


VI-HPS

SOFTWARE



FAST SOLUTIONS


- ☒ PAPI_L1_DCM
- ☒ PAPI_L1_ICM
- ☐ PAPI_L2_DCM
- ☒ PAPI_L2_ICM
- ☒ PAPI_L2_TCM
- ☐ PAPI_L2_TCM

PRODUCTIVITY

Performance Analysis with Vampir

Matthias Weber

ZIH, Technische Universität Dresden



Part I: Welcome to the Vampir Tool Suite


- Mission
- Event Trace Visualization
- Vampir & VampirServer

Part II: Vampir Demo

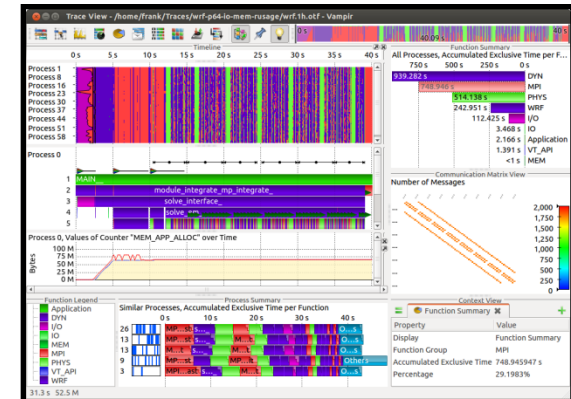
- Visualizing and Analyzing NPB-MZ-MPI / BT

Part III: Use Cases and Summary

- Case Study
- Vampir Use Cases
- Summary



- Visualization of dynamics of complex parallel processes
- Requires two components
 - Monitor/Collector (Score-P)
 - Charts/Browser (Vampir)




Typical questions that Vampir helps to answer:

- What happens in my application execution during a given time in a given process or thread?
- How do the communication patterns of my application execute on a real system?
- Are there any imbalances in computation, I/O or memory usage and how do they affect the parallel execution of my application?

- 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




- **Timeline Charts:**

-  Master Timeline
-  Process Timeline
-  Counter Data Timeline
-  Performance Radar

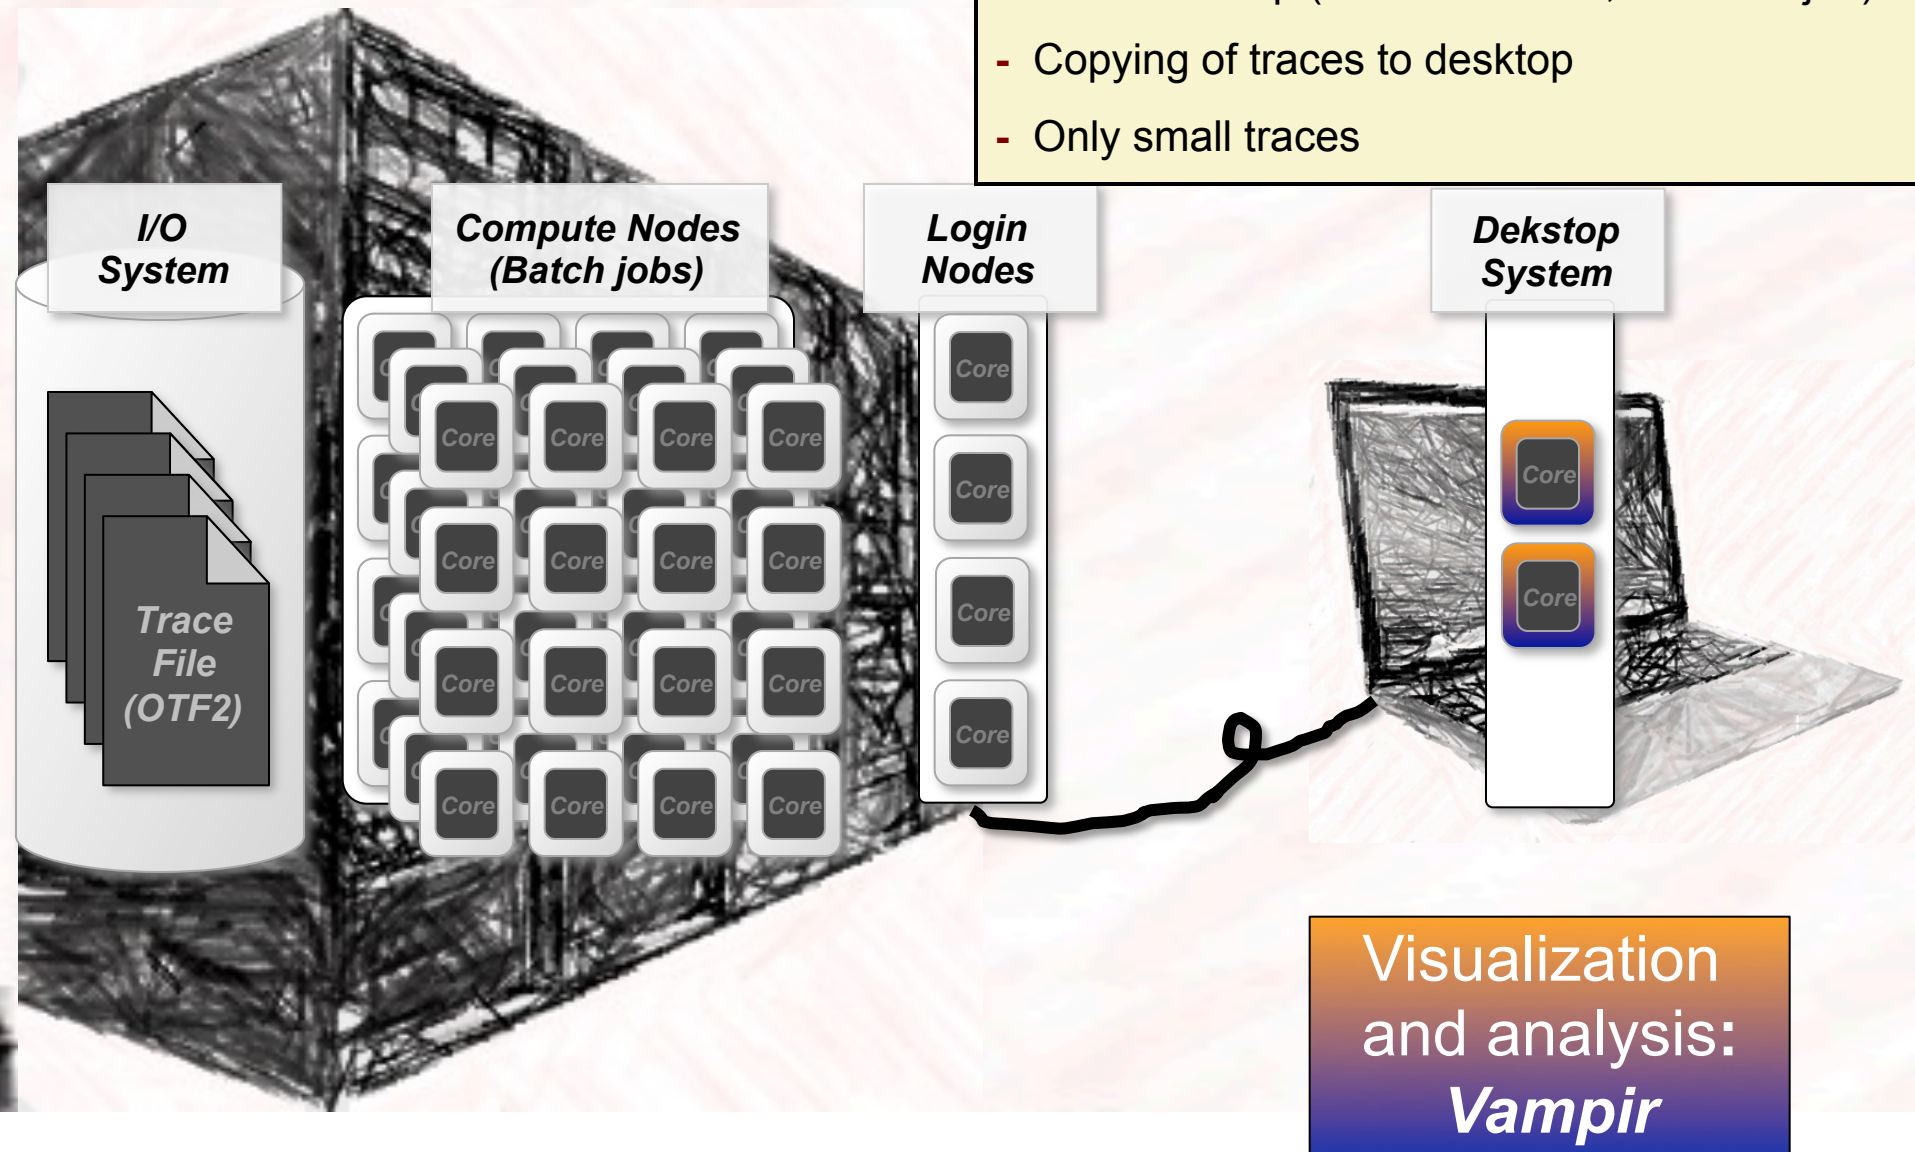
- **Summary Charts:**

-  Function Summary
-  Message Summary
-  Process Summary
-  Communication Matrix View



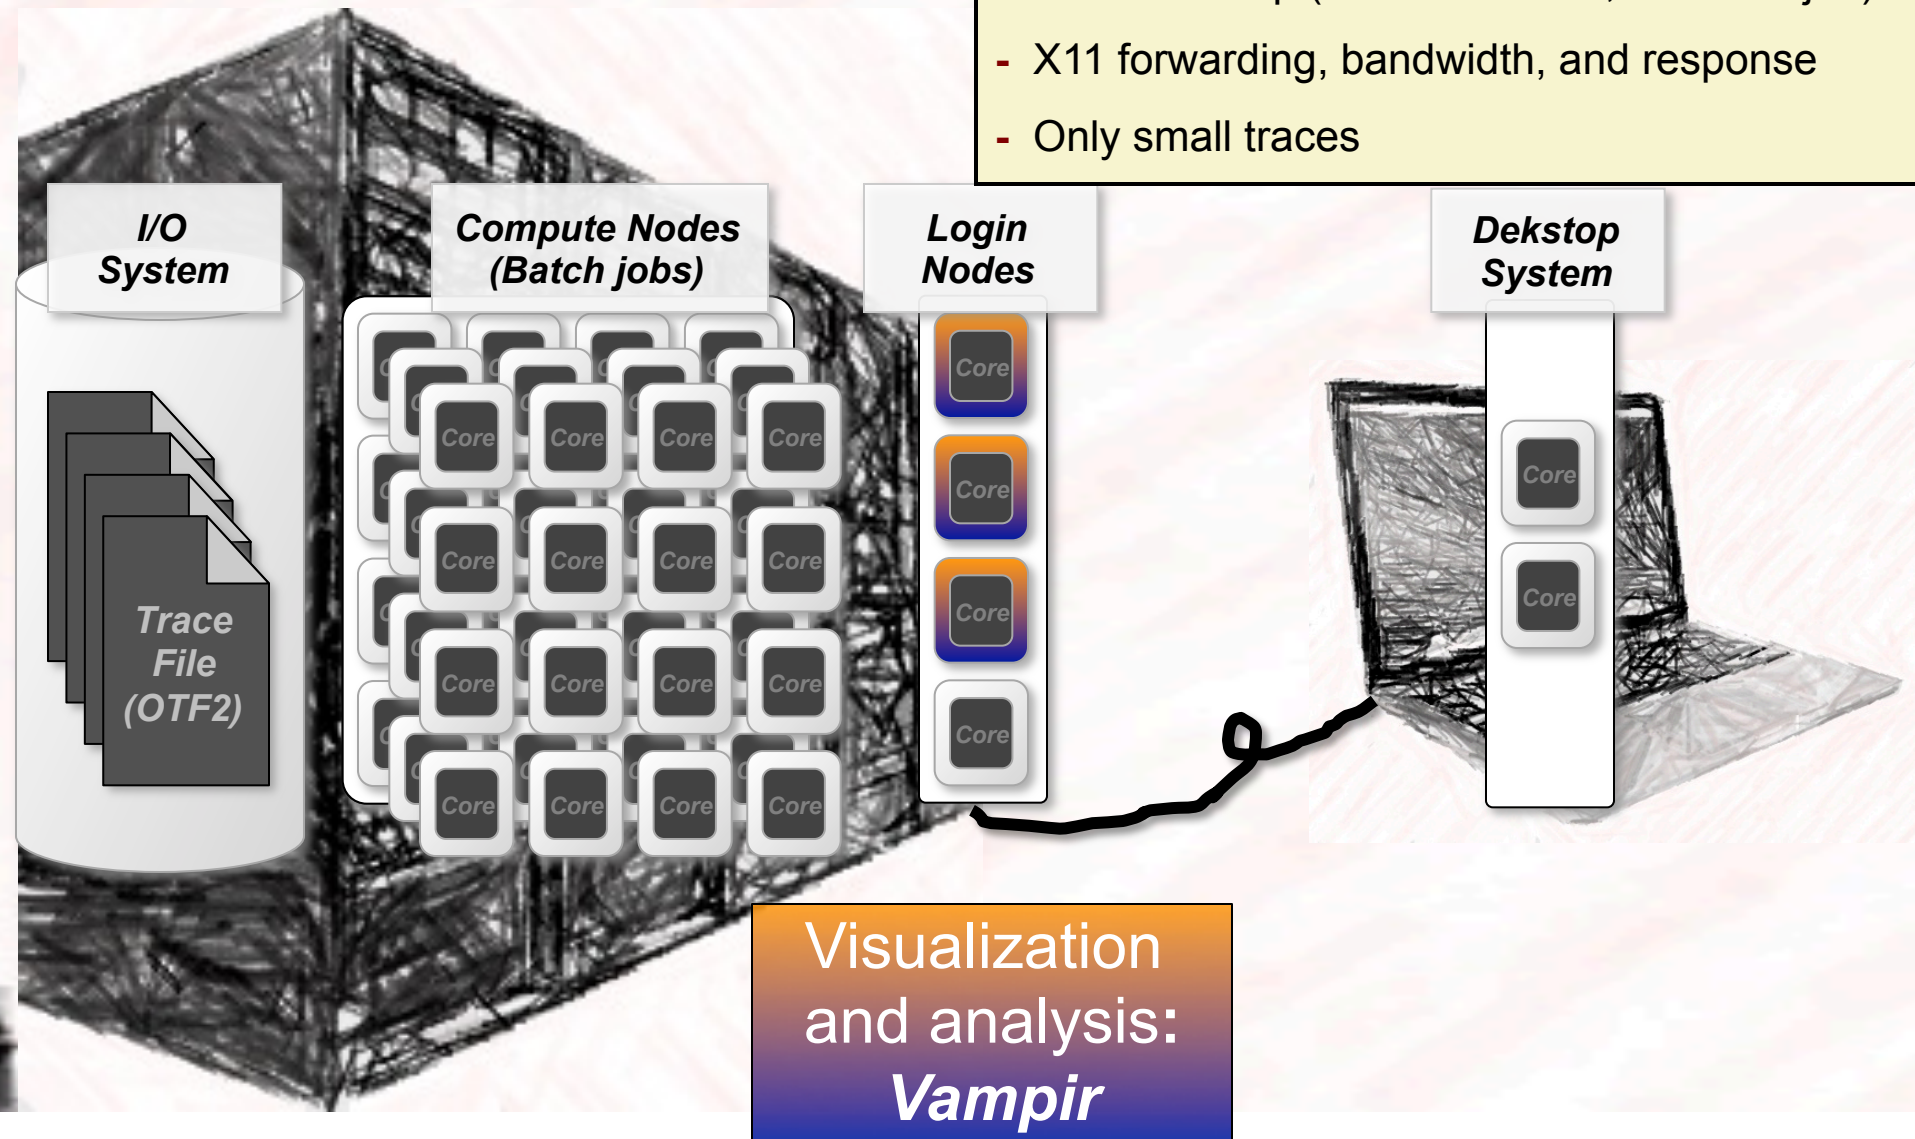
Visualization: Most simple (Analysis on Desktop)

- + Minimal setup (no installations, no batch job)
- Copying of traces to desktop
- Only small traces



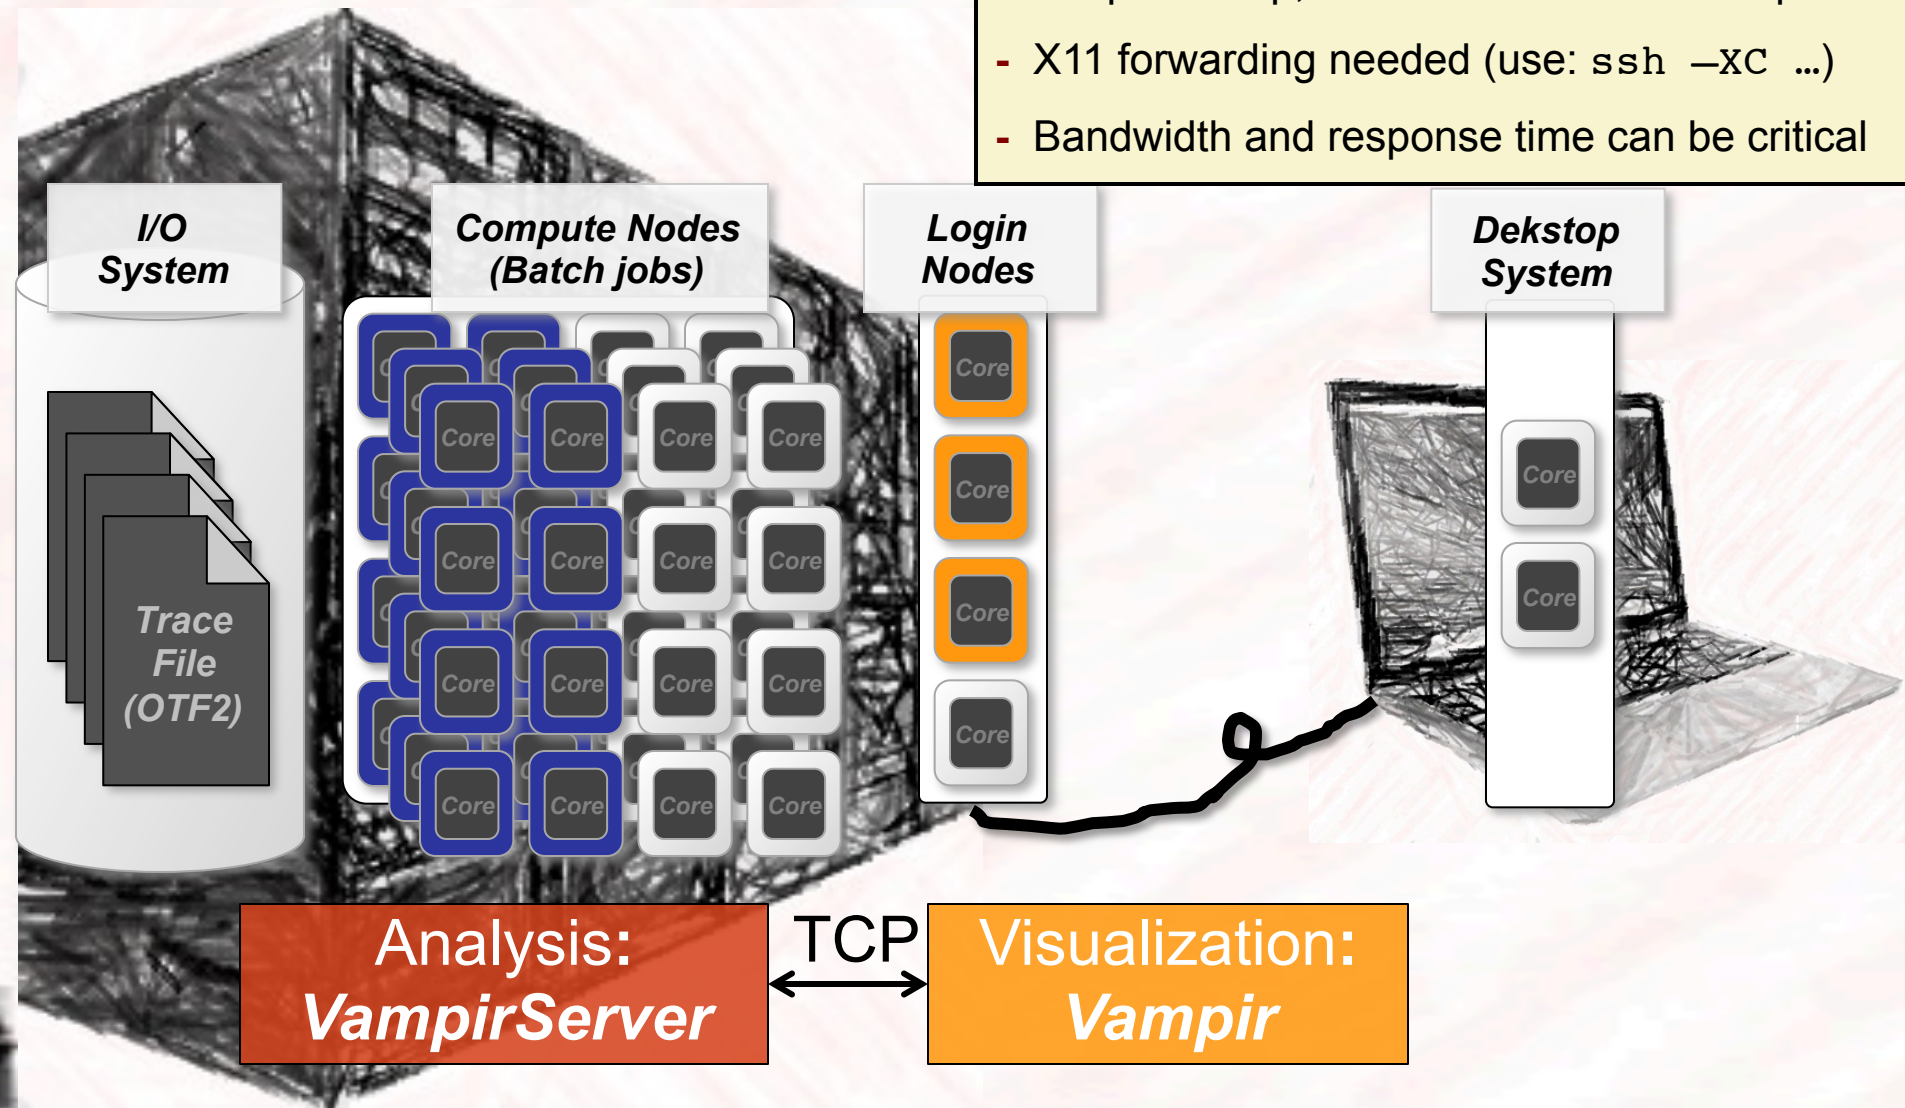
Visualization: Most simple (Analysis on Frontend)

- + Minimal setup (no installations, no batch job)
- X11 forwarding, bandwidth, and response
- Only small traces



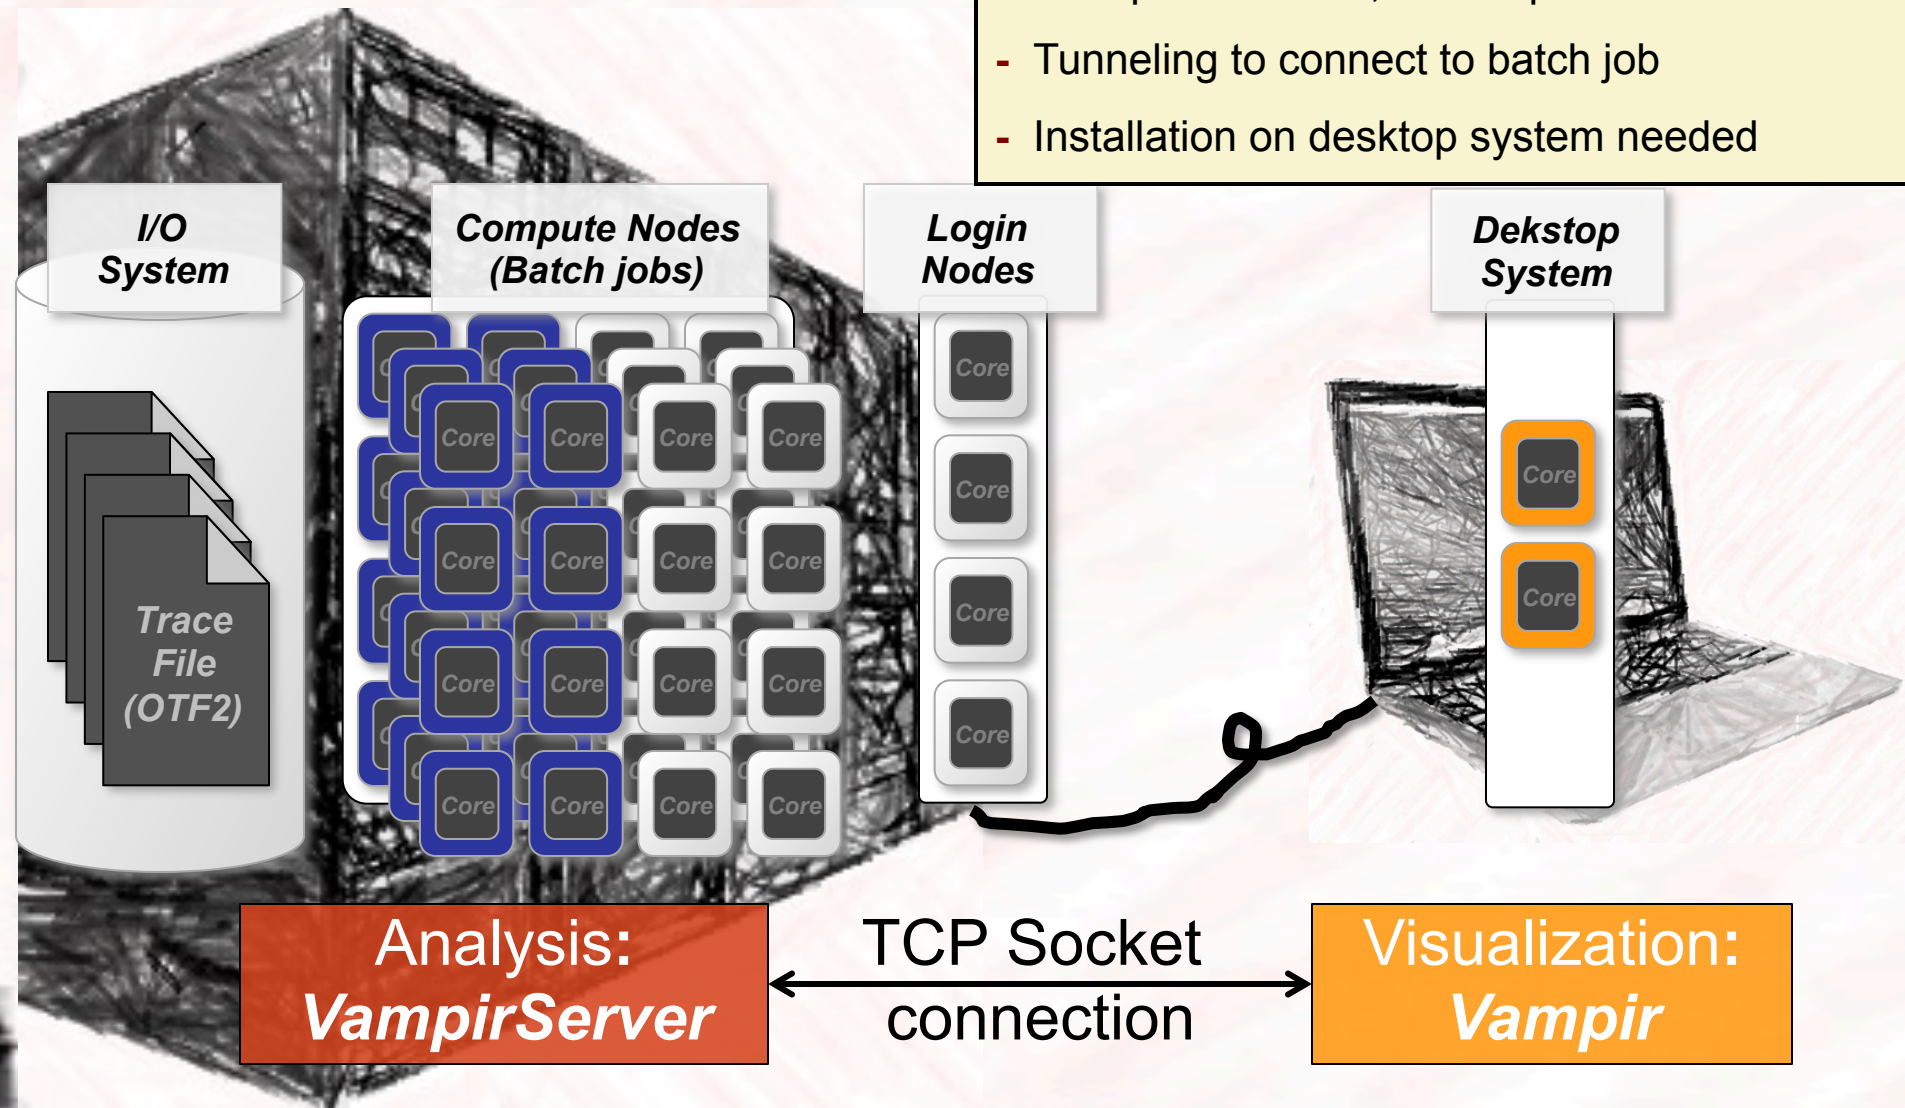
Visualization: Alternative (Analysis on HPC system)

- + Simpler setup, no installation on desktop
- X11 forwarding needed (use: `ssh -XC ...`)
- Bandwidth and response time can be critical



Visualization: Best Option (Analysis on HPC system)



- + Best performance, low response time
- Tunneling to connect to batch job
- Installation on desktop system needed



1. Instrument your application with Score-P
2. Run your application with an appropriate test set
3. Analyze your trace file with Vampir
 - Small trace files can be analyzed on your local workstation
 1. Start your local Vampir
 2. Load trace file from your local disk
 - Large trace files should be stored on the HPC file system
 1. Start VampirServer on your HPC system
 2. Start your local Vampir
 3. Connect local Vampir with the VampirServer on the HPC system
 4. Load trace file from the HPC file system

VI-HPS

SOFTWARE



FAST SOLUTIONS

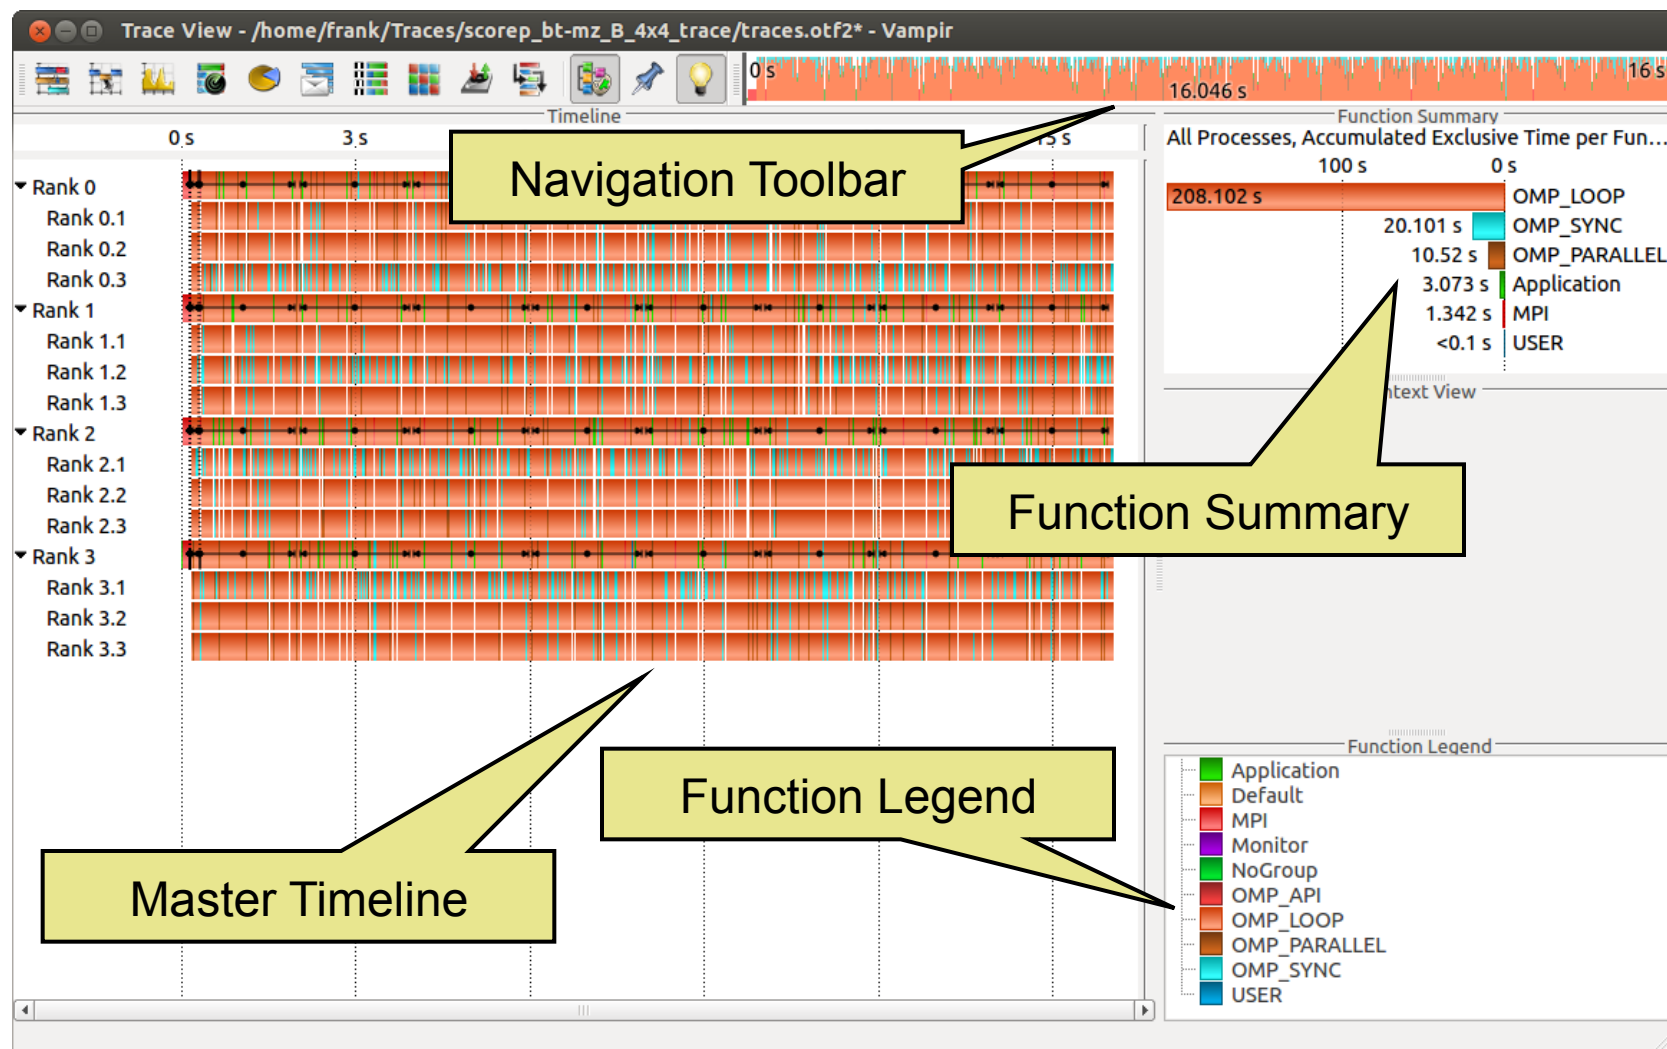
- ☒ PAPI_L1_DCM
- ☒ PAPI_L1_ICM
- ☐ PAPI_L2_DCM
- ☒ PAPI_L2_ICM
- ☒ PAPI_L2_TCM
- ☐ PAPI_L2_TCM

PRODUCTIVITY

Vampir Demo

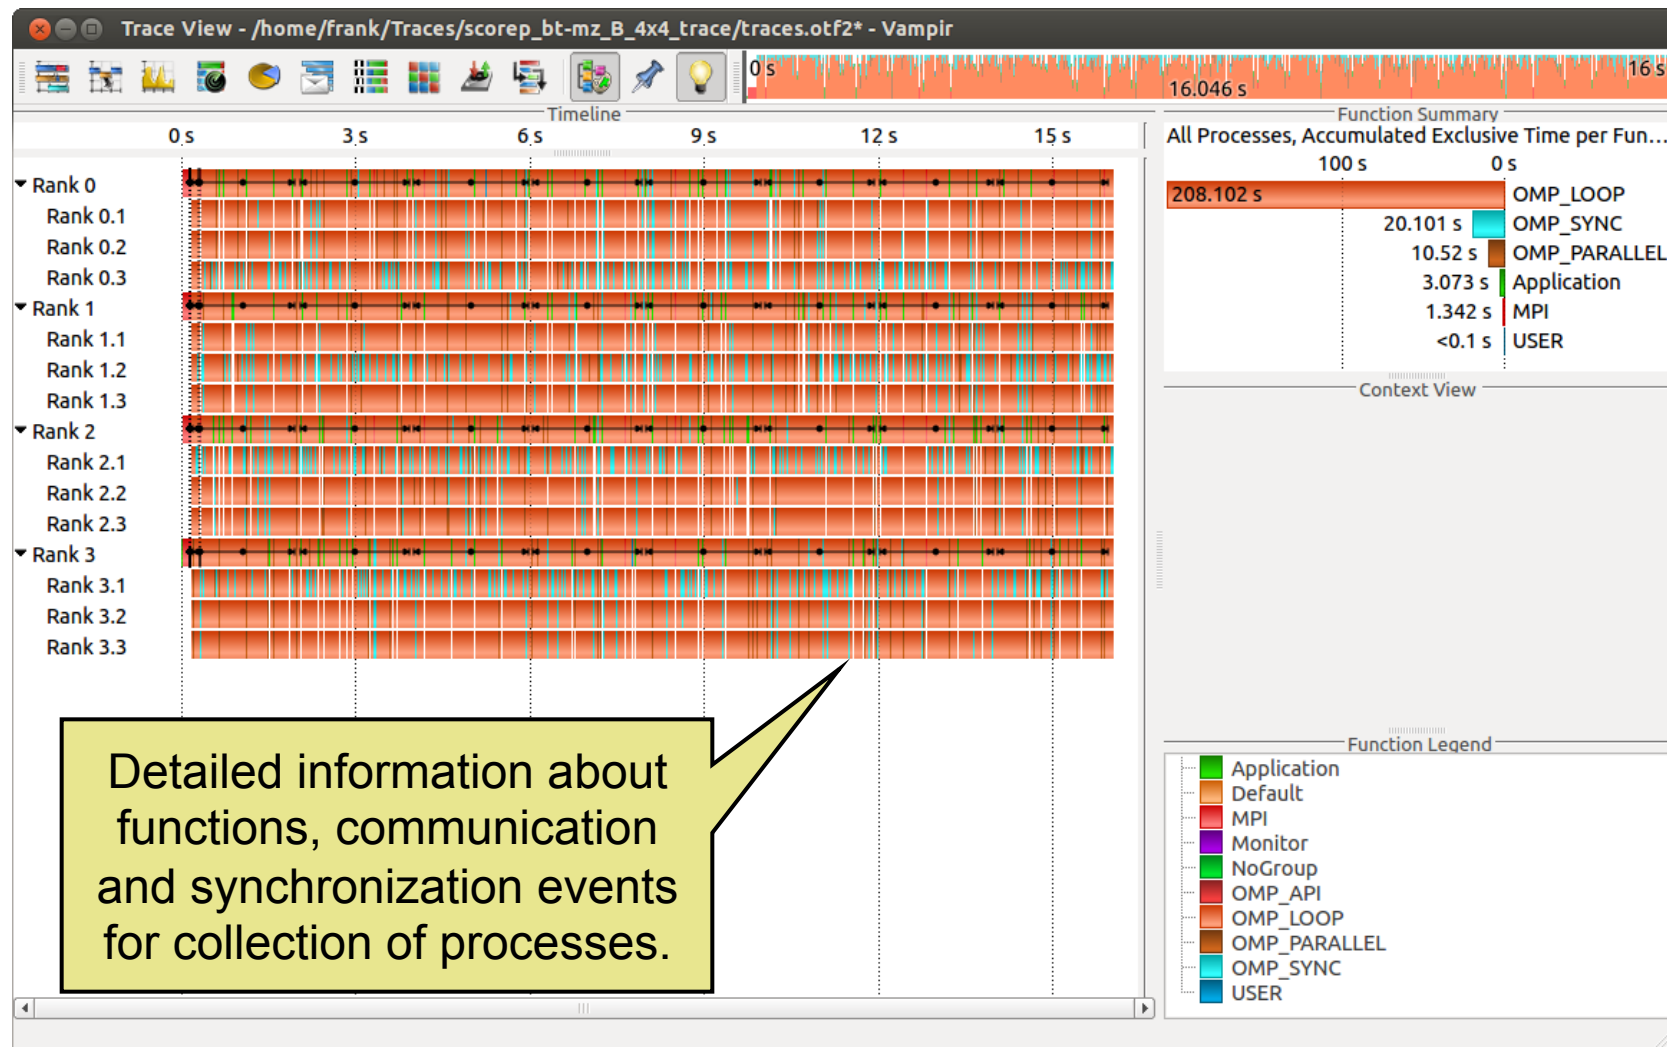
Visualizing and analyzing NPB-MZ-MPI / BT


```
% vampir scorep_bt-mz_B_4x4_trace
```



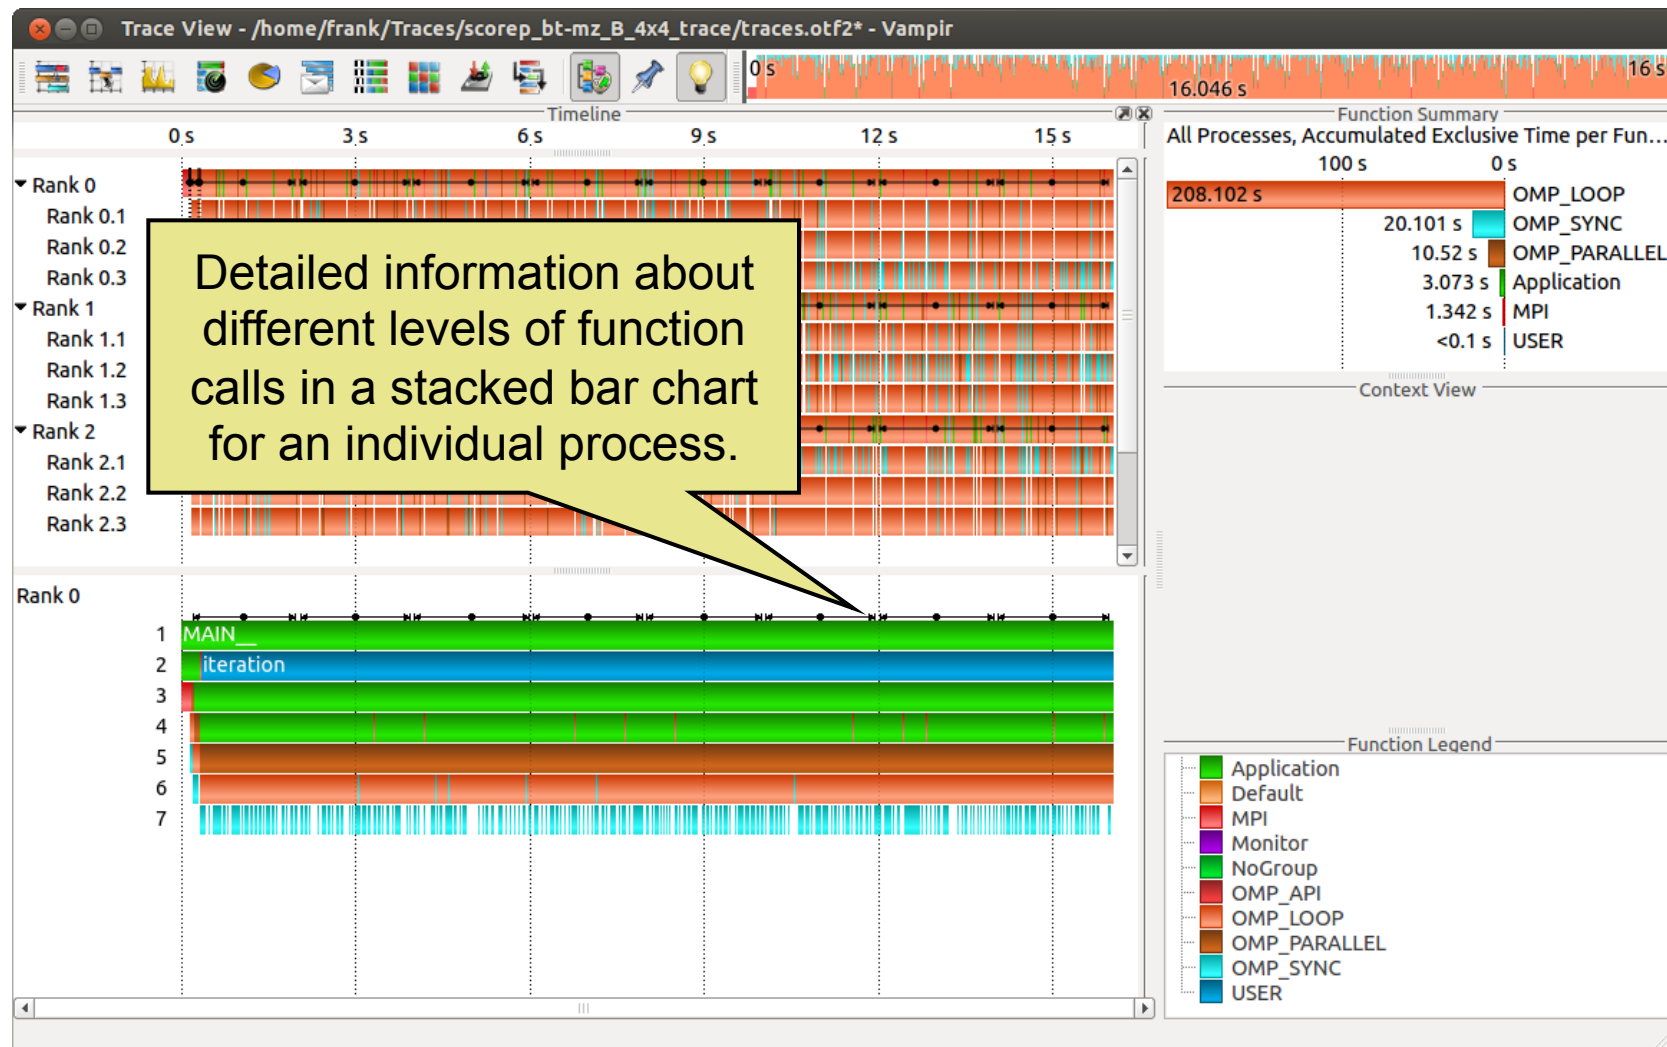


Master Timeline

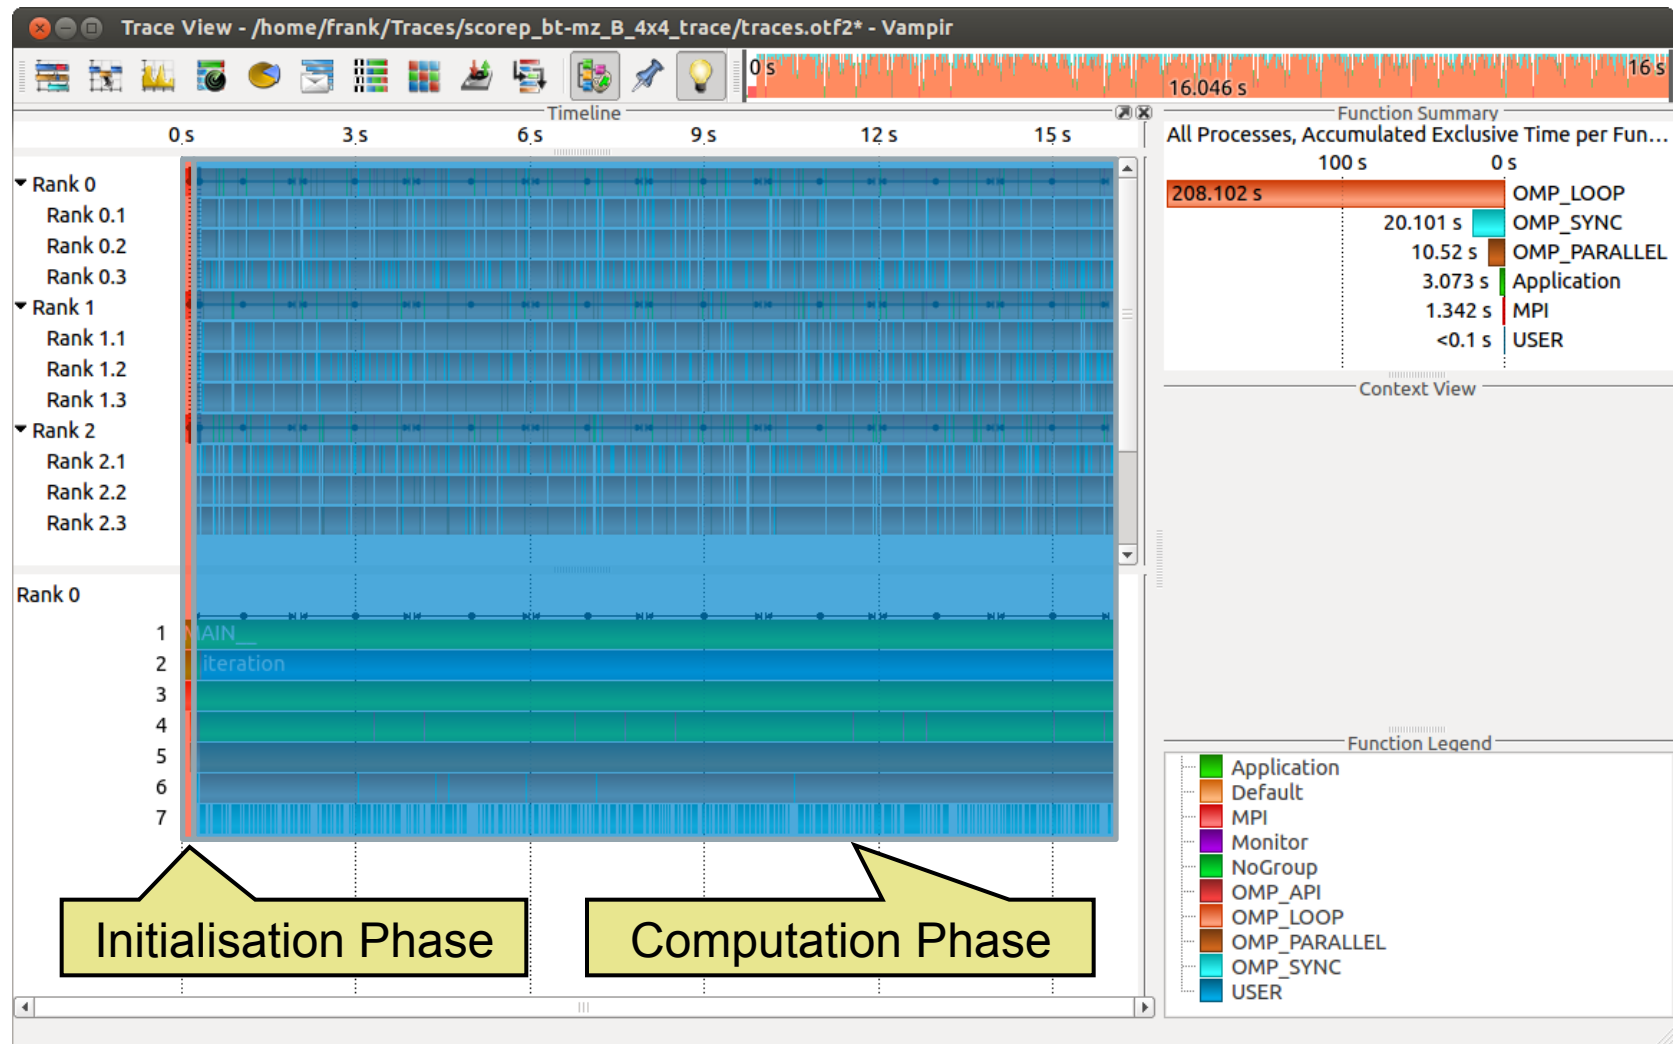




Process Timeline

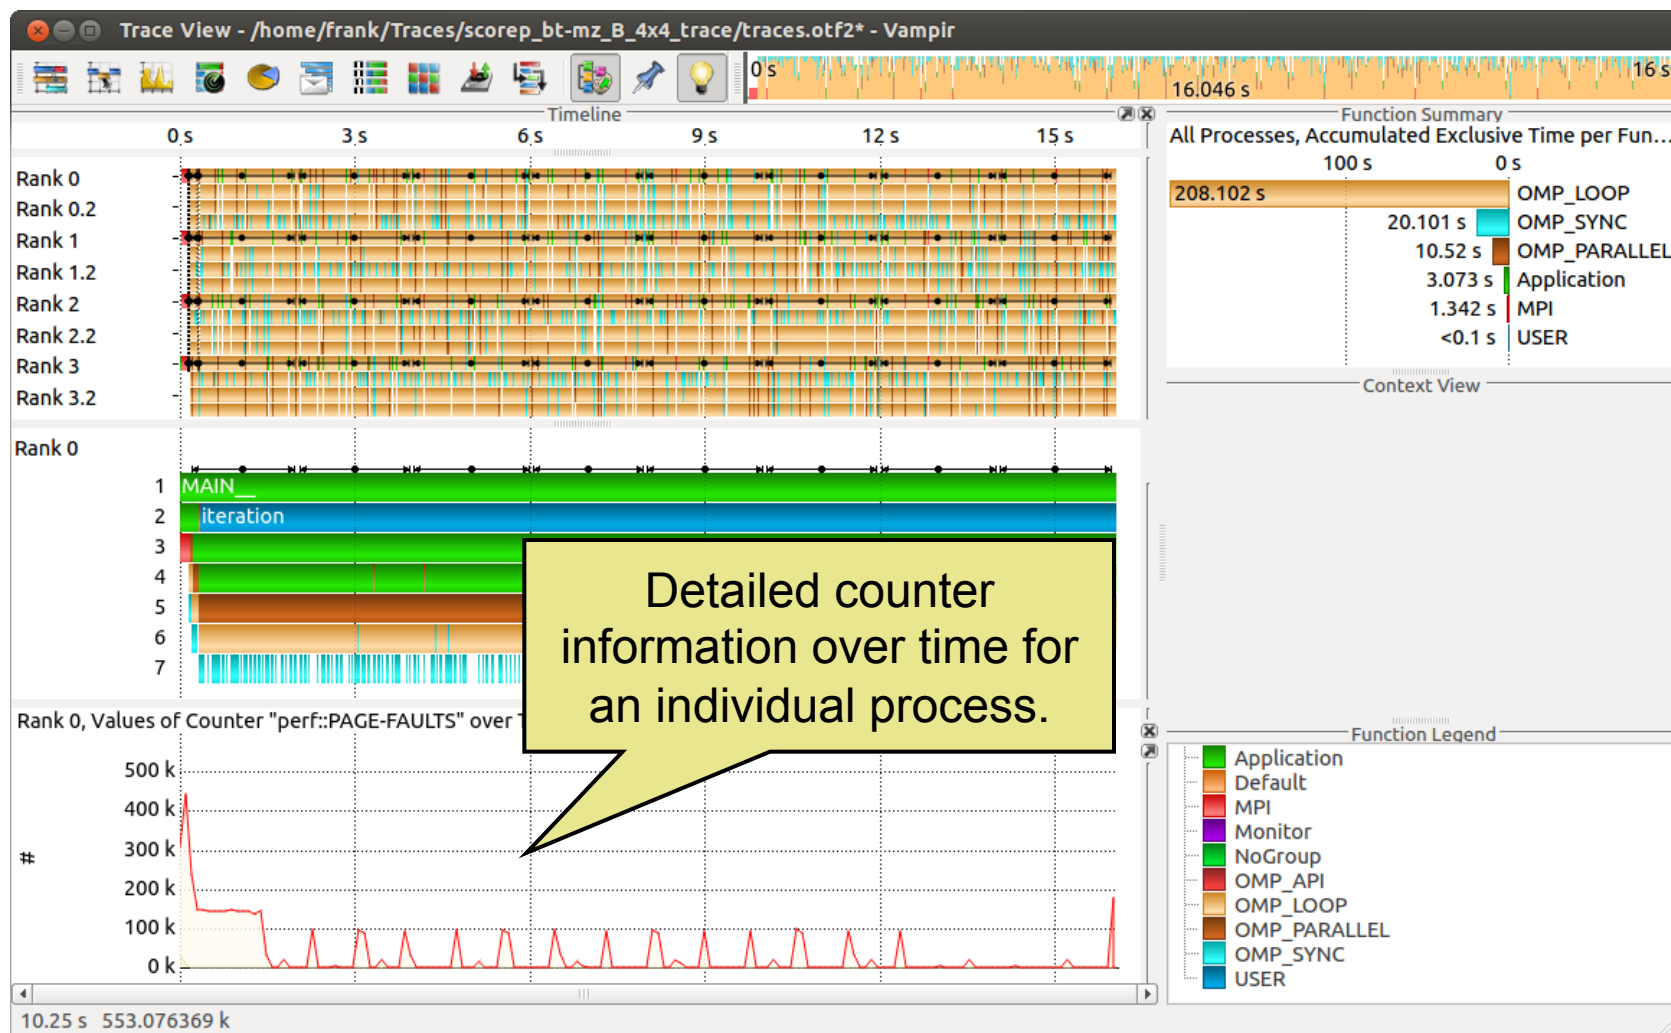


Typical program phases



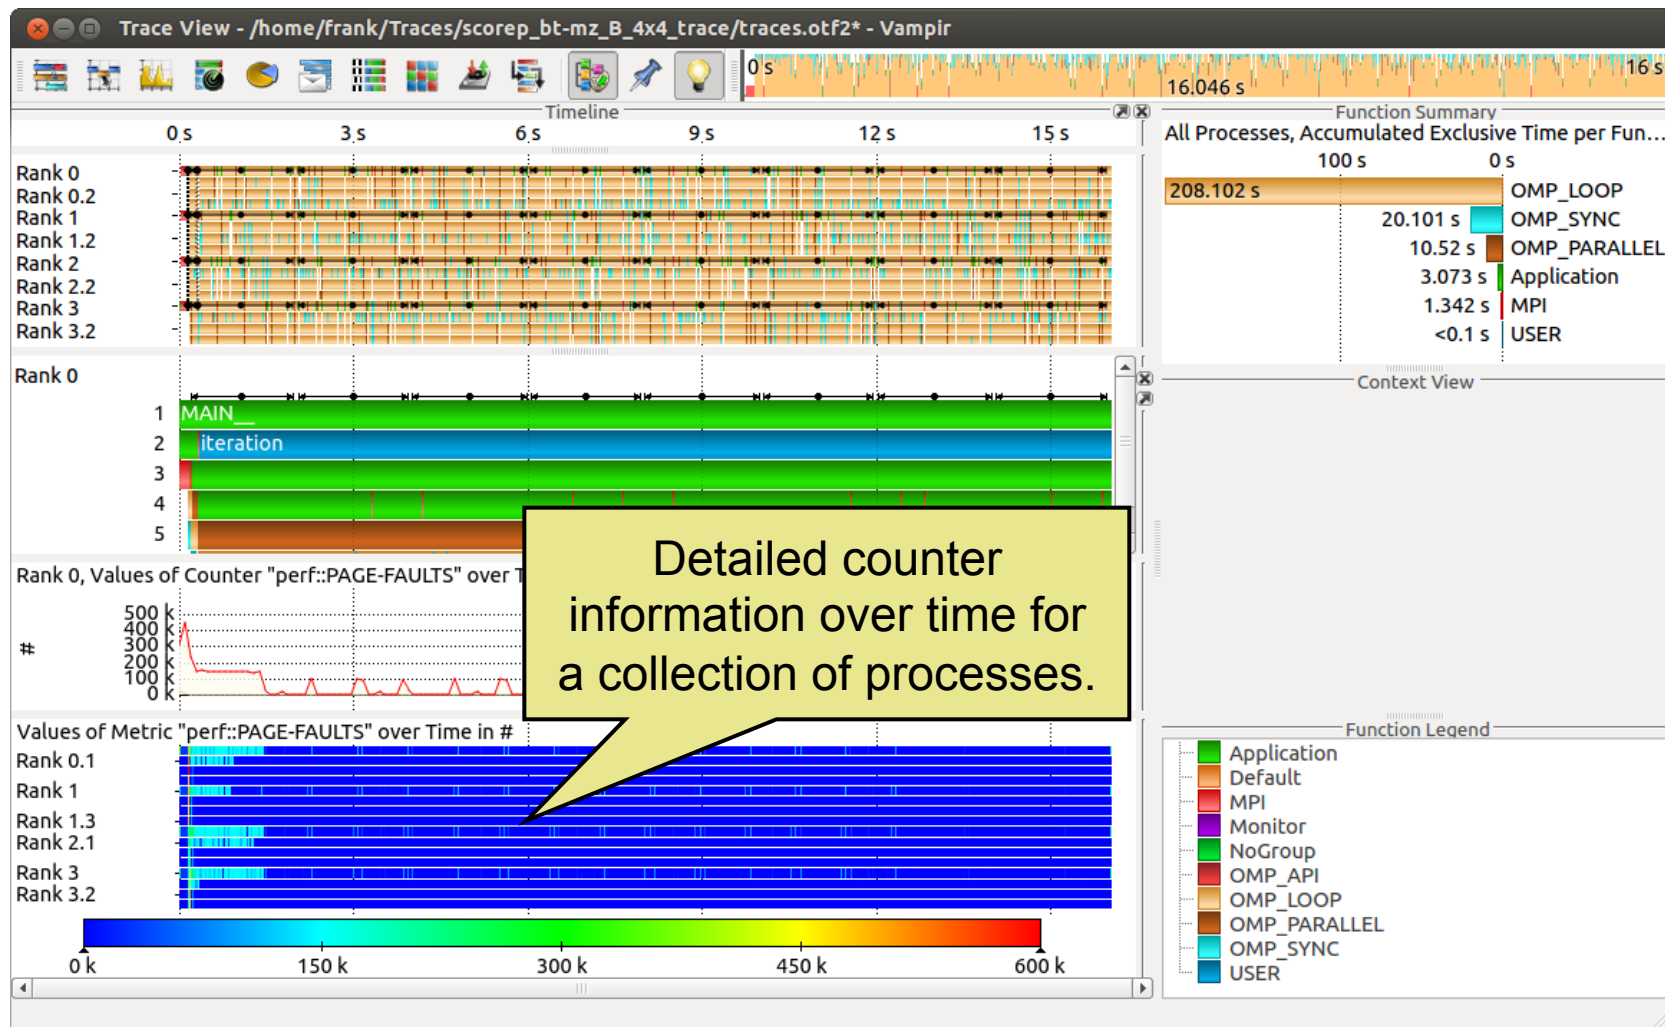


Counter Data Timeline

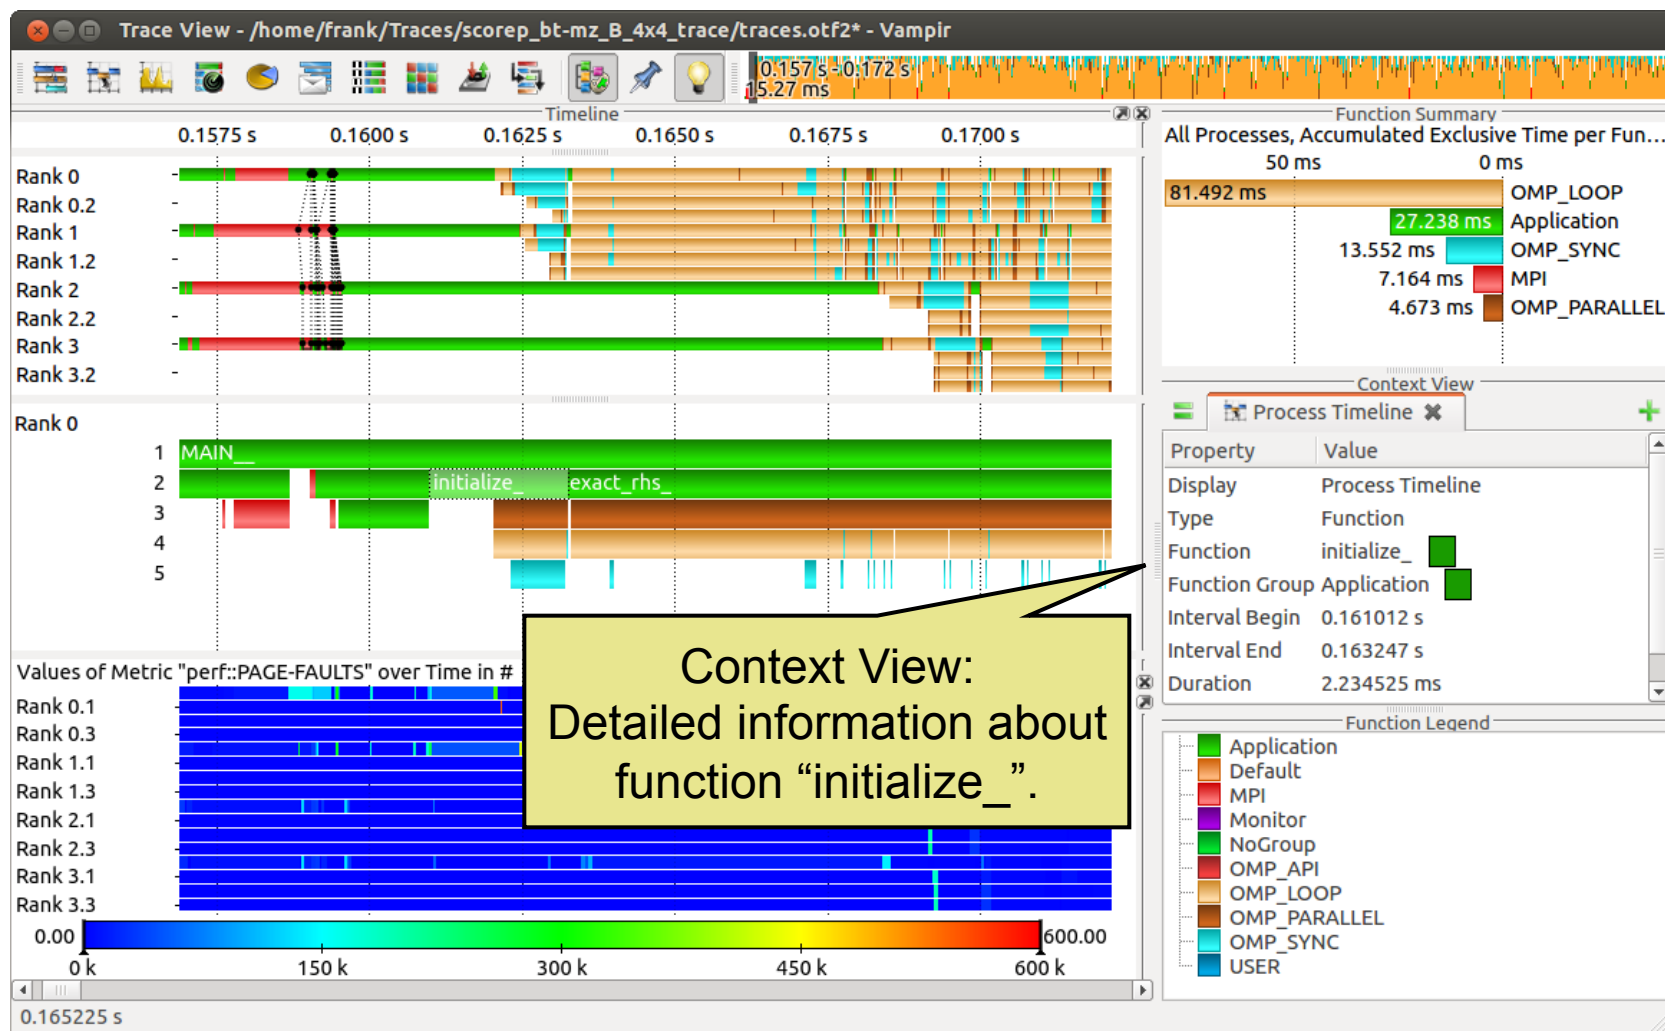




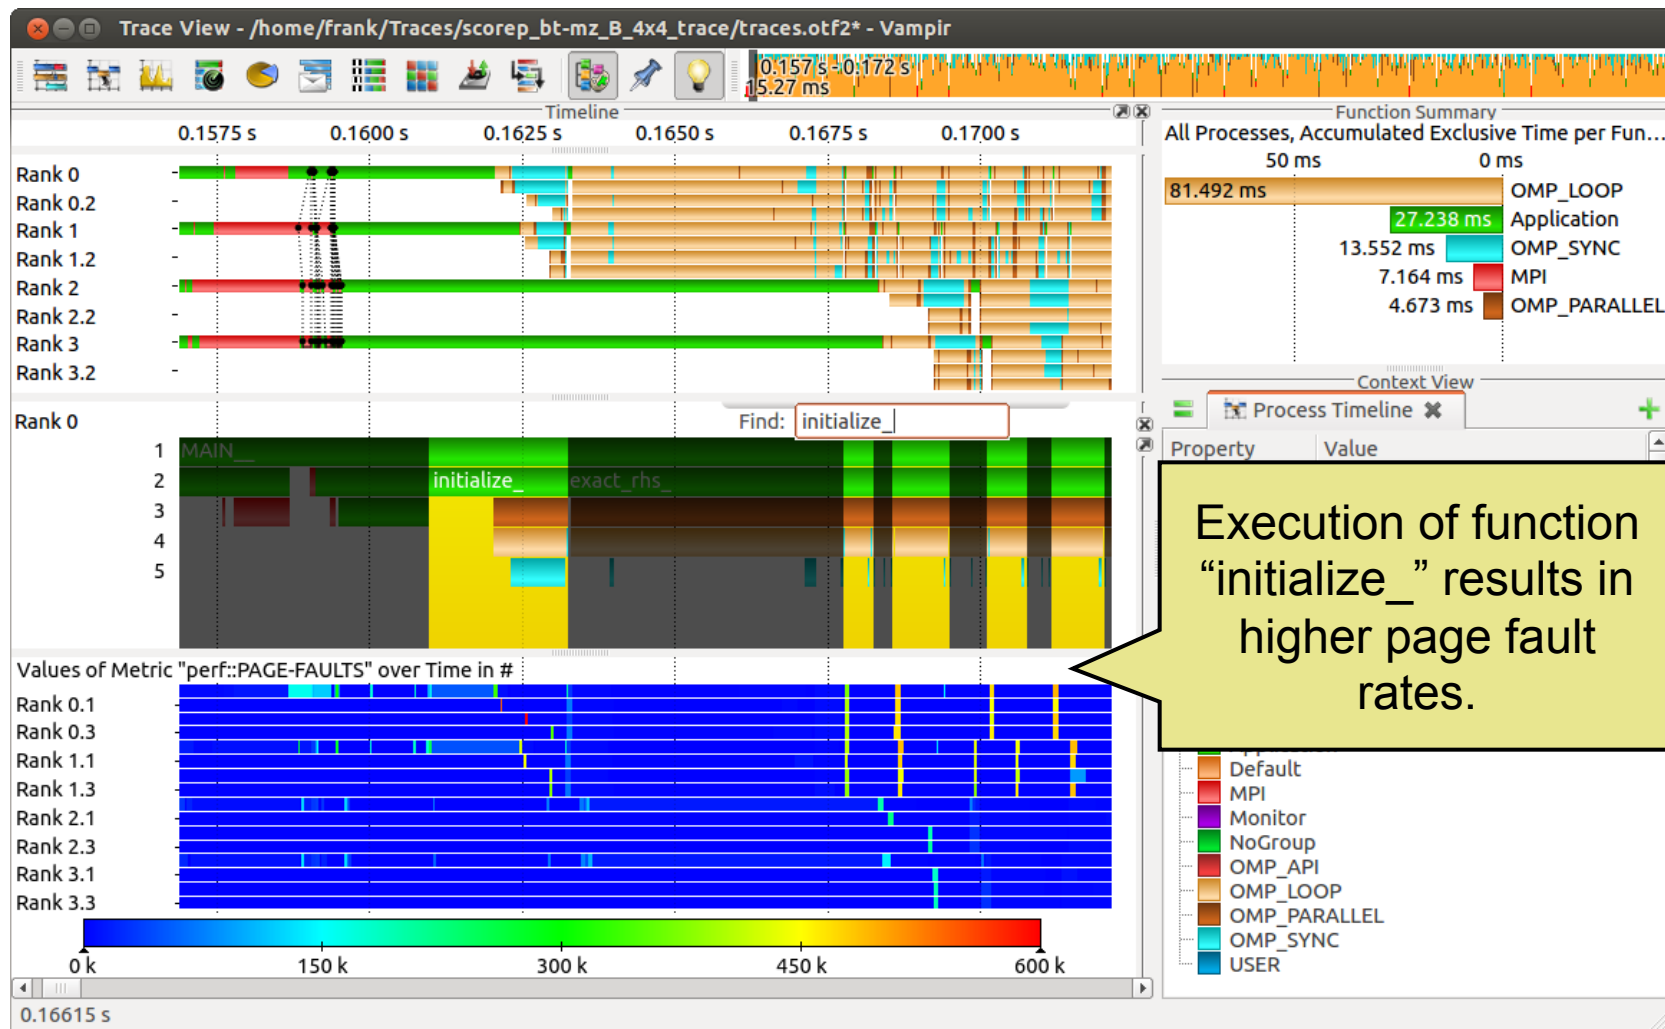
Performance Radar



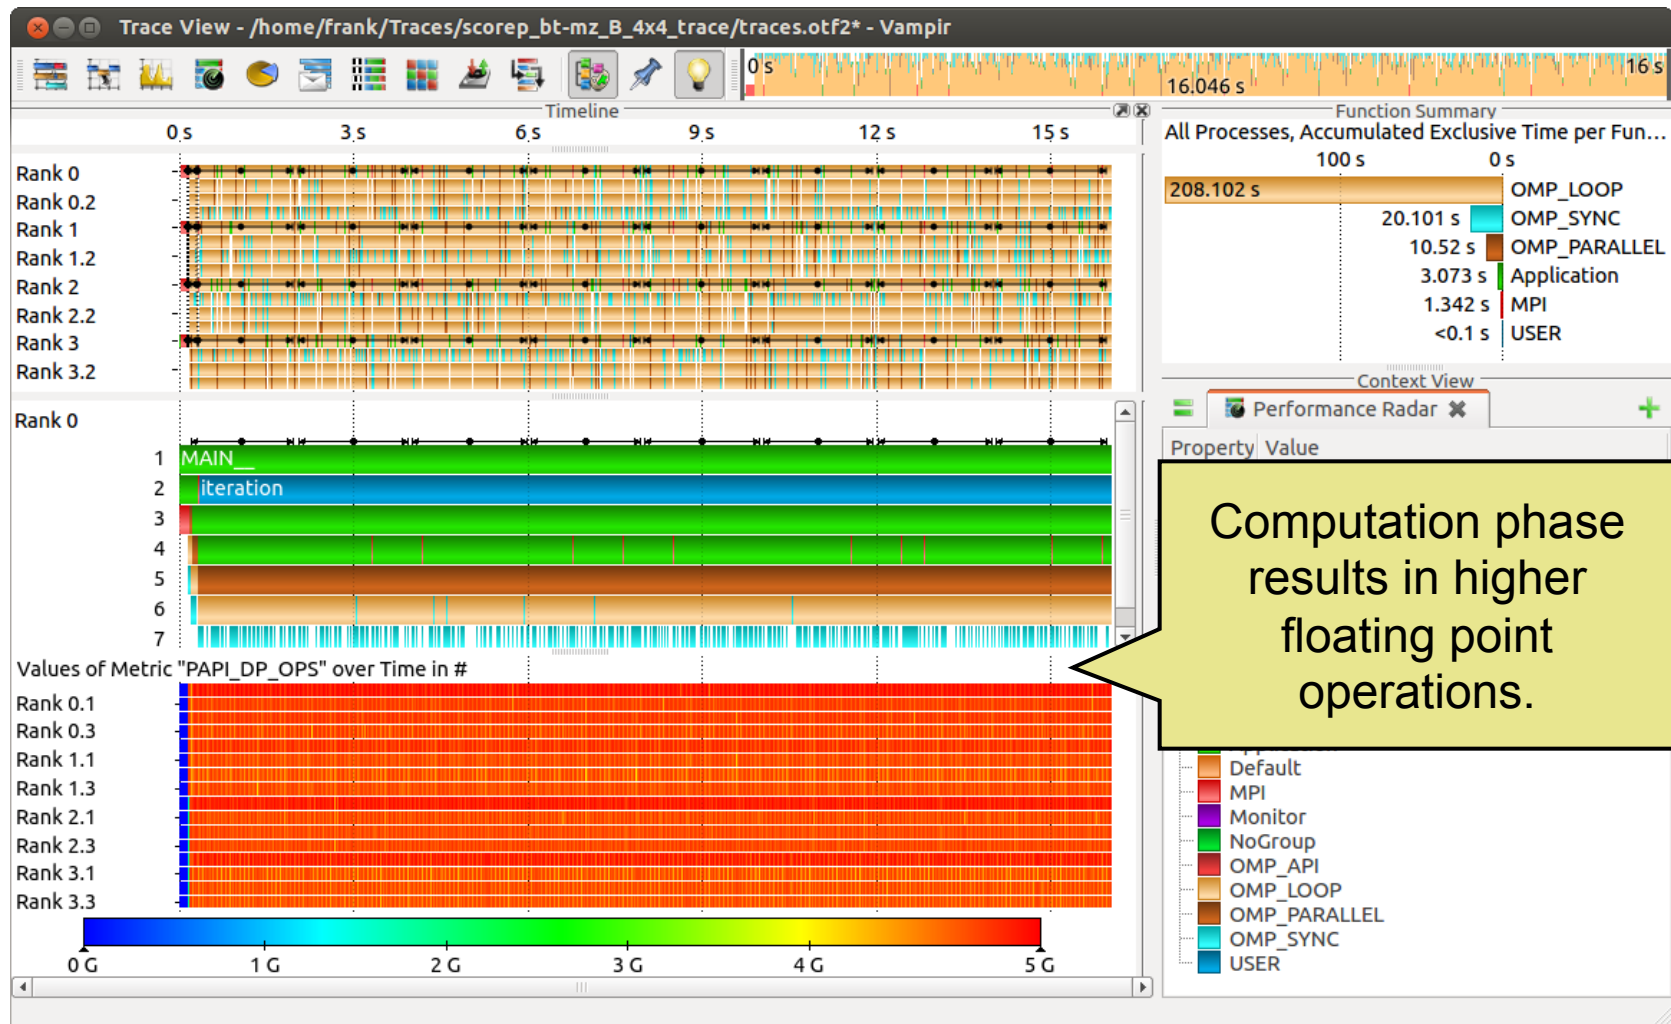
Zoom in: Initialisation Phase



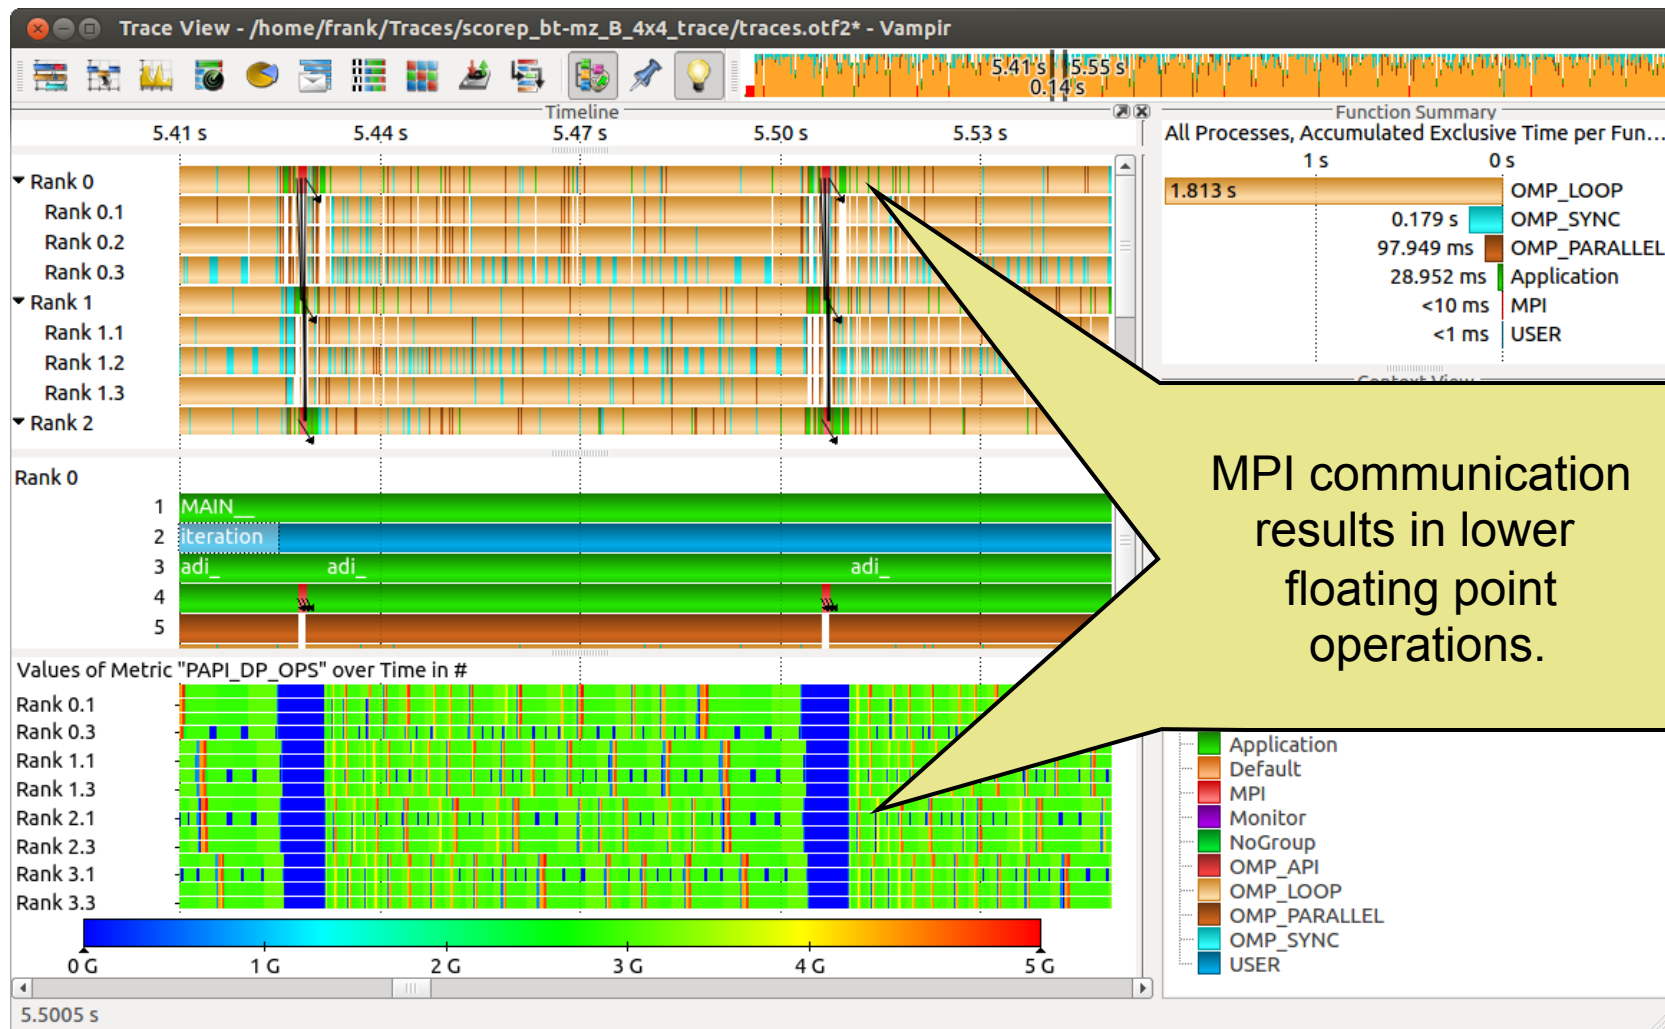
Feature: Find Function



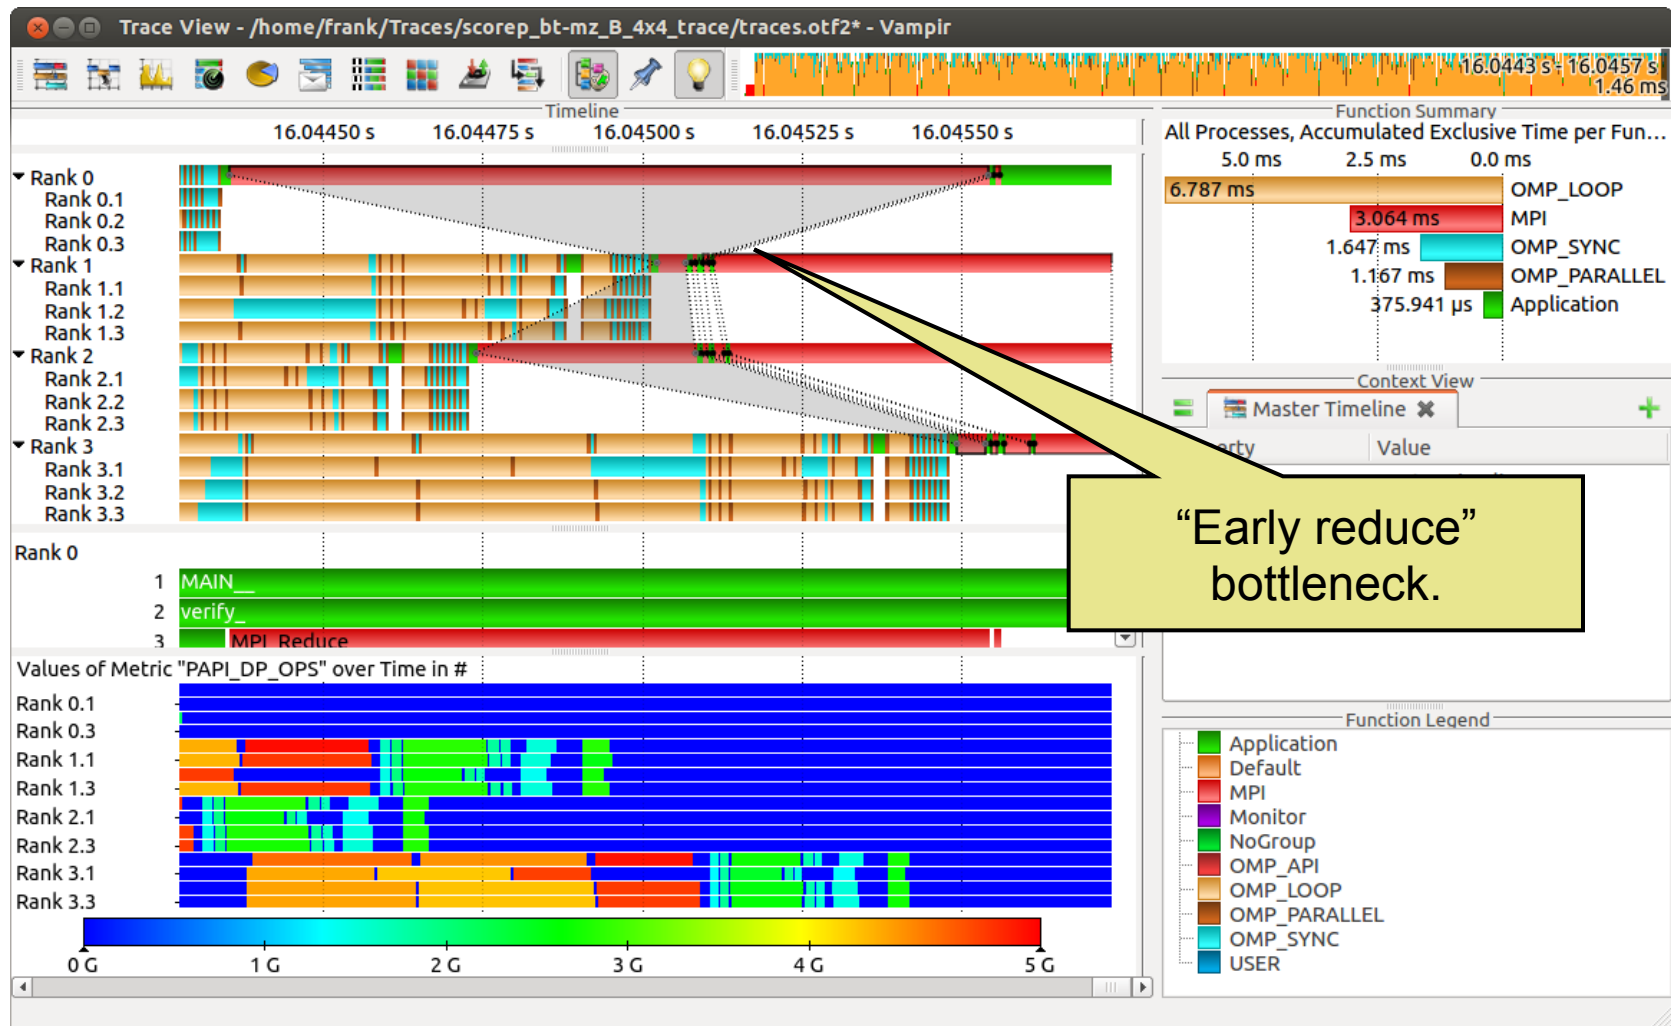
Computation Phase



Zoom in: Computation Phase

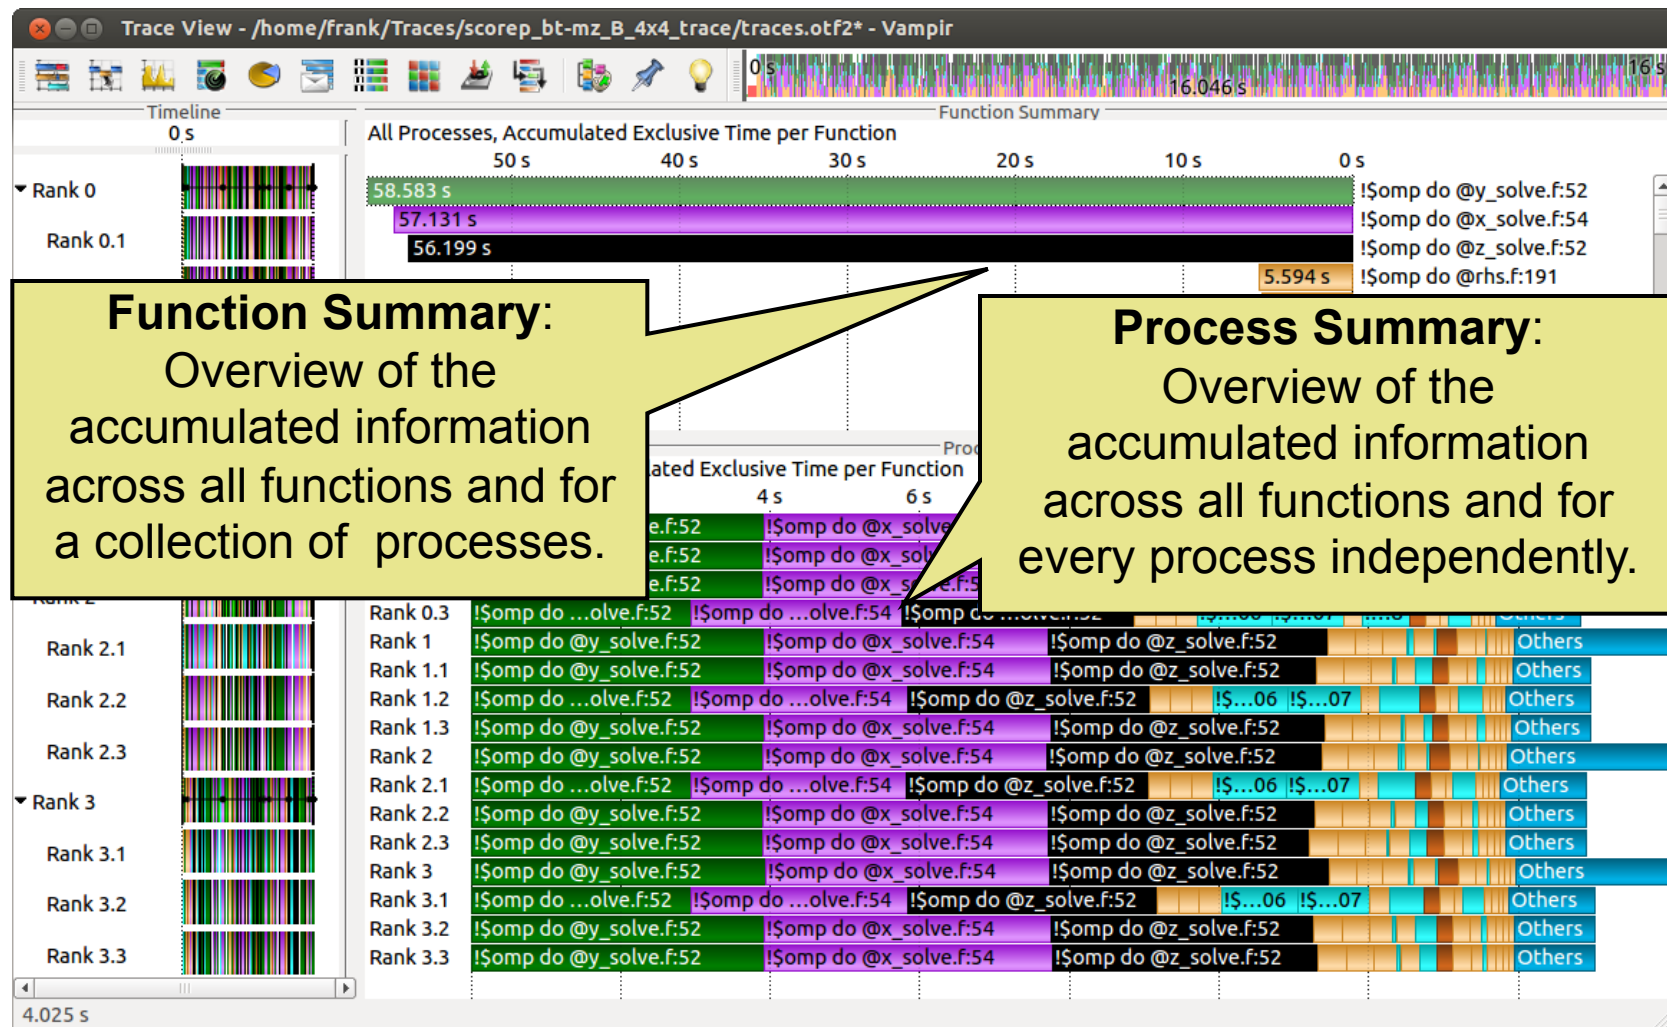


Zoom in: Finalisation Phase



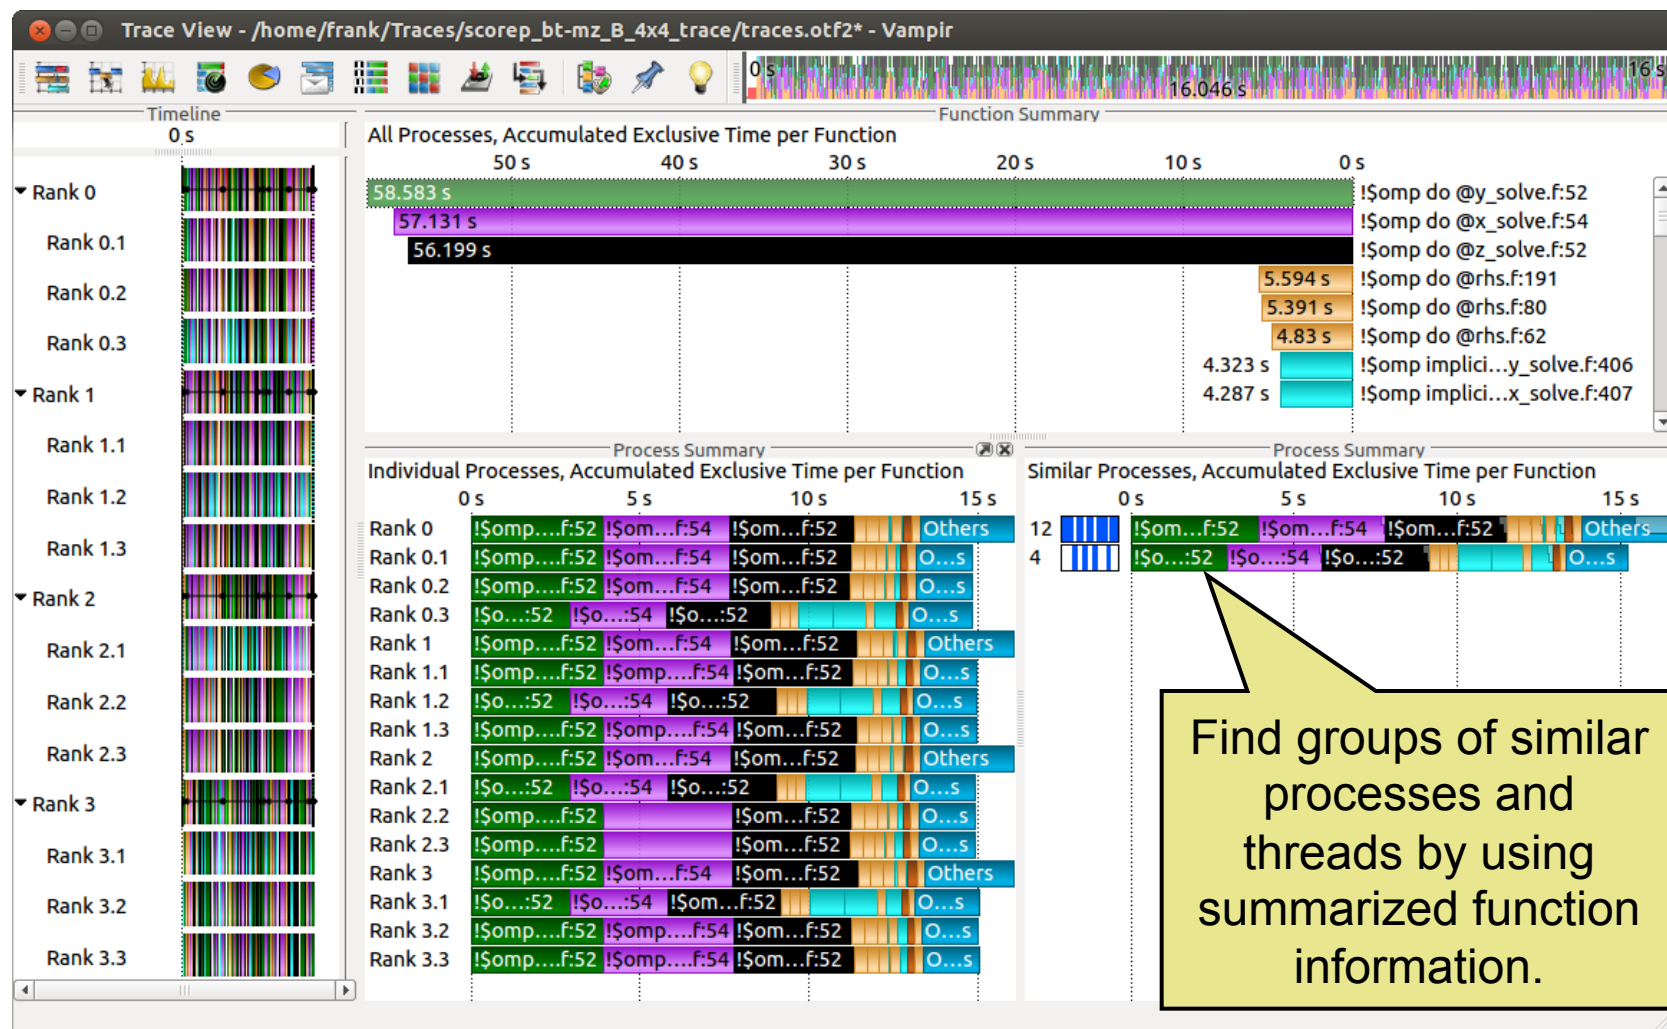


Process Summary






Process Summary



VI-HPS

SOFTWARE



FAST SOLUTIONS


- ☒ PAPI_L1_DCM
- ☒ PAPI_L1_ICM
- ☐ PAPI_L2_DCM
- ☒ PAPI_L2_ICM
- ☒ PAPI_L2_TCM
- ☐ PAPI_L2_TCM

PRODUCTIVITY


Use Cases and Summary

- COSMO-SPECS a coupling of:
 - Weather forecast model
 - Detailed cloud microphysics scheme


COSMO:
Approximation of
cloud droplet size



Cloud droplets



COSMO-SPECS:
Bin-wise discretization
of cloud droplet size



Developer observation:

Runtime per iteration increases over time, why?

Case Study (2)

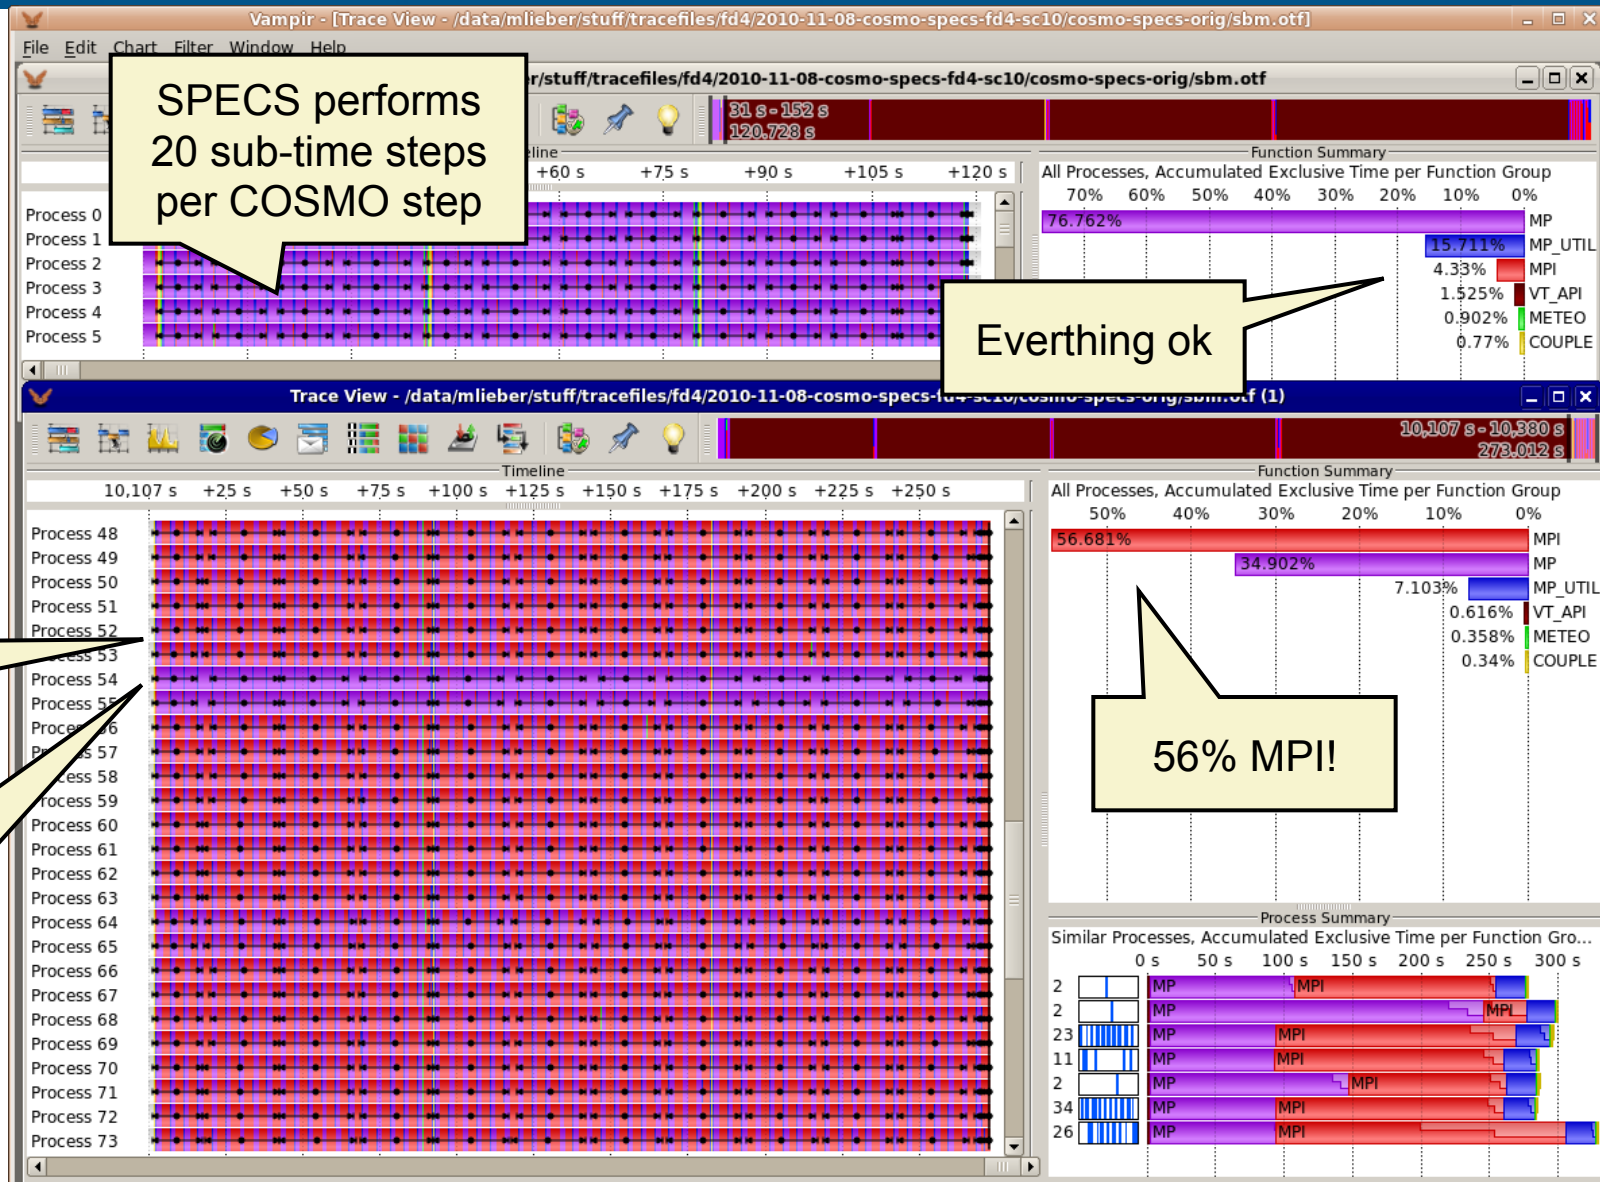
First 3 time steps of COSMO-SPECS run

SPECS performs 20 sub-time steps per COSMO step


Last 3 time steps of COSMO-SPECS run

Heavy load imbalance

Cloud grows in grid cells of these MPI ranks




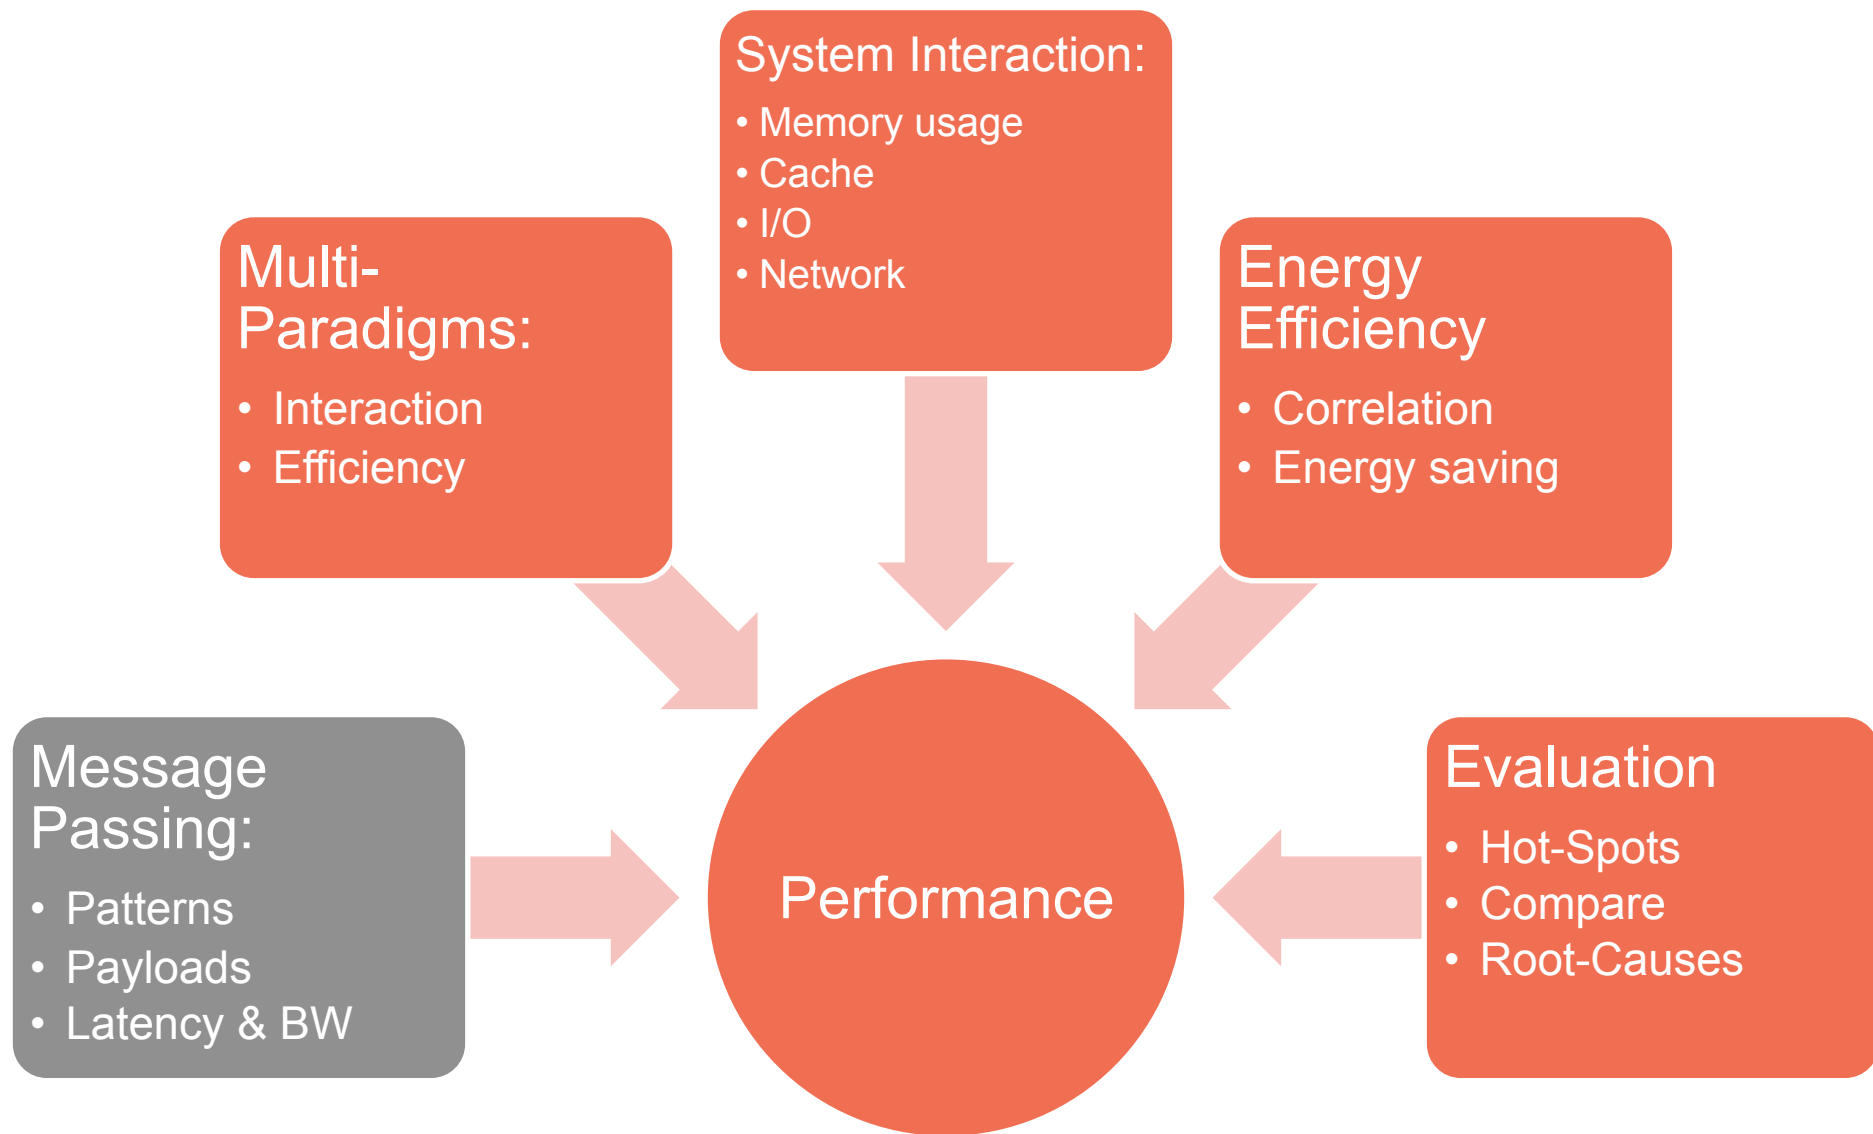
- Domain decomposition:
 - 2D (horizontal) decomposition into $M \times N$ processor domains
 - No dynamic load balancing

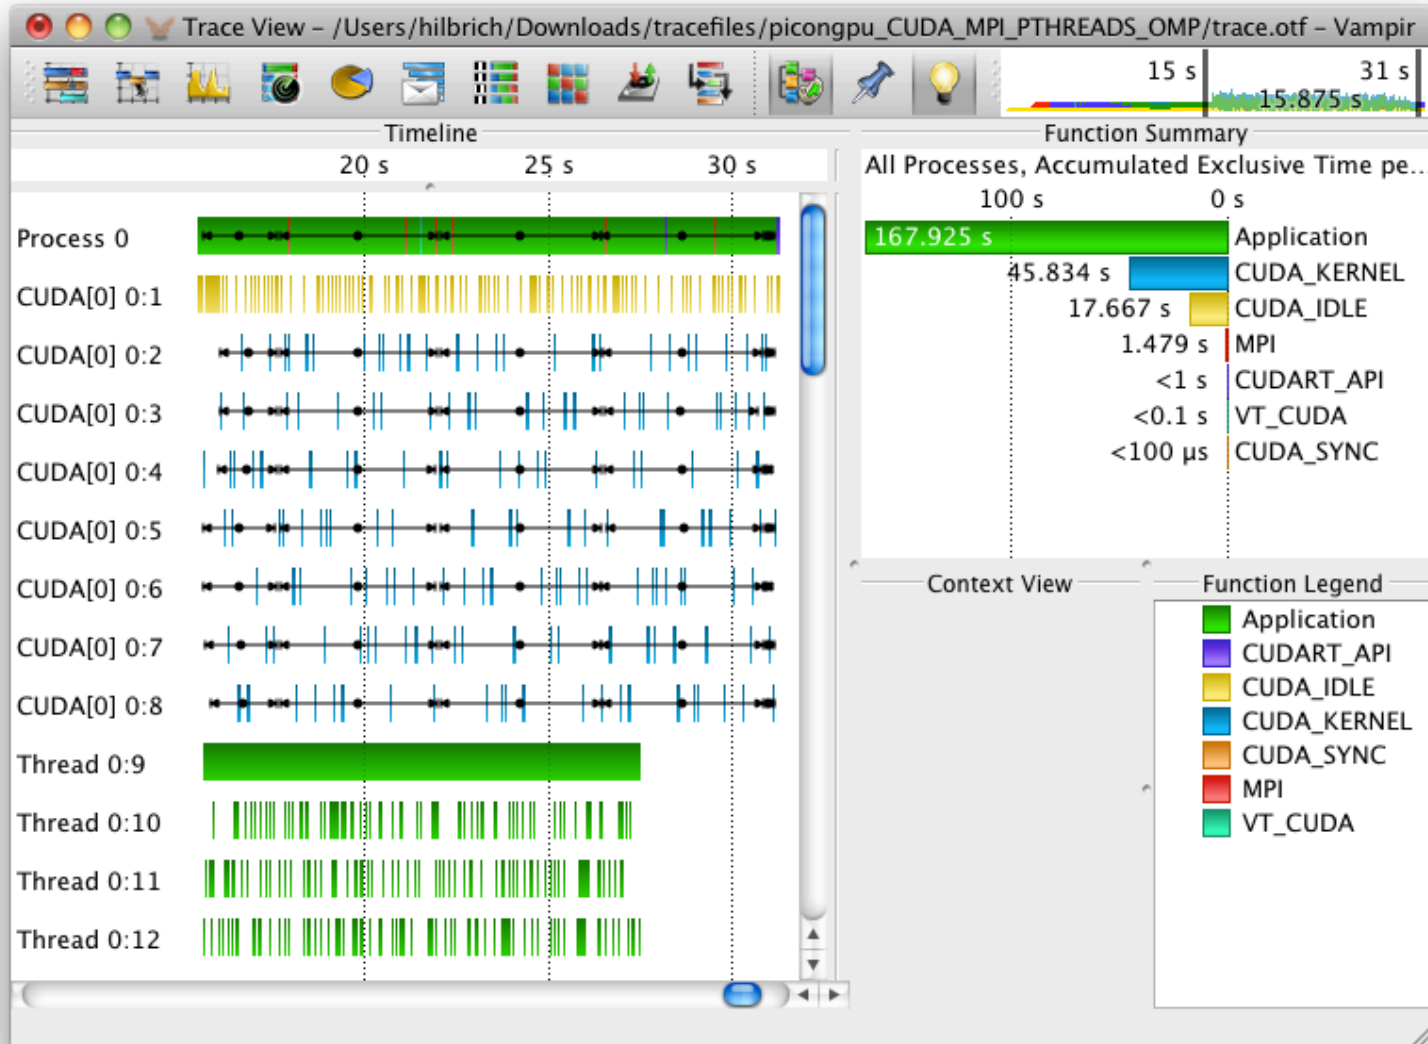


Lesson learned:

Domain decomposition for detailed cloud microphysics must adapt to presence of clouds.





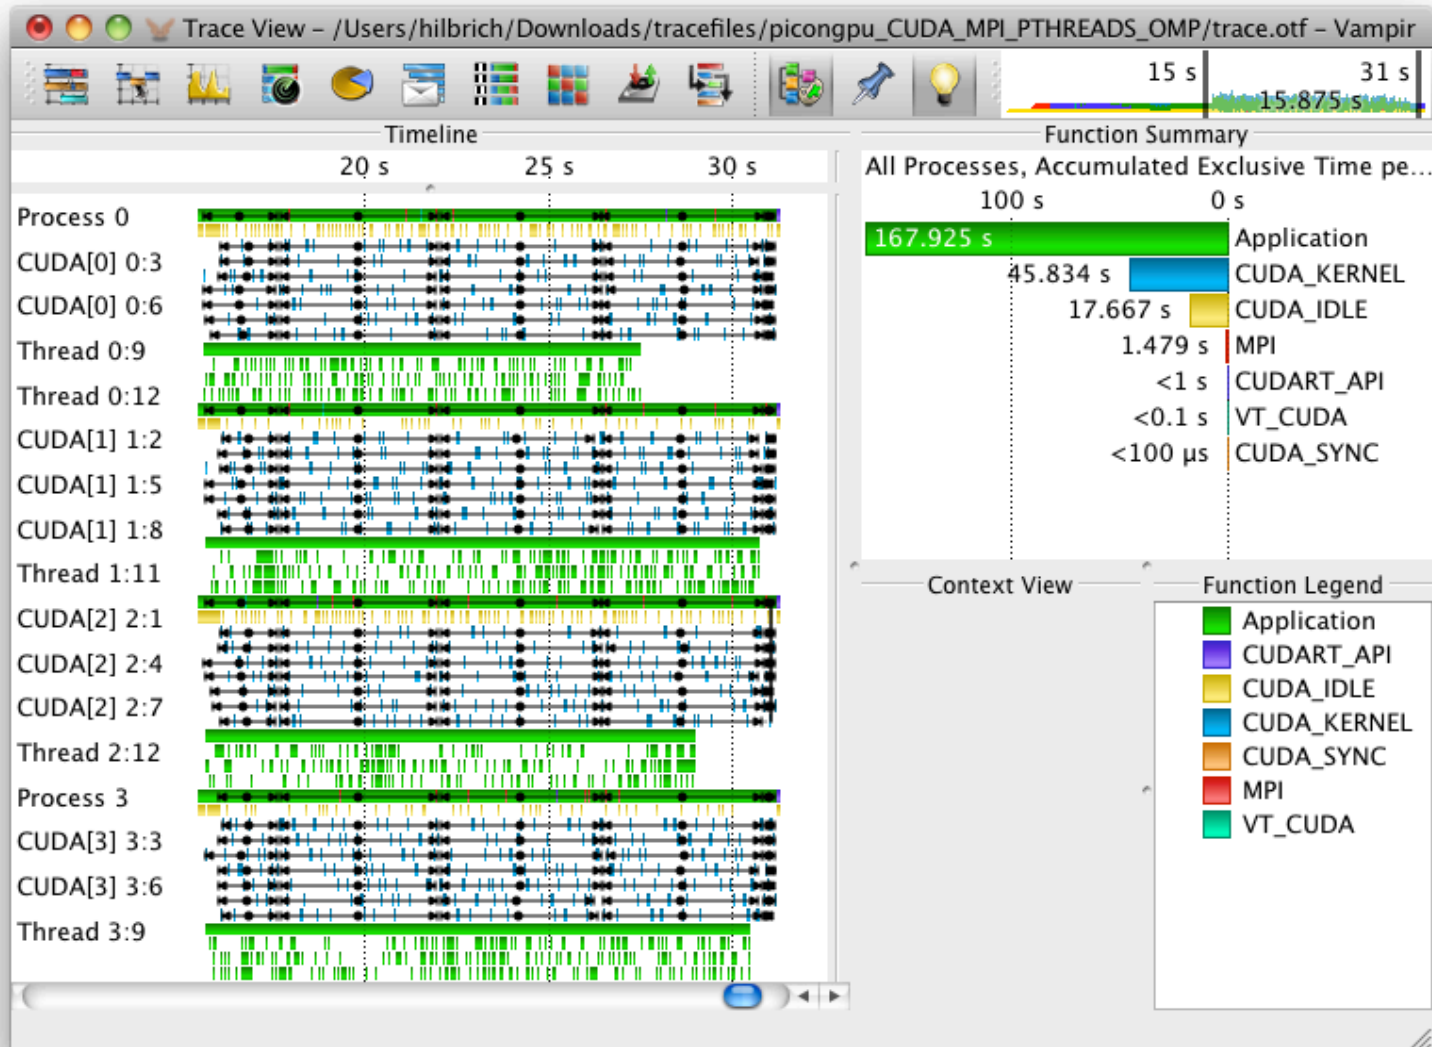


Paradigms:

- MPI
- Pthreads
- CUDA

View:

- Initial

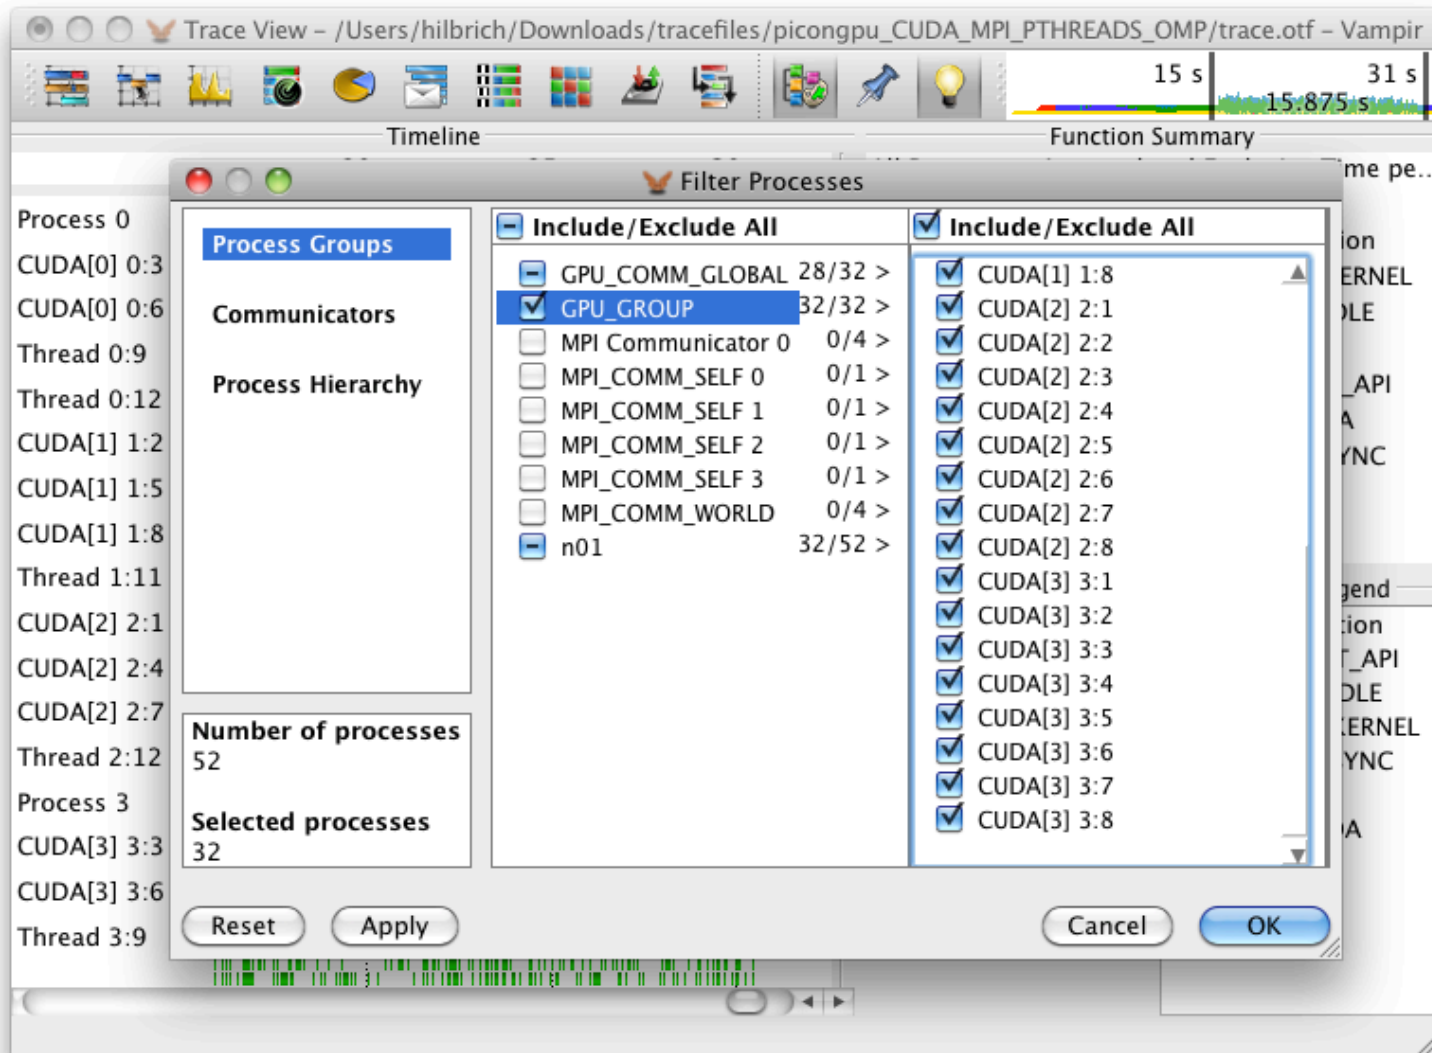


Paradigms:

- MPI
- Pthreads
- CUDA

View:

- All

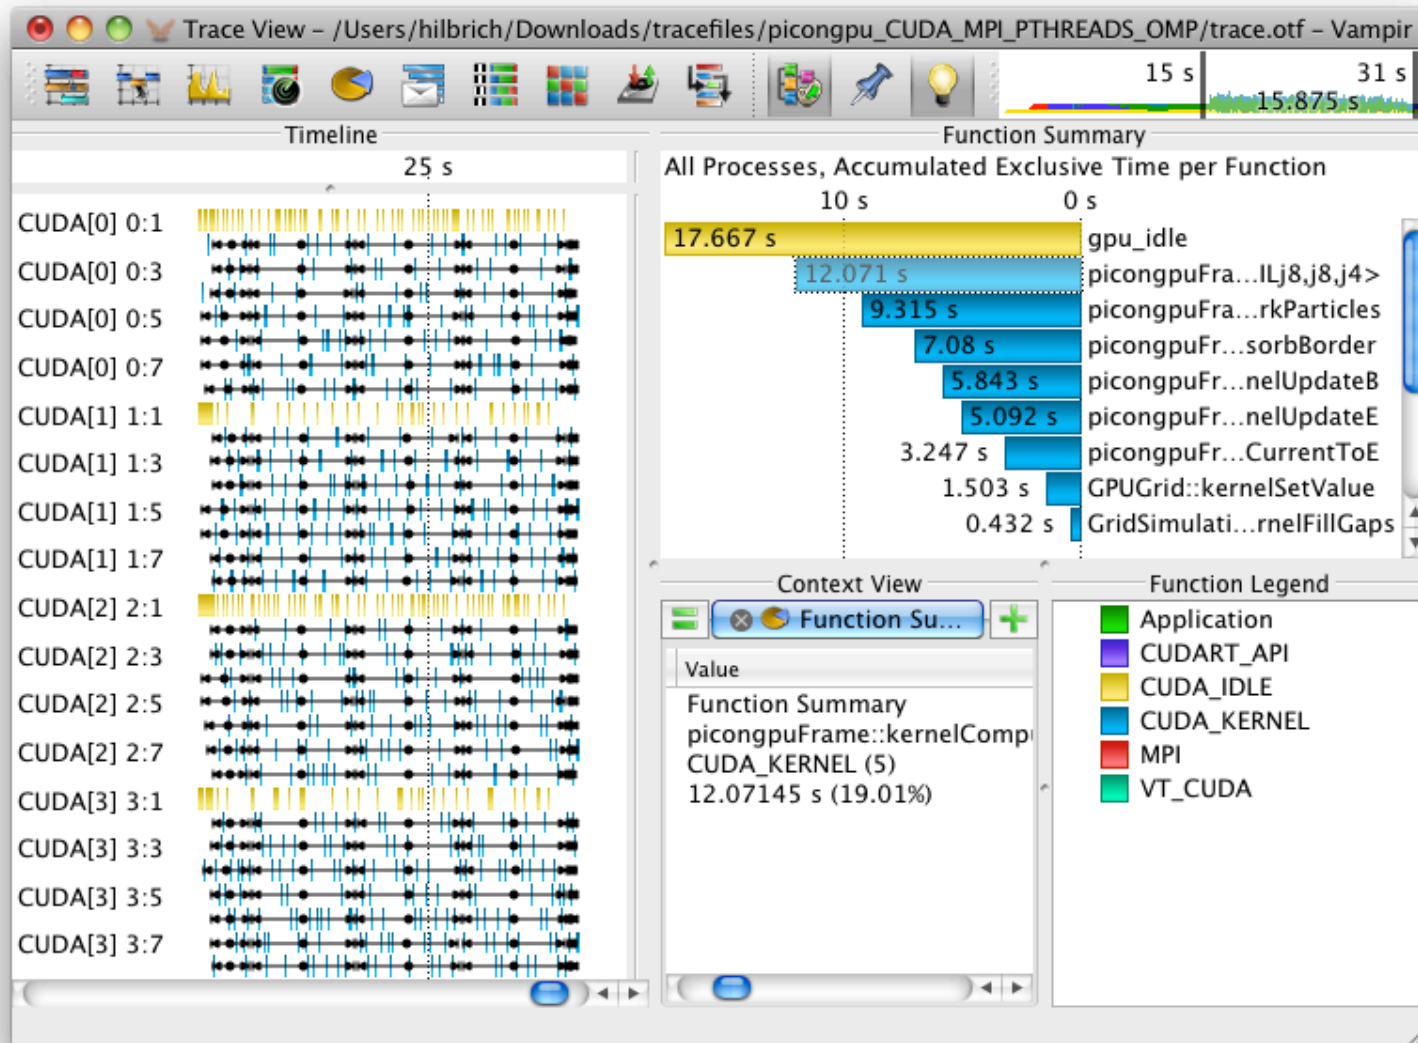


Paradigms:

- MPI
- Pthreads
- CUDA

View:

- Filter

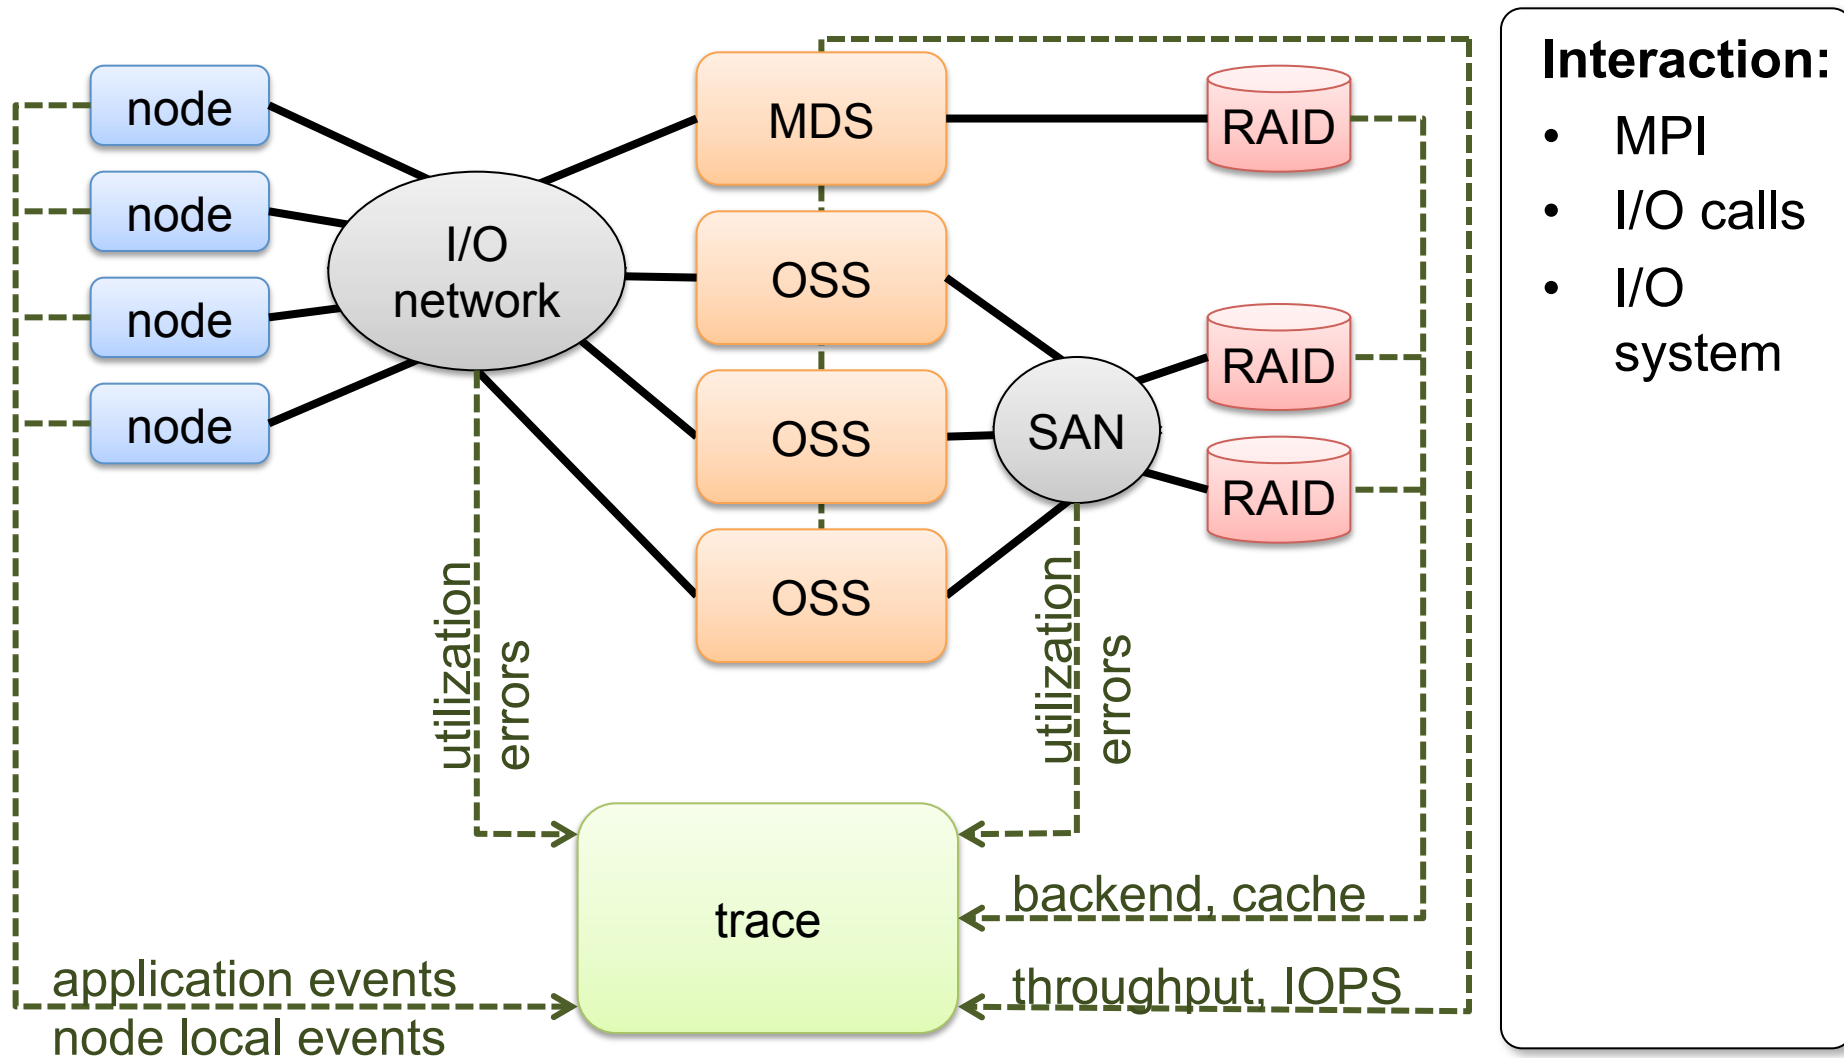


Paradigms:

- MPI
- Pthreads
- CUDA

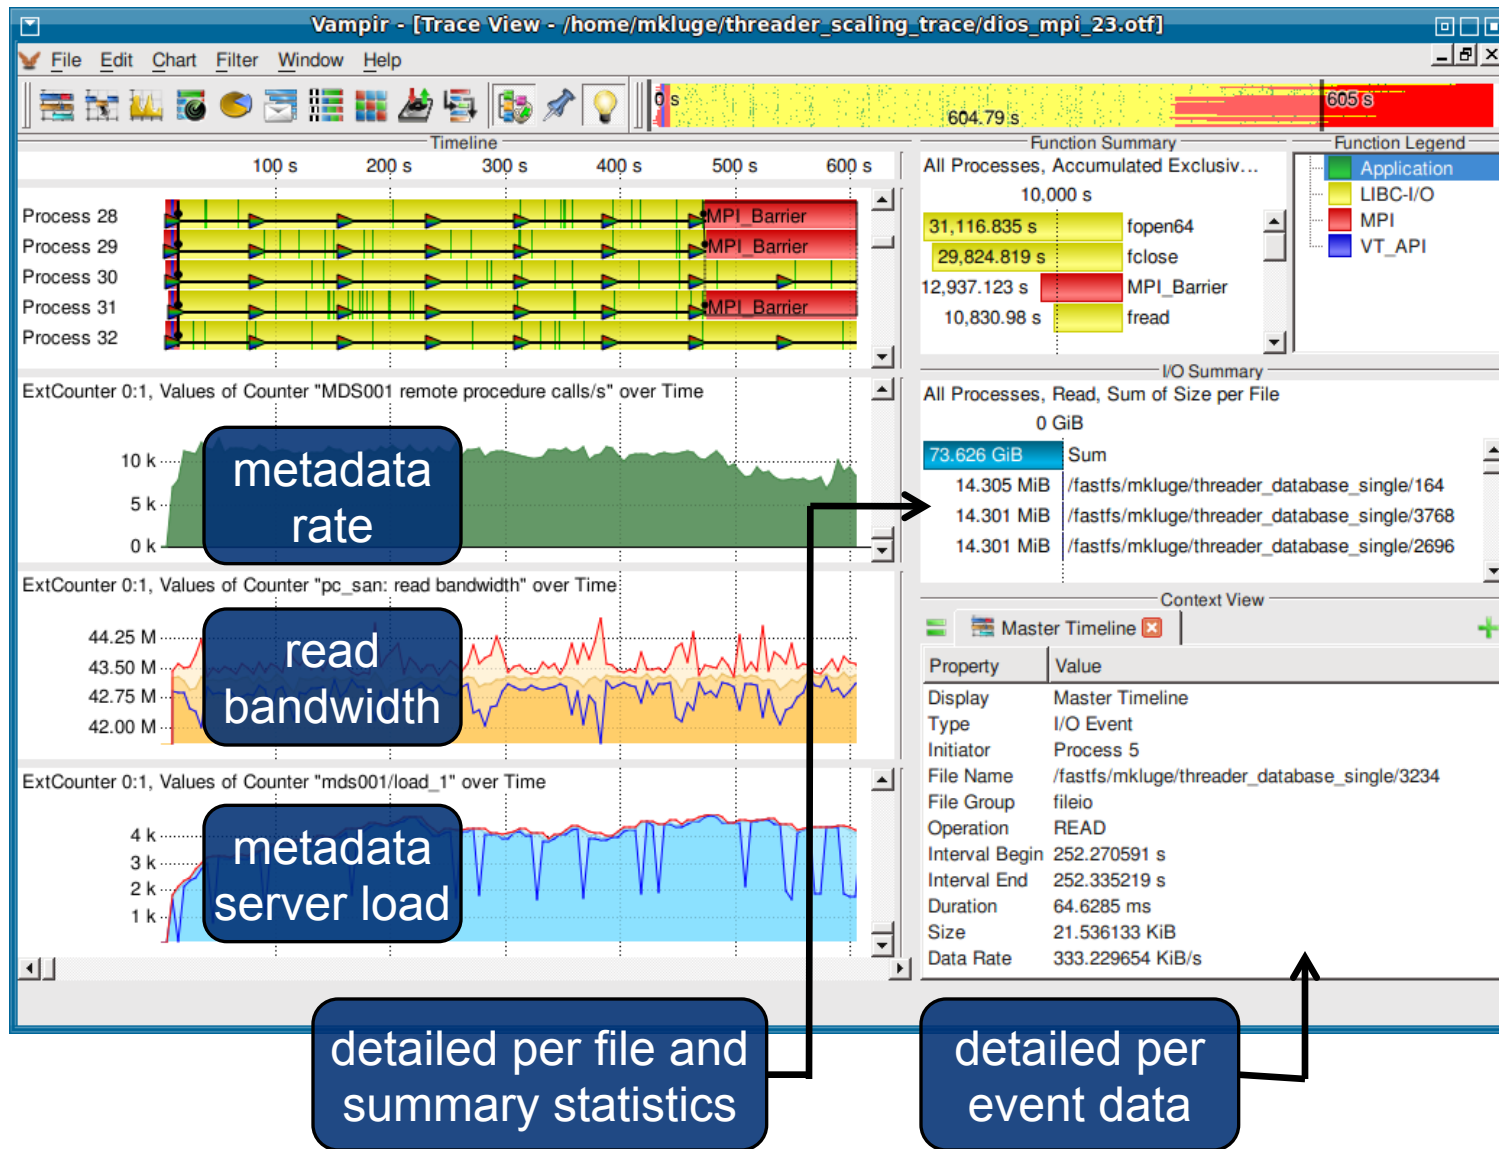
View:

- Just CUDA

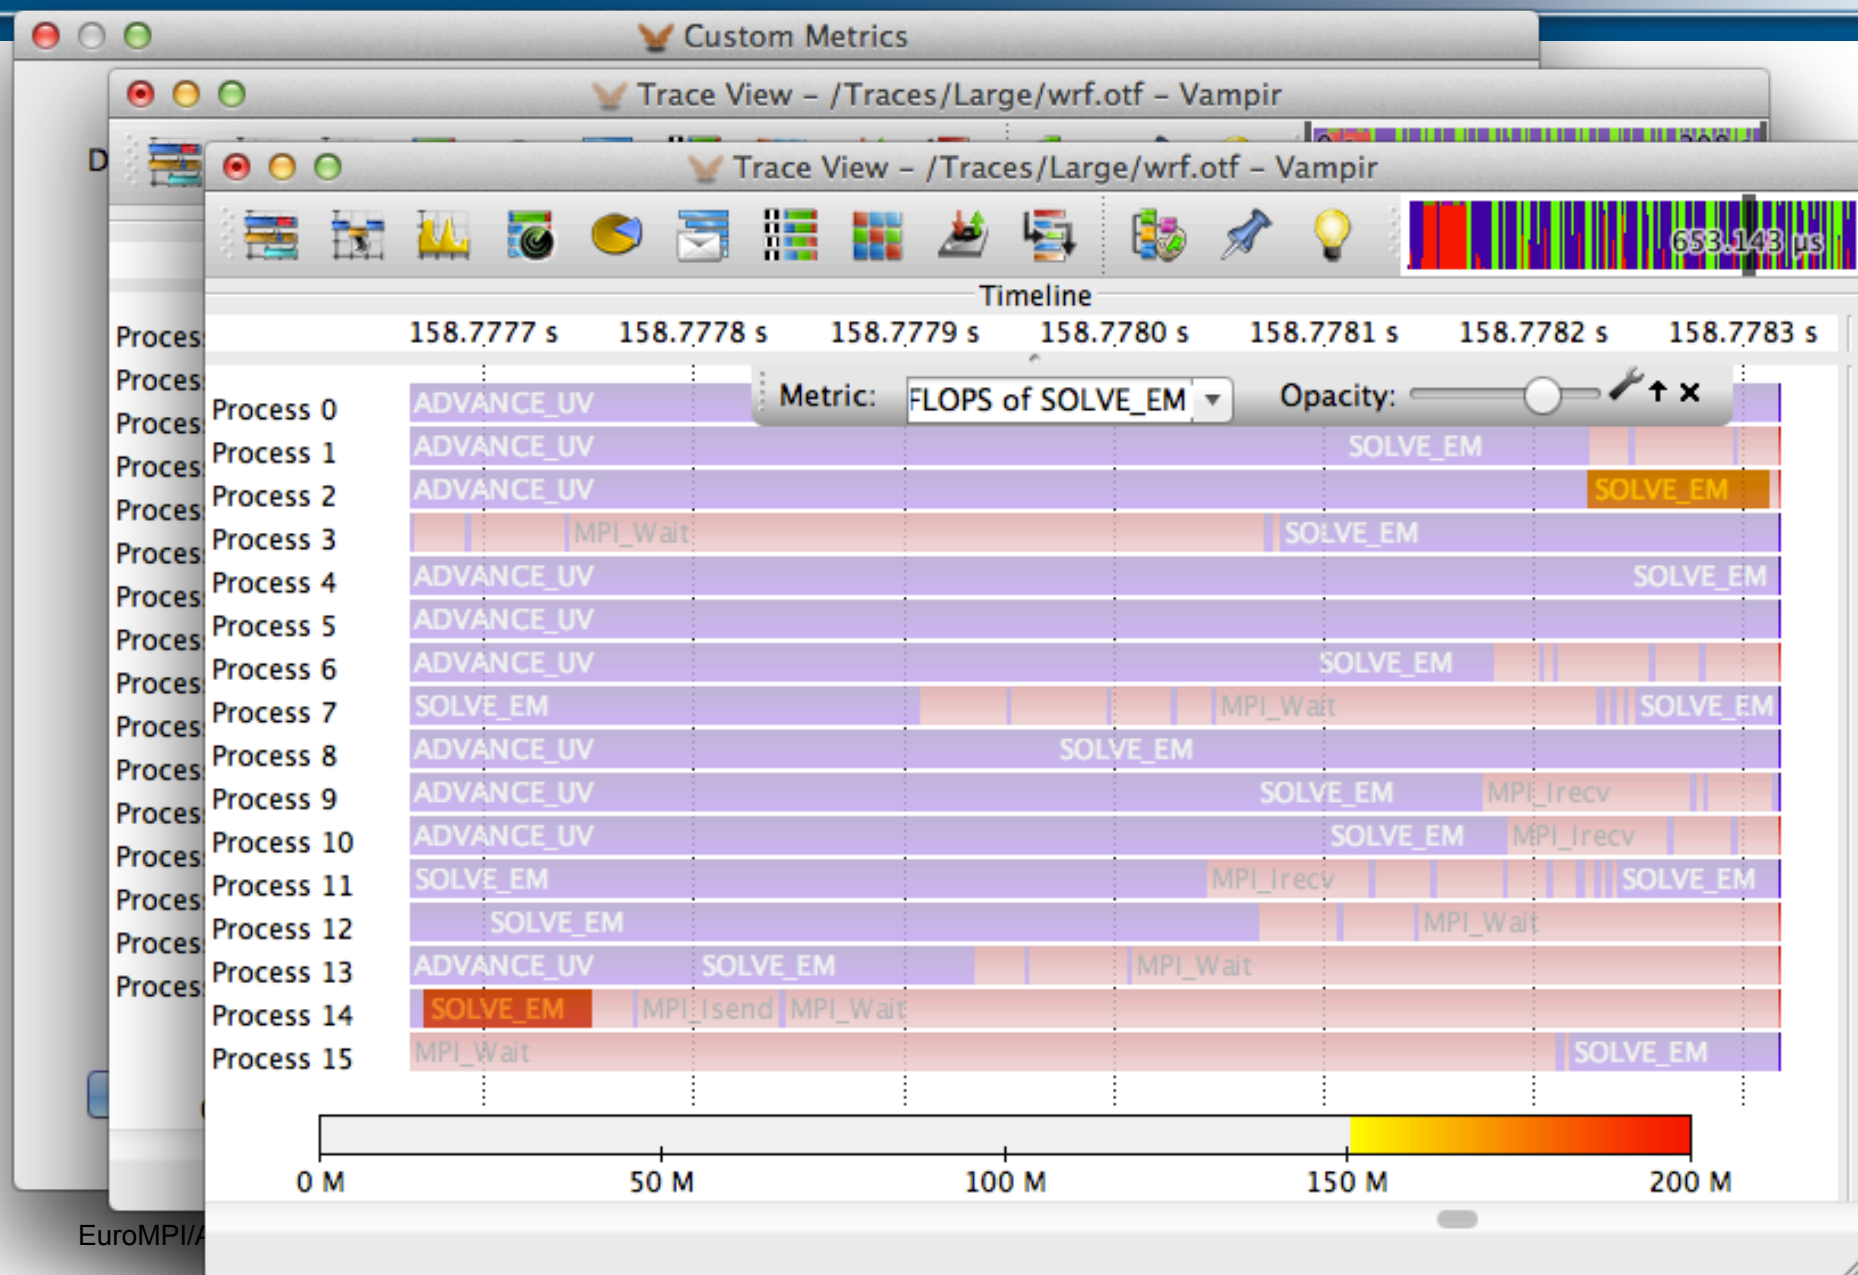


Interaction:

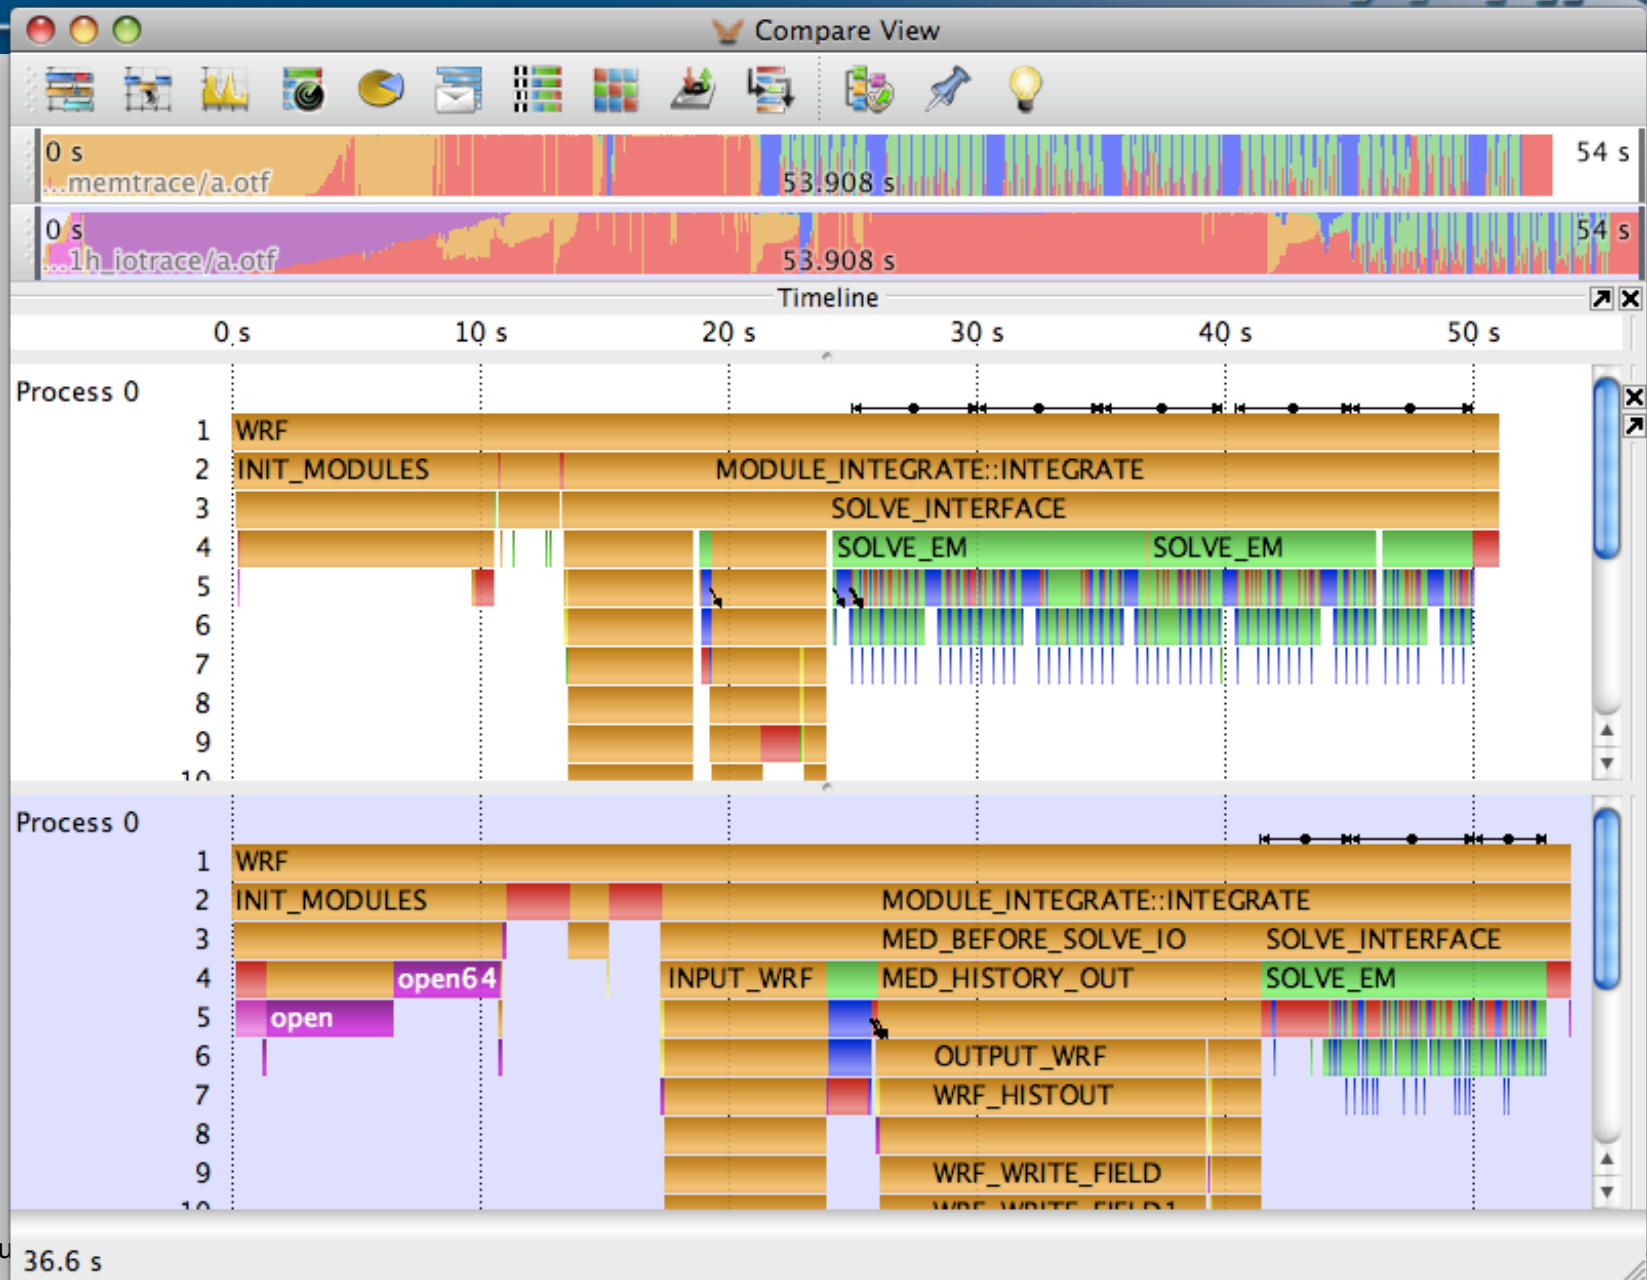
- MPI
- I/O calls
- I/O system



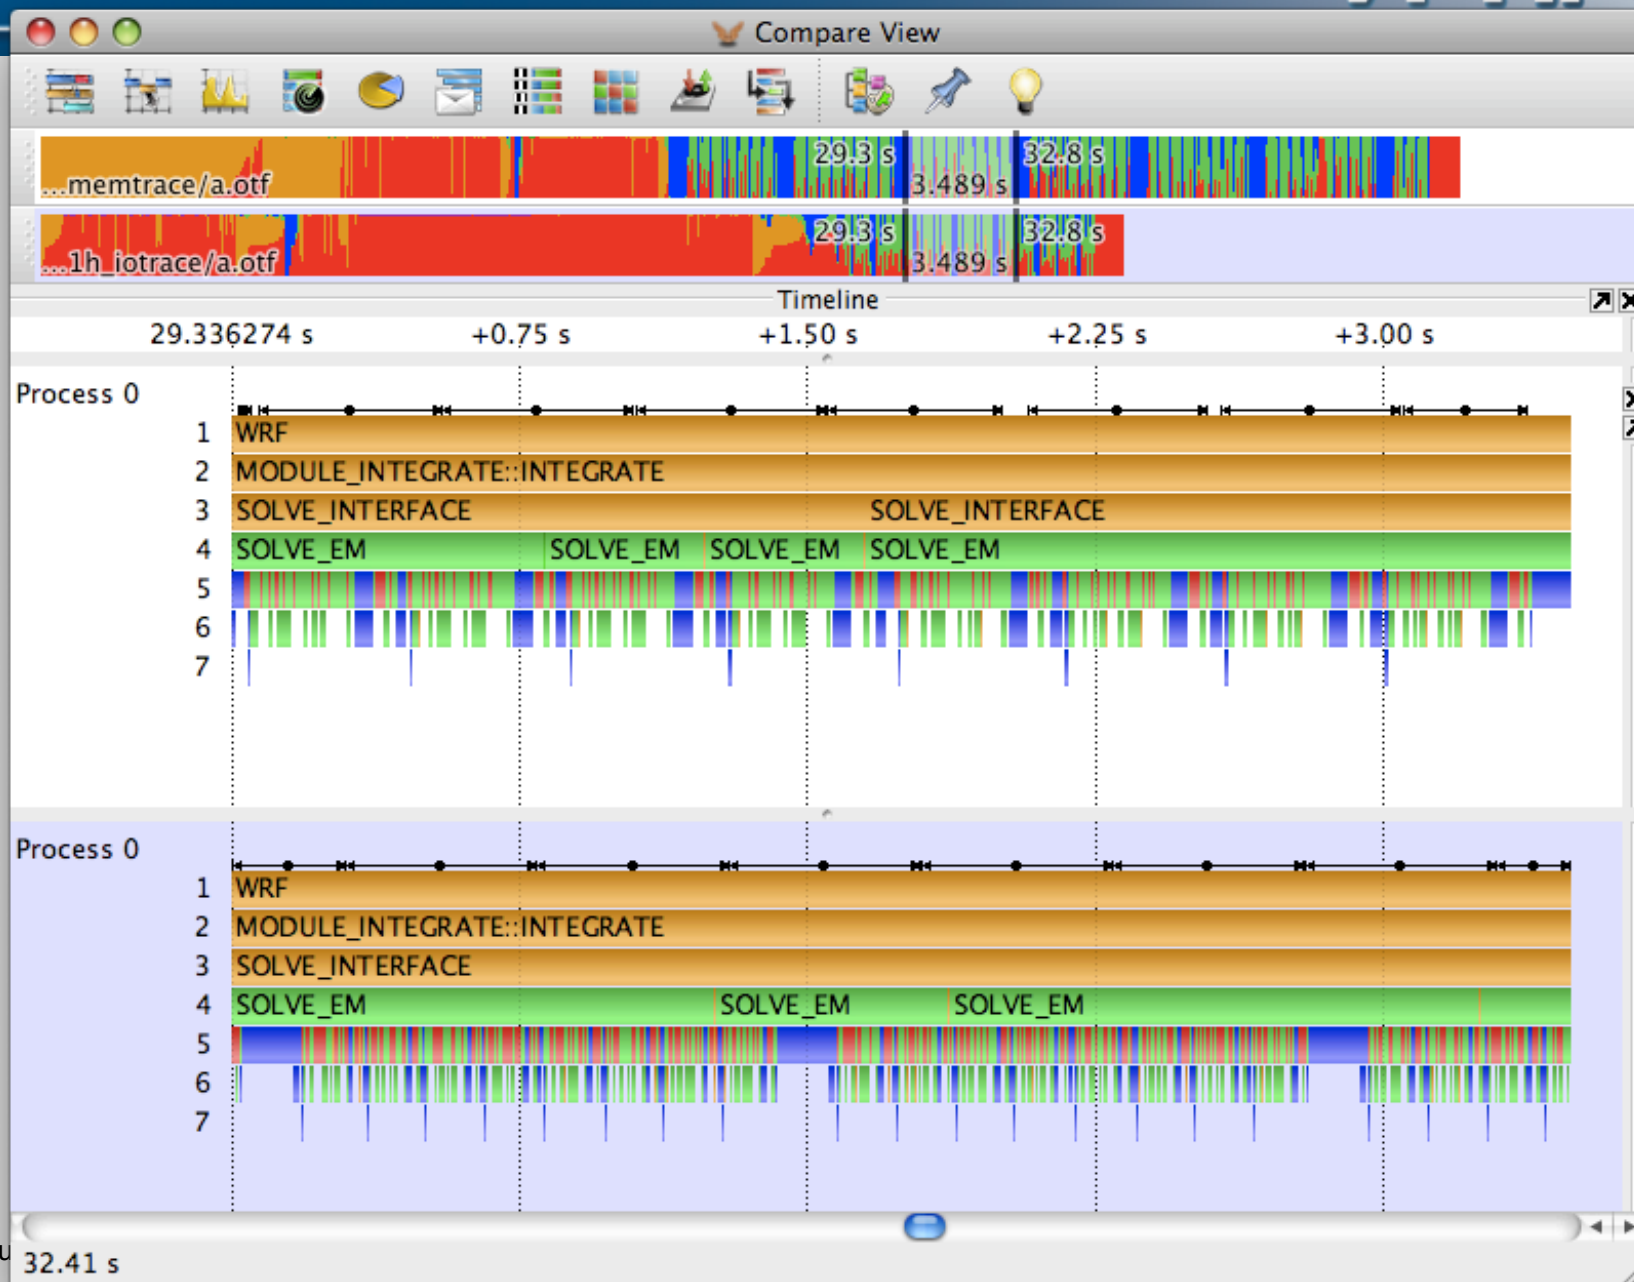
Advanced – Derived Counters



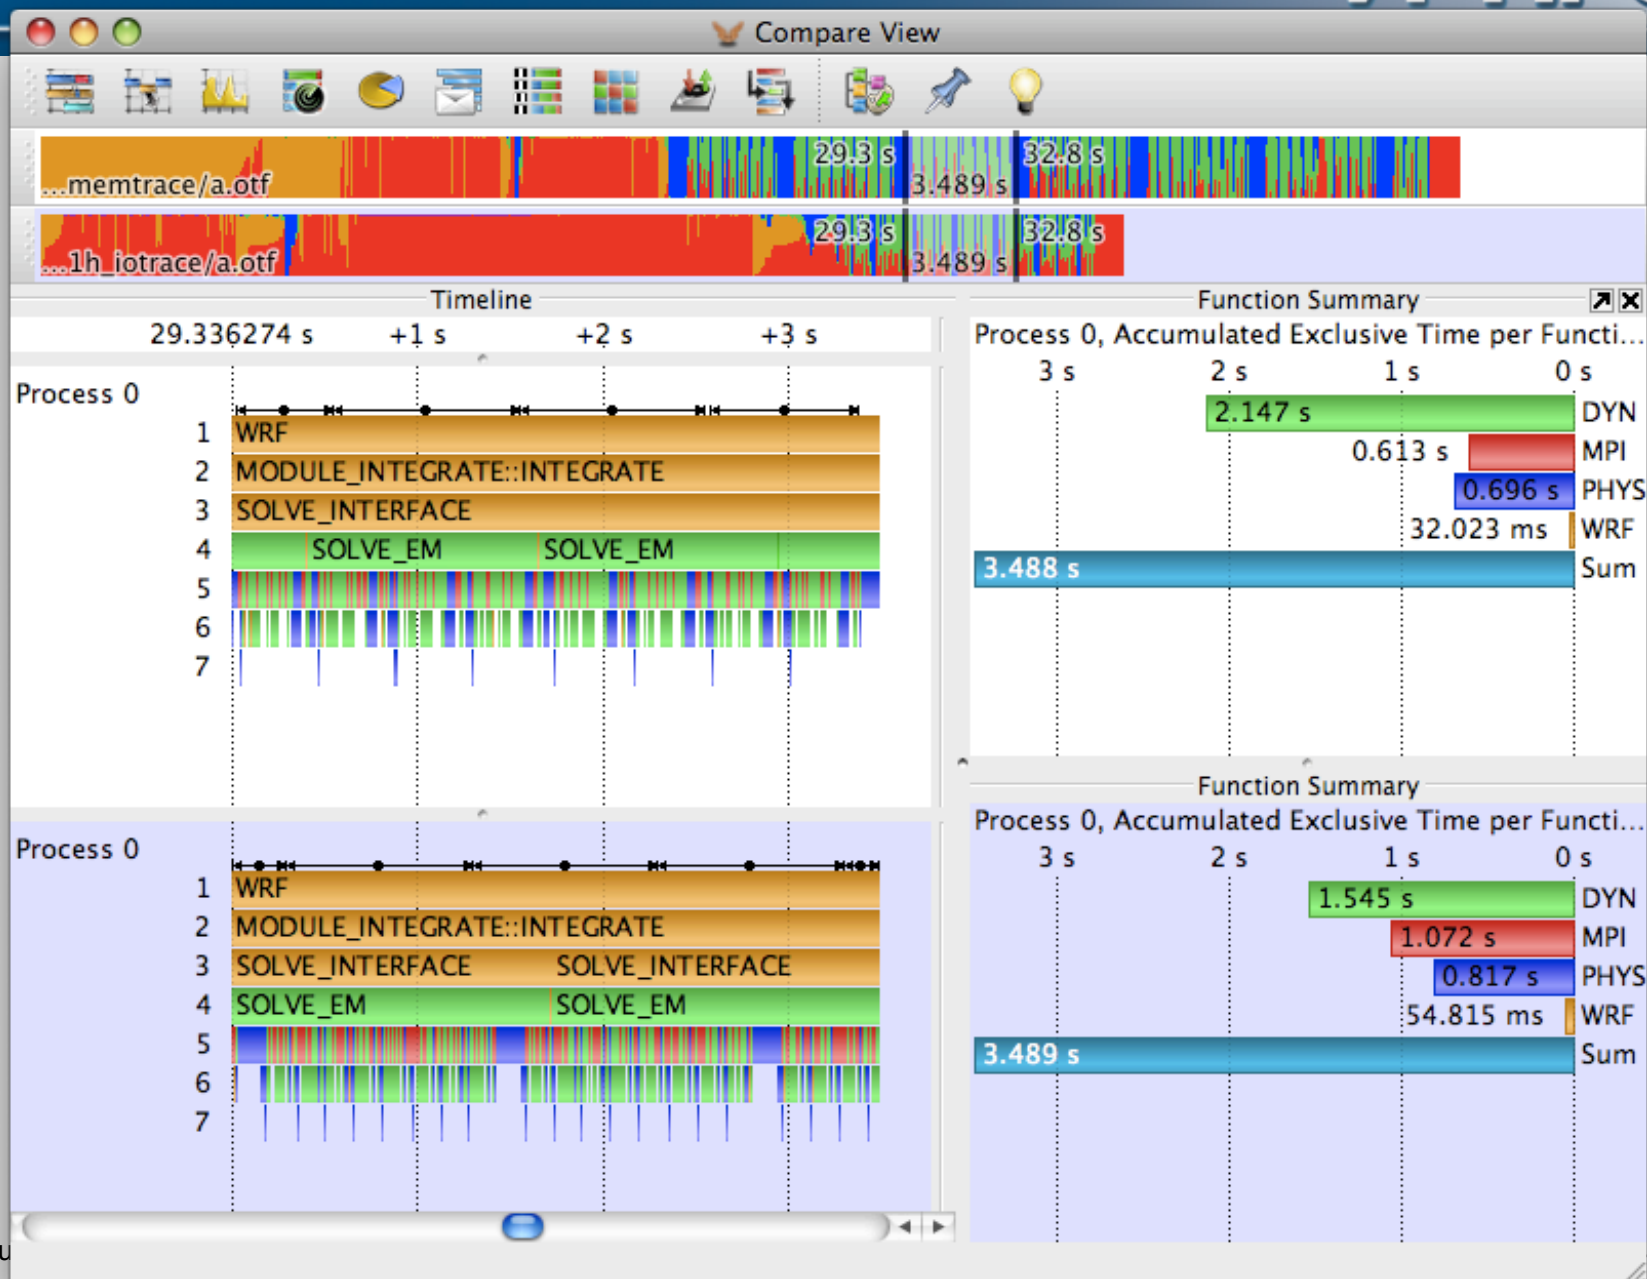
Use Cases – Comparison View




Use Cases – Comparison View (2)



Use Cases – Comparison View (3)



- Vampir & VampirServer
 - Interactive trace visualization and analysis
 - Intuitive browsing and zooming
 - Scalable to large trace data sizes (20 TByte)
 - Scalable to high parallelism (200000 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