



TECHNISCHE  
UNIVERSITÄT  
DRESDEN

Center for Information Services and High Performance Computing (ZIH)

## Performance Analysis with Open MPI and VampirTrace

Ronny Brendel, Holger Brunst, Jens Doleschal, Matthias Jurenz, Andreas Knüpfer,  
Matthias Lieber, Holger Mickler, Hartmut Mix, **Matthias Müller**, Michael Peter,  
Matthias Weber, Thomas William, Wolfgang E. Nagel

Zellescher Weg 12  
Willers-Bau A113  
Tel. +49 351 - 463 - 39835

Matthias S. Mueller

(matthias.mueller@tu-dresden.de)



### Motivation

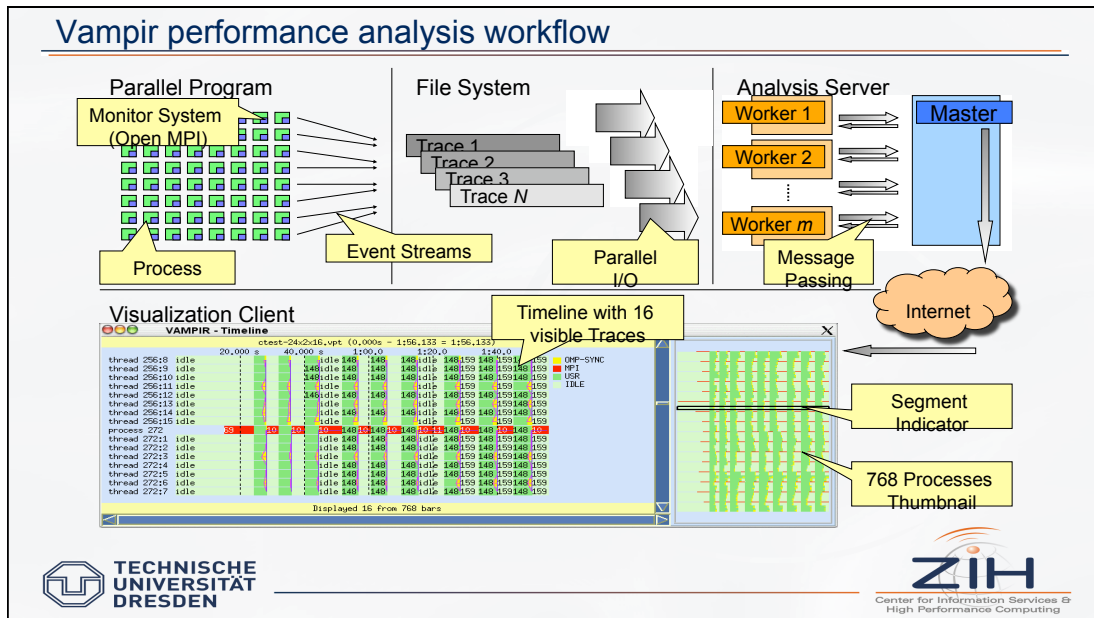
- Parallel programming is about performance!
- Scaling to thousands of cores is required
- You need a decent MPI implementation, e.g. Open MPI
- You also need a ready-to-use performance monitoring and analysis tool



TECHNISCHE  
UNIVERSITÄT  
DRESDEN



## Vampir performance analysis workflow



## VampirTrace

- VampirTrace 5.x.y supports OTF as trace format
- Uses the PMPI interface
- Library implemented in collaboration with FZ Juelich, Univ. Oregon and University of Tennessee
  - OTF (developed with the TAU team at Univ. Oregon)
  - Epilog MPI wrappers (FZ Juelich)
  - Hardware performance counters (PAPI)
- Integrated in Open MPI v1.3 or later
- Available as open source under BSD license:
  - <http://www.tu-dresden.de/zih/vampirtrace>



## Design of Open Trace Format

---

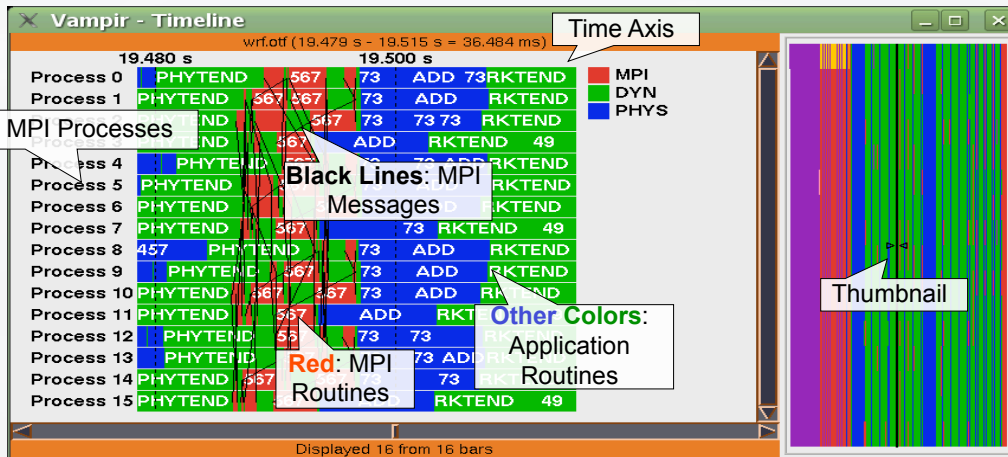
- Design of OTF is directed at 3 objectives:
- **Openness**
  - open format defines record types and file structure so that OTF traces can be generated and read correctly
  - external wishes will be considered .. just talk to us!
- **Flexibility**
  - efficiently selective access is supported
- **Performance**
  - is determined by how efficient & fast OTF trace query and manipulation can be done
  - parallel I/O

## Tools using the Open Trace Format

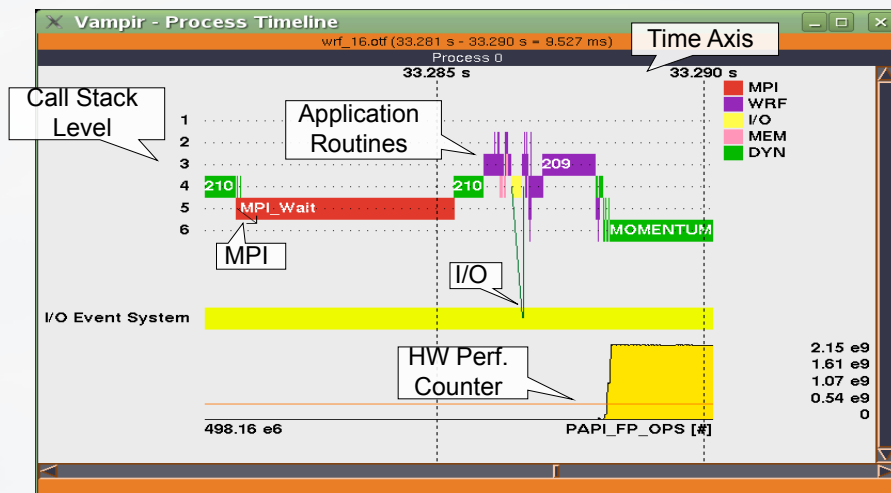
---

- OTF profiler (comes with latest VampirTrace)
- Vampir
- Scalasca
- TAU
- Open SpeedShop
- Microsoft Windows Compute Cluster Environment

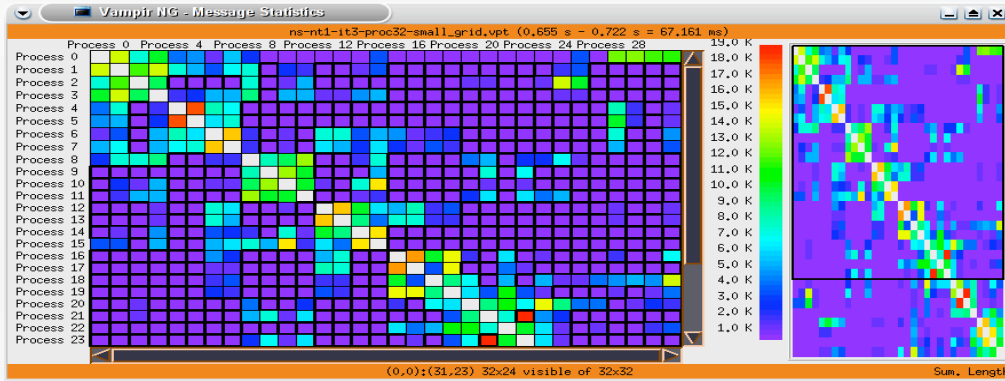
## Most prominent displays: Global Timeline



## Most prominent displays: Single Process Timeline

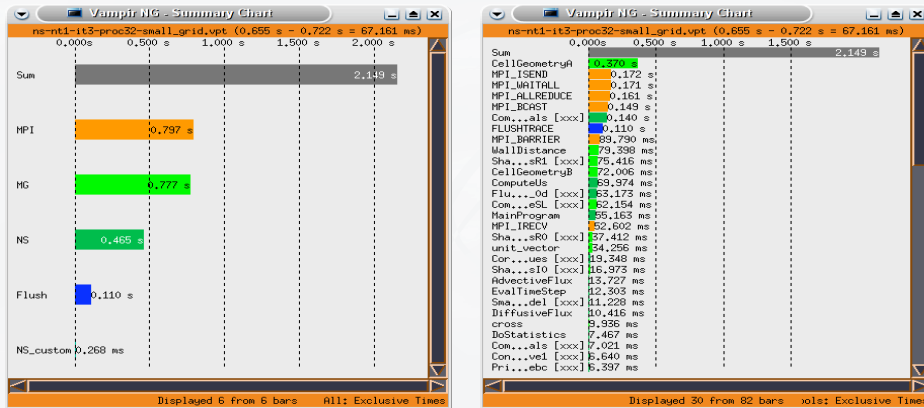


## Vampir Displays - Message Statistics




Message Statistics with Thumbnails

## Vampir Displays - Summary Chart



Grouped / Comprehensive Function Statistics

## Vampir NG Displays - Summary Timeline



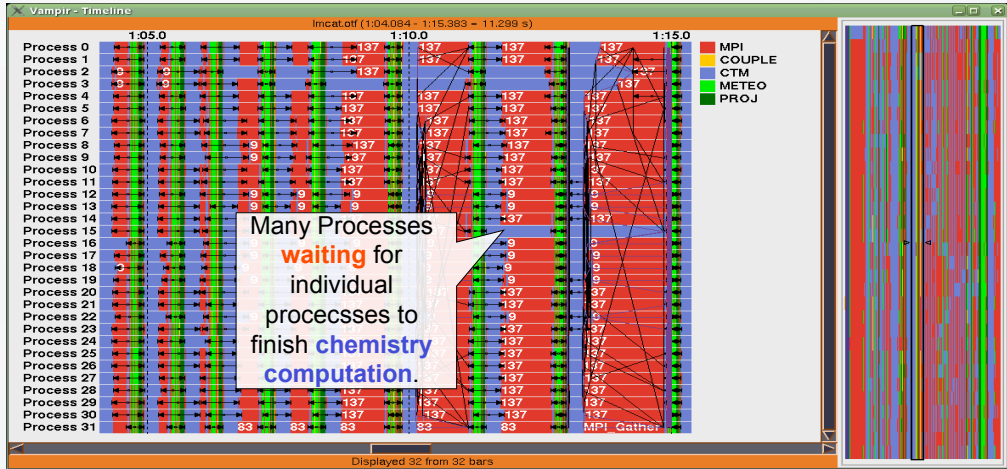
## A Tour Through Typical Performance Phenomena Seen With Vampir

Zellescher Weg 12  
Willers-Bau A113  
Tel. +49 351 - 463 - 39835

Matthias S. Mueller

(matthias.mueller@tu-dresden.de)

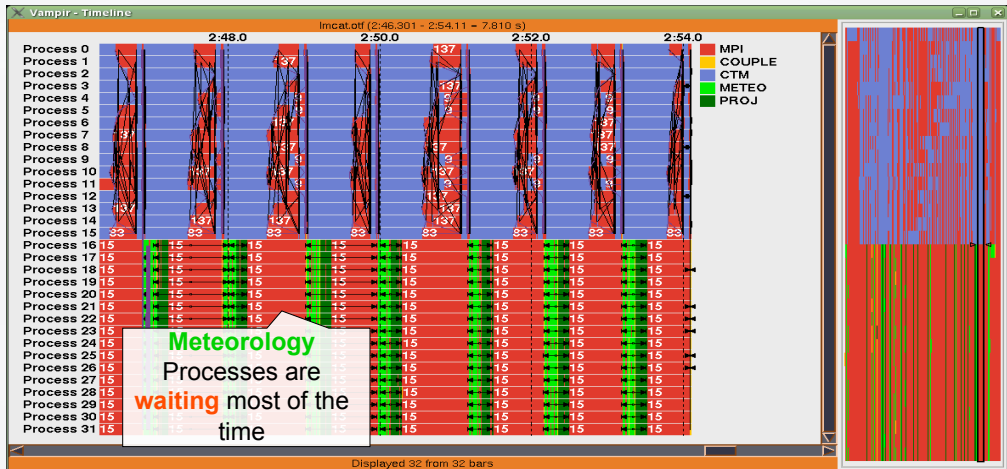
## LM-MUSCAT Air Quality Model: Emerging Load Imbalance



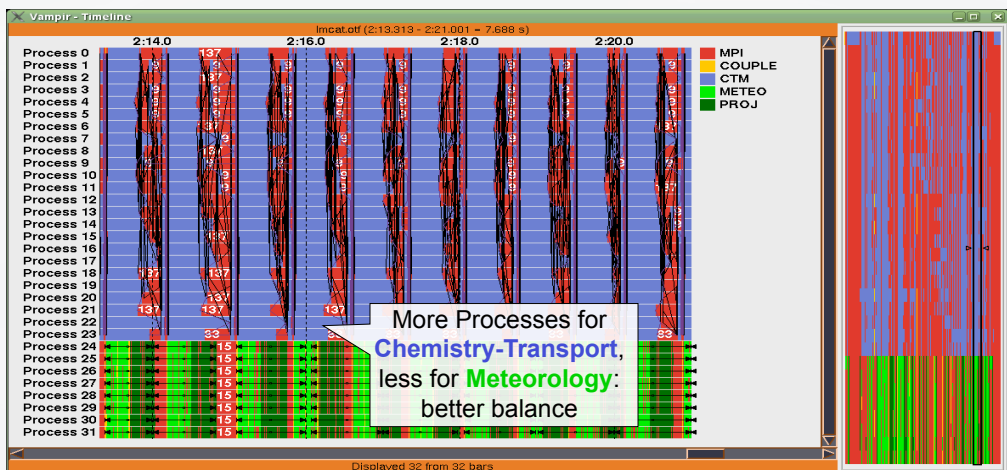
## LM-MUSCAT Air Quality Model: Emerging Load Imbalance



## LM-MUSCAT Air Quality Model: Bad Processor Partitioning

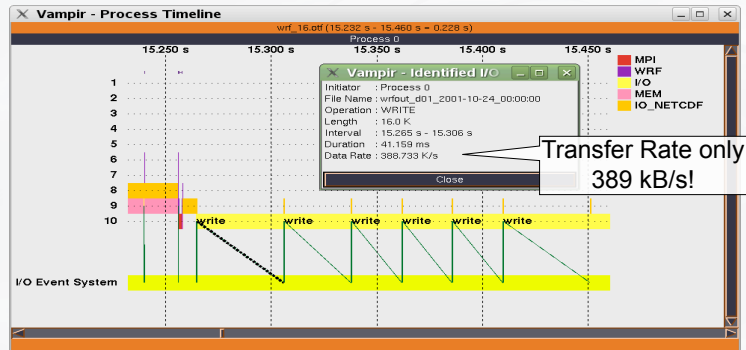


## LM-MUSCAT Air Quality Model: Good Processor Partitioning

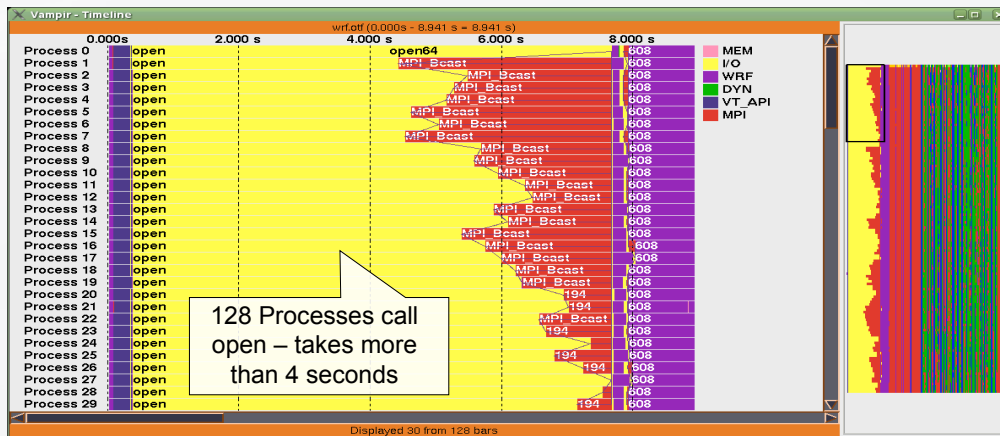




## WRF Weather Model: Low I/O Performance



## WRF Weather Model: Slow Parallel Metadata Operations



## Semtex CFD Application: Serial I/O Leads to Waiting Times



## WRF Weather Model: Hybrid MPI/OpenMP - Idle Threads

