Performance Analysis with Vampir

Bill Williams
TU Dresden

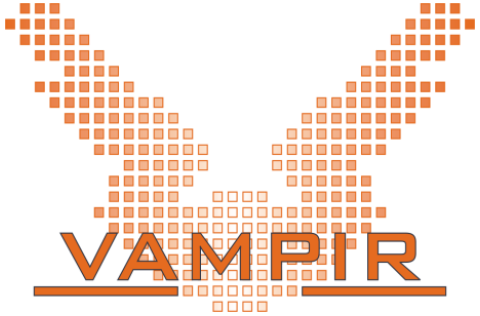

Outline

- **Part I: Welcome to the Vampir Tool Suite**
  - Mission
  - Event Trace Visualization
  - Vampir & VampirServer
  - The Vampir Displays
- **Part II: Vampir Demo**
  - Analyzing load imbalance over time in COSMO

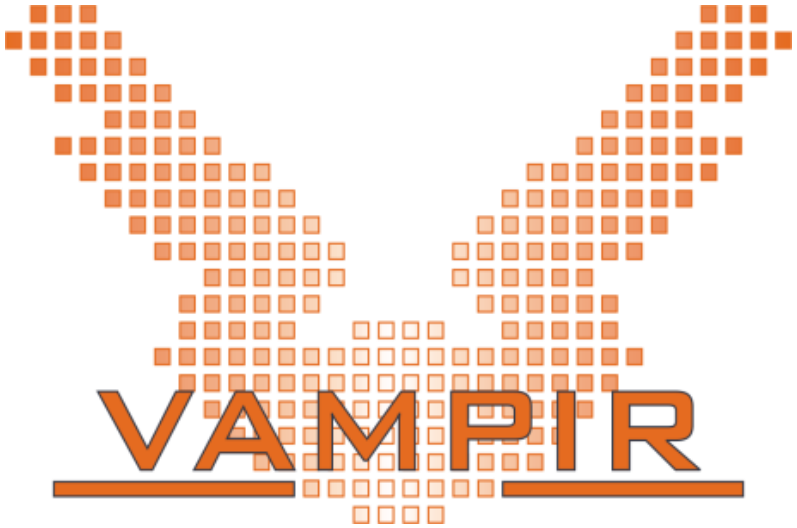

Event Trace Visualization with Vampir

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

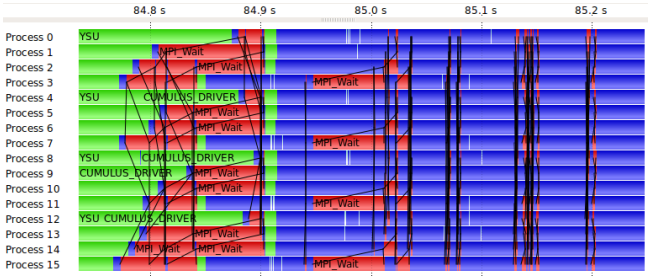

- **Timeline charts**
  - Show application activities and communication along a time axis

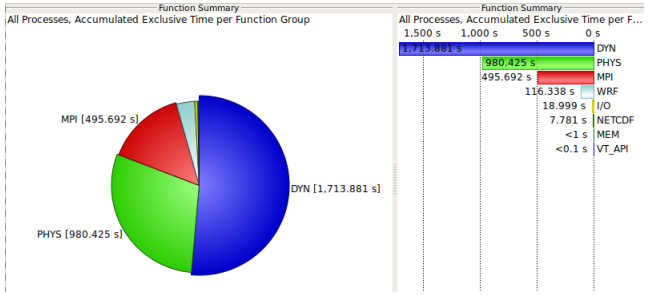

- **Summary charts**
  - Provide quantitative results for the currently selected time interval

When is Vampir the right tool?

- Performance problems depend on information that changes over time or location
- Fine grained information is necessary:
  - What events occur simultaneously?
  - Which messages are being delayed?
  - Which threads are idle and why?

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

```
% vampir
```

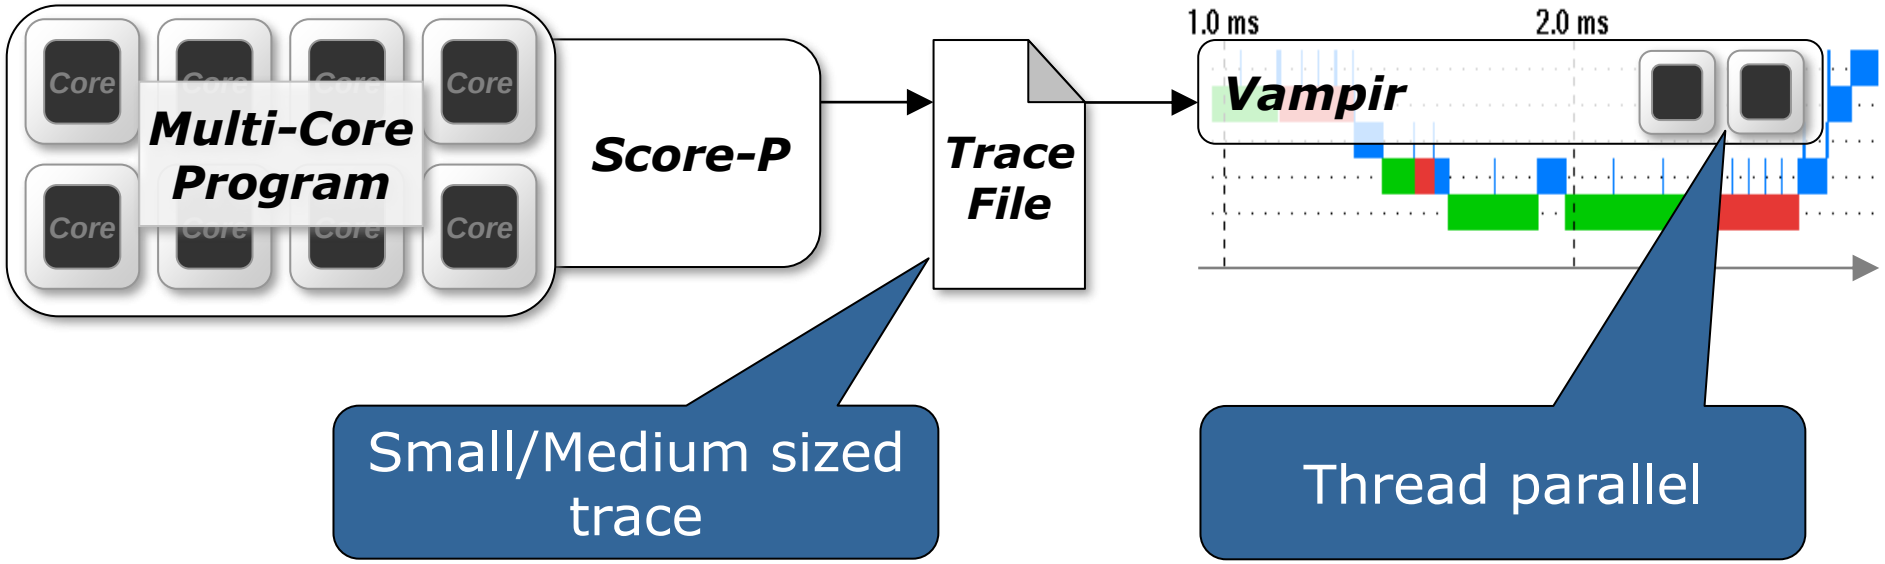

Visualization Modes (2)
On local machine with remote VampirServer

```
% vampirserver start
```

```
% vampir
```

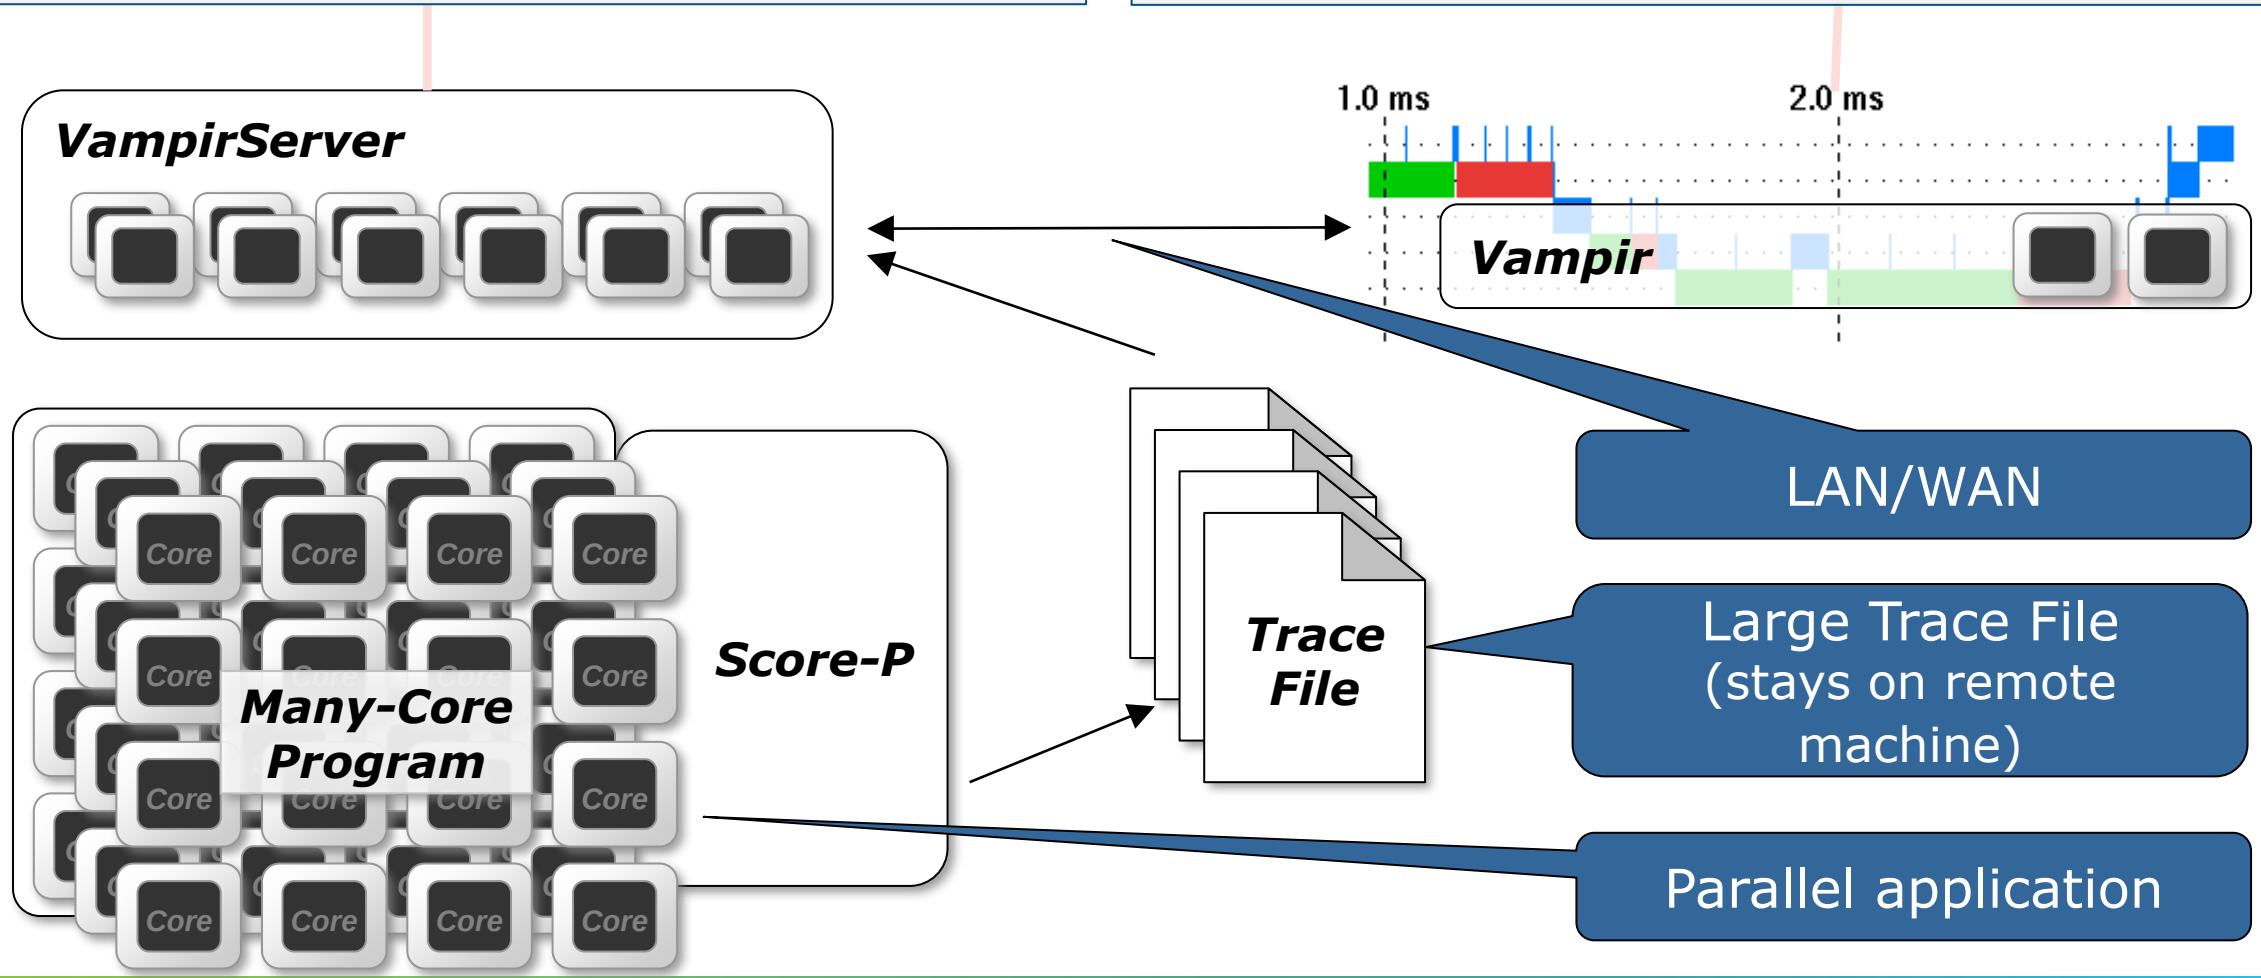

Local setup for Vampirserver
Explicit SSH tunneling

```
% module use /home/vihps/software/modulefiles/
% module load vampir
% module load mpi/openmpi/3.1.2-gcc-8.2.0
% vampirserver start
Launching VampirServer...
VampirServer 9.9.0  (14232279)
Licensed to Goethe-Universitaet Frankfurt am Main
Running 4 analysis processes... (abort with vampirserver stop 18299)
VampirServer <18299> listens on: login02.cm.cluster:30018
```

```
% vampir &
% ssh -L 30000:login02.cm.cluster:30018 <user>@goethe.hhlr-gu.de
```

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

Vampir Case Study: Optimizing COSMO-SPECS

COSMO-SPECS Original

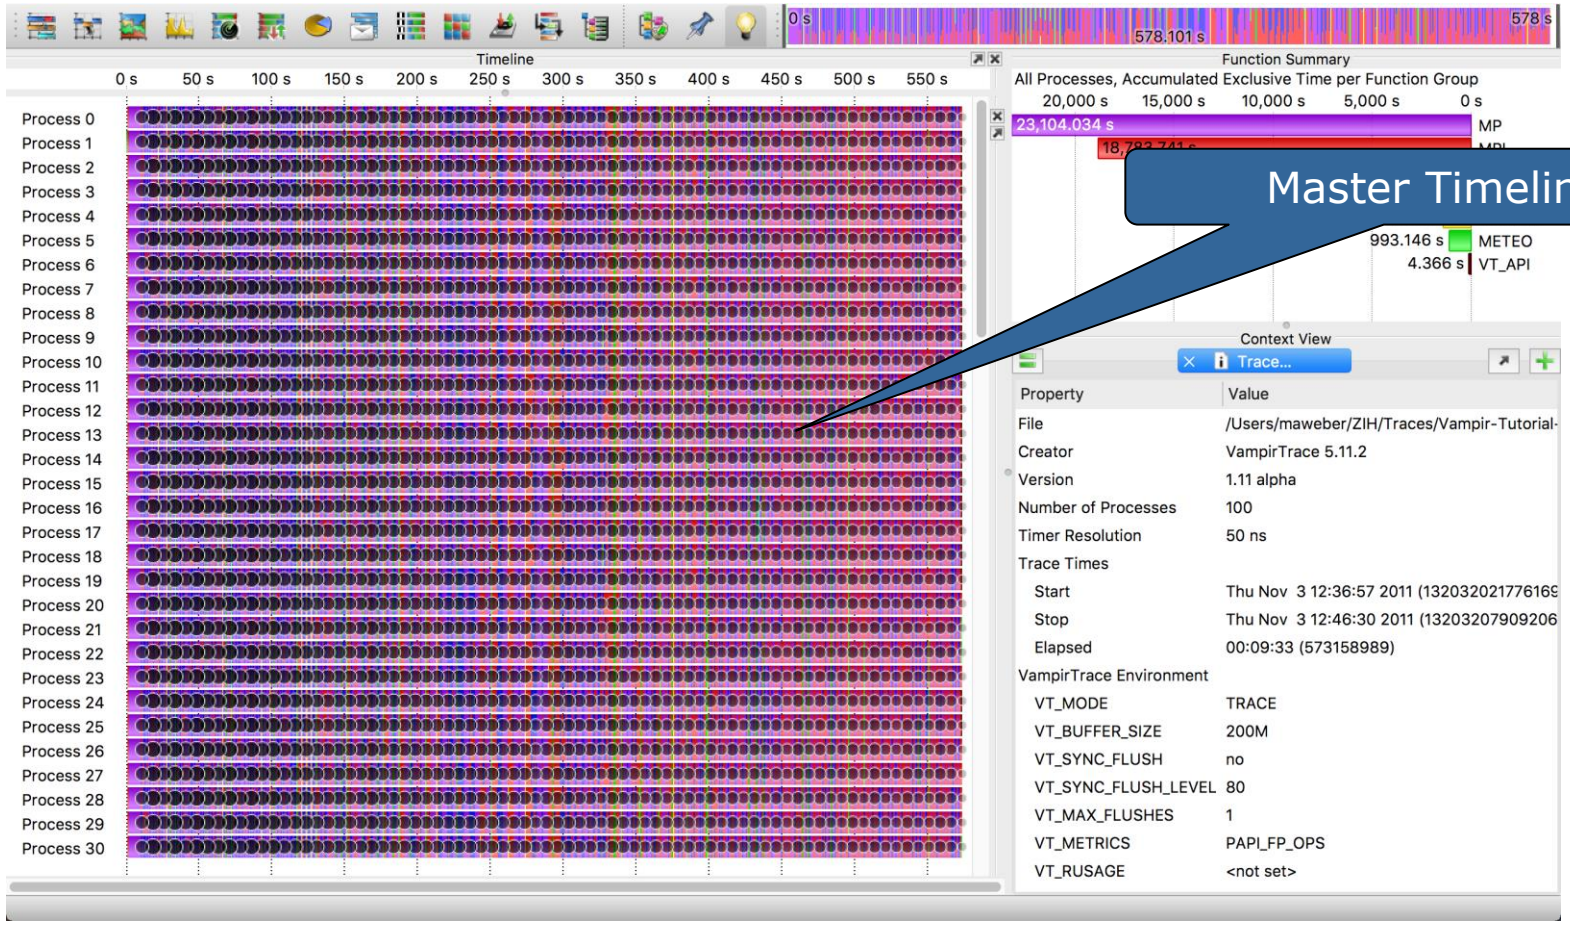

- Weather forecast code COSMO-SPECS
  - with 100 processes
- COSMO: weather model (METEO group)
- SPECS: microphysics for accurate cloud calculation (MP and MP_UTIL group)
- Coupling of both models done in COUPLE group

COSMO-SPECS Original

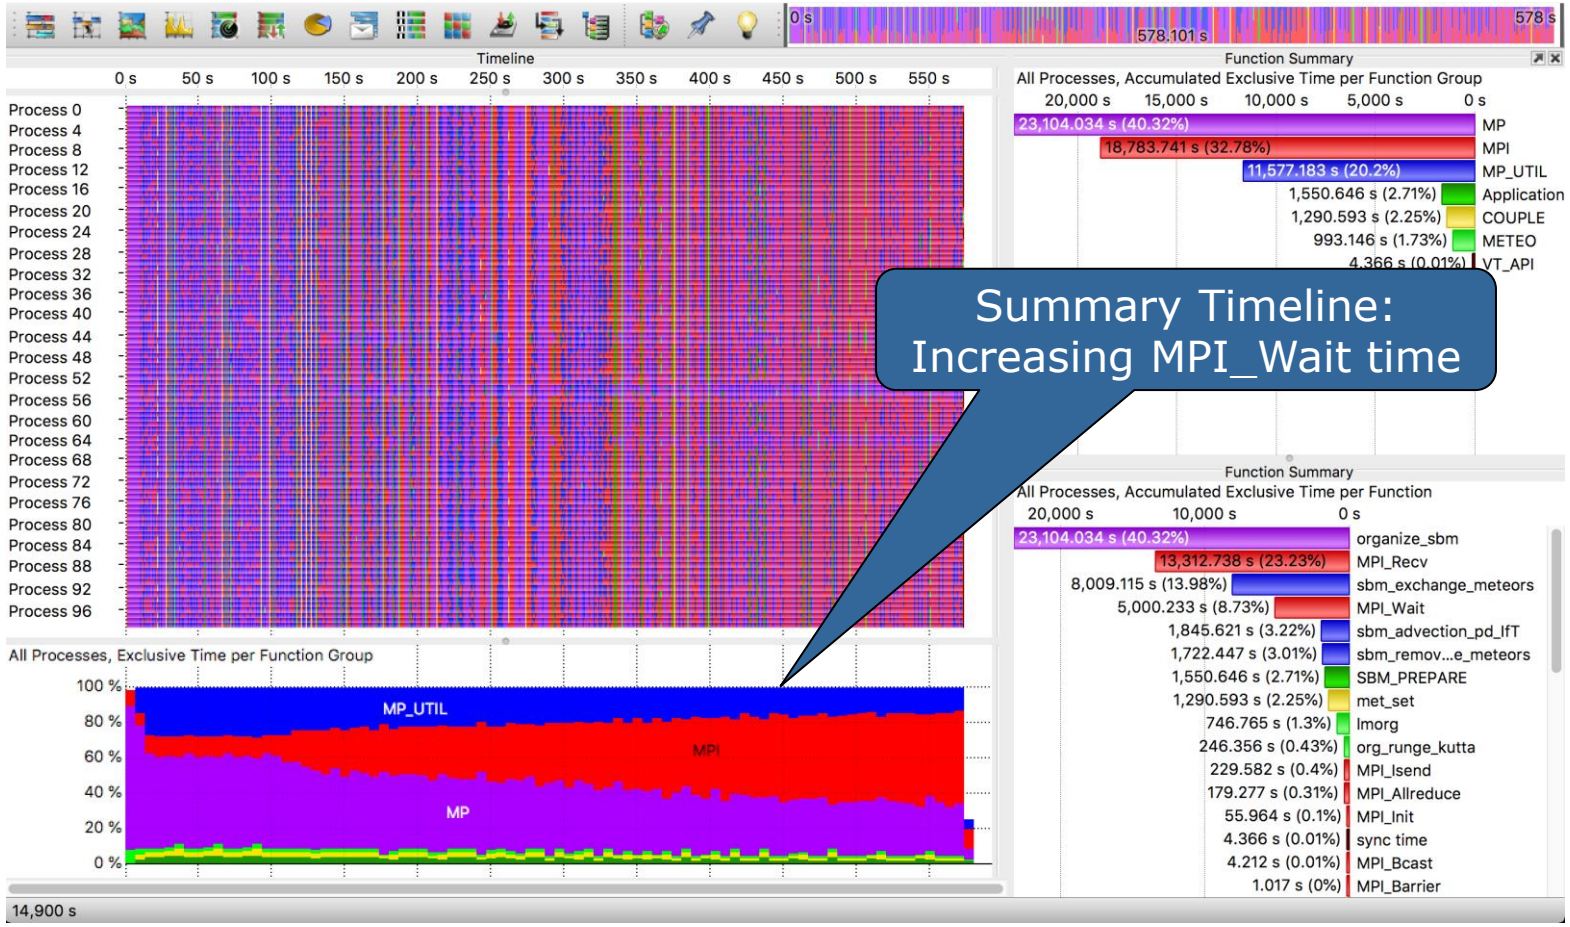

- Compared to METEO, MP and MP_UTIL are very compute intensive, however this is due to more complex calculations and no performance issue
- Problem: >32% of time spent in MPI
- MPI runtime share increases throughout the application run

COSMO-SPECS Original

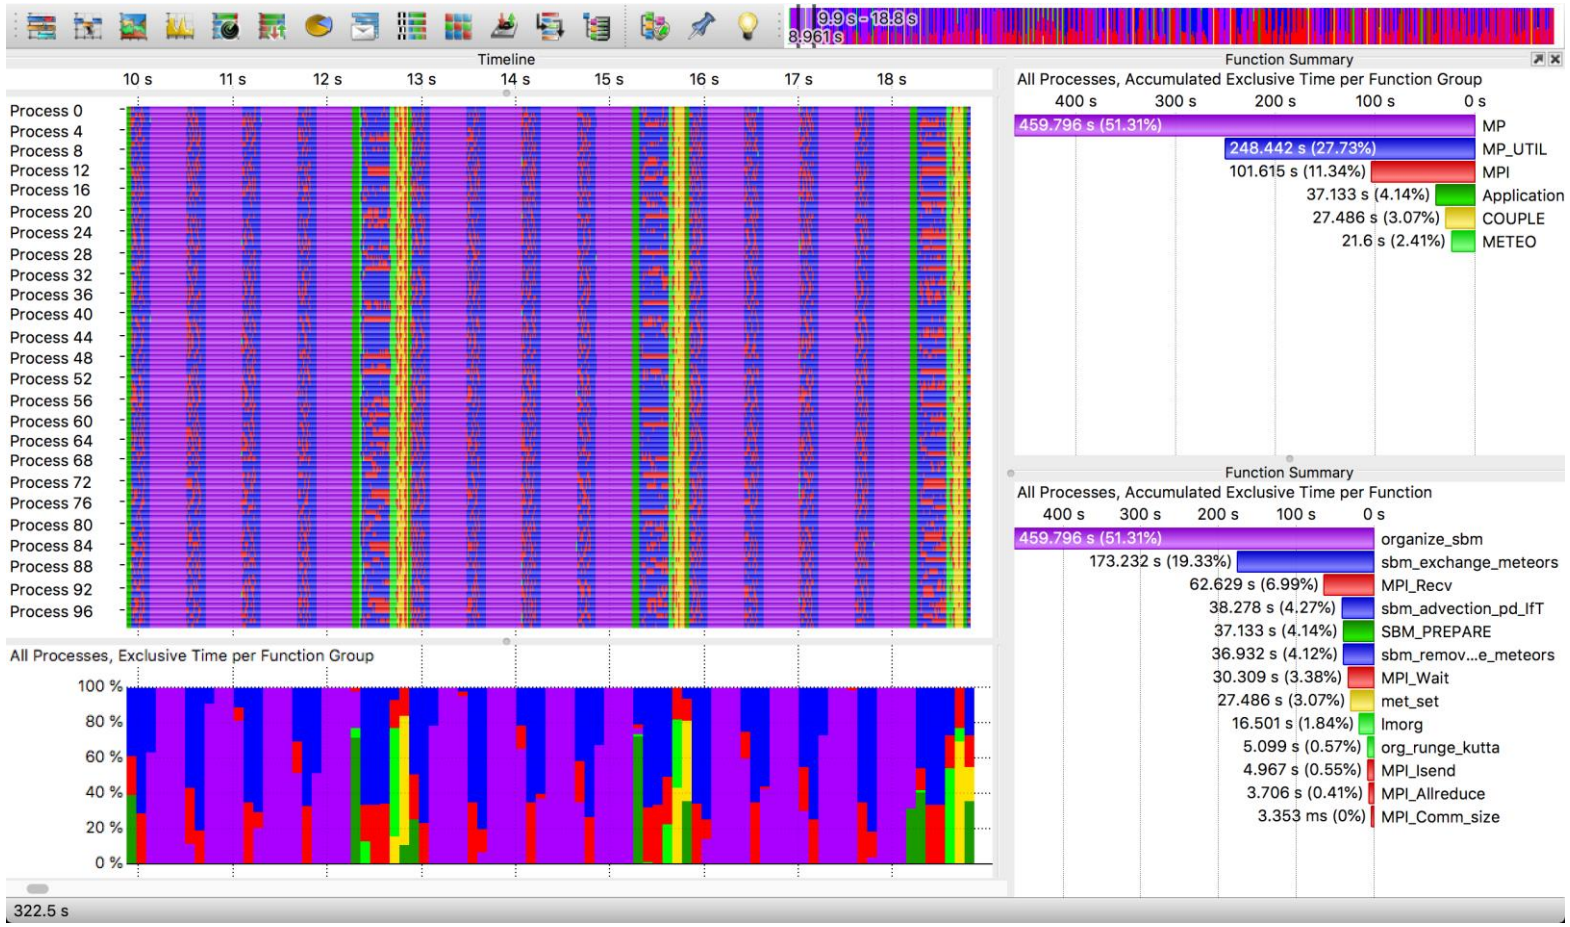

- Zoom into the first three iterations
- MP/MP_UTIL perform four sub-steps in one iteration
- Low MPI time share
- Everything is balanced and looks okay

COSMO-SPECS Original

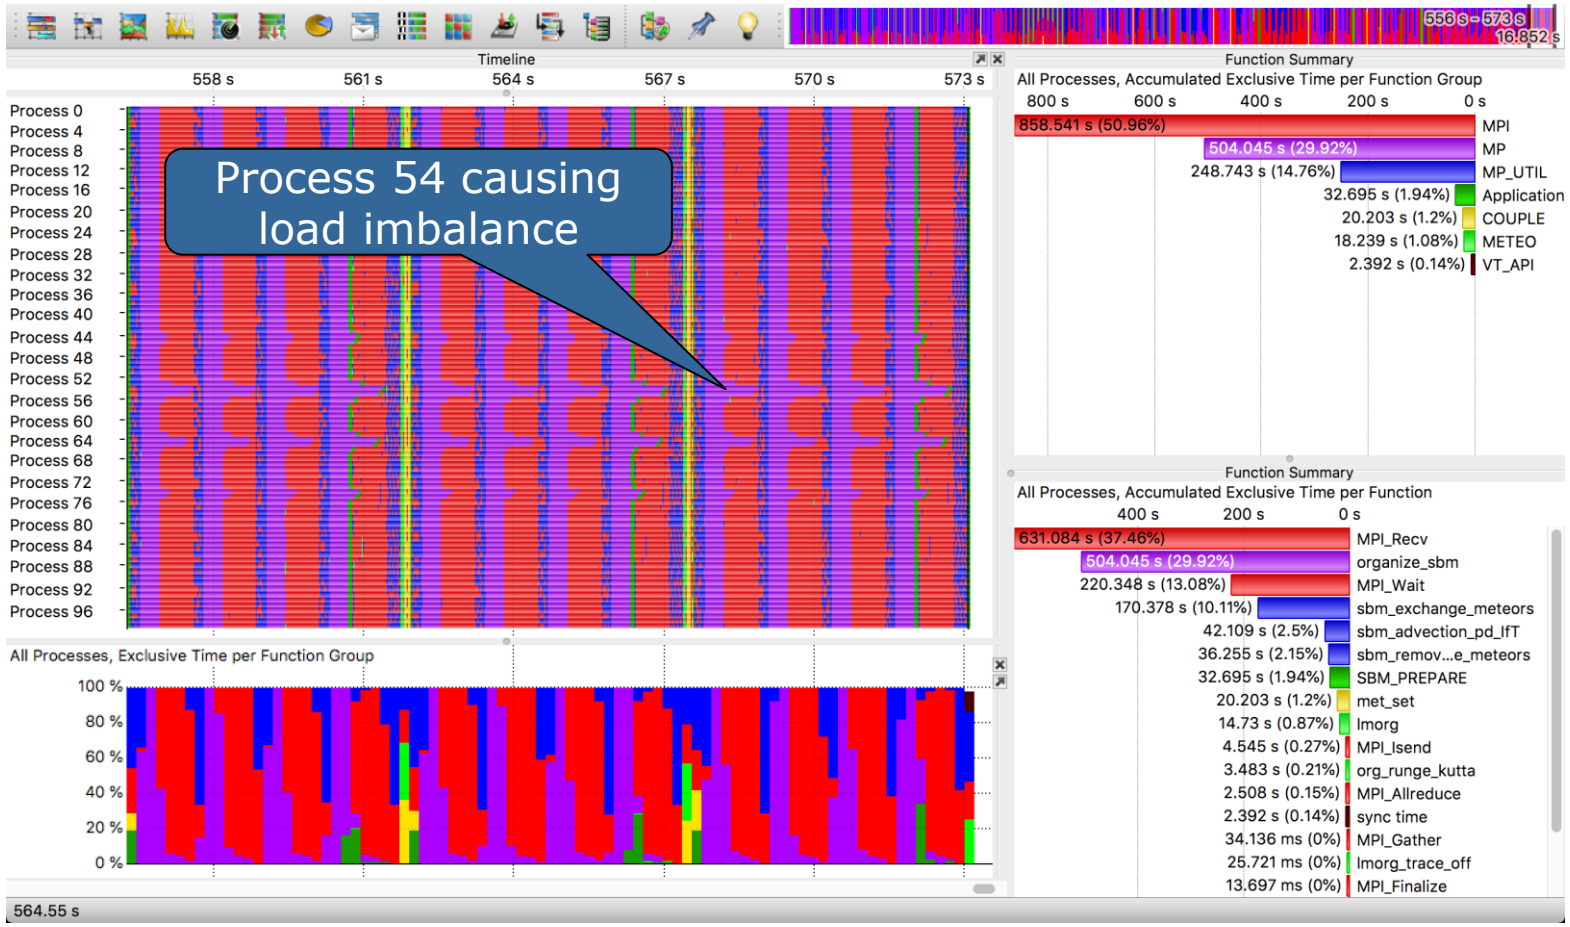

- Zoom into the last three iterations
- Very high MPI time share (>50%)
- Large load imbalance caused by MP functions around **`Process 54`** and **`Process 64`**

COSMO-SPECS Original

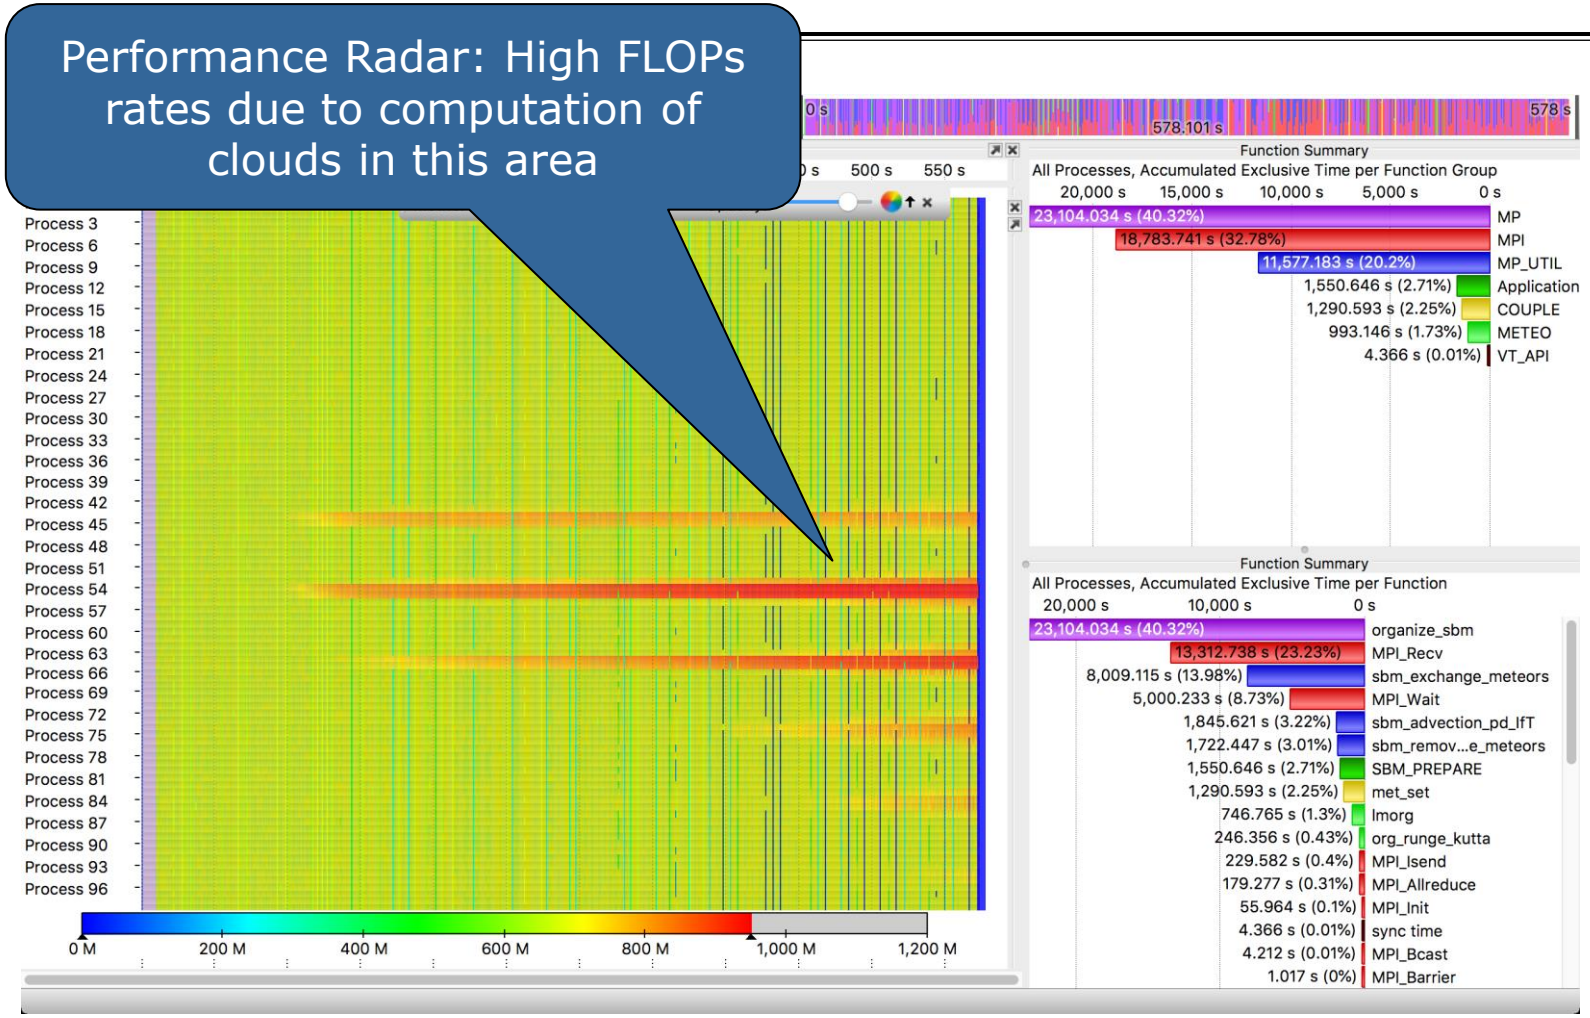

- **`PAPI_FP_OPS`** counter showing higher FLOPs rates on processes causing the imbalance
- Reason for imbalance: Static grid used for distribution of processes. Depending on the weather, expensive cloud computations (MP group) may be only necessary on some processes

COSMO-SPECS FD4

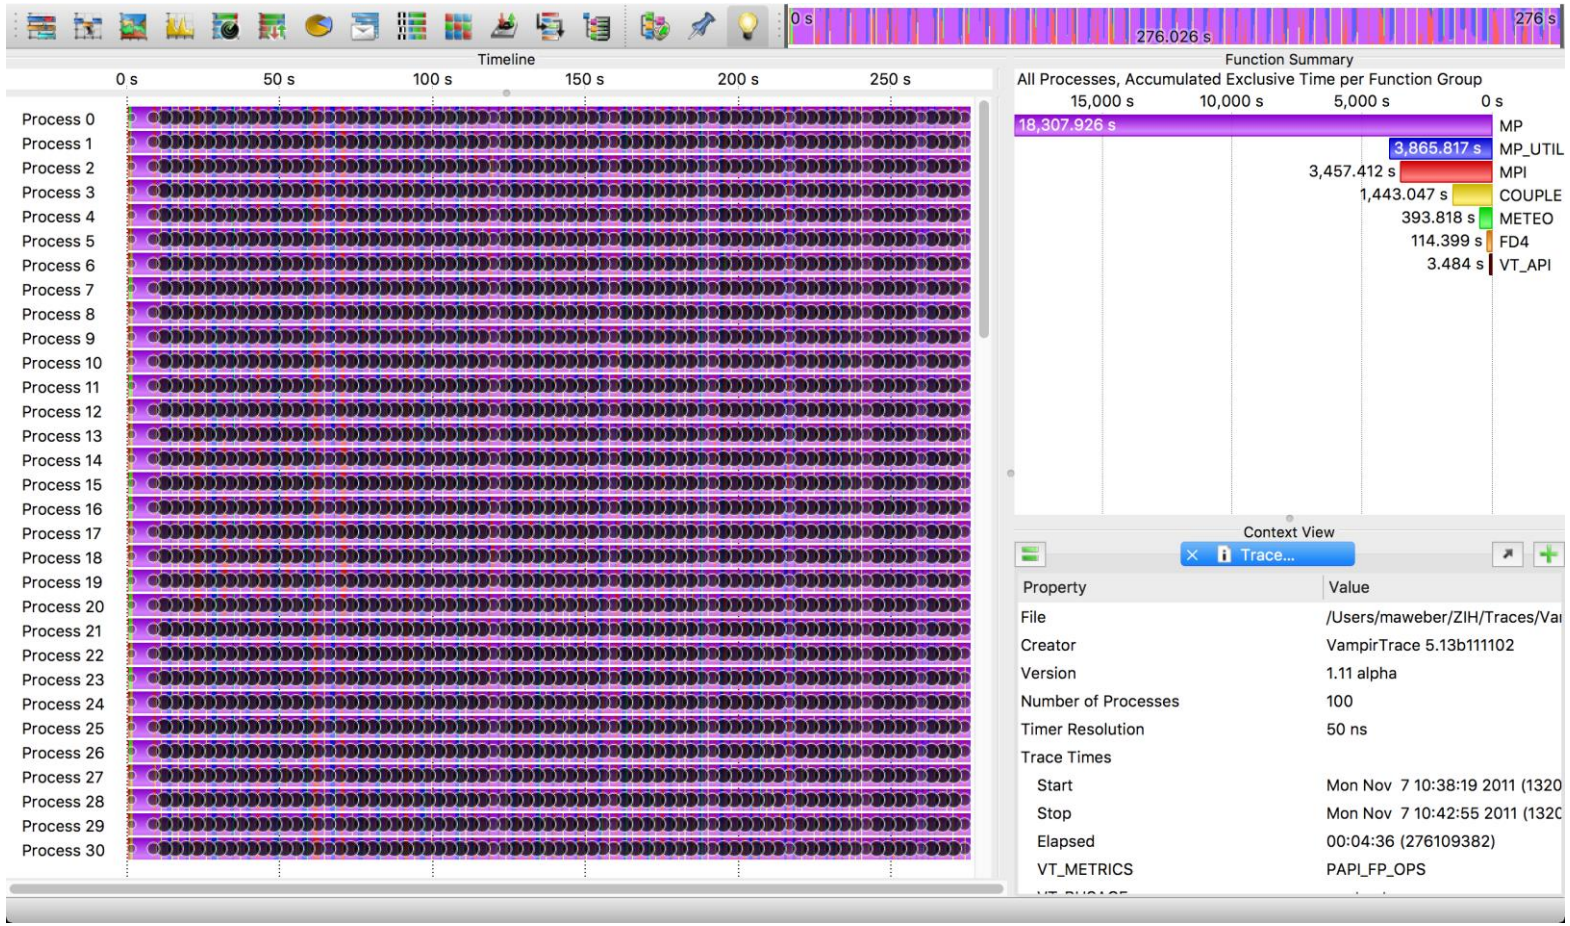

- Weather forecast code COSMO-SPECS
- Run with 100 processes
- COSMO: weather model (METEO group)
- SPECS: microphysics for accurate cloud calculation (MP and MP_UTIL group)
- Coupling of both models done in COUPLE group
- Dynamic load balancing (FD4 group)

COSMO-SPECS FD4

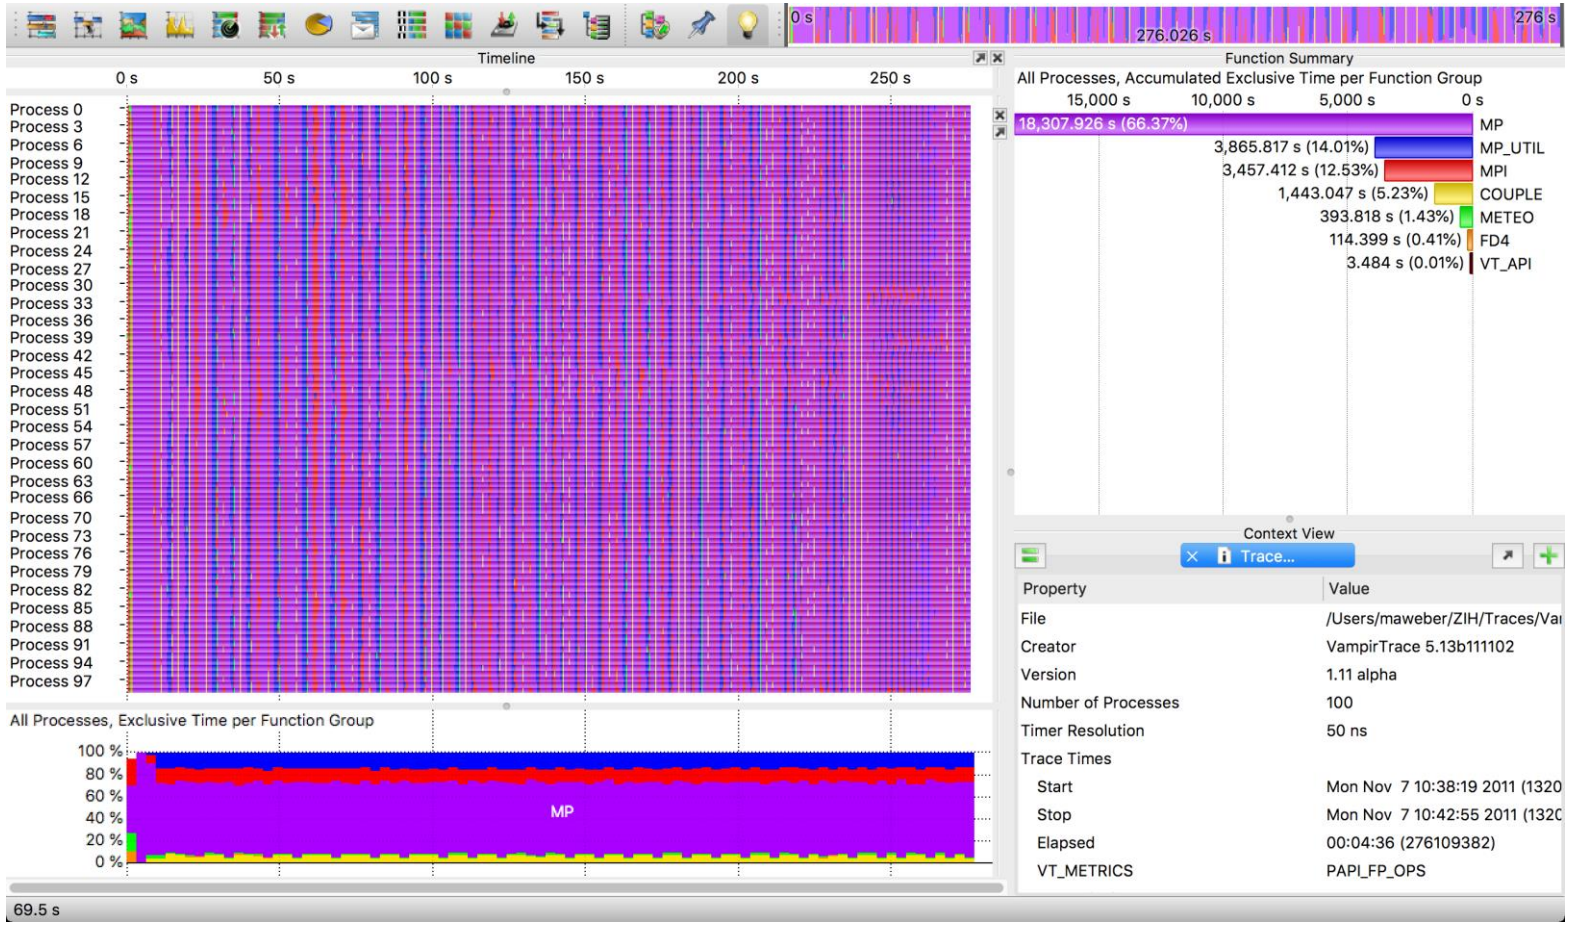

- Dynamic load balancing mitigates the balance problems of the original COSMO-SPECS version
- MPI time share is reduced to <13%
- MPI time share stays constant throughout the application runtime
- Runtime reduced by factor of 2.1, from initially 578s to 276s

COSMO-SPECS FD4

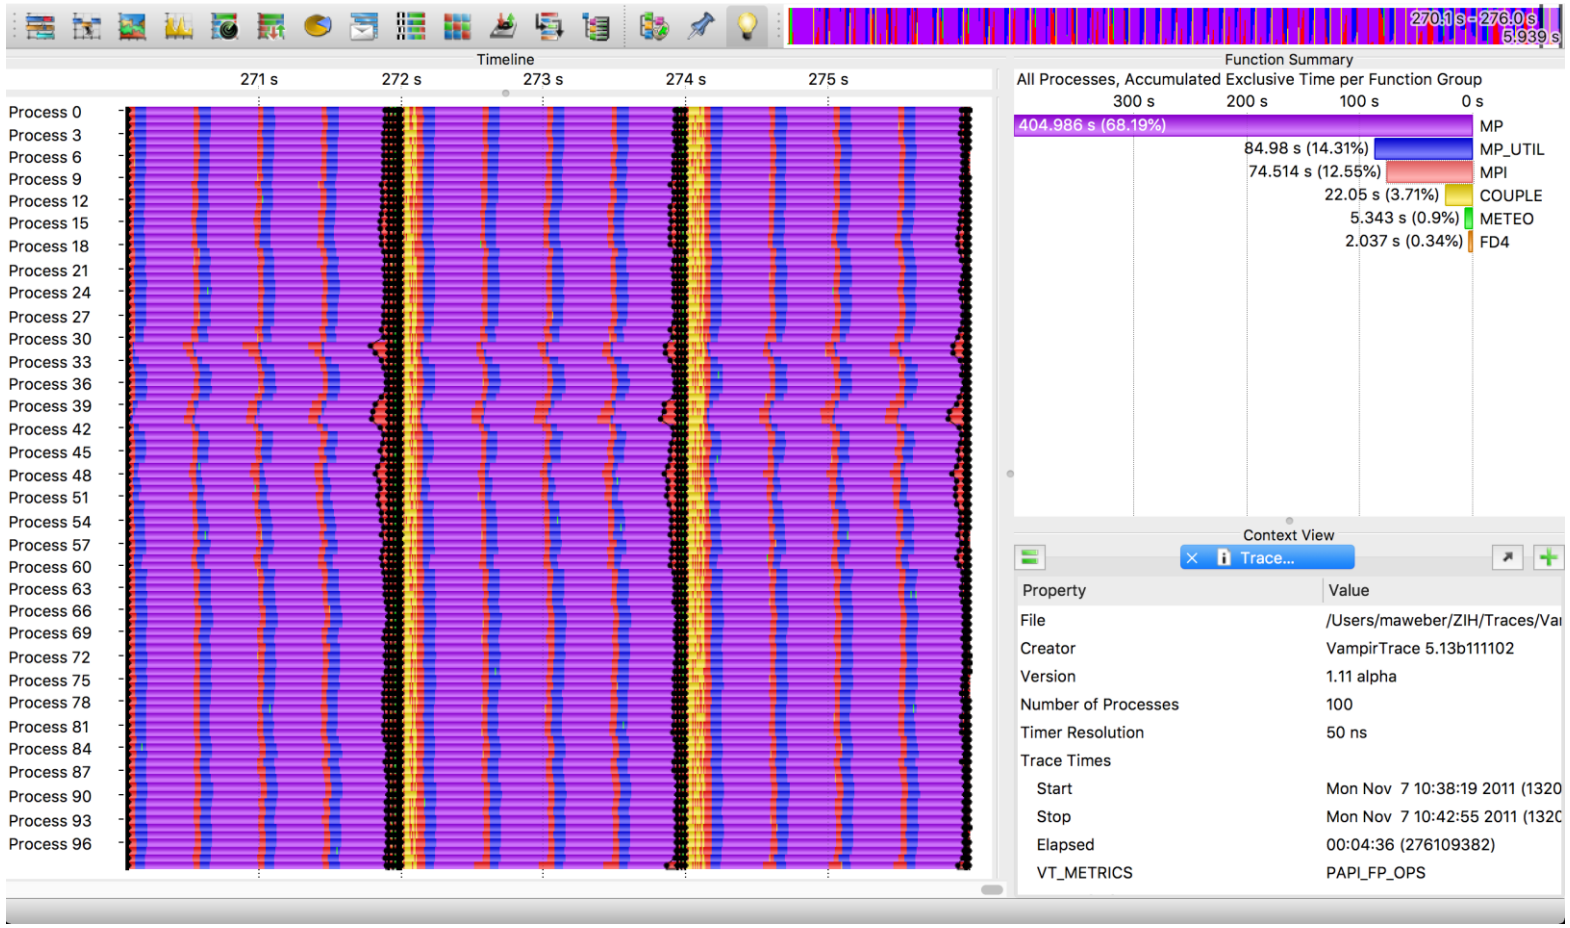

- Zoom into last three iterations
- FD4 balances MP load (precipitation processes in clouds) across all available processes

Summary and Conclusion

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

http://www.vampir.eu

vampirsupport@zih.tu-dresden.de