Performance Analysis with Vampir

Ronald Geisler, Holger Brunst, Bert Wesarg, Matthias Weber,
Hartmut Mix, Ronny Tschüter, Robert Dietrich, and Andreas Knüpfer

Technische Universität Dresden

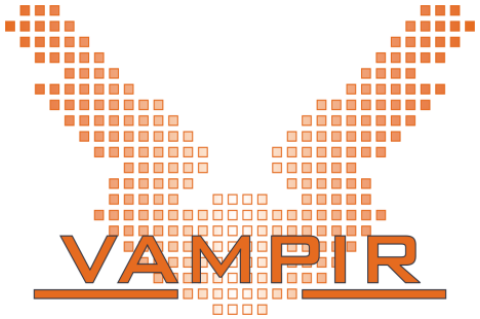

Outline

- **Part I: Welcome to the Vampir Tool Suite**
  - Event Trace Visualization
  - Vampir & VampirServer
  - The Vampir Displays
- **Part II: Vampir Hands-On**
  - Visualizing and analyzing NPB-MZ-MPI / BT

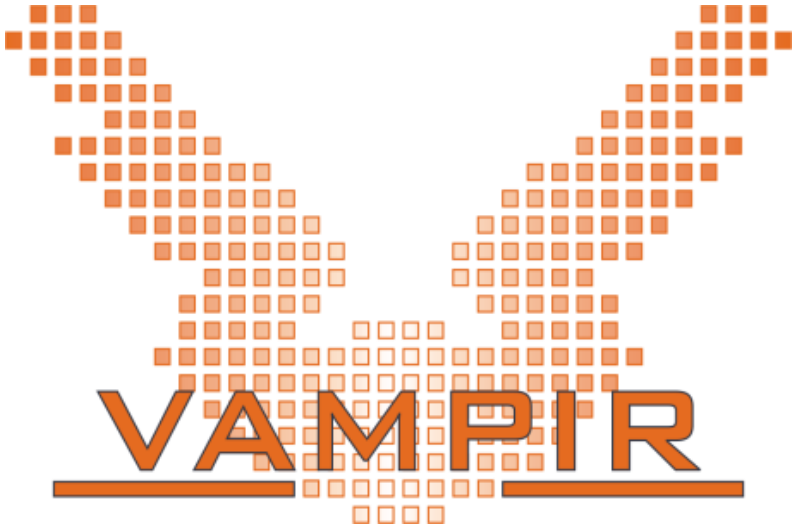

Event Trace Visualization with Vampir

- 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

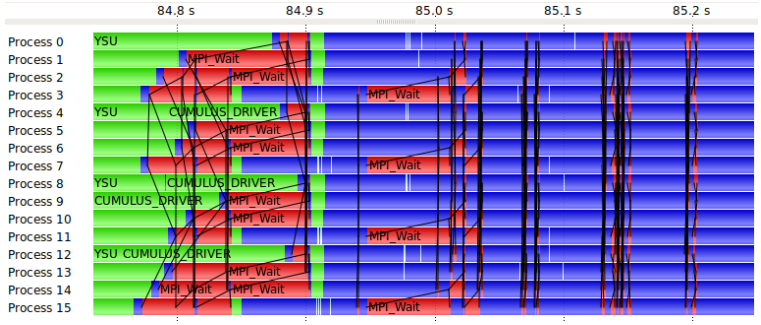

- **Summary charts**
  Provide quantitative results for the currently selected time interval

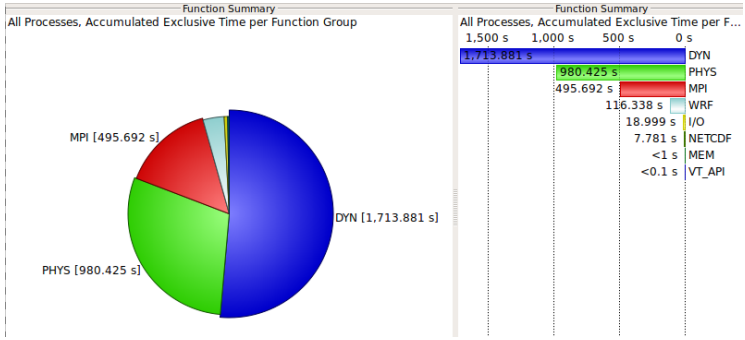

Visualization Modes (1)
On front end or local machine

```
% module load vampir
% vampir
```

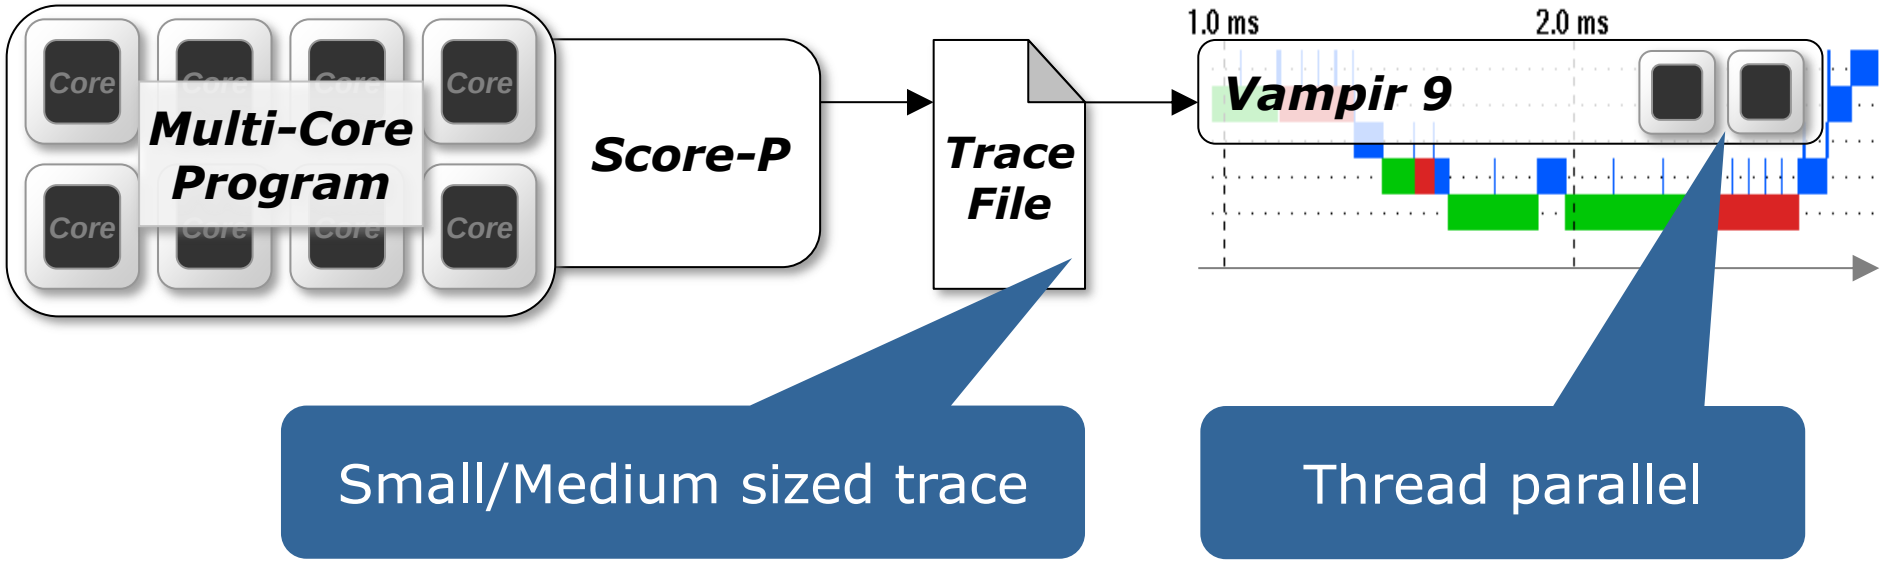

Visualization Modes (2)
On local machine with remote VampirServer

```
% module load vampirserver
% vampirserver start
```

```
% module load vampir
% vampir
```

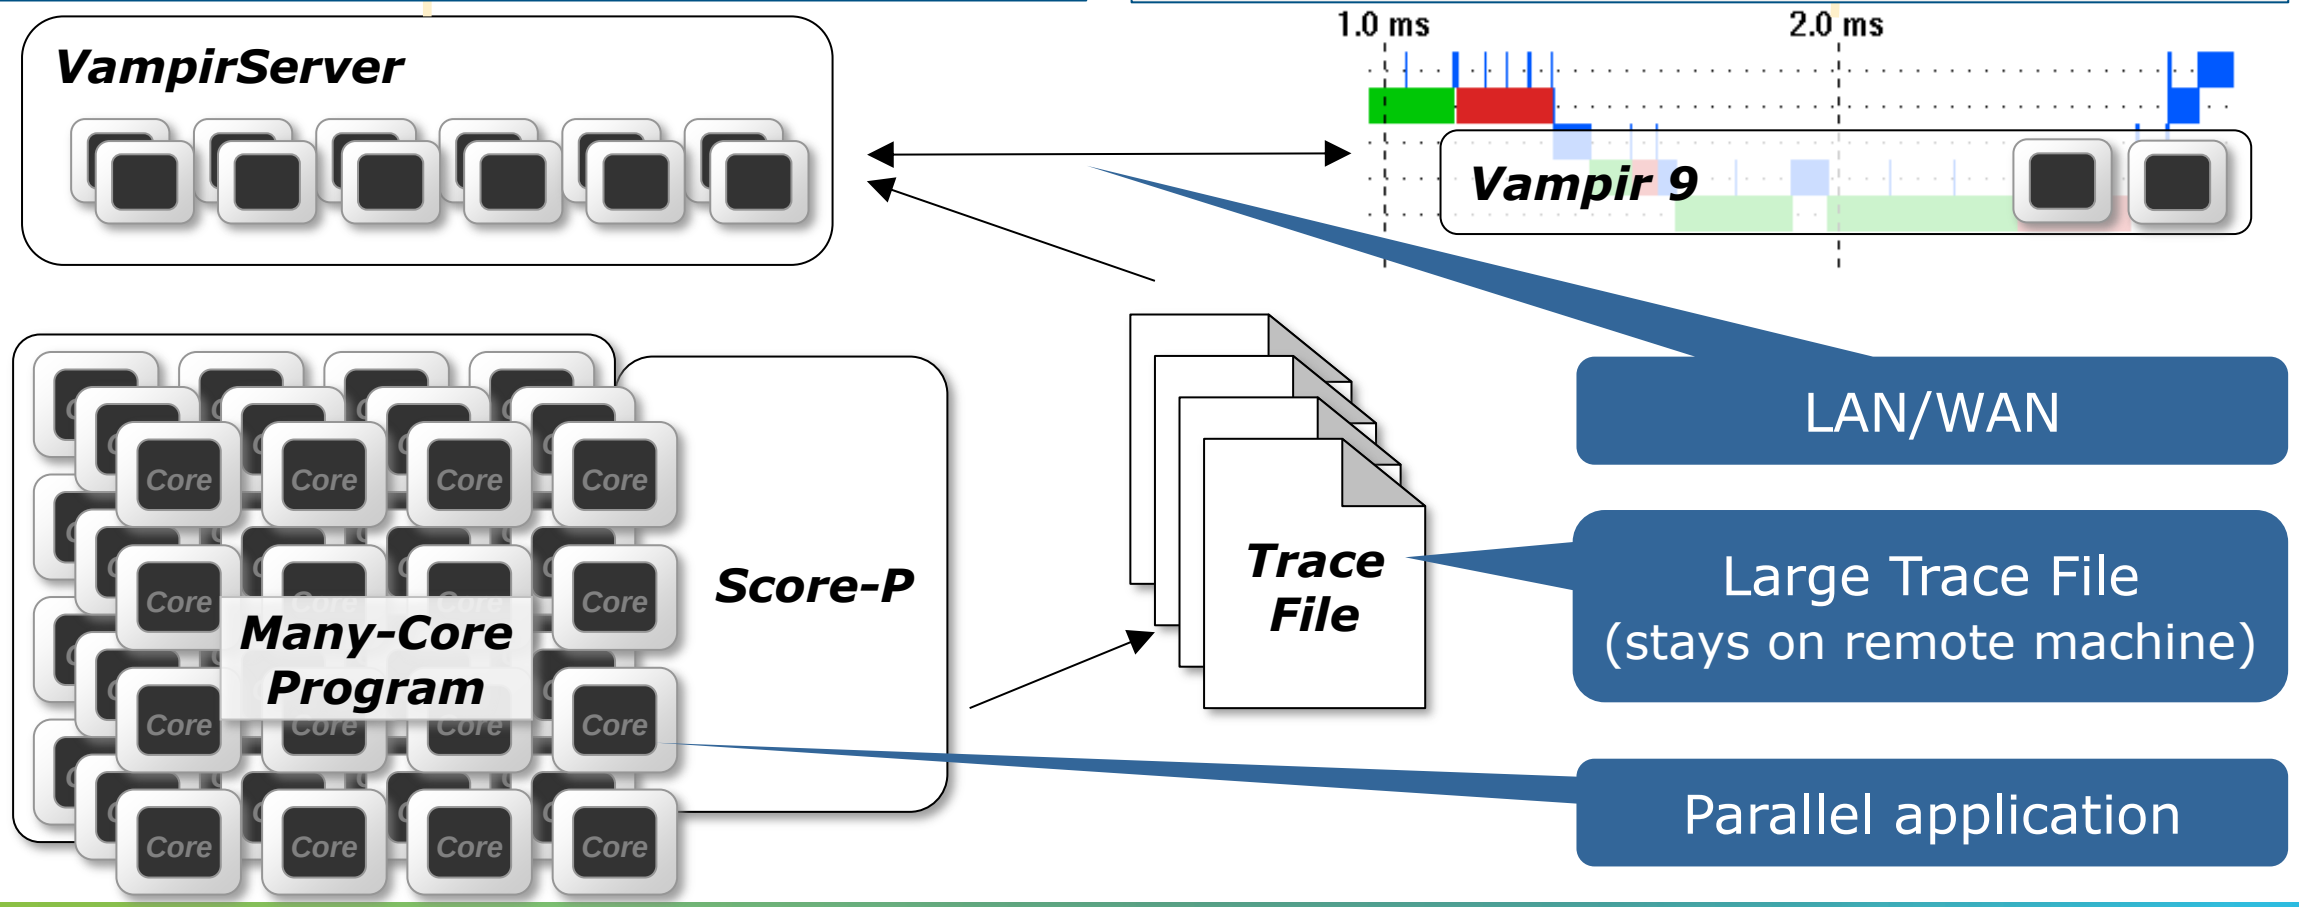

The main displays of Vampir

- 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

Help! Where is my trace file?

- If you followed the Score-P hands-on up to the trace experiment

```
% ls /home/dkrz/$USER/NPB3.3-MZ-MPI/bin.scorep/scorep_bt-mz_B_4x12_trace
scorep.cfg    traces/   traces.def  traces.otf2
```

- If you did not follow to that point, take a prepared trace

```
% ls /work/kg0166/PATworkshop2016/experiments/scorep_bt-mz_B_4x12_trace
profile.cubex  scorep.filt  scorep.score  scout.err  summary.cubex  traces      traces.otf2
scorep.cfg     scorep.log   scout.cubex   scout.log  trace.cubex    traces.def  trace.stat
```

Starting VampirServer on Mistral

- Load VampirServer module

```
% module load vampirserver
```

- Start VampirServer with 8 processes on the compute partition for 60 minutes

```
% vampirserver start -n 8 -- -p compute --account=kg0166 -t 60
Launching VampirServer...
Submitting slurm 60 minutes job (this might take a while)...
```

Install and start Vampir on local computer

- Start a new terminal on your local computer
- Copy the appropriate Vampir package and license file from Mistral

```
% scp -r user@mistral.dkrz.de:/work/kg0166/PATworkshop2016/vampir .
```

- Install Vampir

```
% sh vampir-9.1.0-demo_plus-linux-x86_64-setup.bin [--instdir=]
```

- Start Vampir

```
% /vampir/install/path/bin/vampir
```

Activate Vampir with provided license file

- Select the vampir.licence file

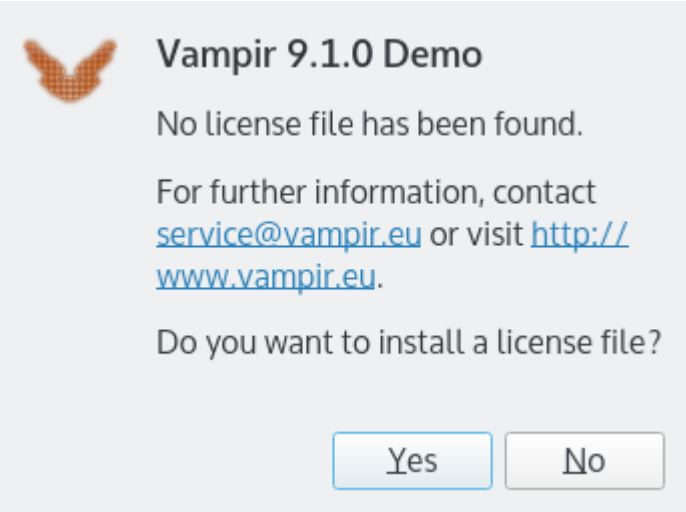


File Help
VAMPIR 9
Recent Files:
Open
Open Other...
Cancel
Use the “Open Other” option

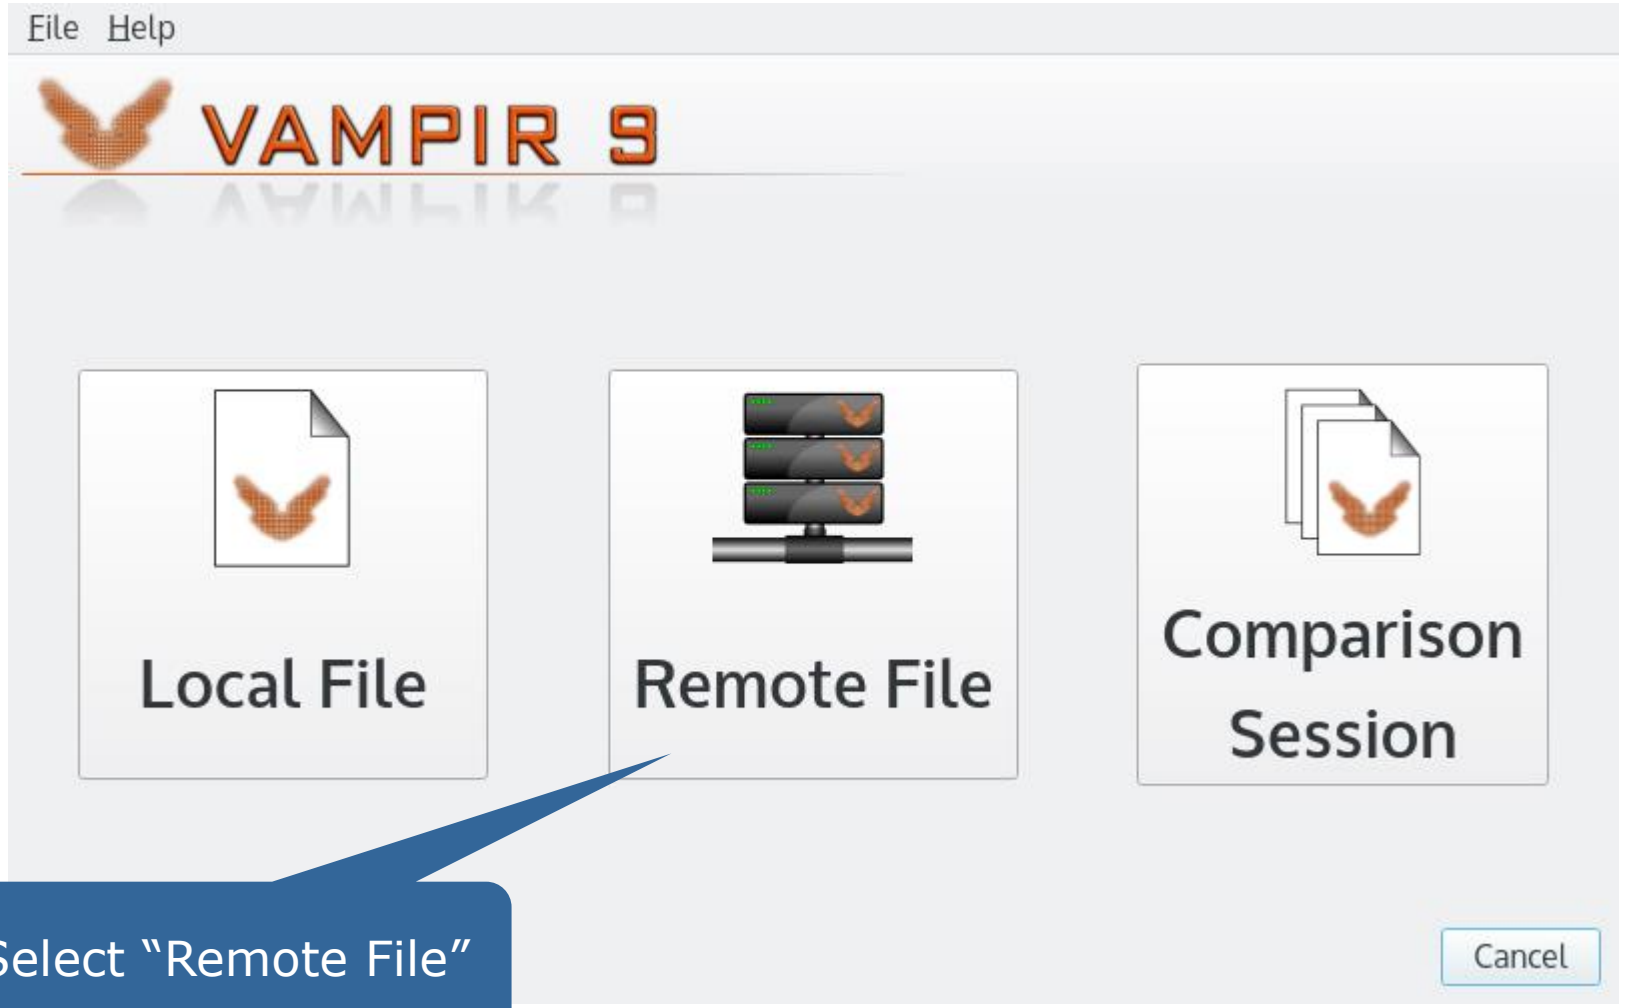
File Help
VAMPIR 9
Local File
Remote File
Comparison Session
Cancel
Select “Remote File”

Back to: Starting VampirServer on Mistral + Port Forwarding

- Start VampirServer with 8 processes on the compute partition for 30 minutes

```
% vampirserver start -n 8 -- -p compute --account=kg0166 -t 30
Launching VampirServer...
Submitting slurm 60 minutes job (this might take a while)...
salloc: Granted job allocation 4302669
… some warnings that can be ignored …
VampirServer 9.1.0  (r10418)
Licensed to DKRZ
Running 8 analysis processes... (abort with vampirserver stop 3695)
VampirServer <3695> listens on: m11285:30026
```

On Mistral
Start VampirServer

- Port forwarding to Mistral to access the VampirServer on the compute node

```
% ssh -N -L 30000:m11285:30026 user@mistral.dkrz.de
```

On local machine
Port forwarding

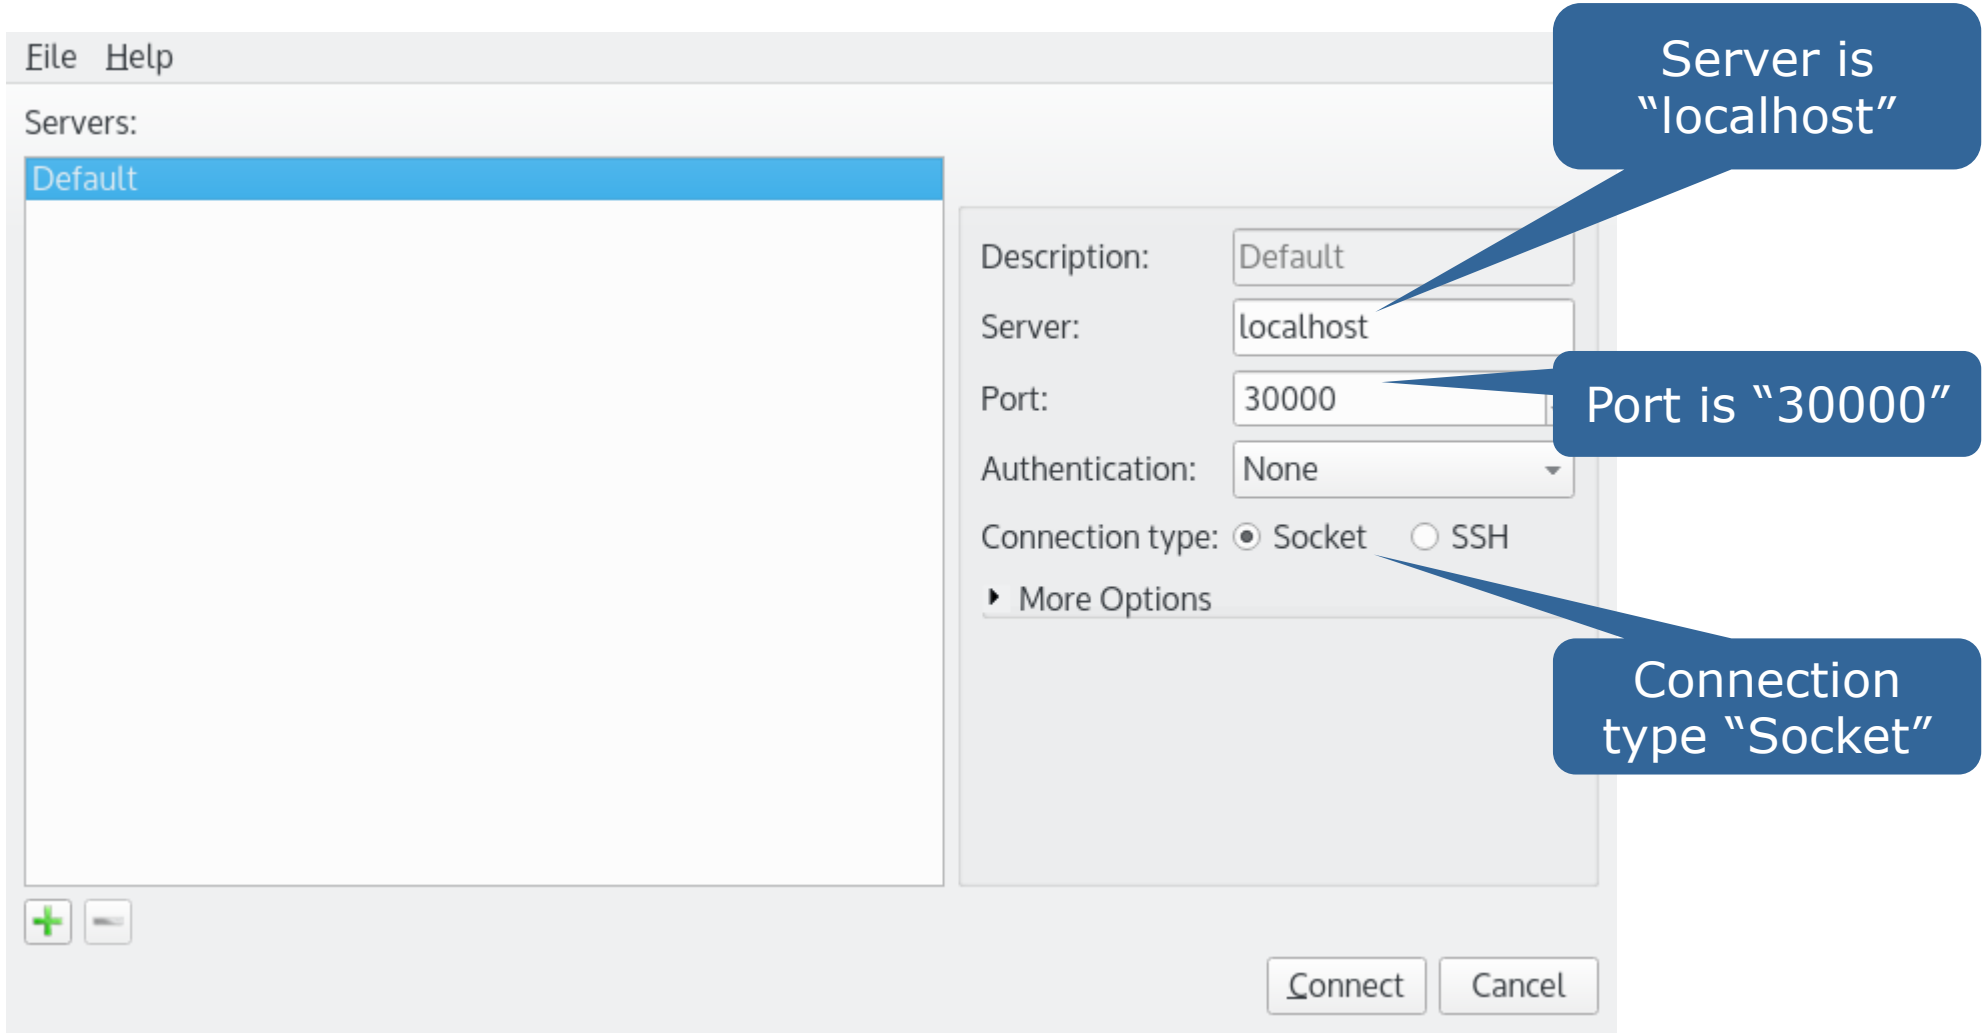
File Help
Servers:
Default
Description:
Default
Server:
localhost
Port:
30000
Authentication:
None
Connection type:
Socket
SSH
More Options
Connect
Cancel
Server is "localhost"
Port is "30000"
Connection type "Socket"

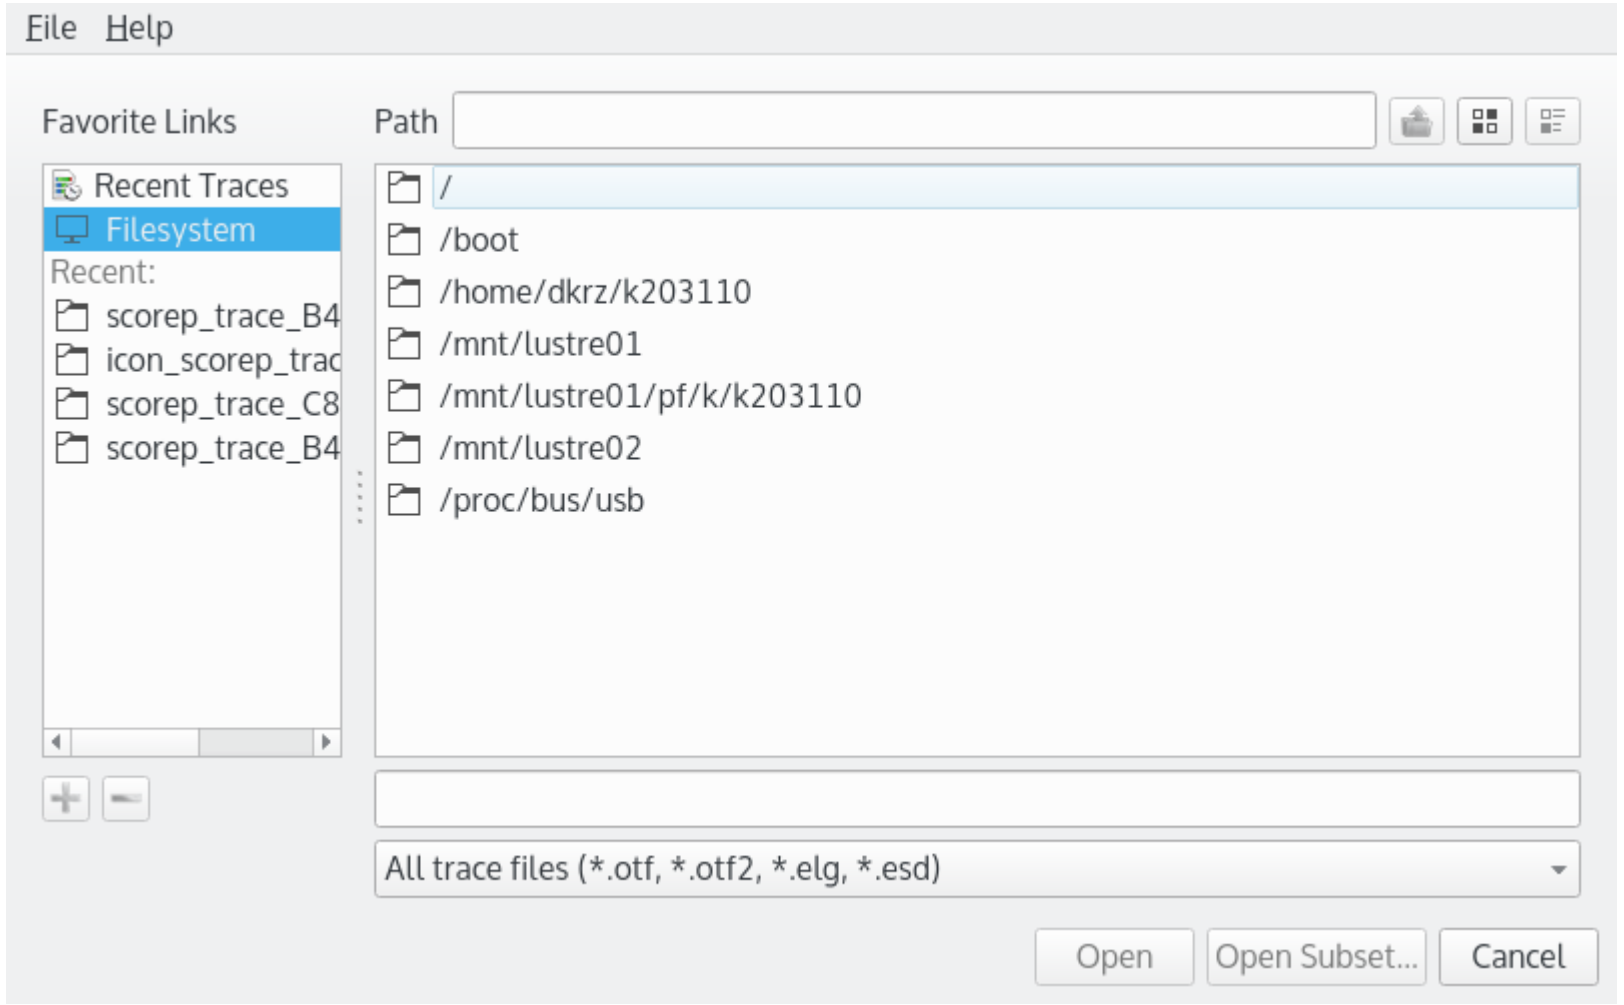
File Help
Favorite Links
Path
Recent Traces
Filesystem
Recent:
scorep_trace_B4
icon_scorep_trac
scorep_trace_C8
scorep_trace_B4
/
/boot
/home/dkrz/k203110
/mnt/lustre01
/mnt/lustre01/pf/k/k203110
/mnt/lustre02
/proc/bus/usb
All trace files (*.otf, *.otf2, *.elg, *.esd)
Open
Open Subset...
Cancel

Summary: Setup local Vampir and start VampirServer on Mistral

```
% module load vampirserver
% vampirserver start -n 8 -- -p compute --account=kg0166 -t 30
  …
  VampirServer <3695> listens on: m11285:30026
```

On Mistral
Start VampirServer

```
% scp -r mistral.dkrz.de:/work/kg0166/PATworkshop2016/vampir .
% sh vampir-9.1.0-demo_plus-linux-x86_64-setup.bin [--instdir=]
% /vampir/install/path/bin/vampir
```

On local machine
1) Copy, install and start Vampir

- Activate Vampir with provided license
- Choose “Open Other” → “Remote File”
- Set server to “localhost”, port to “30000”, connection to “Socket”

```
% ssh -N -L 30000: m11285:30026 user@mistral.dkrz.de
```

2) Port forwarding

- Connect to VampirServer
- Select *traces.oft2* file in the file dialog

Alternative: Start Vampir on Mistral

Don‘t do this this, unless everything else fails! X11 forwarding for all participants via WLAN won‘t work.

- Start a VampirServer

```
% module load vampirserver
% vampirserver start -- -p compute --account=kg0166 -t 60
  …
  VampirServer <3688> listens on m11386:30026
```

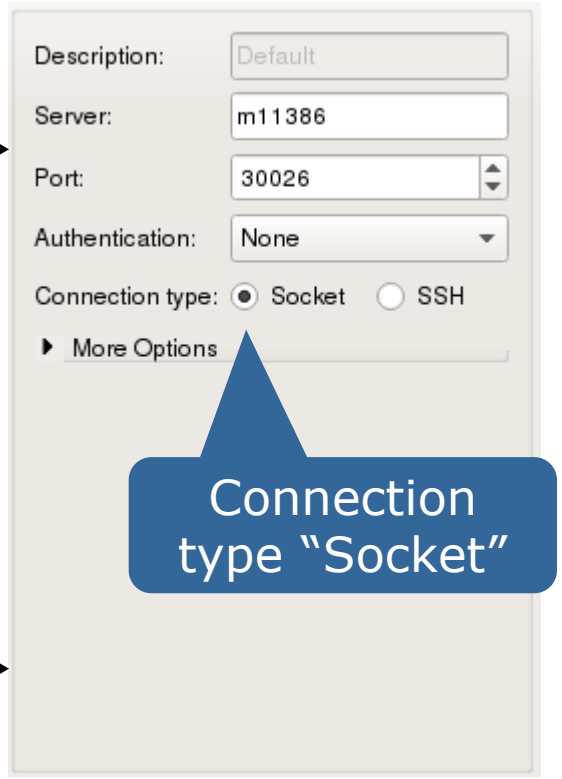

- Use X11 forwarding and start Vampir on Mistral (login node)

```
% ssh -XC user@mistral.dkrz.de
% module load vampir
% vampir
```

This enables X11 forwarding AND compression

- File → Open Remote
- Enter server and port as reported by VampirServer
- Choose trace file in the file dialog

Visualization of the NPB-MZ-MPI / BT trace

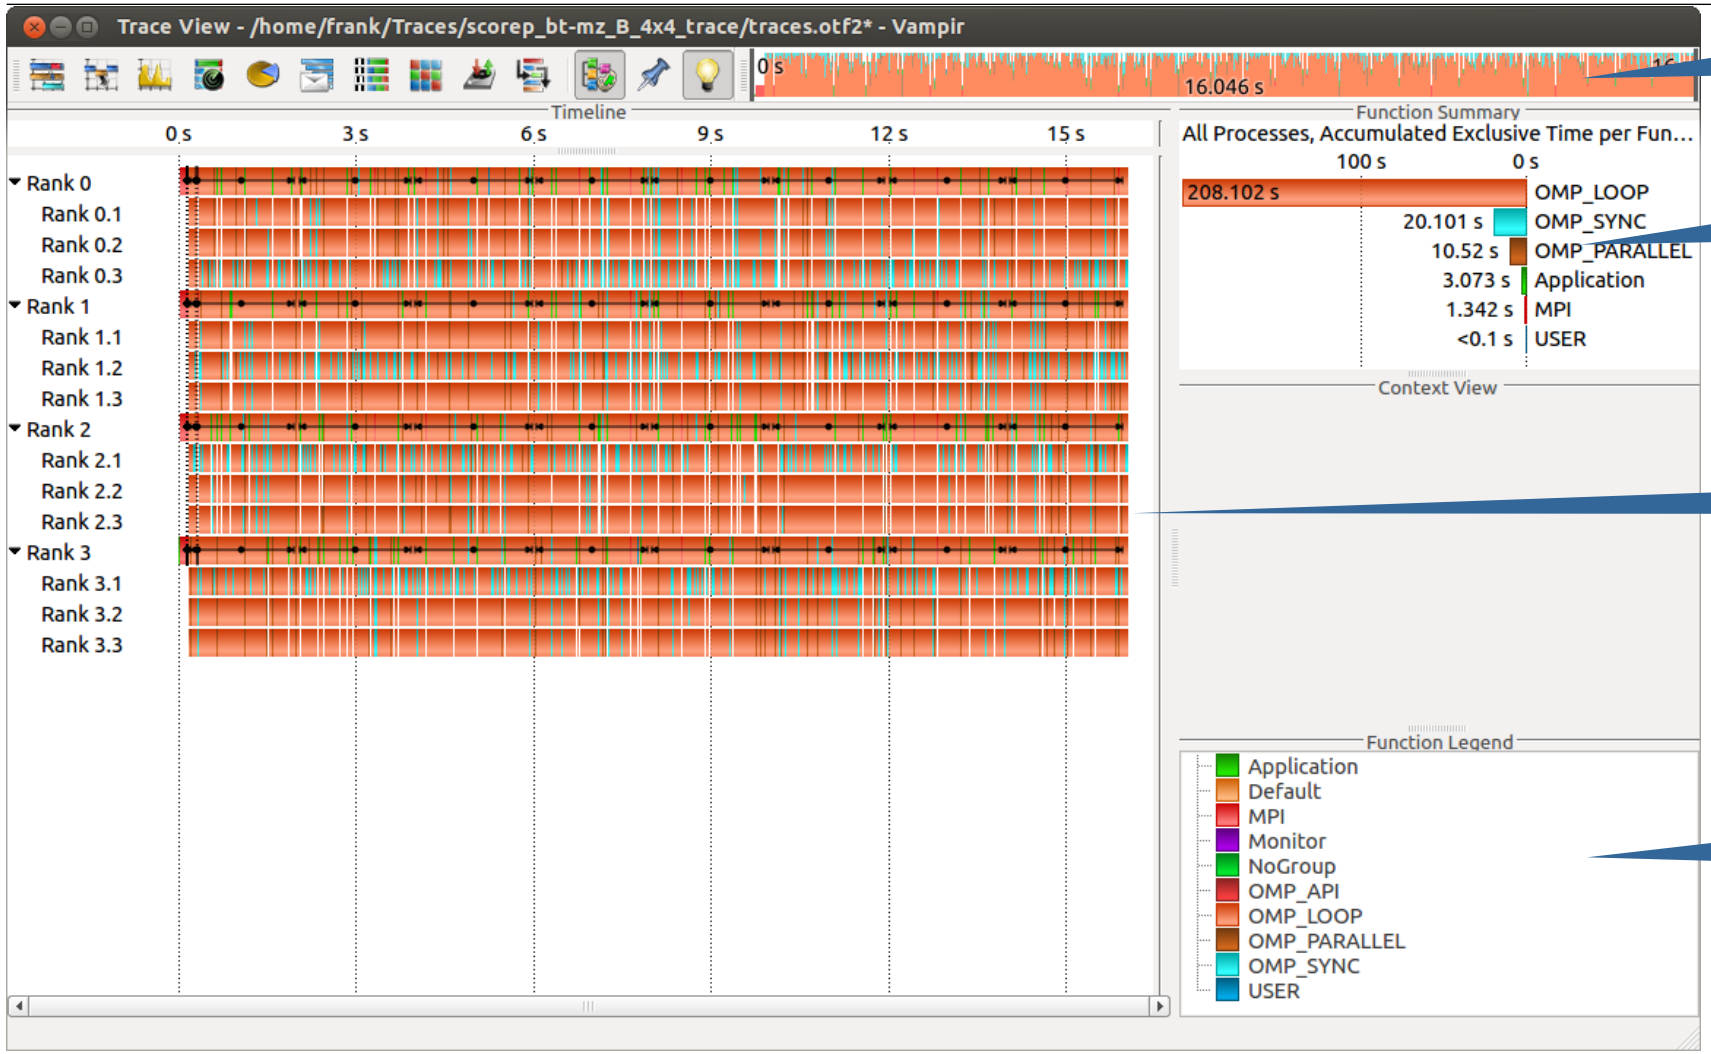

Navigation Toolbar

Function Summary

Master Timeline

Function Legend

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

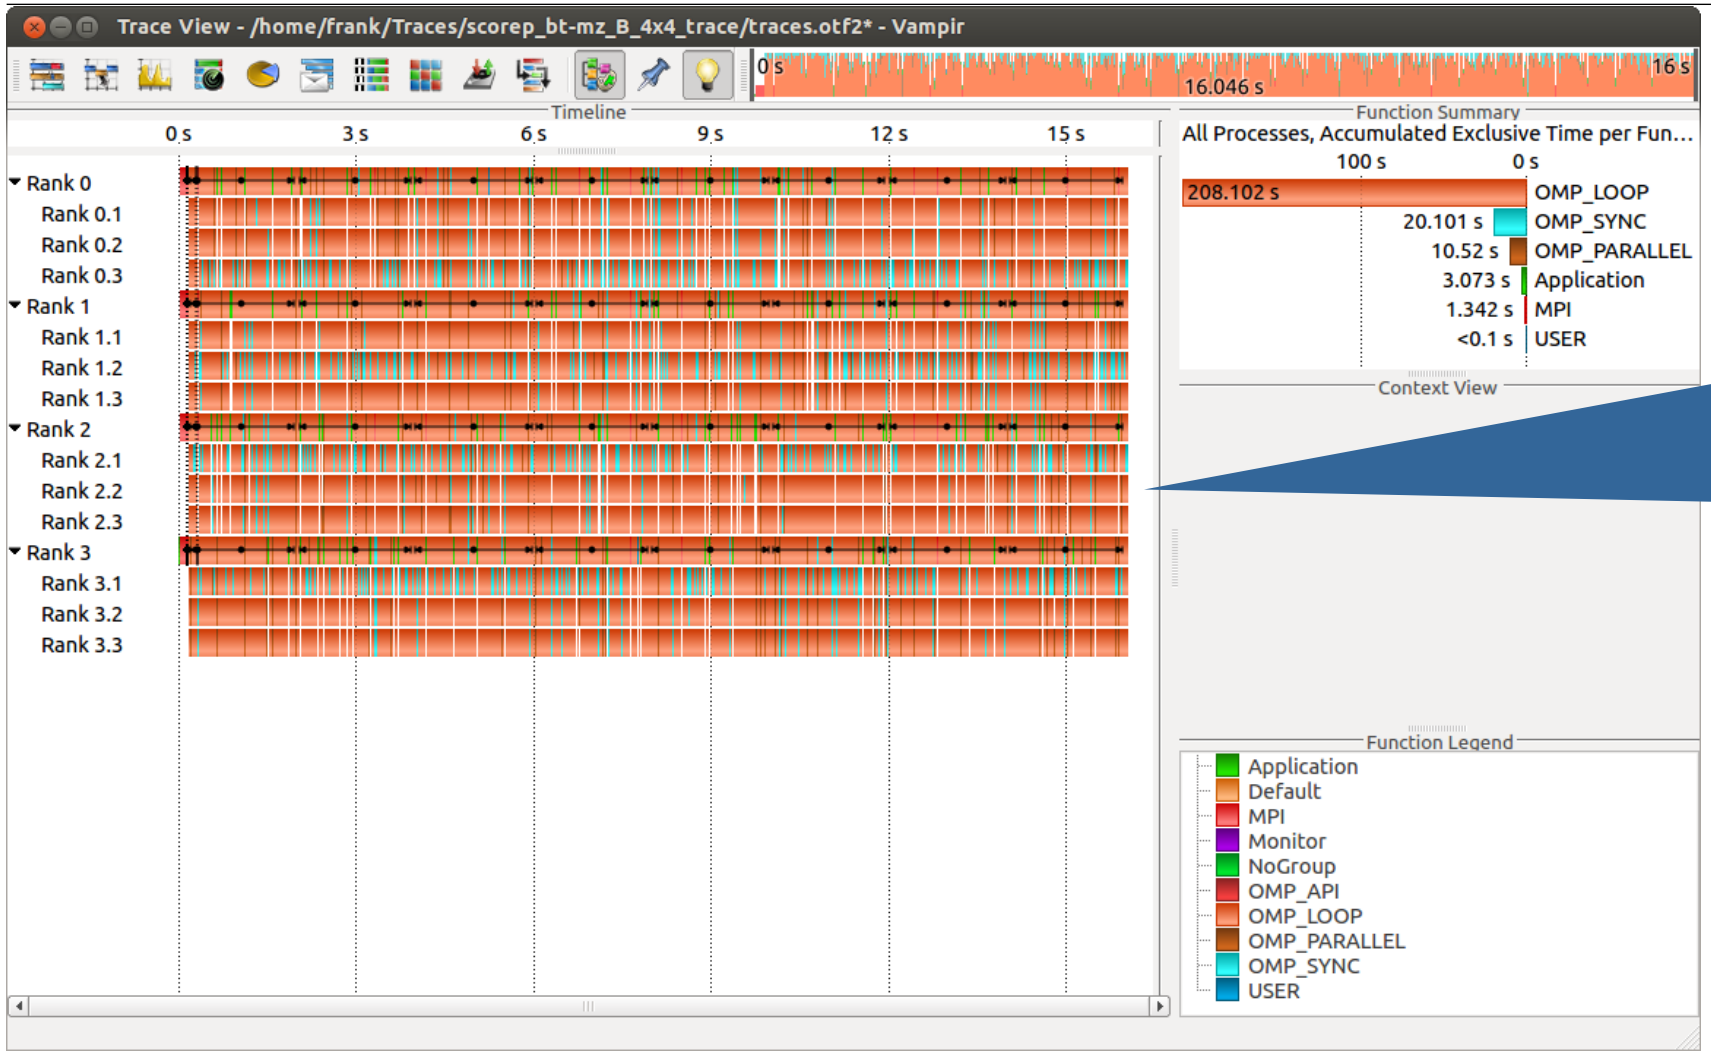

Detailed information about functions, communication and synchronization events for collection of processes.

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

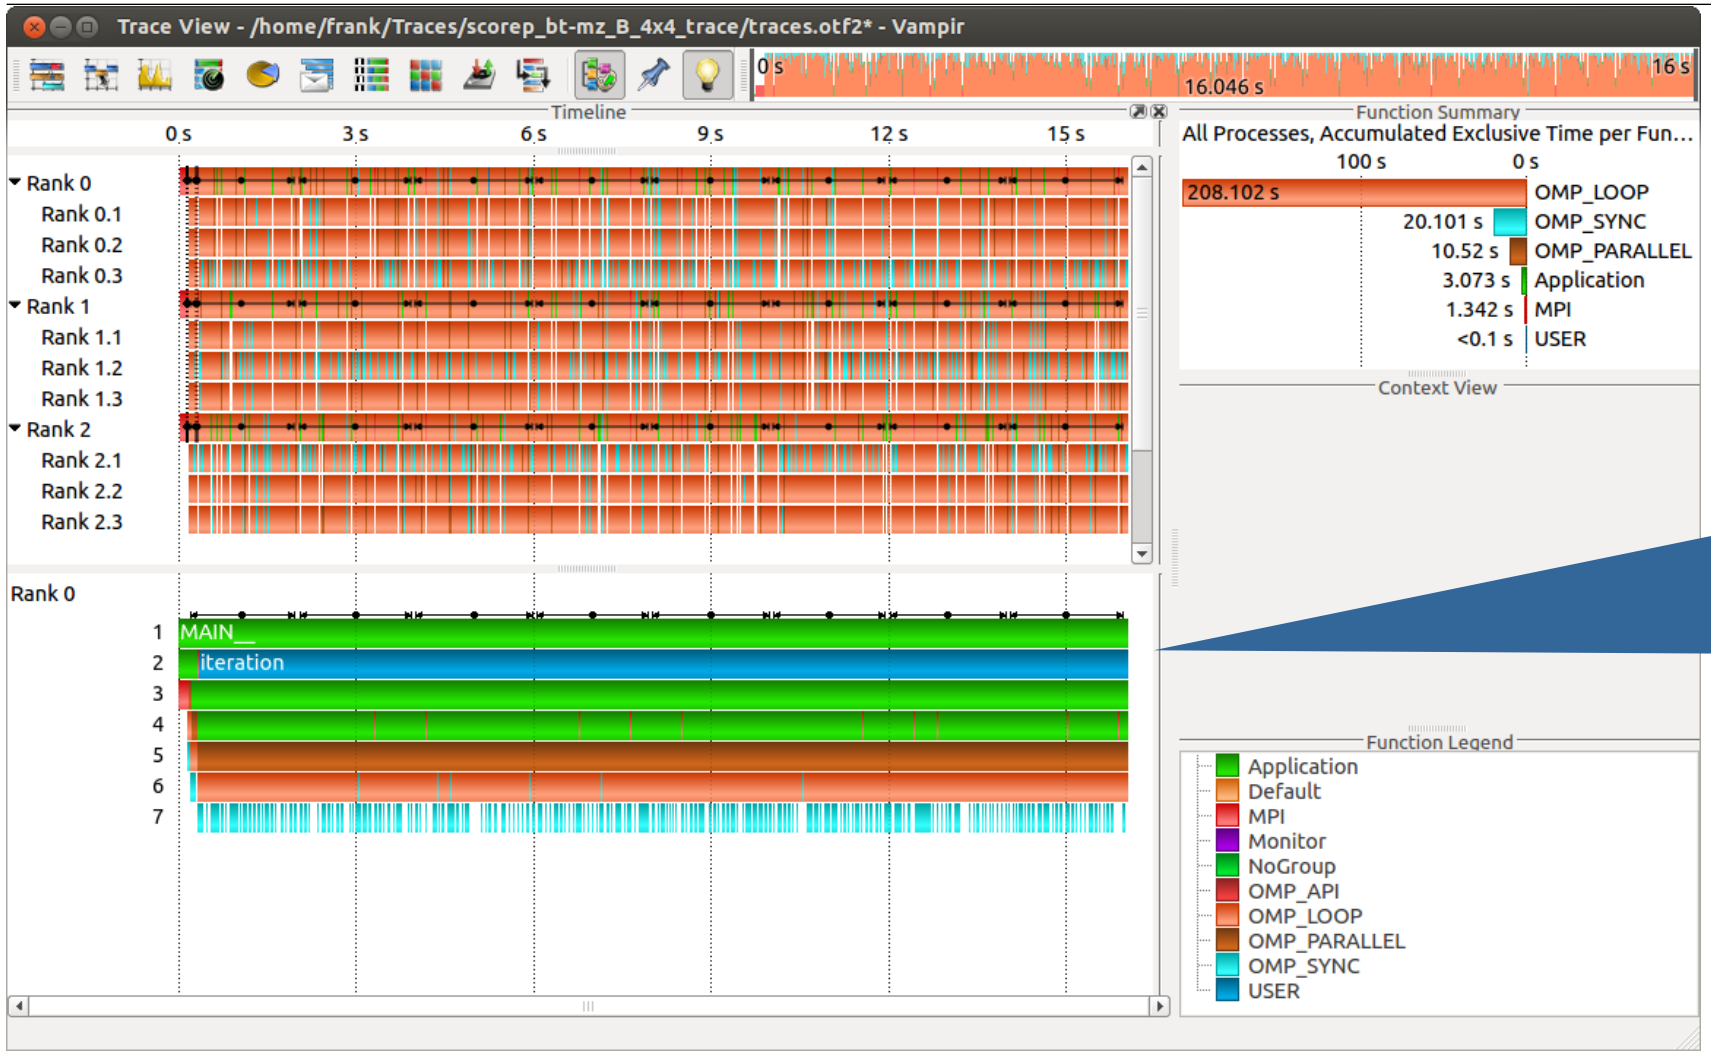

Detailed information about different levels of function calls in a stacked bar chart for an individual process.

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

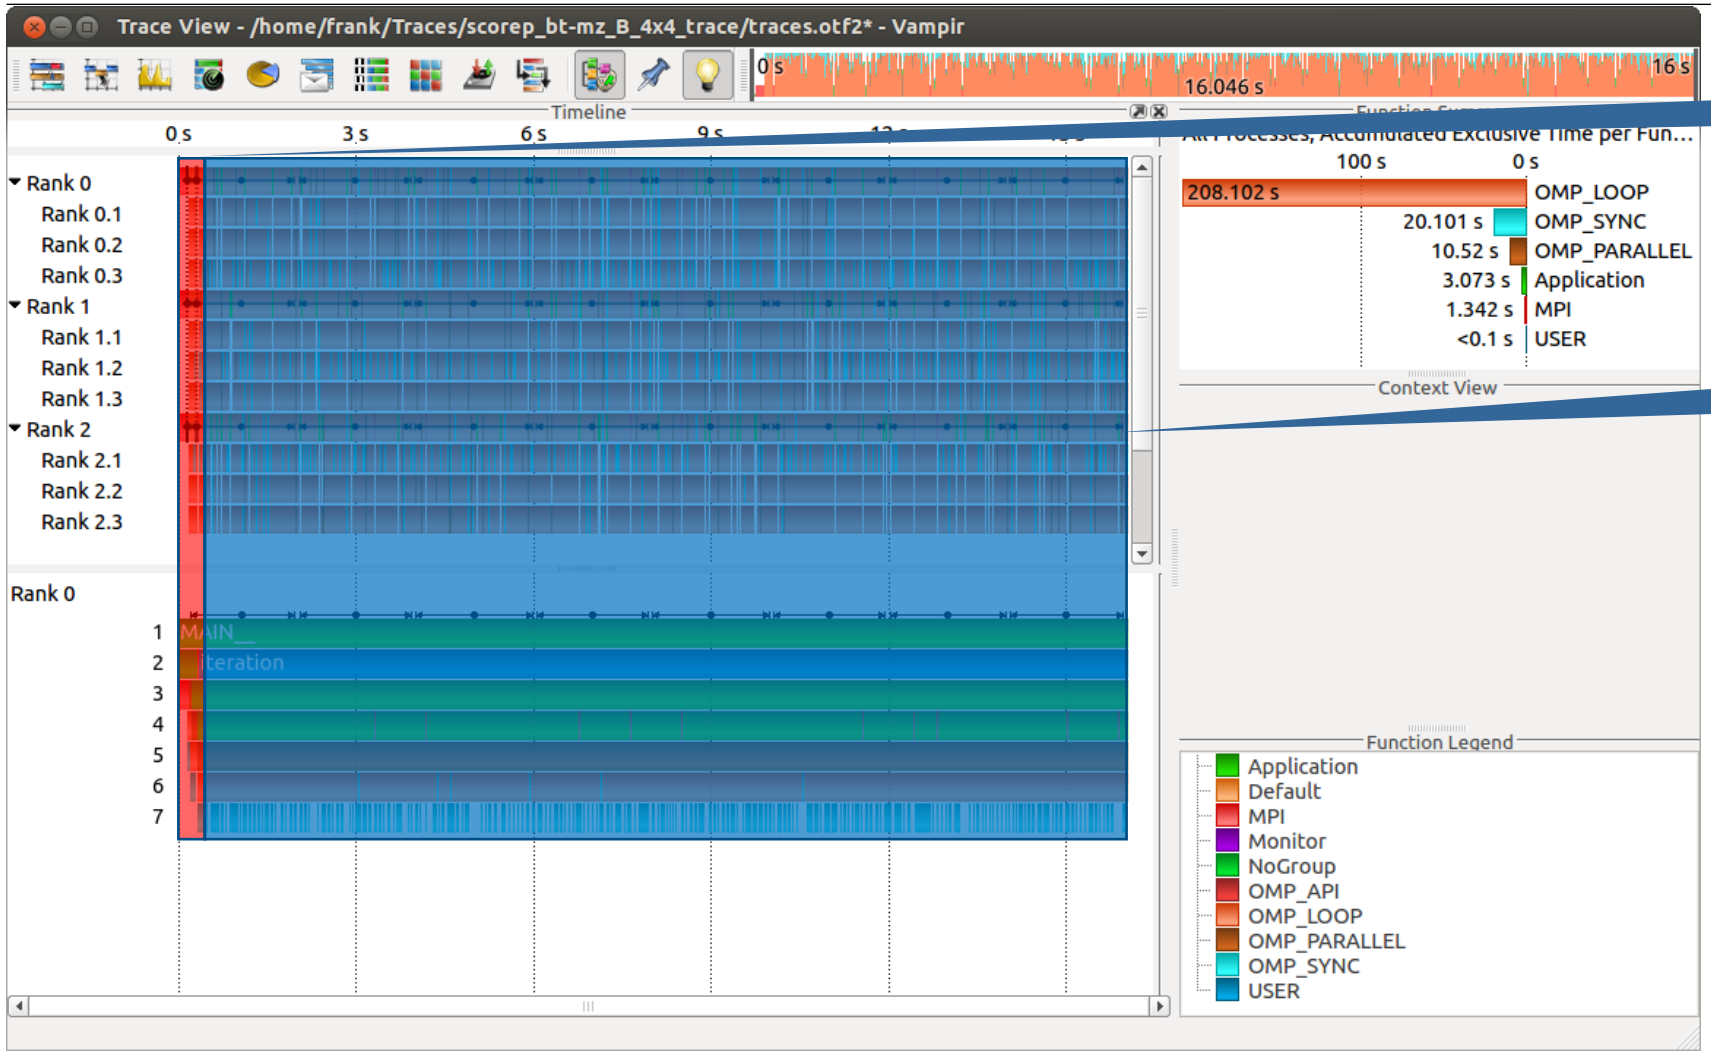

Initialisation Phase

Computation Phase

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

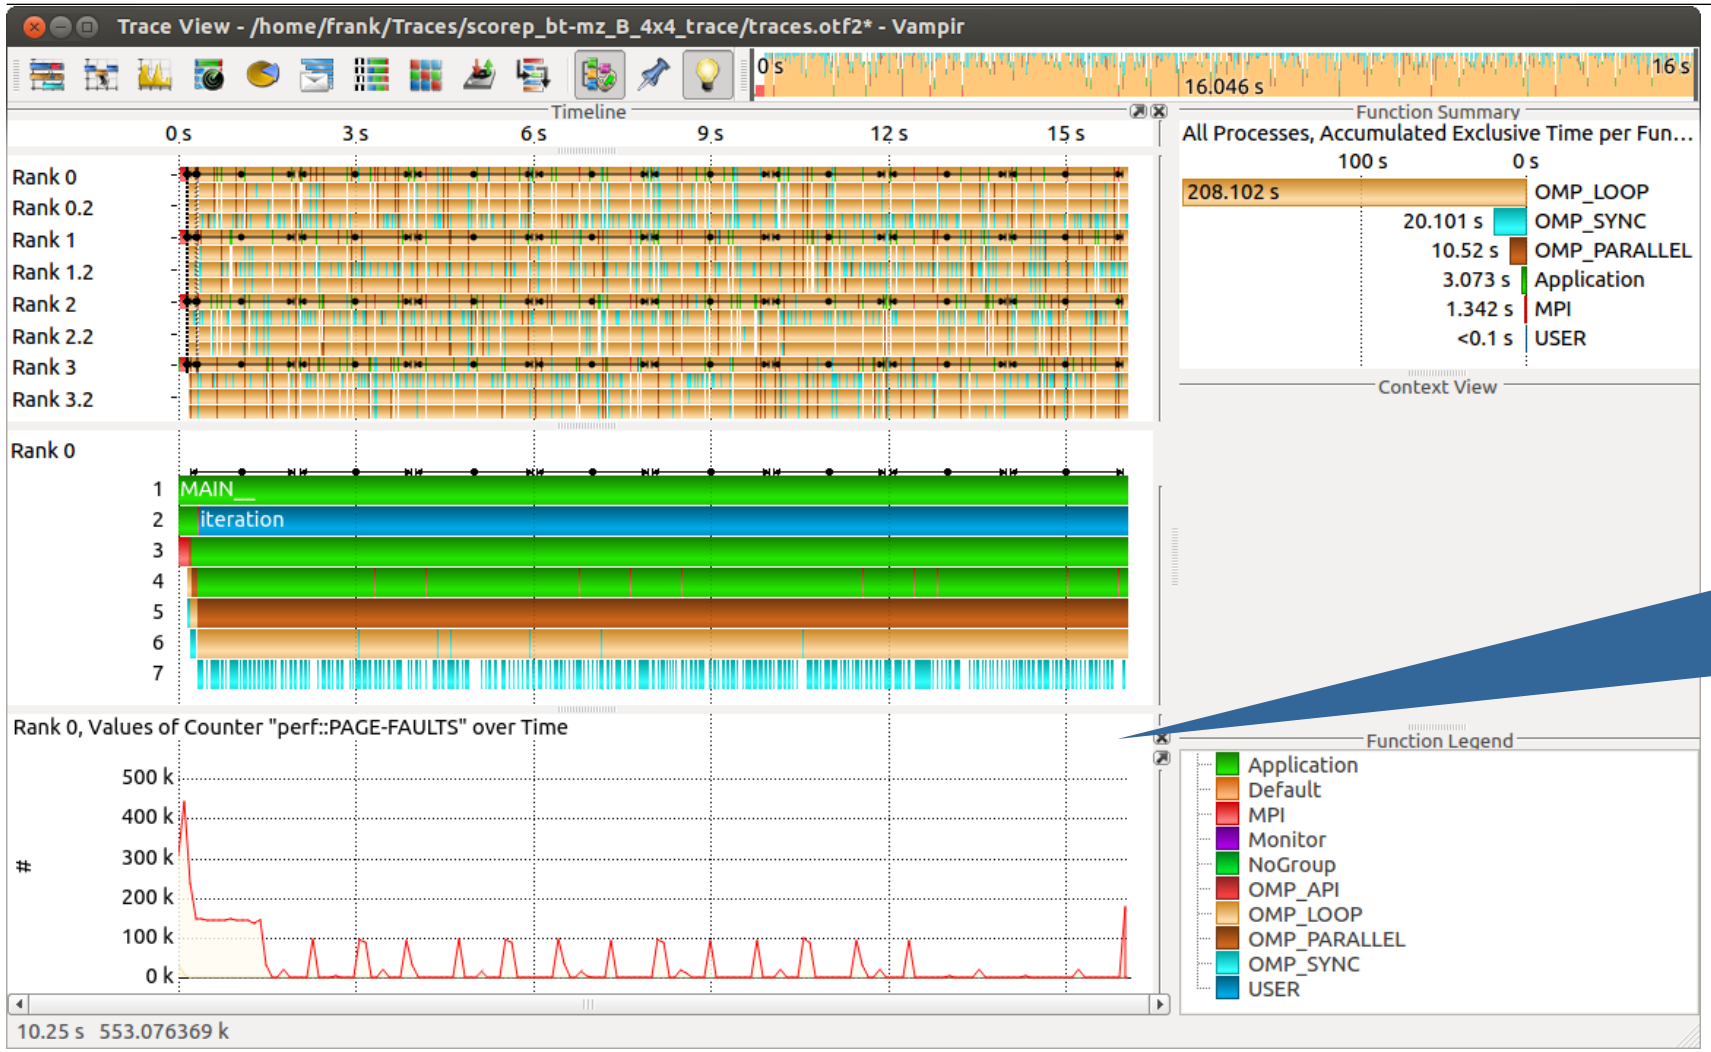

Detailed counter information over time for an individual process.

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

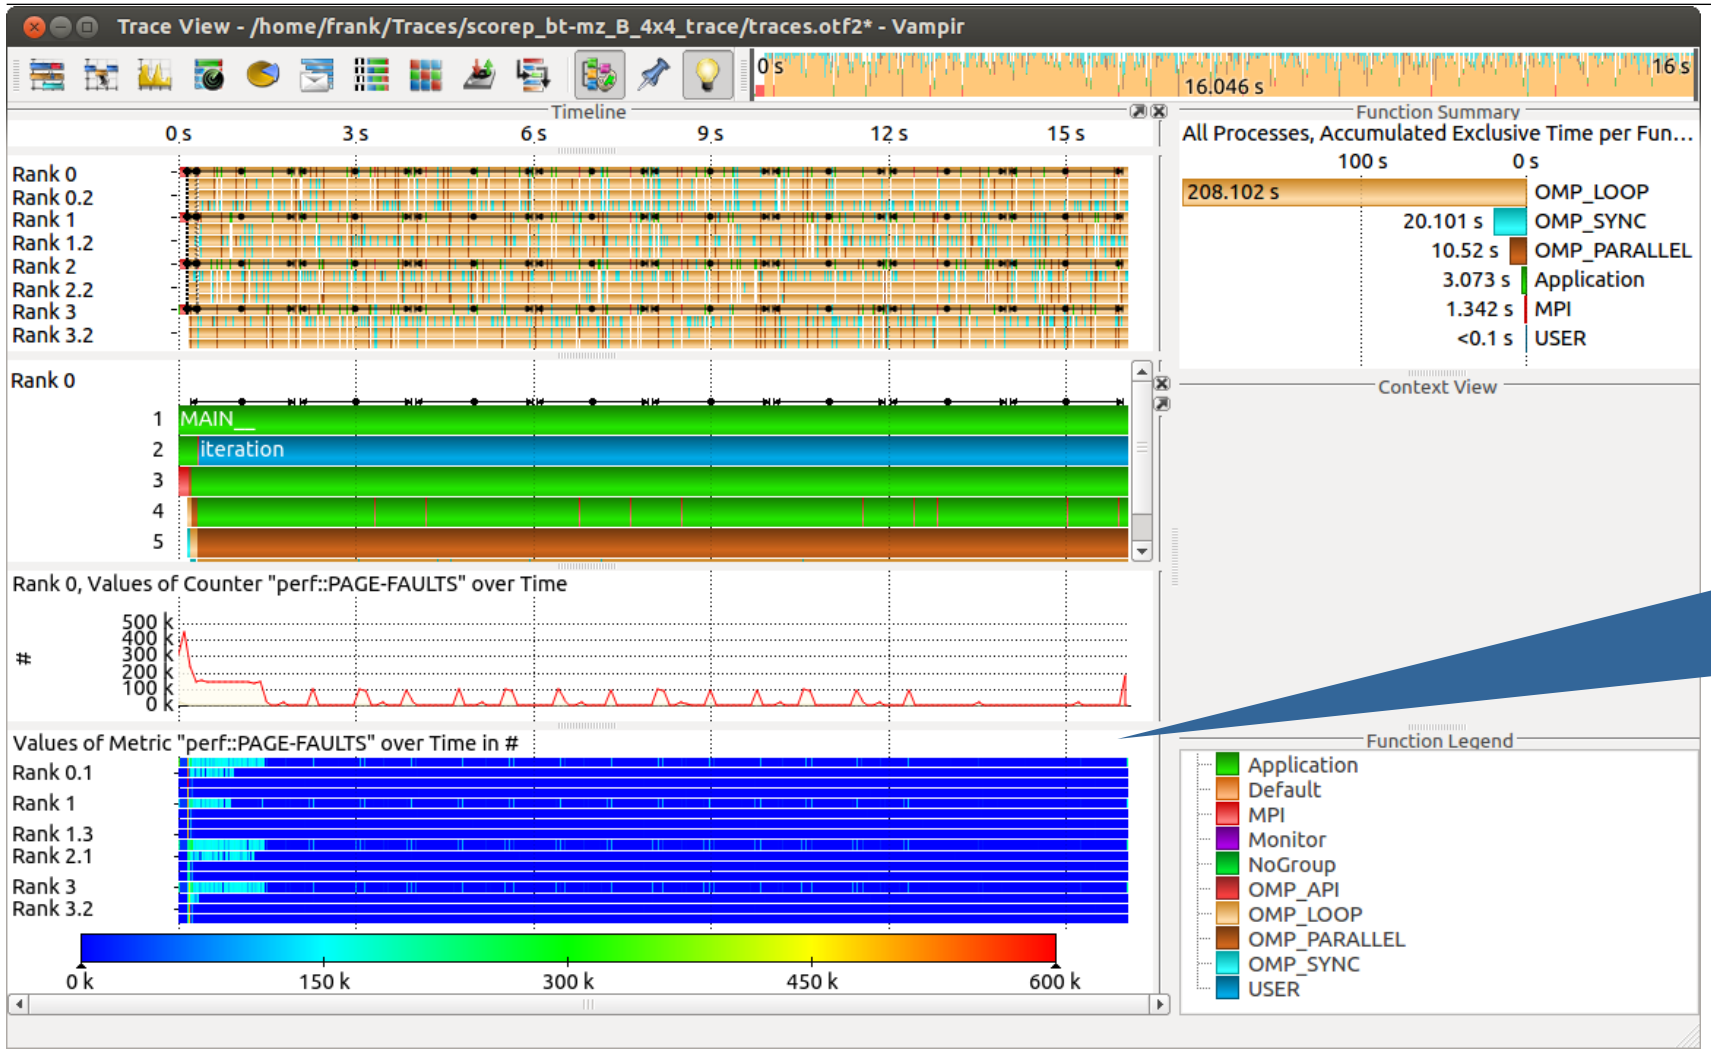

Detailed counter information over time for a collection of processes.

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

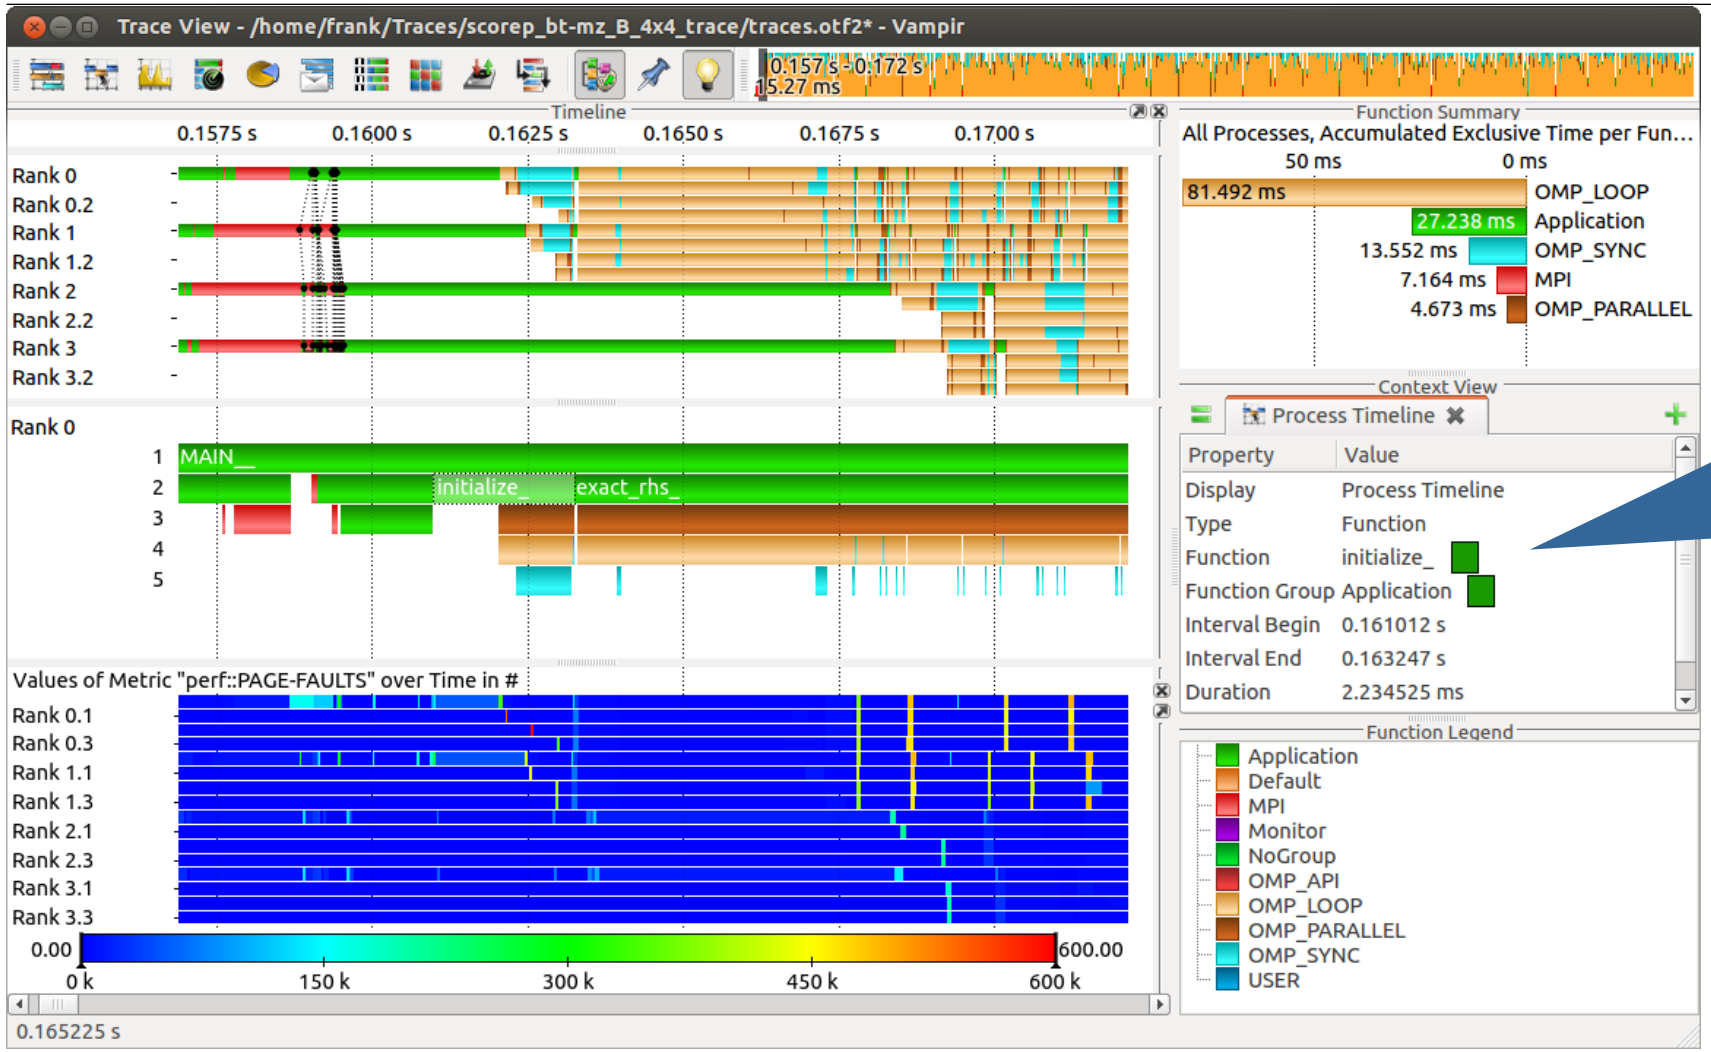

Context View: Detailed information about function “initialize_”.

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

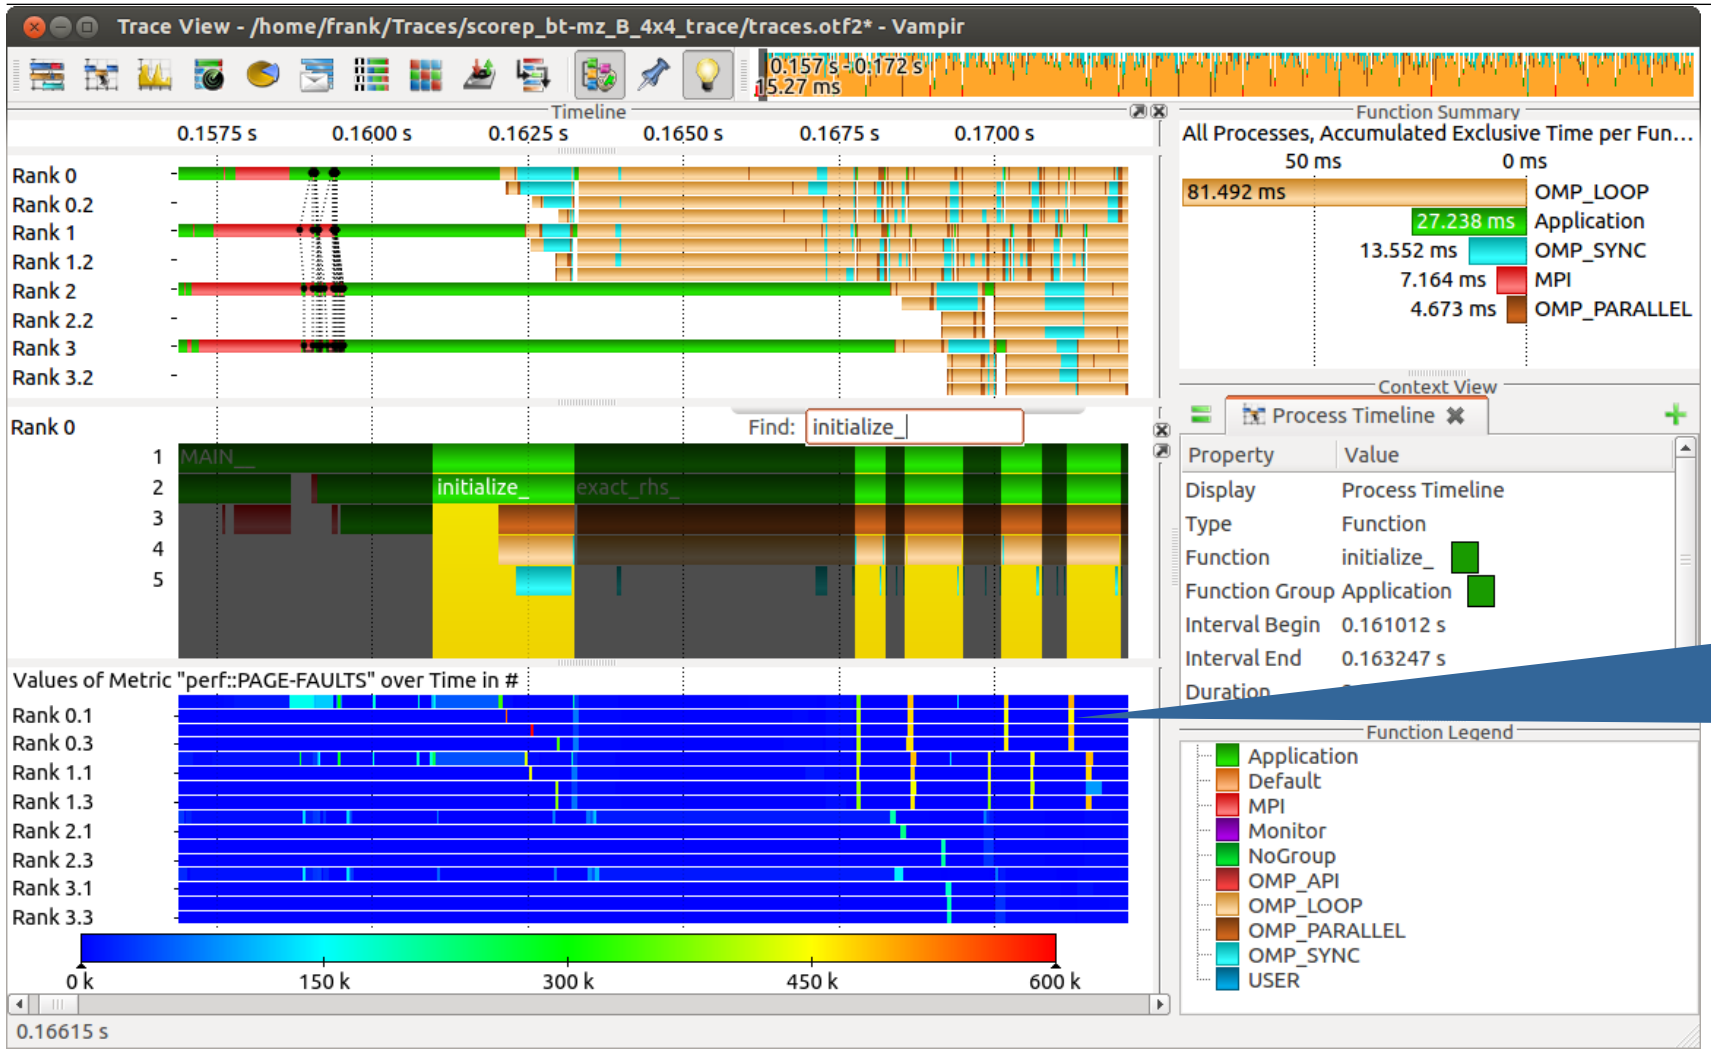

Execution of function "initialize_" results in higher page fault rates.

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

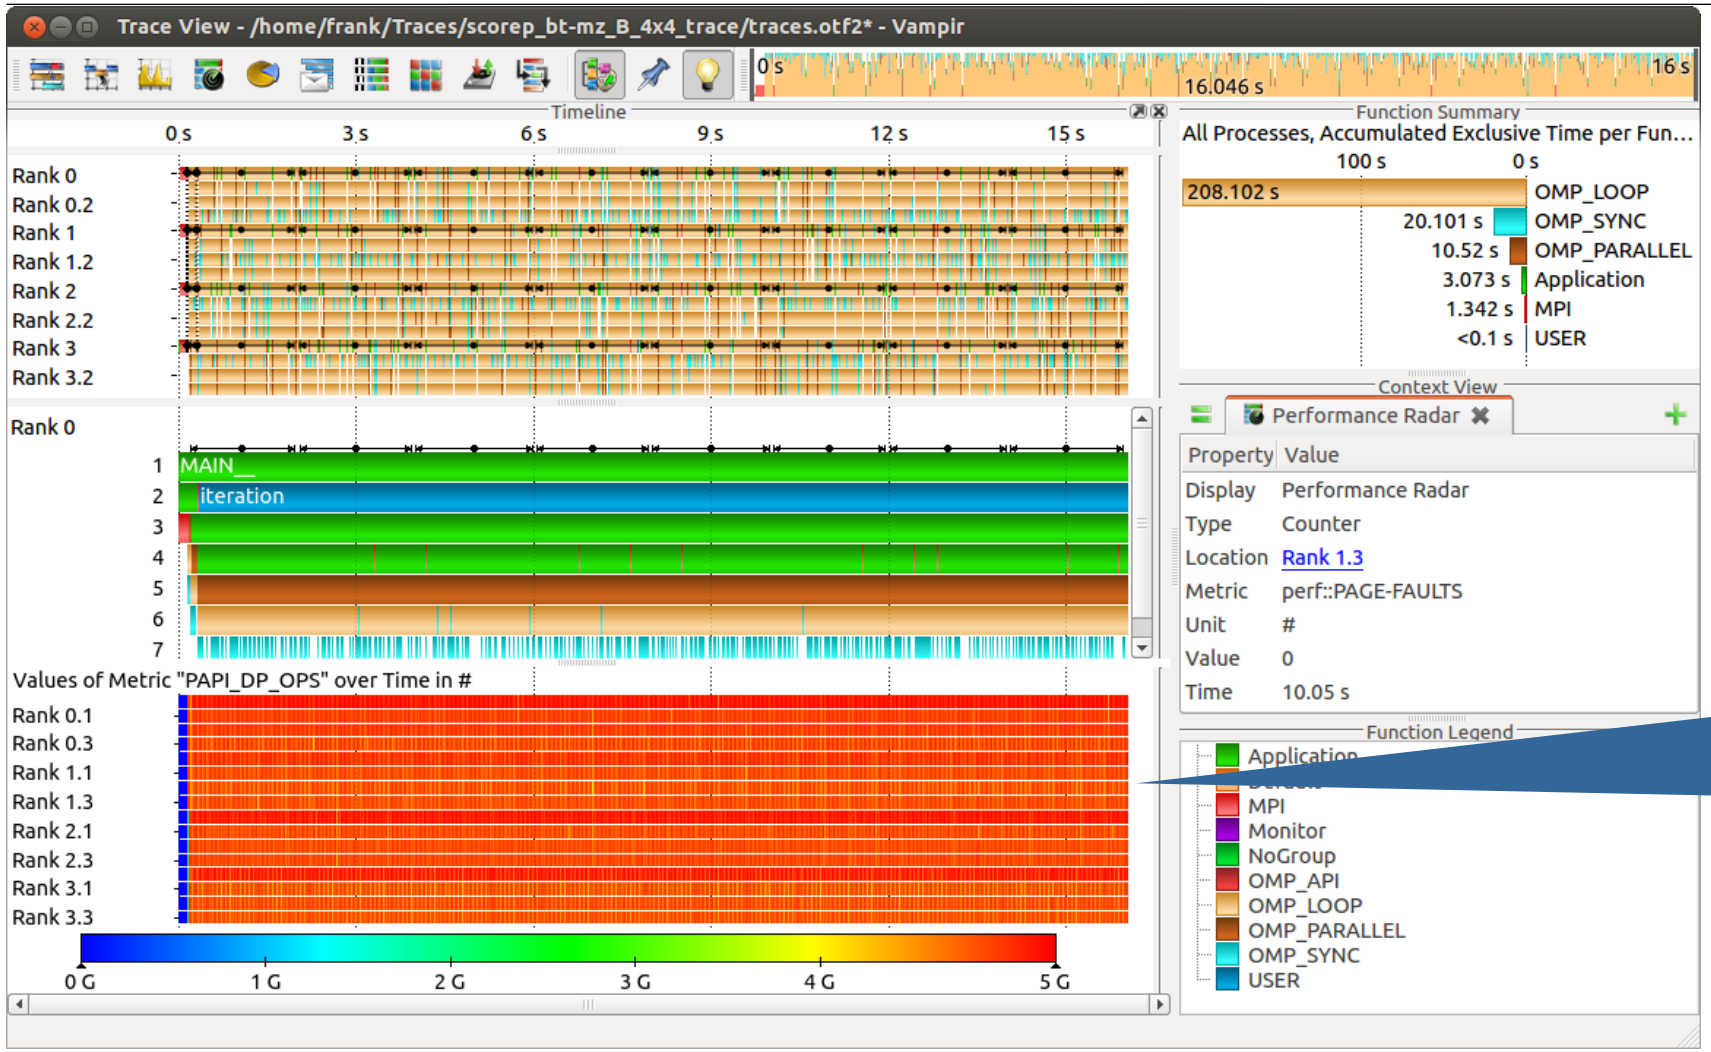

Computation phase results in higher floating point operations.

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

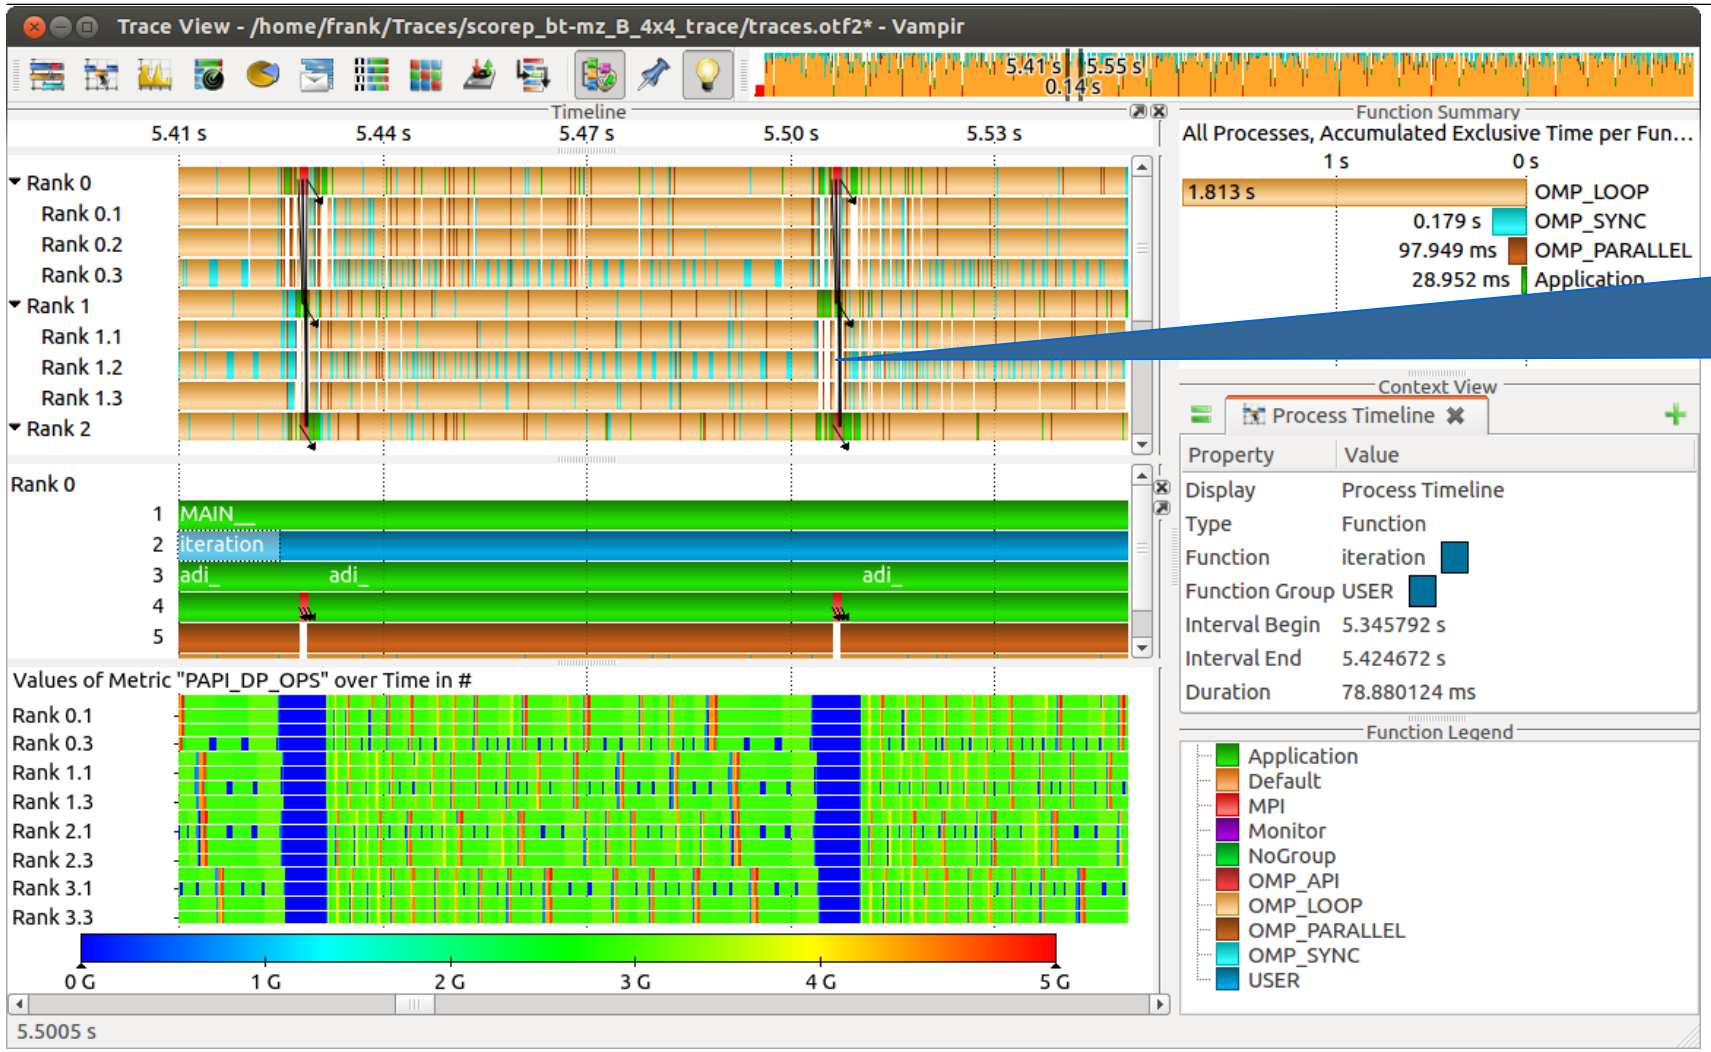

MPI communication results in lower floating point operations.

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

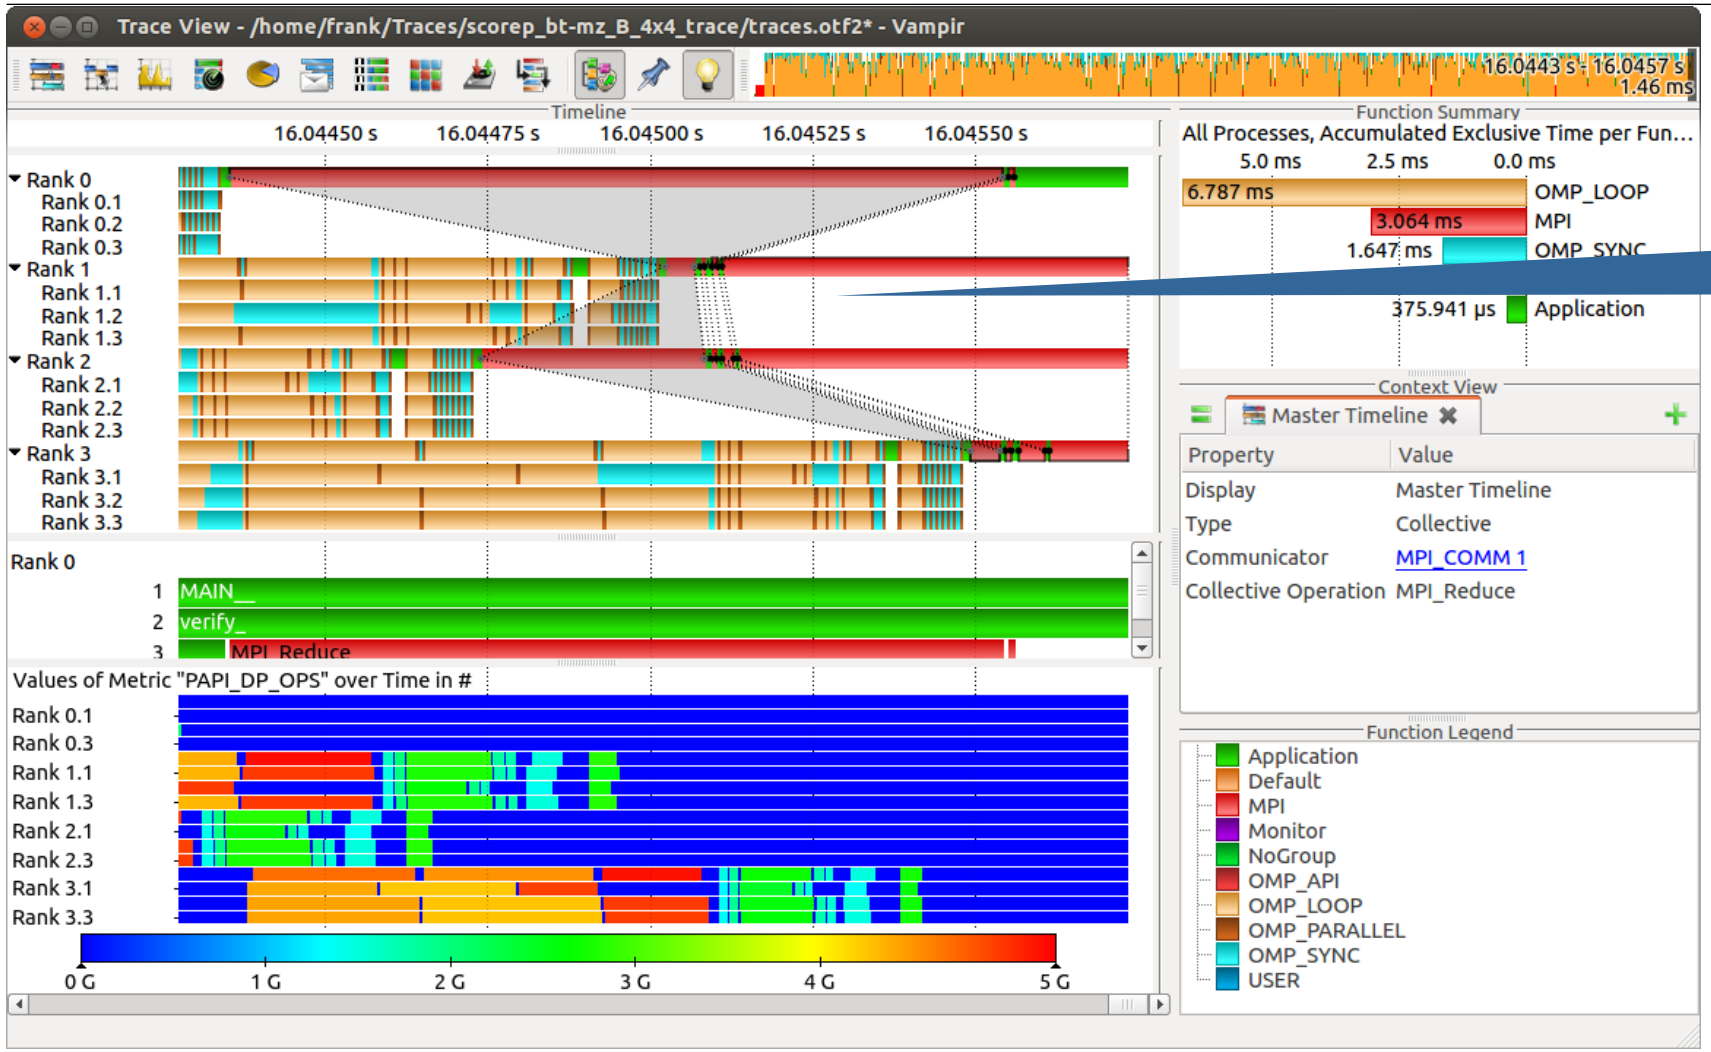

"Early reduce" bottleneck.

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

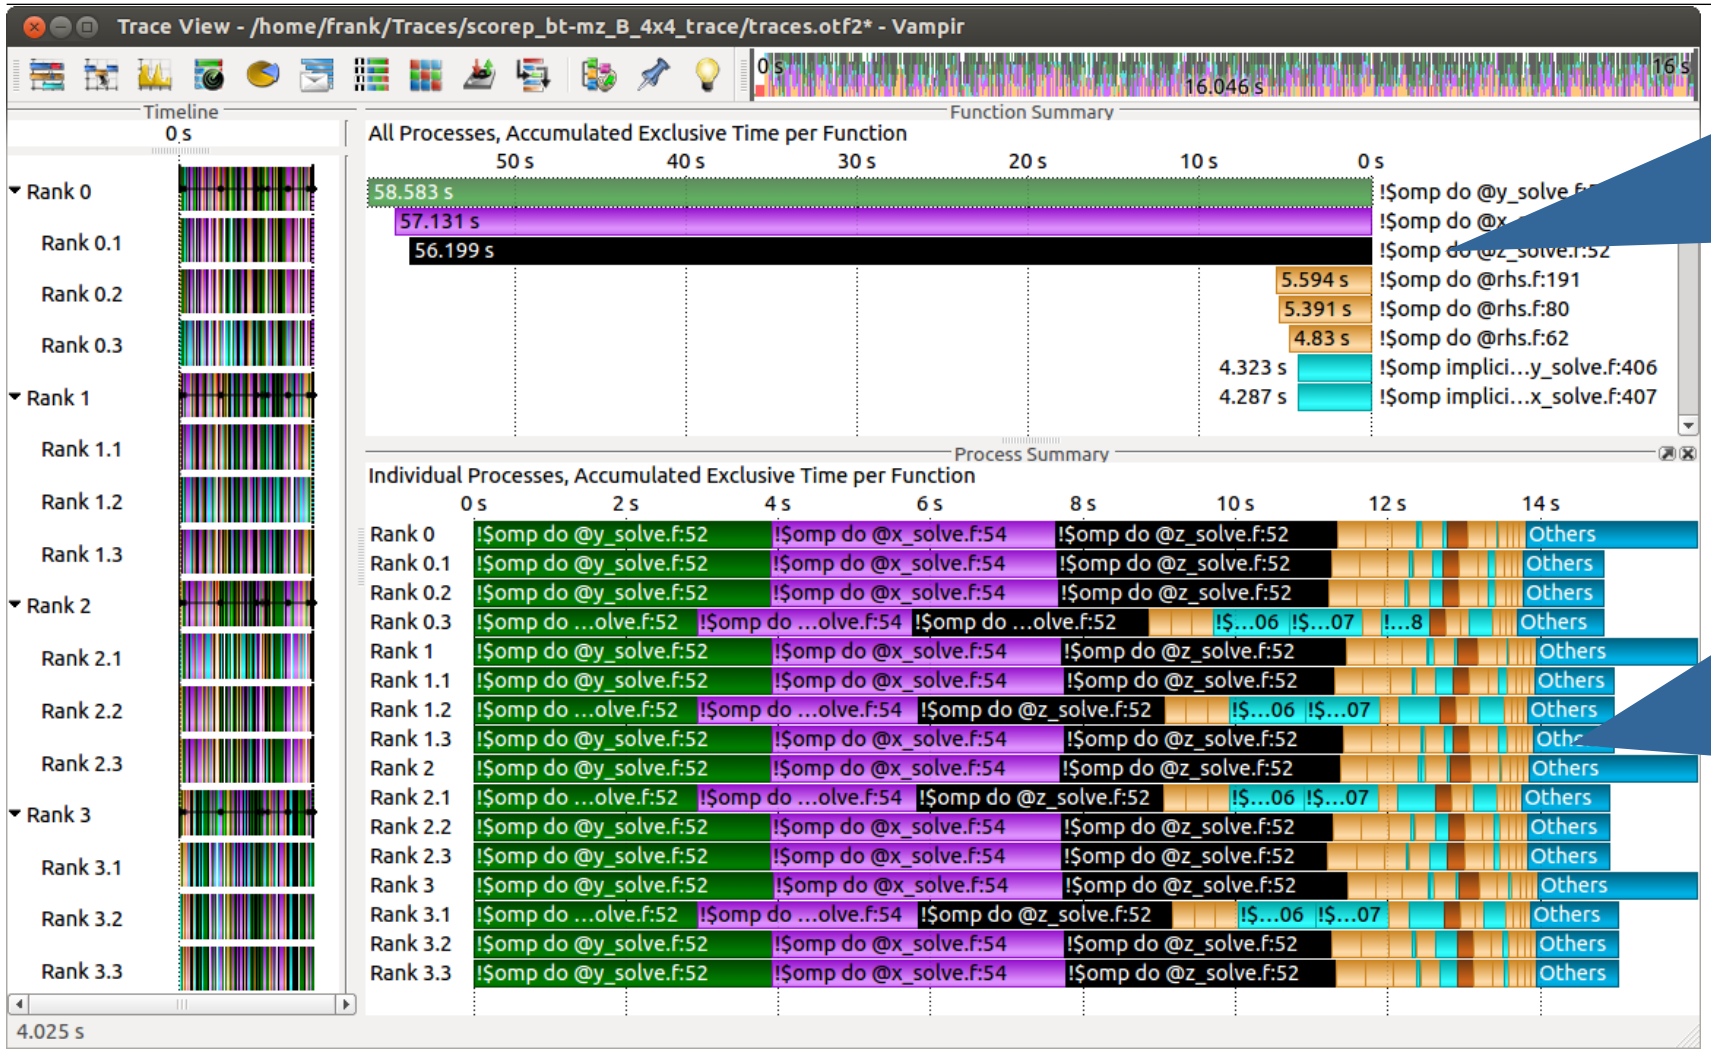

Function Summary: Overview of the accumulated information across all functions and for a collection of processes.

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

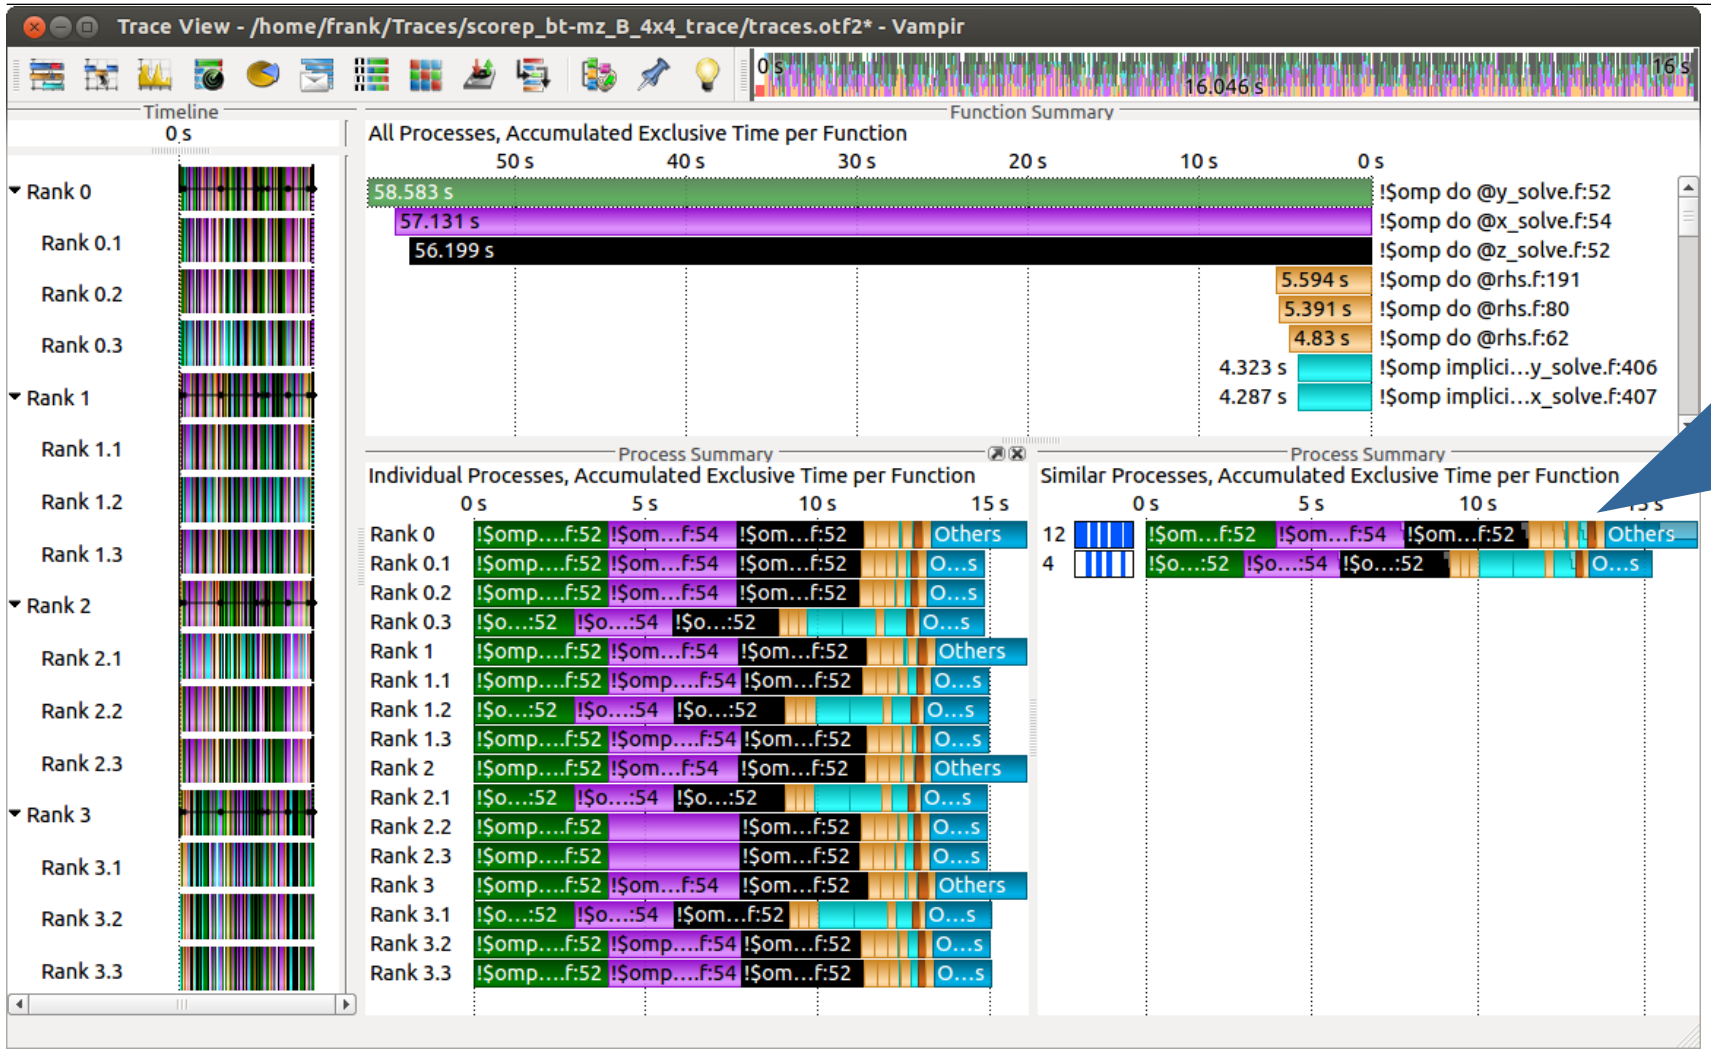

Find groups of similar processes and threads by using summarized function information.

Summary

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

- Note: Vampir does neither solve your problems automatically, nor point you directly at them. It does, however, give you FULL insight into the execution of your application.

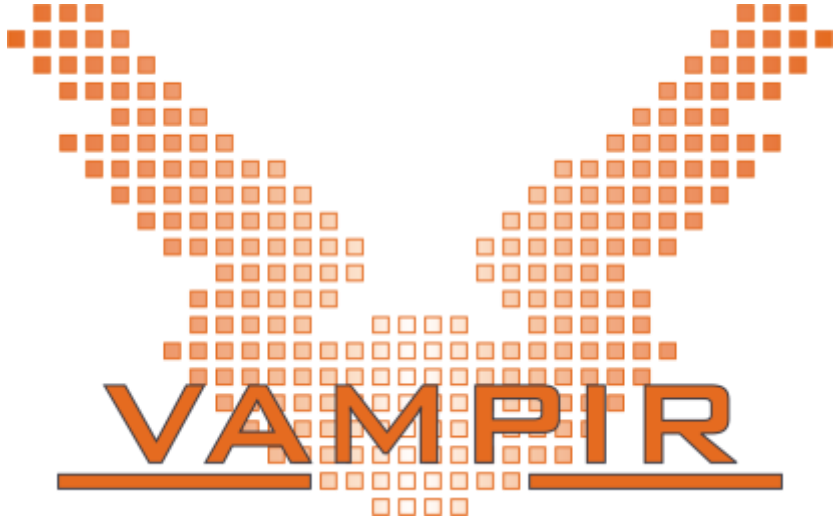

Vampir is available at http://www.vampir.eu

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