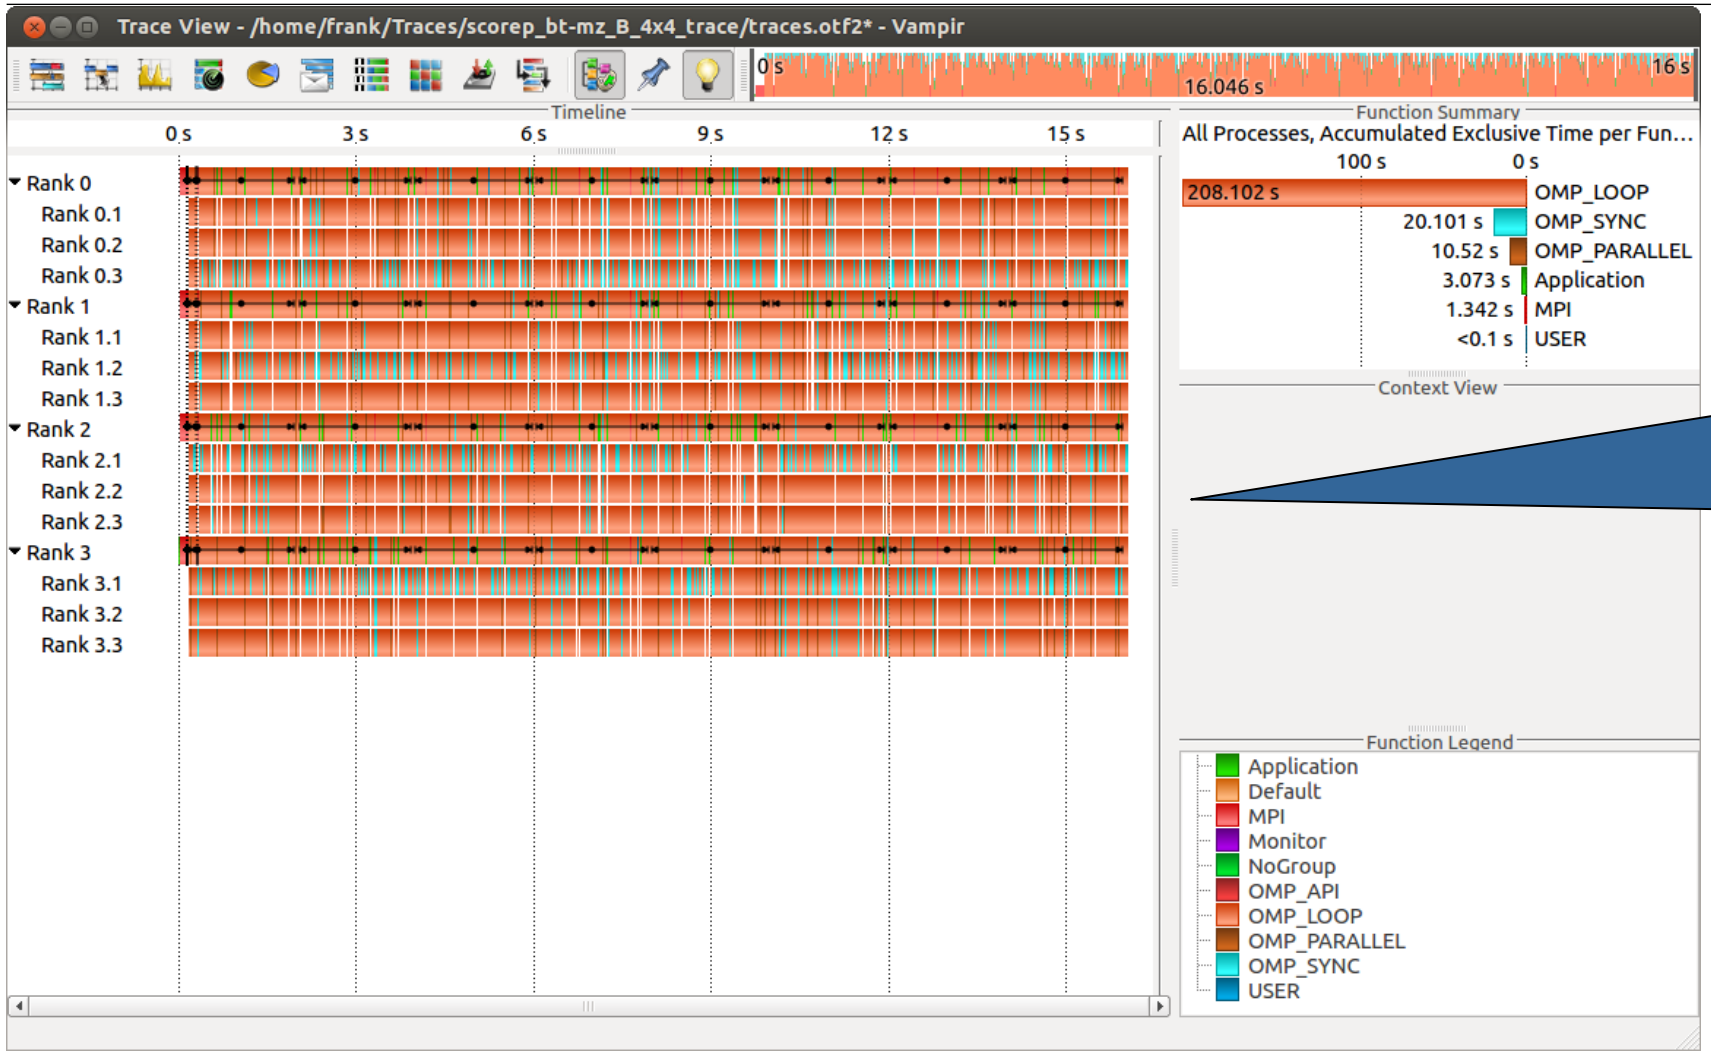

Detailed information about functions, communication and synchronization events for collection of processes.

Visualization of the NPB-MZ-MPI / BT trace
Process Timeline

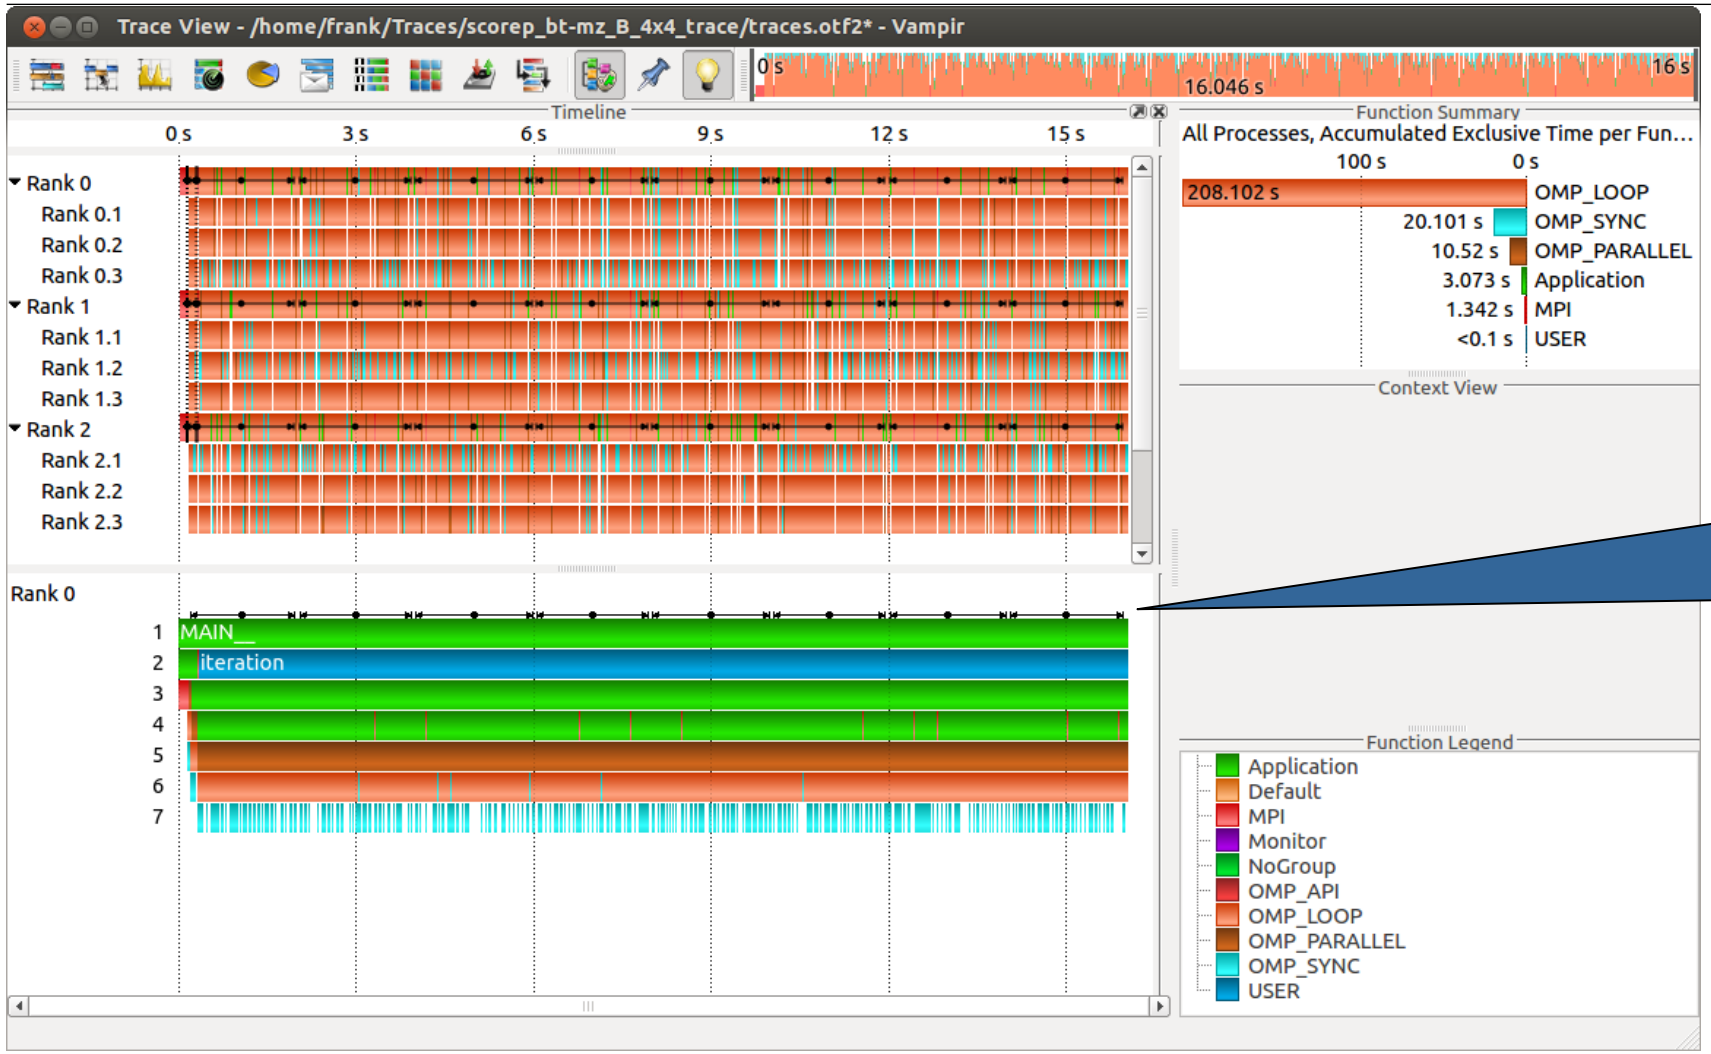

Detailed information about different levels of function calls in a stacked bar chart for an individual process.

Visualization of the NPB-MZ-MPI / BT trace
Typical program phases

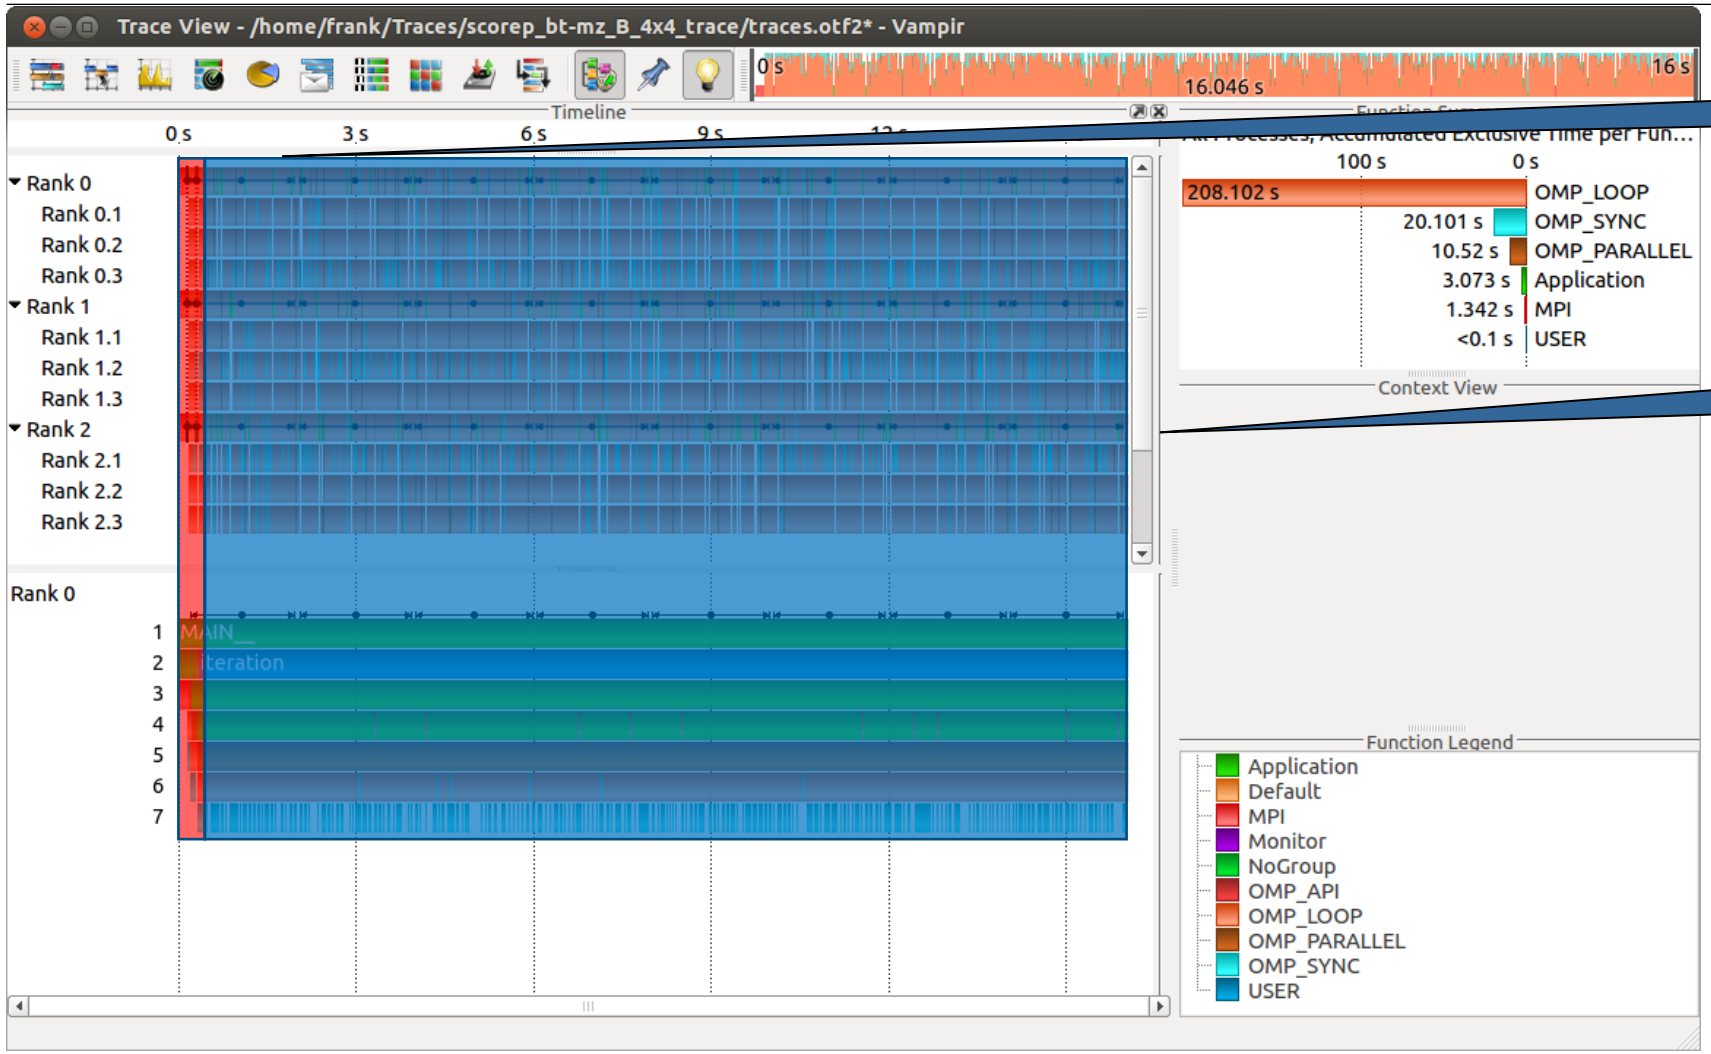

Initialisation Phase

Computation Phase

Visualization of the NPB-MZ-MPI / BT trace
Counter Data Timeline

Detailed counter information over time for an individual process.

Visualization of the NPB-MZ-MPI / BT trace
Performance Radar

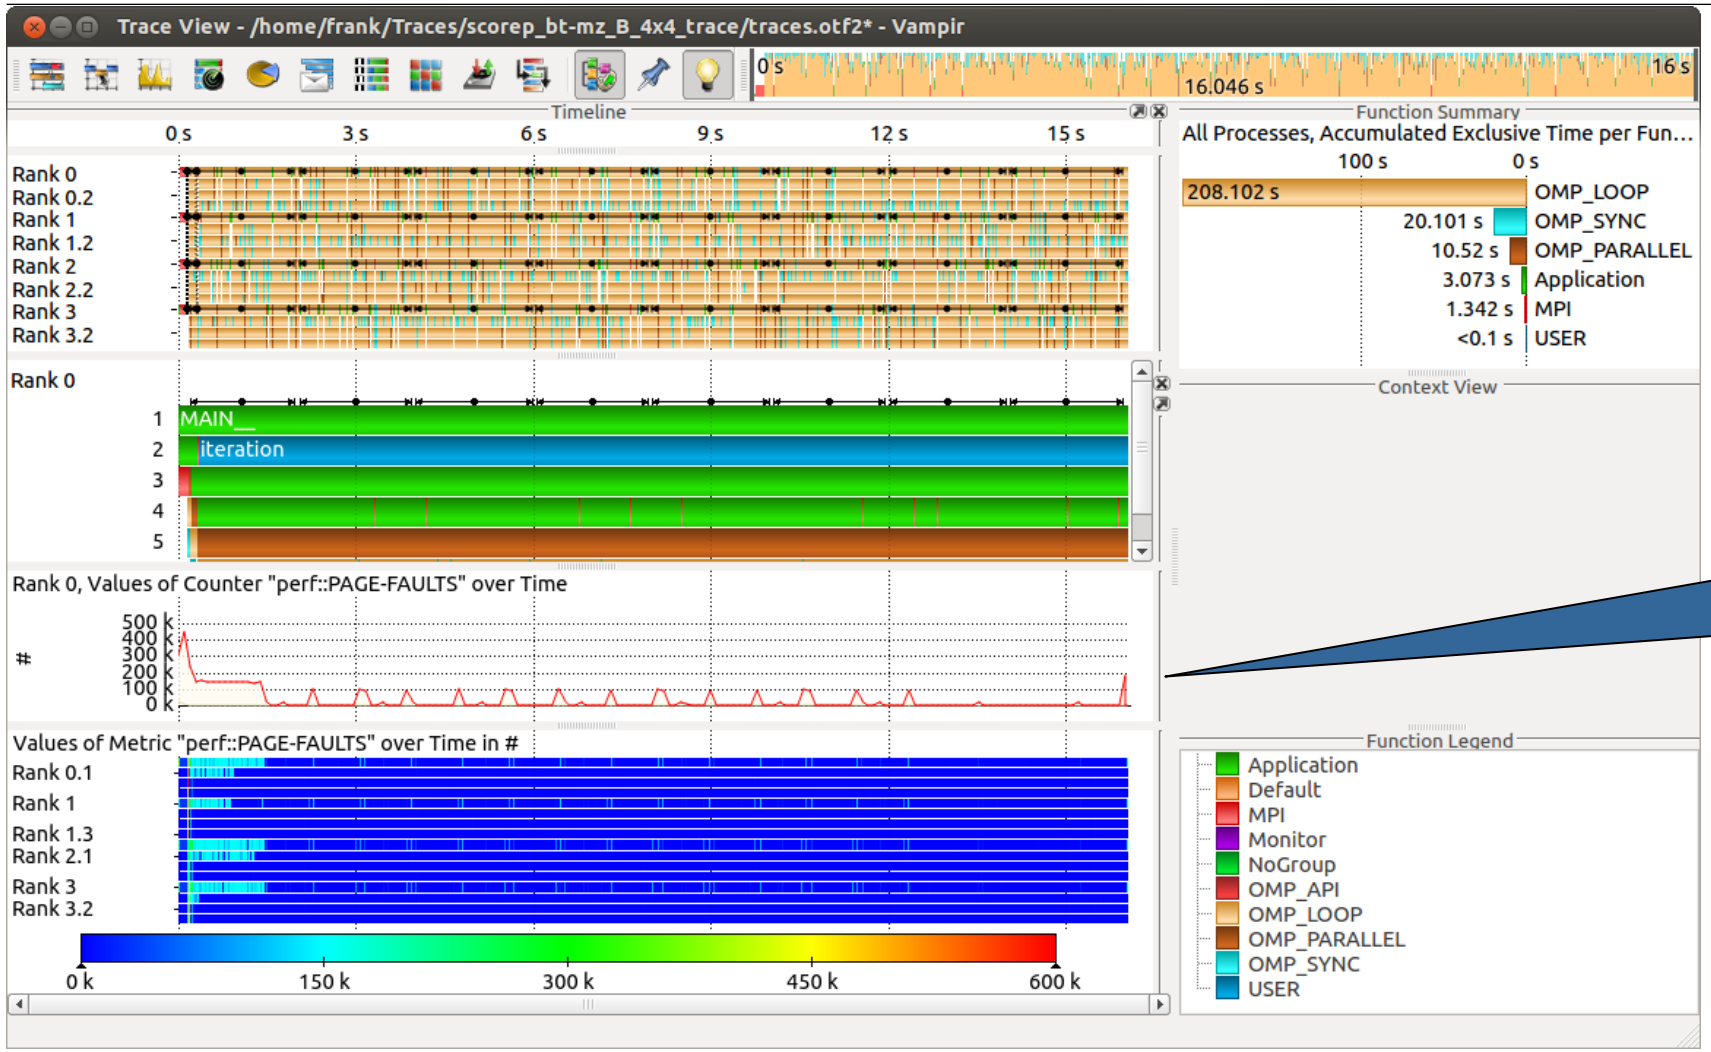

Detailed counter information over time for a collection of processes.

Visualization of the NPB-MZ-MPI / BT trace
Zoom in: Inititialisation Phase

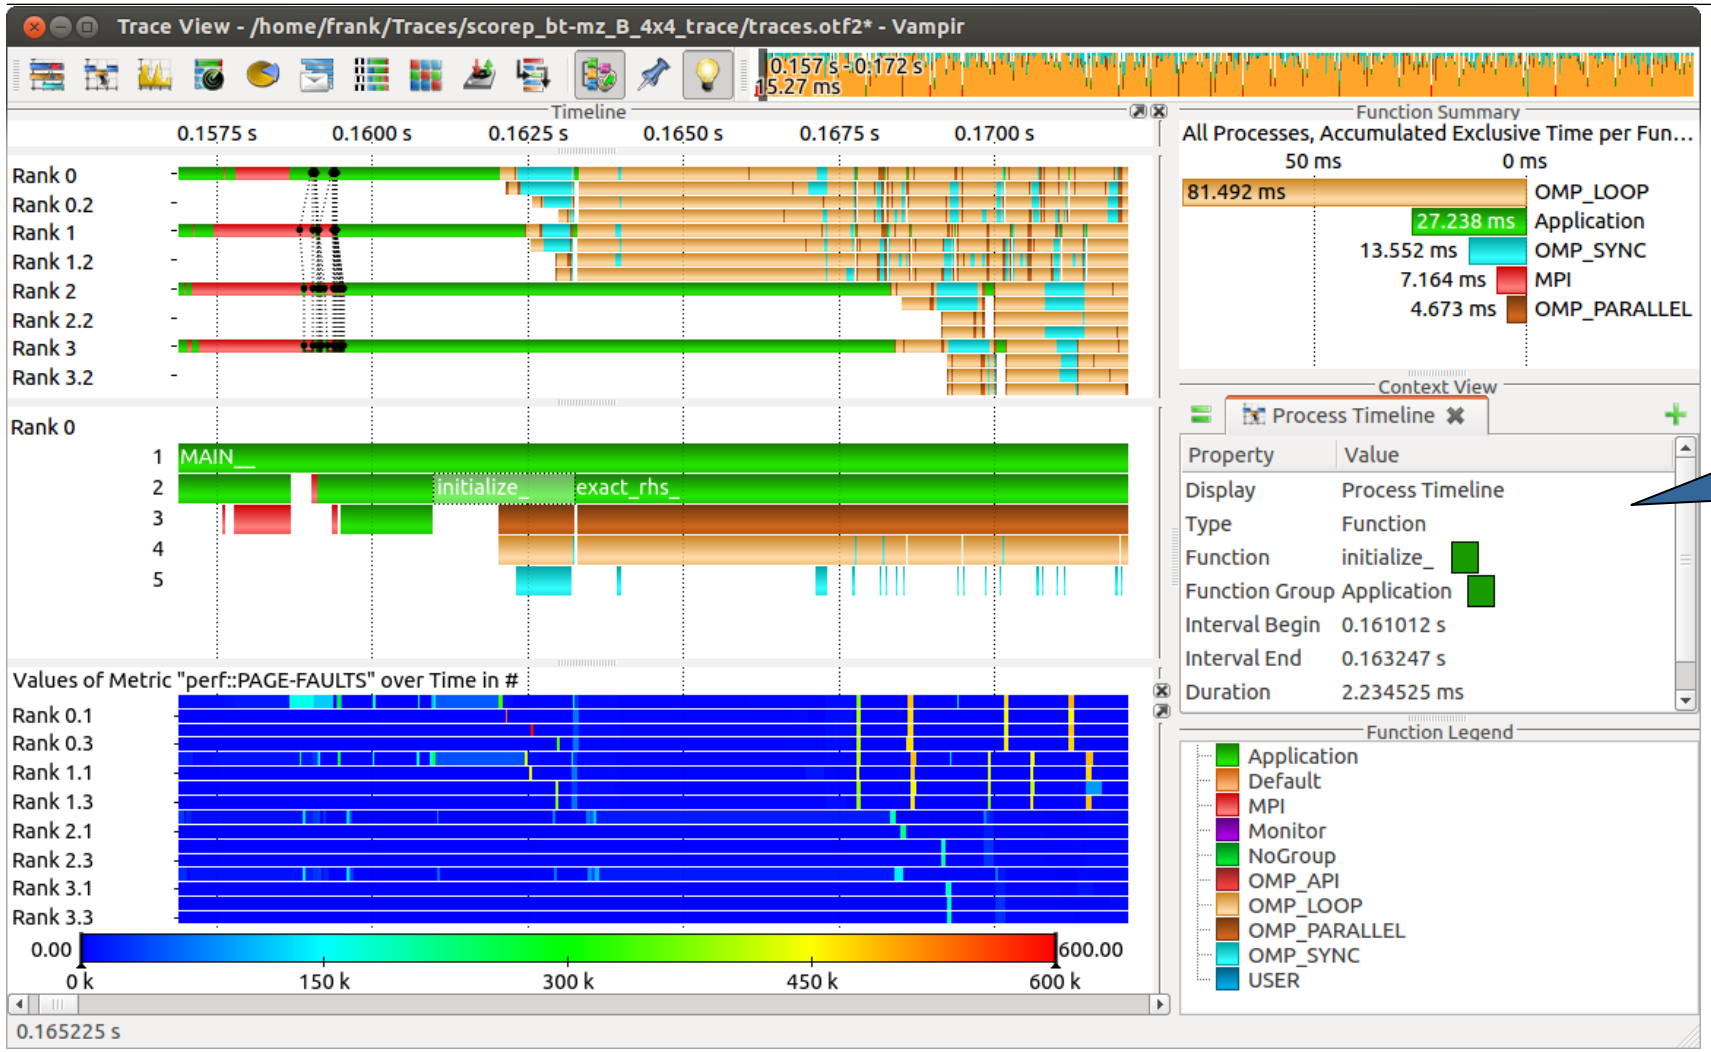

Context View: Detailed information about function “initialize_”.

Visualization of the NPB-MZ-MPI / BT trace
Find Function

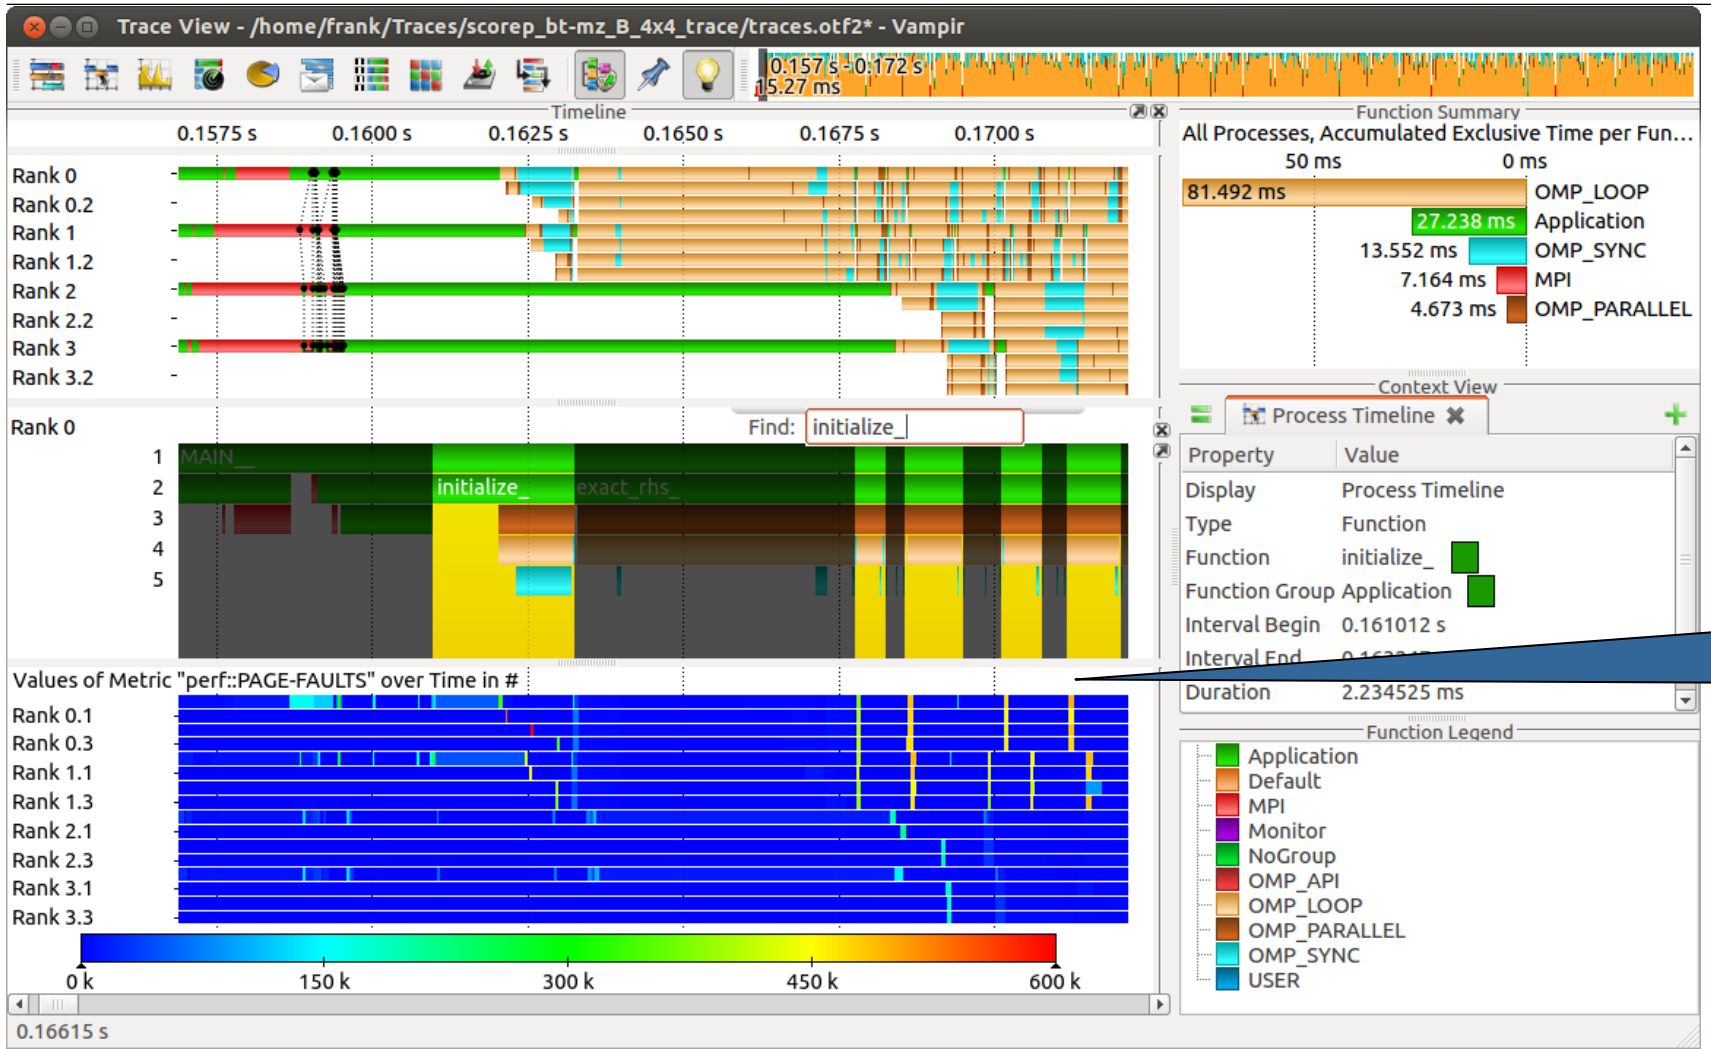

Execution of function “initialize_” results in higher page fault rates.

Visualization of the NPB-MZ-MPI / BT trace
Computation Phase

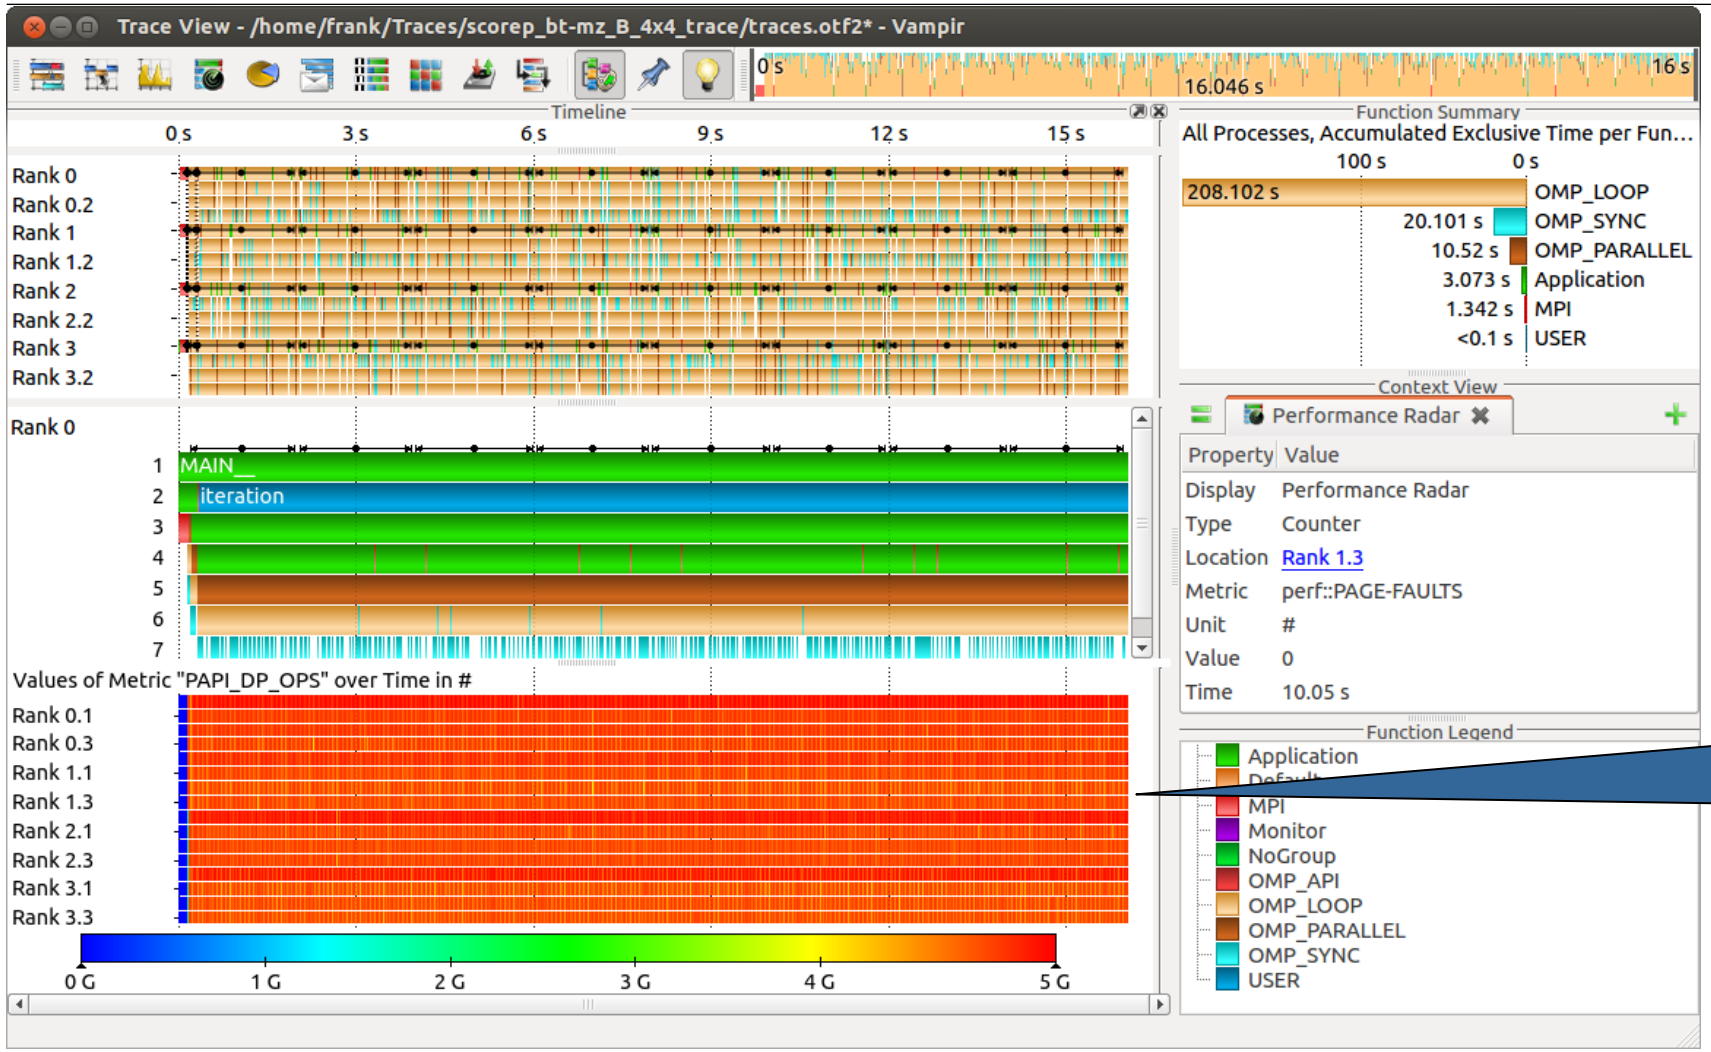

Computation phase results in higher floating point operations.

Visualization of the NPB-MZ-MPI / BT trace
Zoom in: Computation Phase

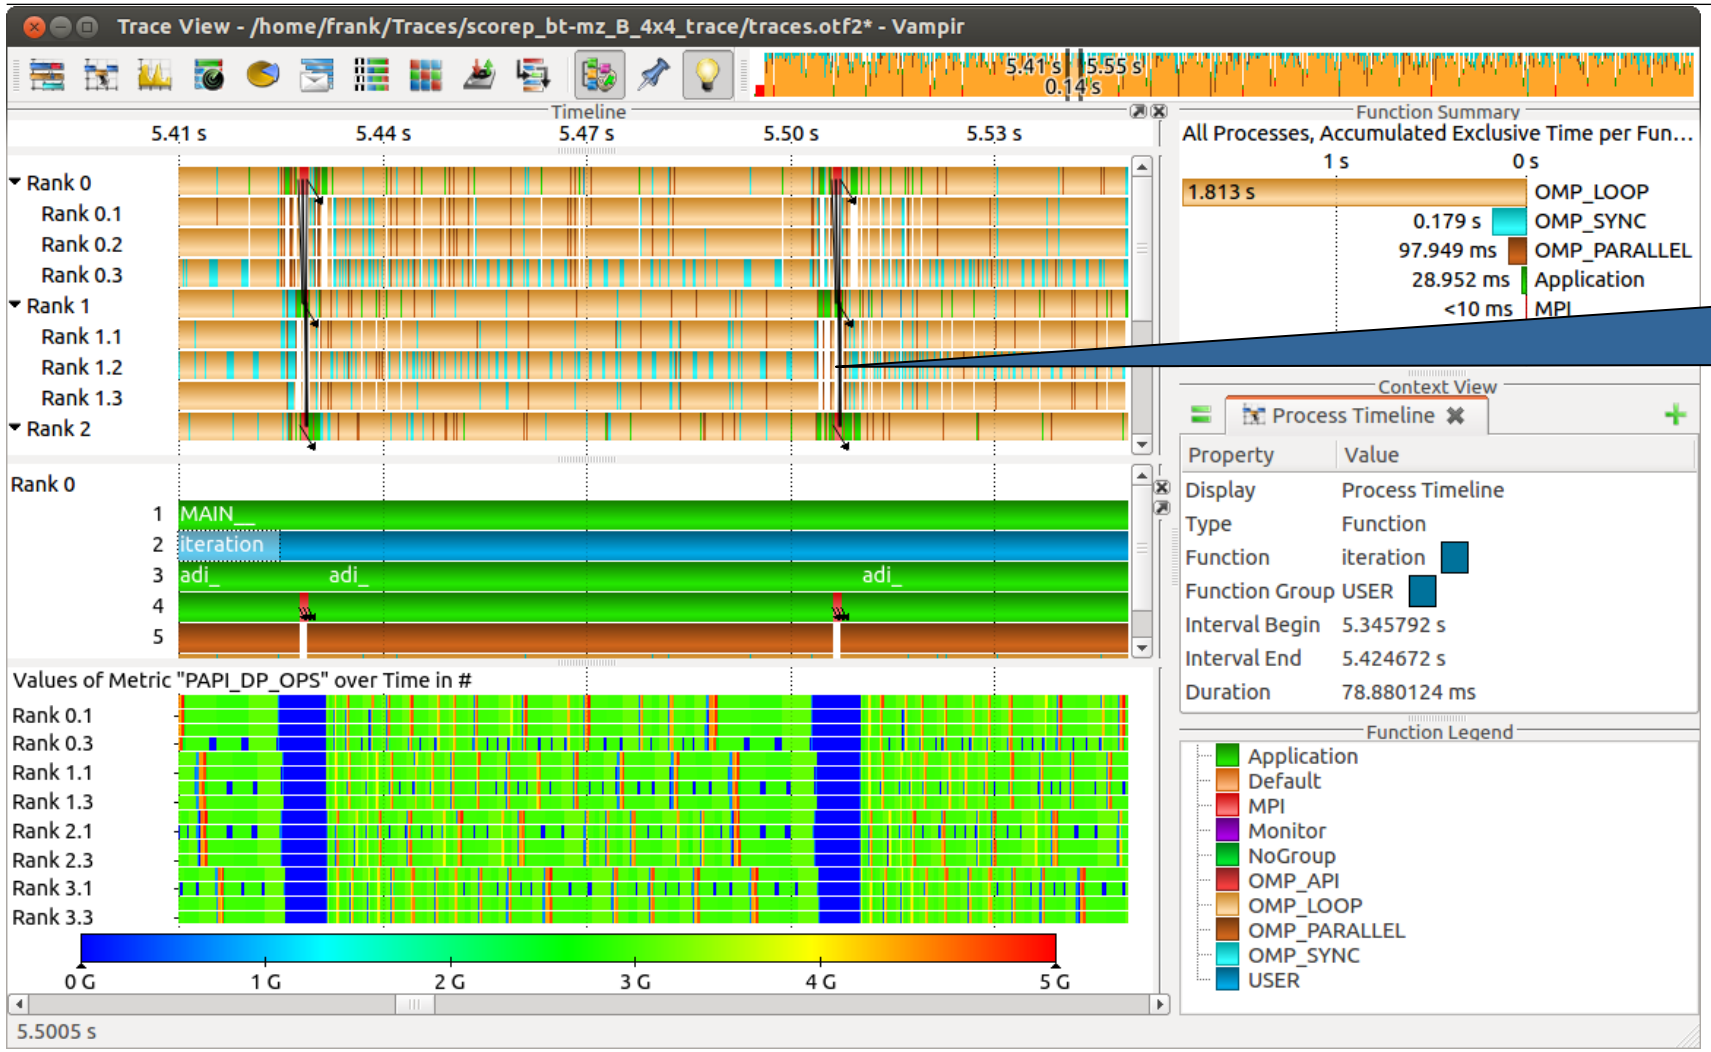

MPI communication results in lower floating point operations.

Visualization of the NPB-MZ-MPI / BT trace
Zoom in: Finalisation Phase

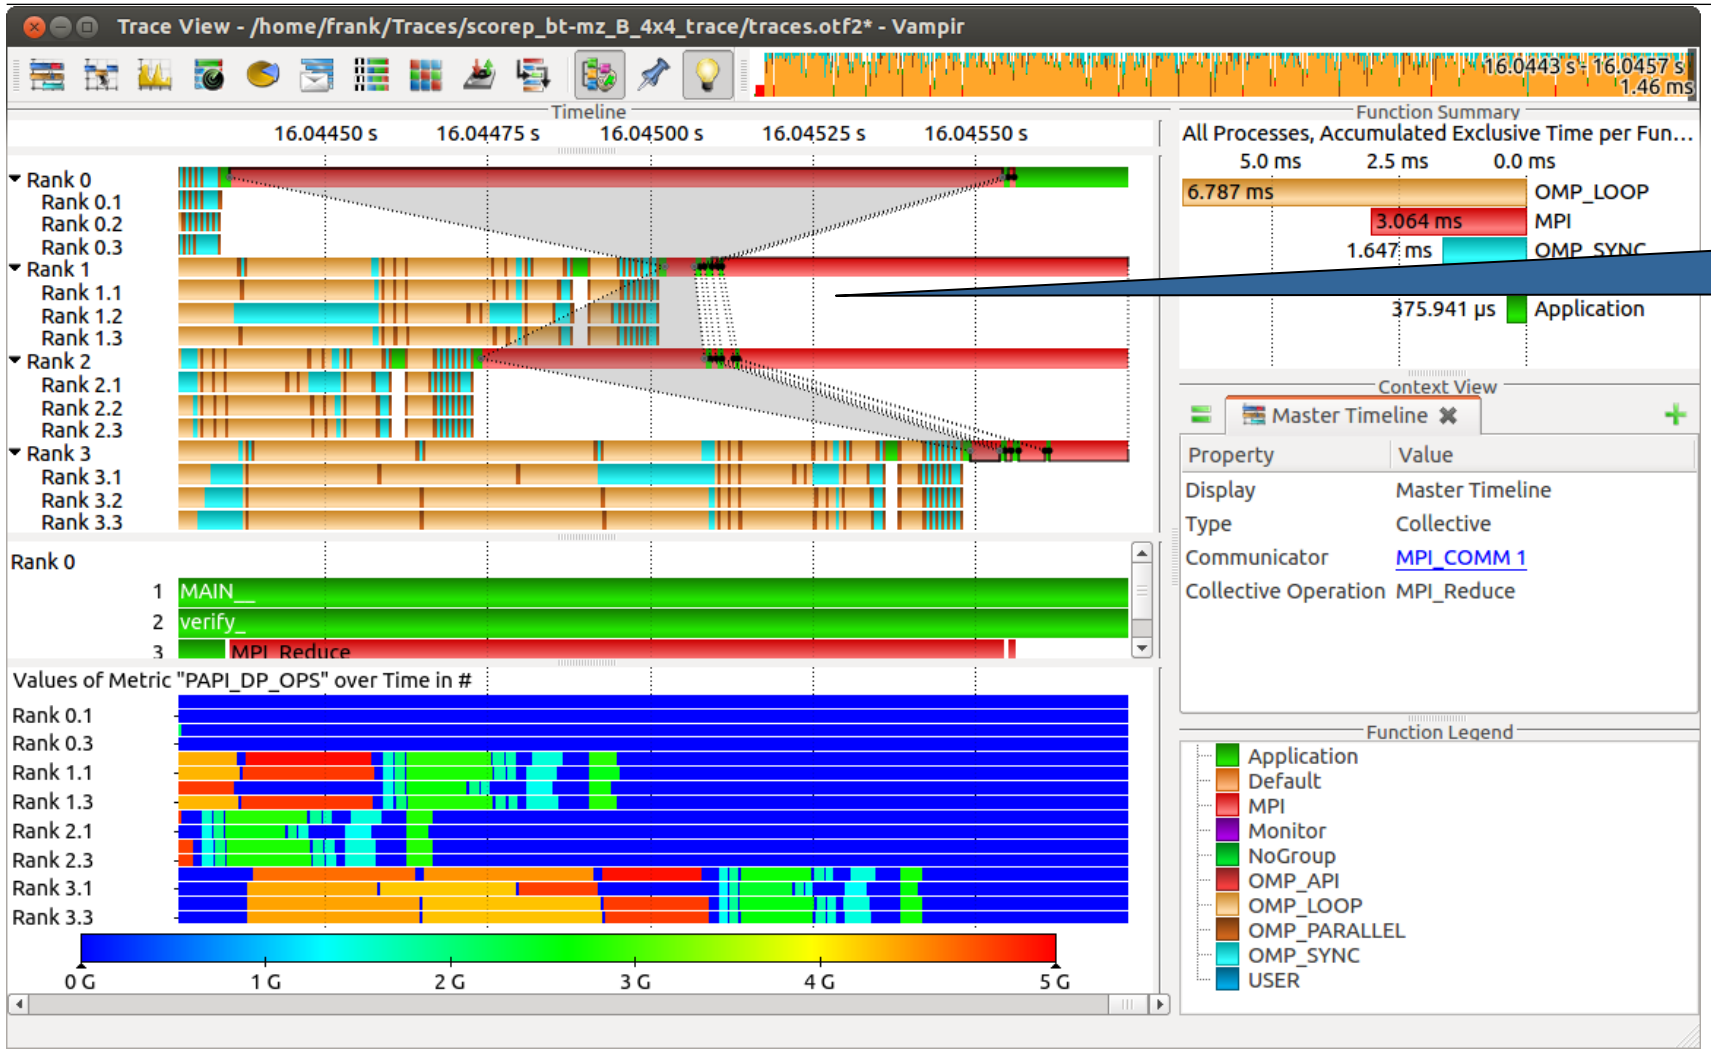

"Early reduce" bottleneck.

Visualization of the NPB-MZ-MPI / BT trace
Process Summary

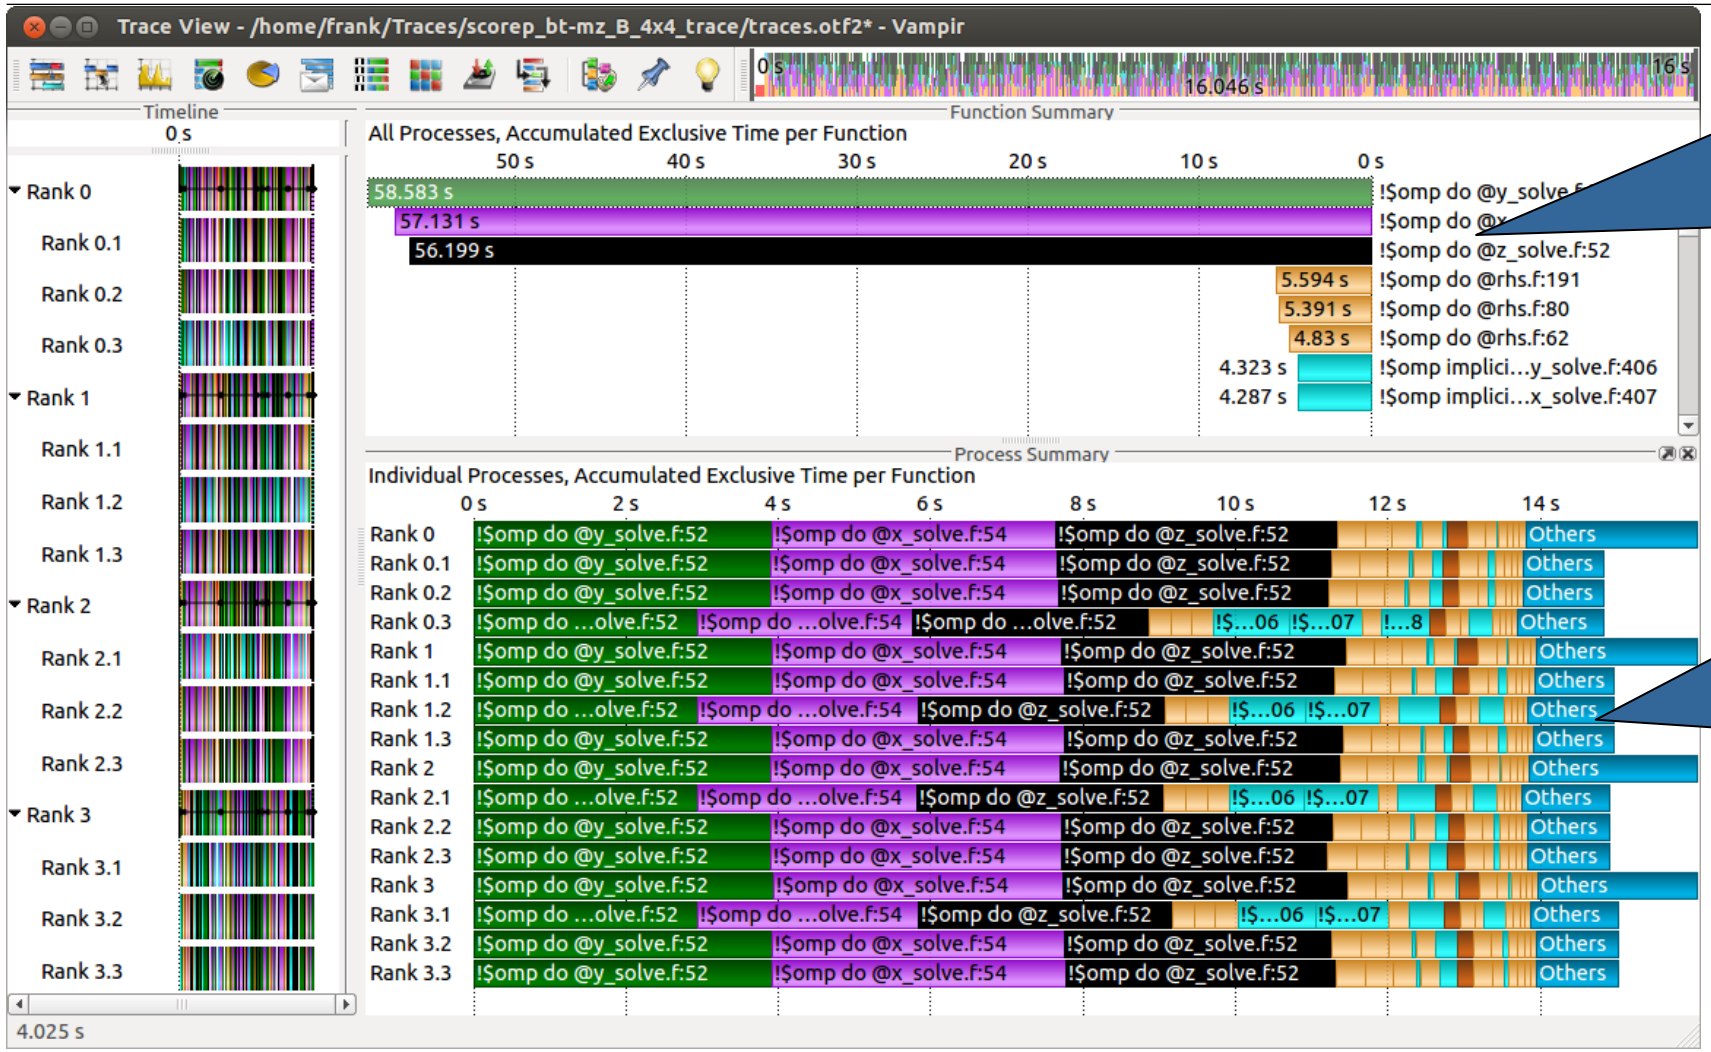

Function Summary: Overview of the accumulated information across all functions and for a collection of processes.

Process Summary: Overview of the accumulated information across all functions and for every process independently.

Visualization of the NPB-MZ-MPI / BT trace
Process Summary

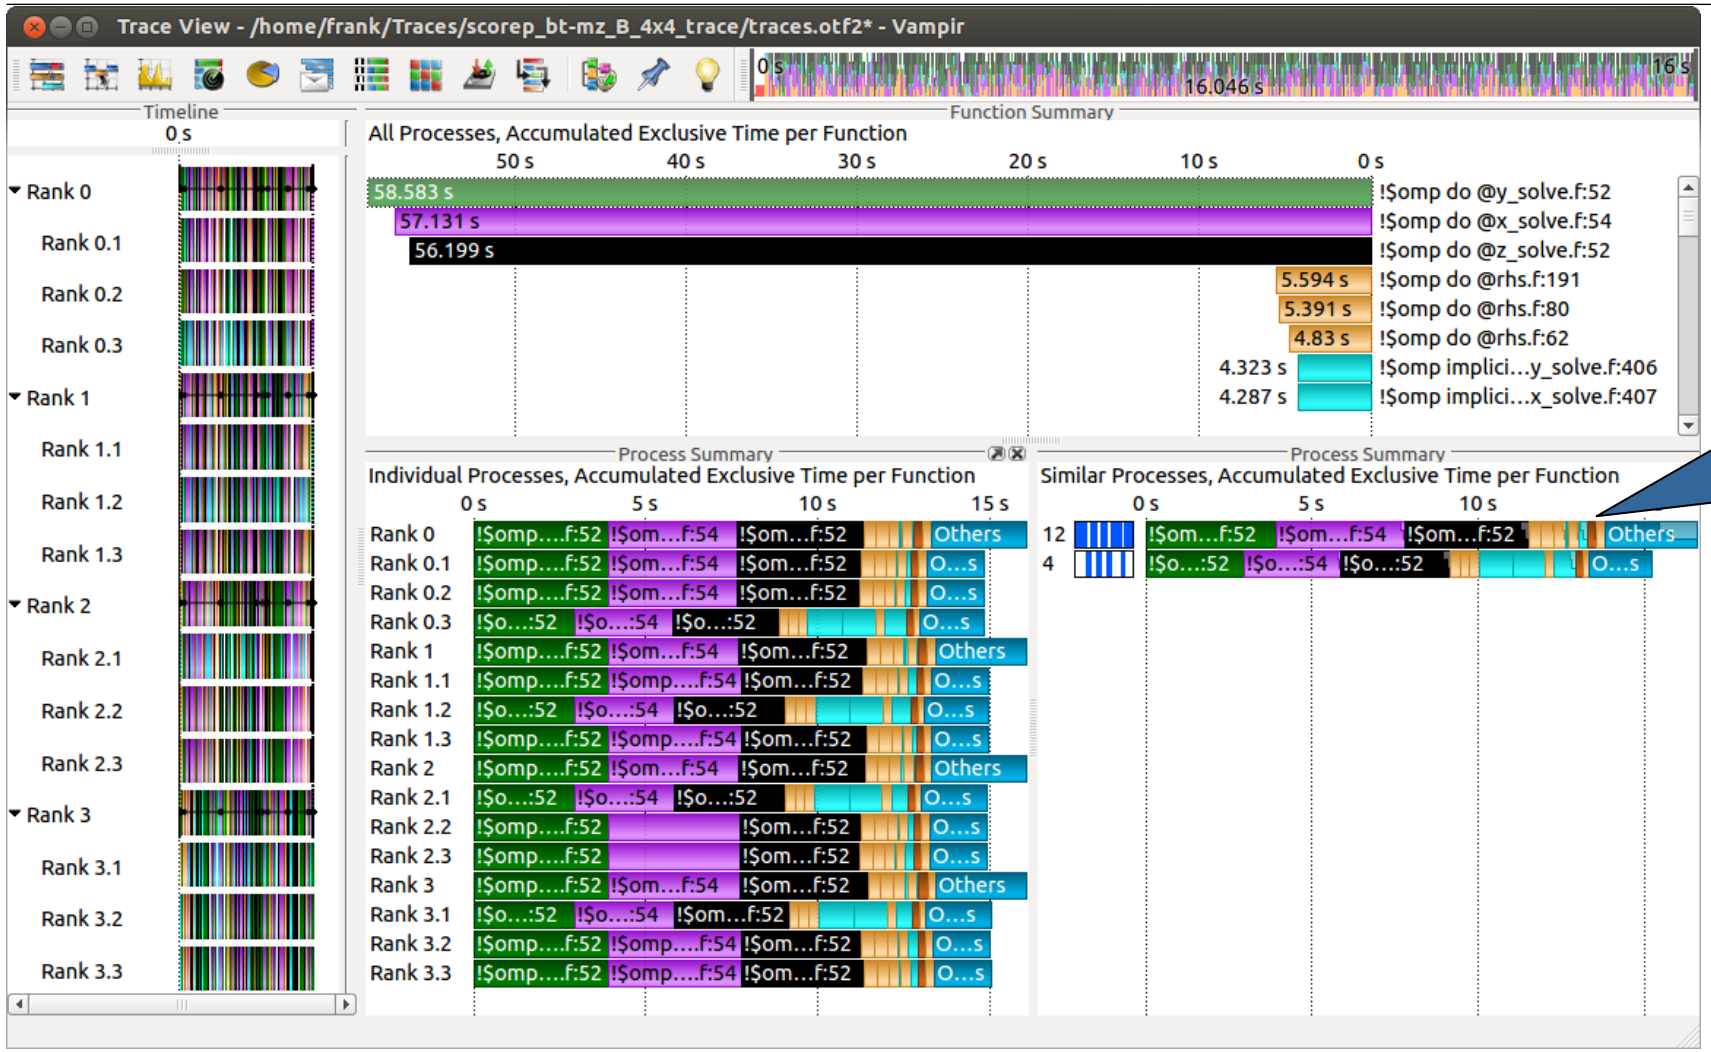

Find groups of similar processes and threads by using summarized function information.

Trace Comparison: Time spent in OpenMP barriers

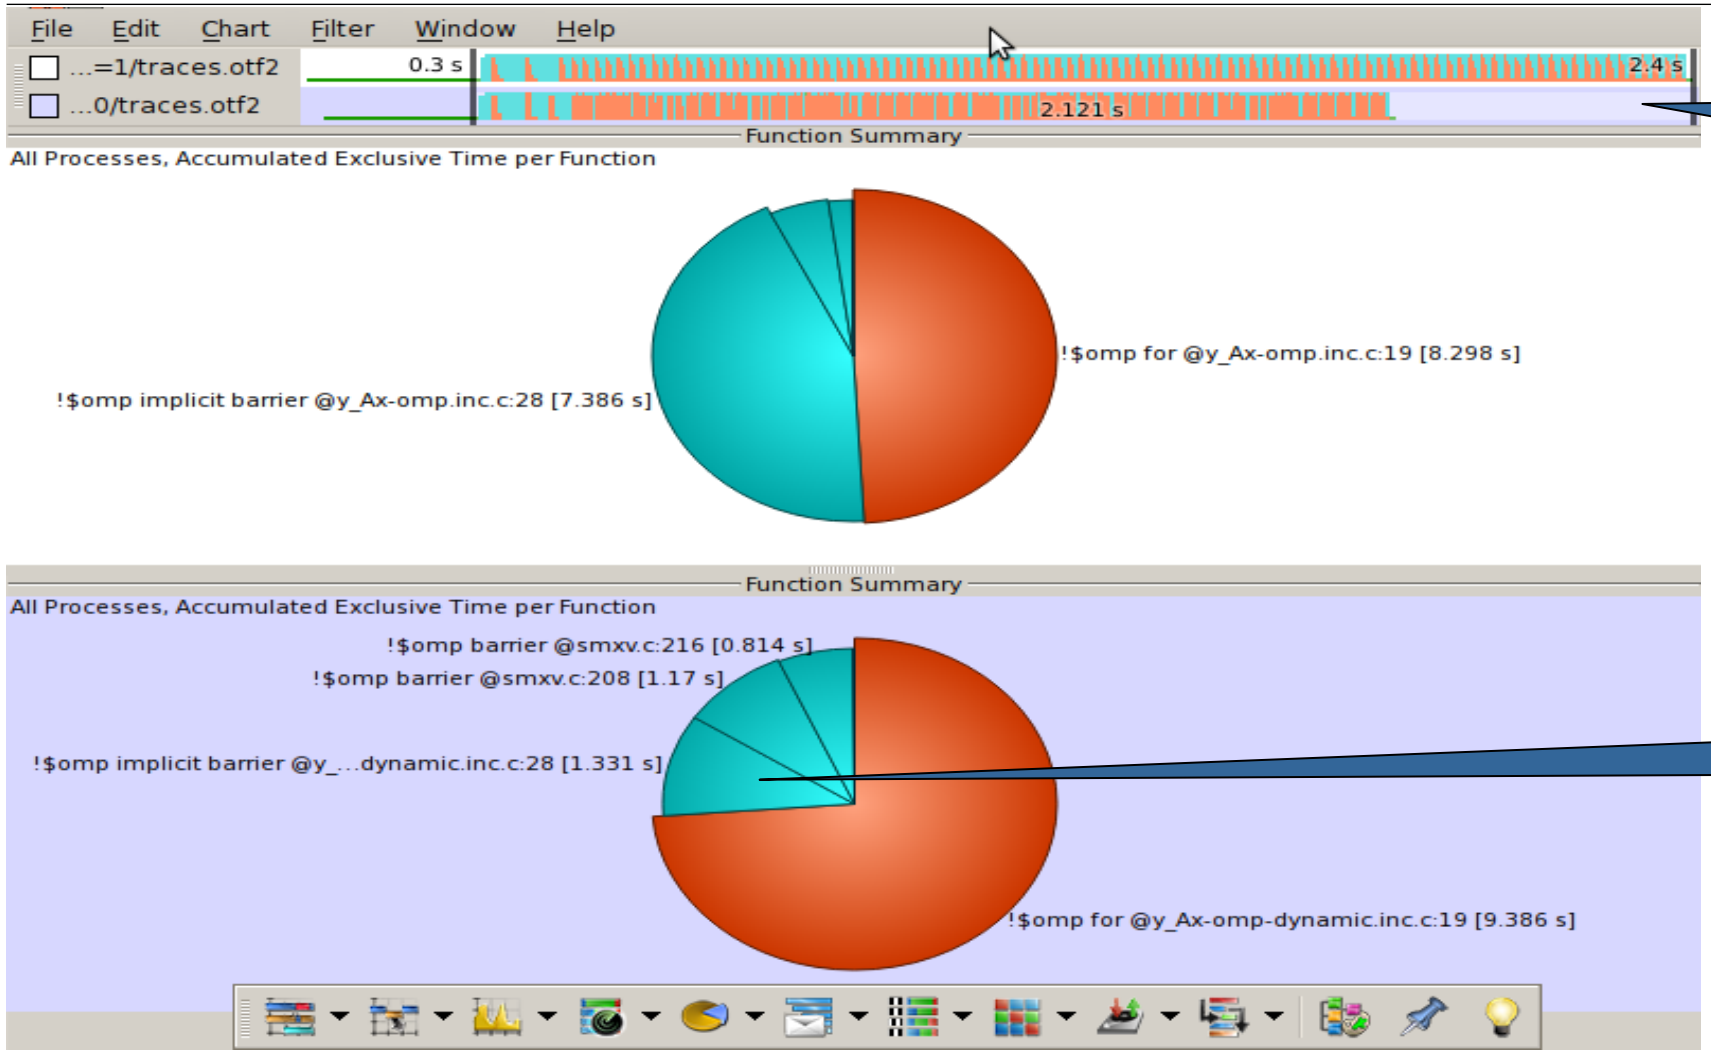

Improved runtime

Less time in OpenMP barrier

Trace Comparison: Computational imbalance

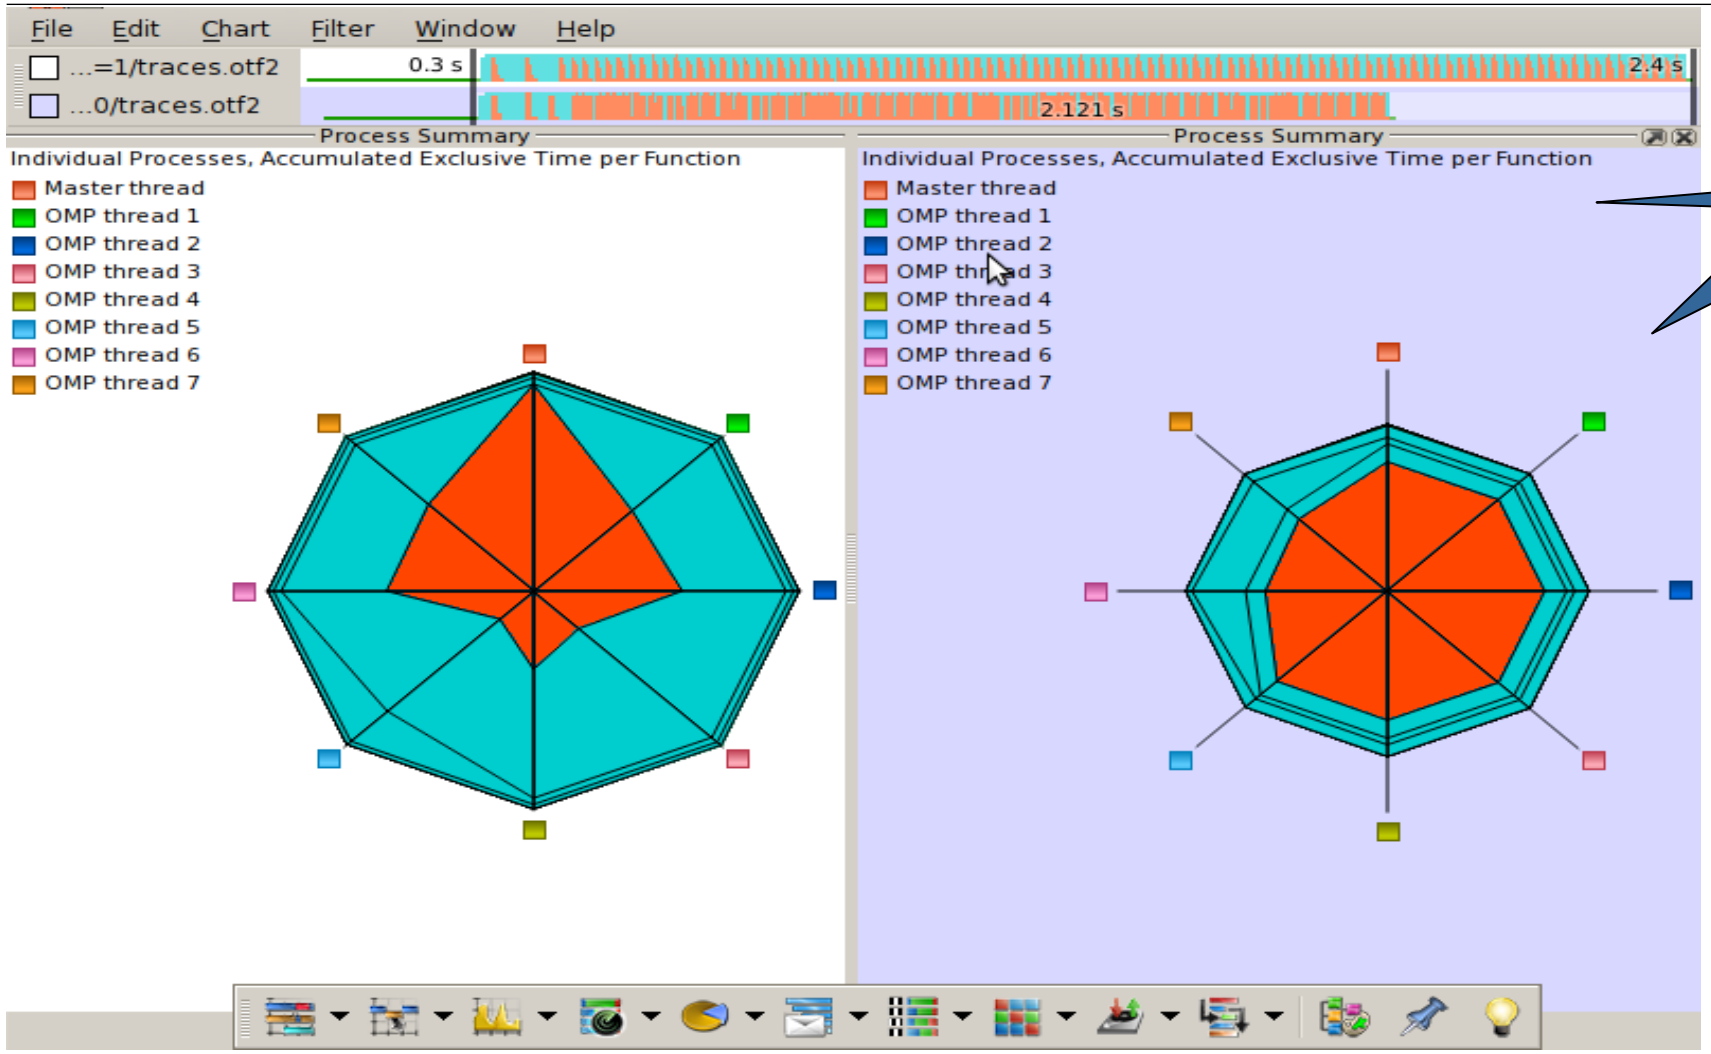

Great imbalance for time spent in computational code

Summary and Conclusion

Summary

- Vampir & VampirServer
  - Interactive trace visualization and analysis
  - Intuitive browsing and zooming
  - Scalable to large trace data sizes (20 TiB)
  - Scalable to high parallelism (200,000 processes)
- Vampir is available for Linux, Windows, and Mac OS X

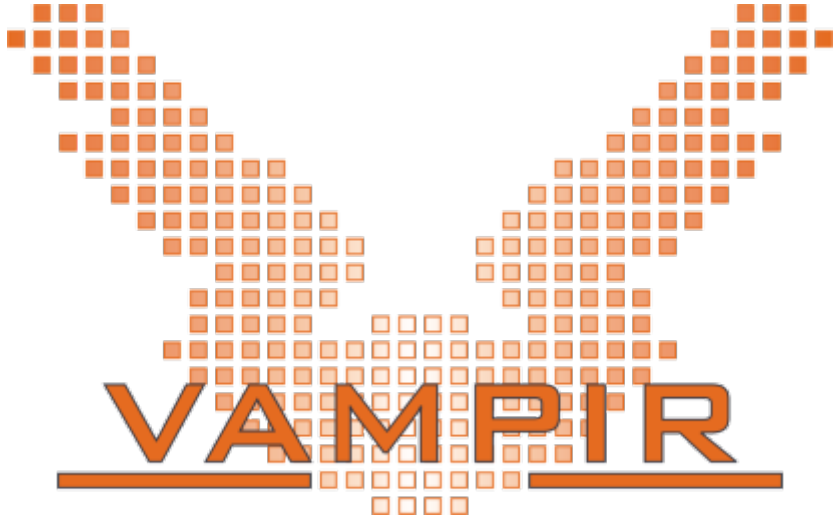

Vampir is available at http://www.vampir.eu

Get support via vampirsupport@zih.tu-dresden.de