TECHNISCHE UNIVERSITÄT DRESDEN

Center for Information Services and High Performance Computing (ZIH)

High performance computing methods for spectral cloud microphysics in the COSMO model

Leibniz Institute for Tropospheric Research
11 October 2012

LEIBNIZ INSTITUTE FOR TROPOSPHERIC RESEARCH

Matthias Lieber [a], Verena Grützun [b], Ralf Wolke [c], Matthias S. Müller [a], Wolfgang E. Nagel [a]

matthias.lieber@tu-dresden.de

[a] Center for Information Services and High Performance Computing (ZIH), TU Dresden, Germany

[b] Max Planck Institute for Meteorology (MPI-M), Hamburg, Germany

[c] Leibniz Institute for Tropospheric Research (IfT), Leipzig, Germany

Max-Planck-Institut für Meteorologie

ZIH
Center for Information Services & High Performance Computing

Outline

- Short ZIH intro
- Short cloud modeling and spectral bin microphysics intro
- Performance analysis of COSMO-SPECS
- Performance optimizations
- Performance comparison
- Summary

Center for Information Services and HPC (ZIH)

- Central scientific unit of Technische Universität Dresden
- IT infrastructure and service provider for the university
- Research and teaching
- High performance computing center for researchers in saxony
- ~70 TFlops installed HPC power
- New HPC system in 2013

Vampir: Performance Analysis of Parallel Programs

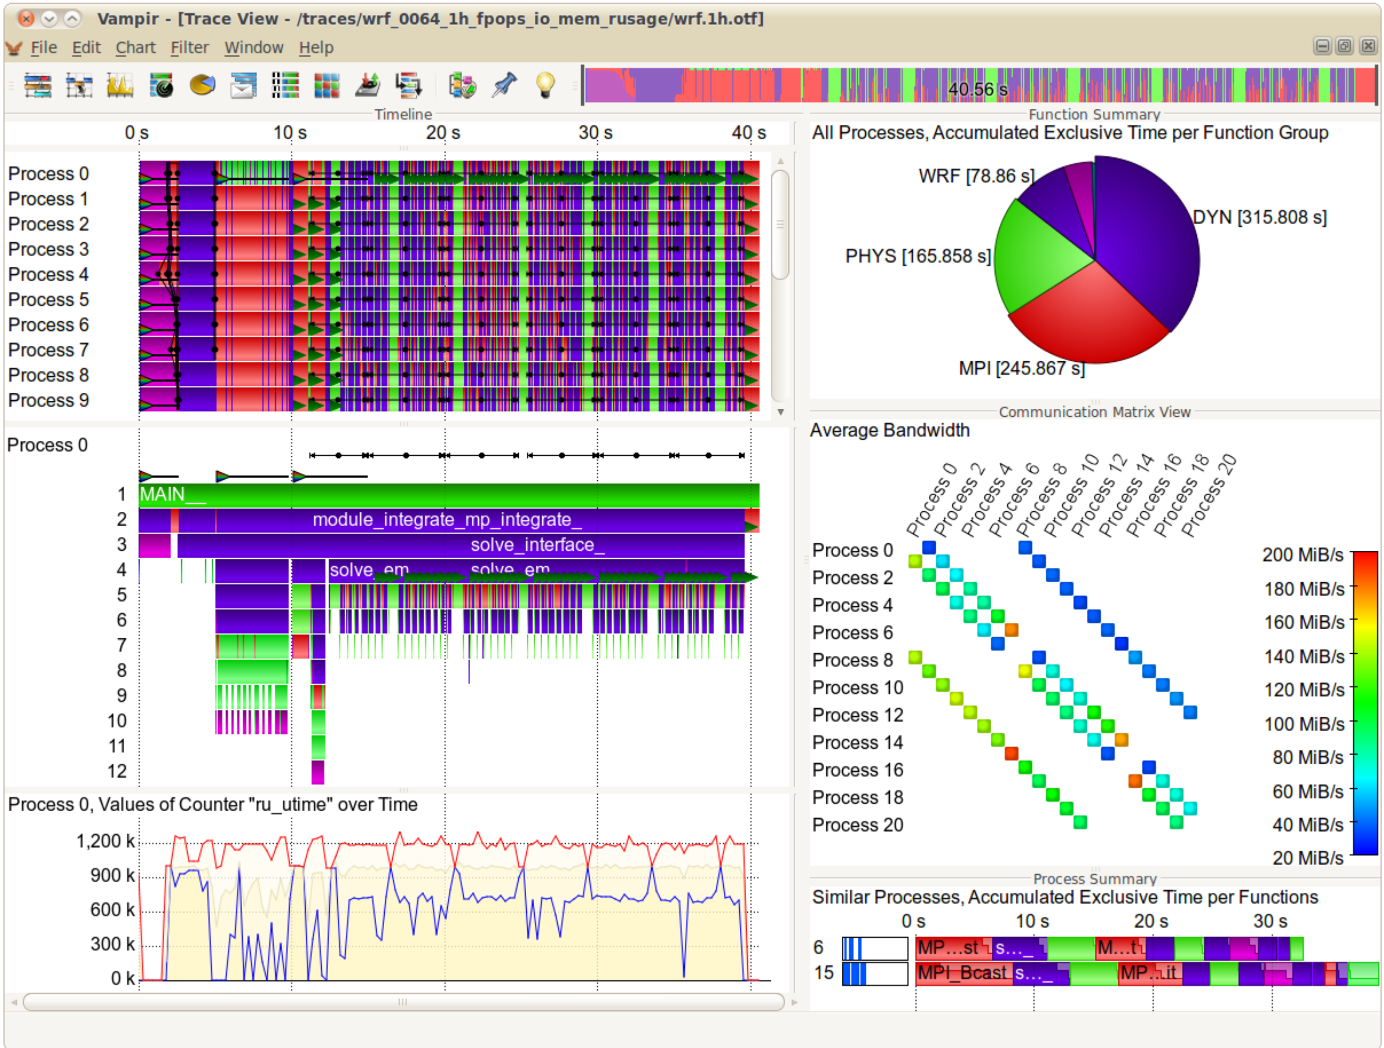

http://www.vampir.eu

TECHNISCHE UNIVERSITÄT DRESDEN

ZIH
Center for Information Services & High Performance Computing

Cloud Modeling

- Clouds play a major role for climate and weather
- Clouds represent one of the major uncertainties in climate and weather models [IPCC07, CCSP09]
- Interaction of aerosols with clouds and precipitation needs more attention

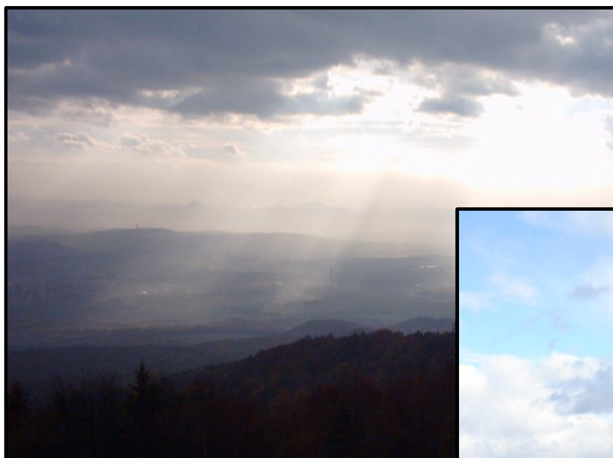

Cloud Modeling: Bulk Parameterization Schemes

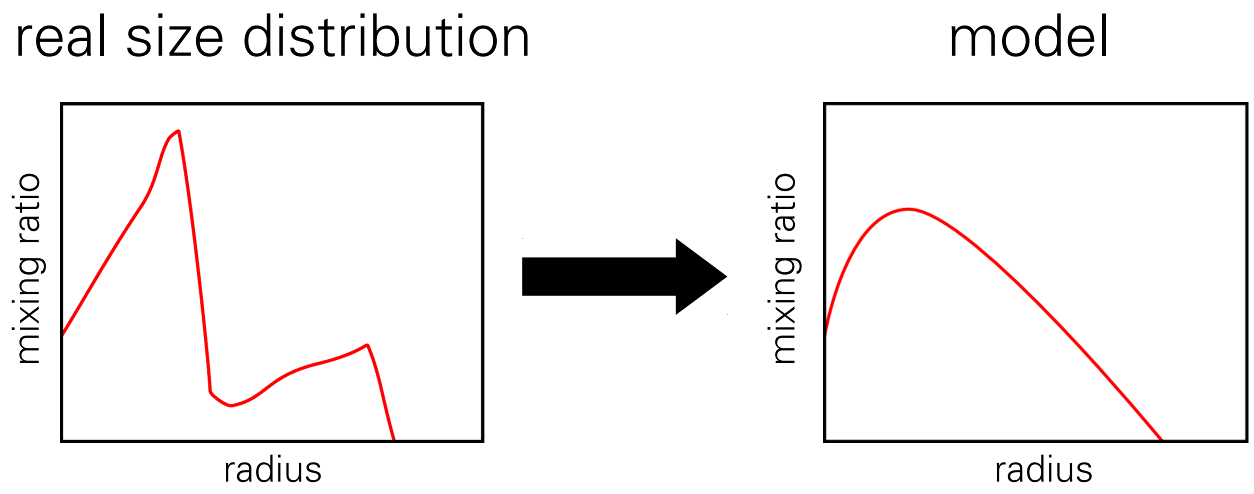

- Cloud droplets are described by their bulk mass only
- Shape of size distribution is prescribed
- Computationally fast
- Used in most weather models

Cloud Modeling: Spectral Bin Microphysics Schemes

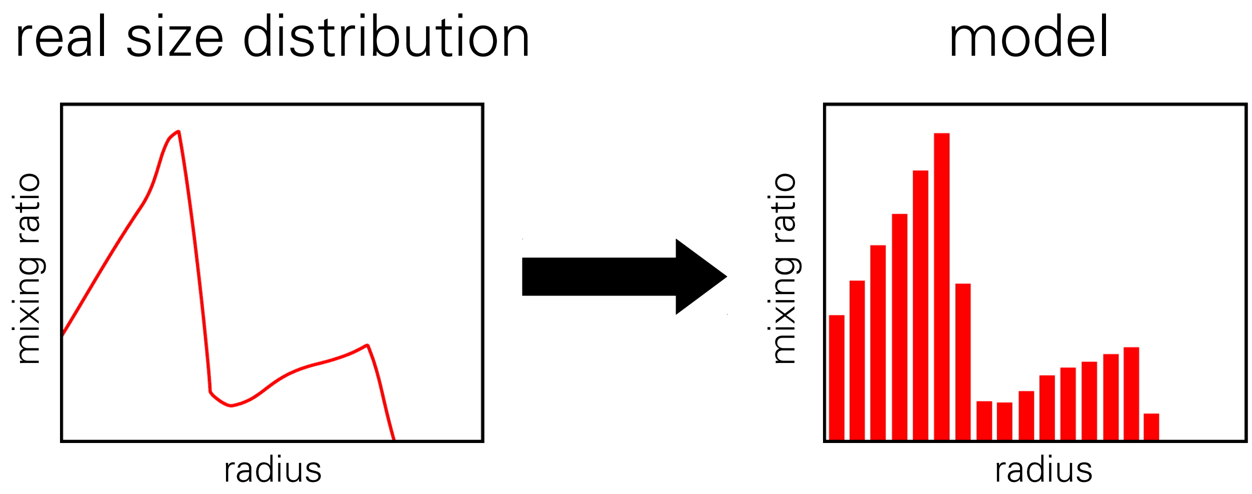

- Bin discretization of size distribution
- Allows more detailed modeling of interaction between aerosols, clouds, and precipitation
- Computationally very expensive (runtime & memory)
- Only used for process studies up to now

Cloud Modeling: 3D Models with Spectral Bin Microphysics

Model	Grid, resolution	#CPUs, system	Forecast and comp. time
MM5-SBM Lynn et al., 2005	100 x 90 x 35 3 km	32 SGI Origin	15 h 8-9 d
COSMO-SPECS Grützun et al., 2008	80 x 80 x 48 1 km	100 SGI Altix 4700	2 h ~1 d
WRF-SBM Khain et al., 2010	133 x 133 x 31 3 km	8 "PC-Cluster"	72 h 10 d
DESCAM-3D Planche et al., 2010	256 x 256 x 80 250 m	64, Altix ICE / IBM Power 6	?

- Relatively small grid sizes
- Small number of CPUs
- No real-time forecast

Cloud Modeling: Tropical Cyclone Forecast?

Horizontal grid:
1000 x 1000

Real-time forecast requires ~10 000 CPU cores

Model systems must be tuned for efficient usage of large machines

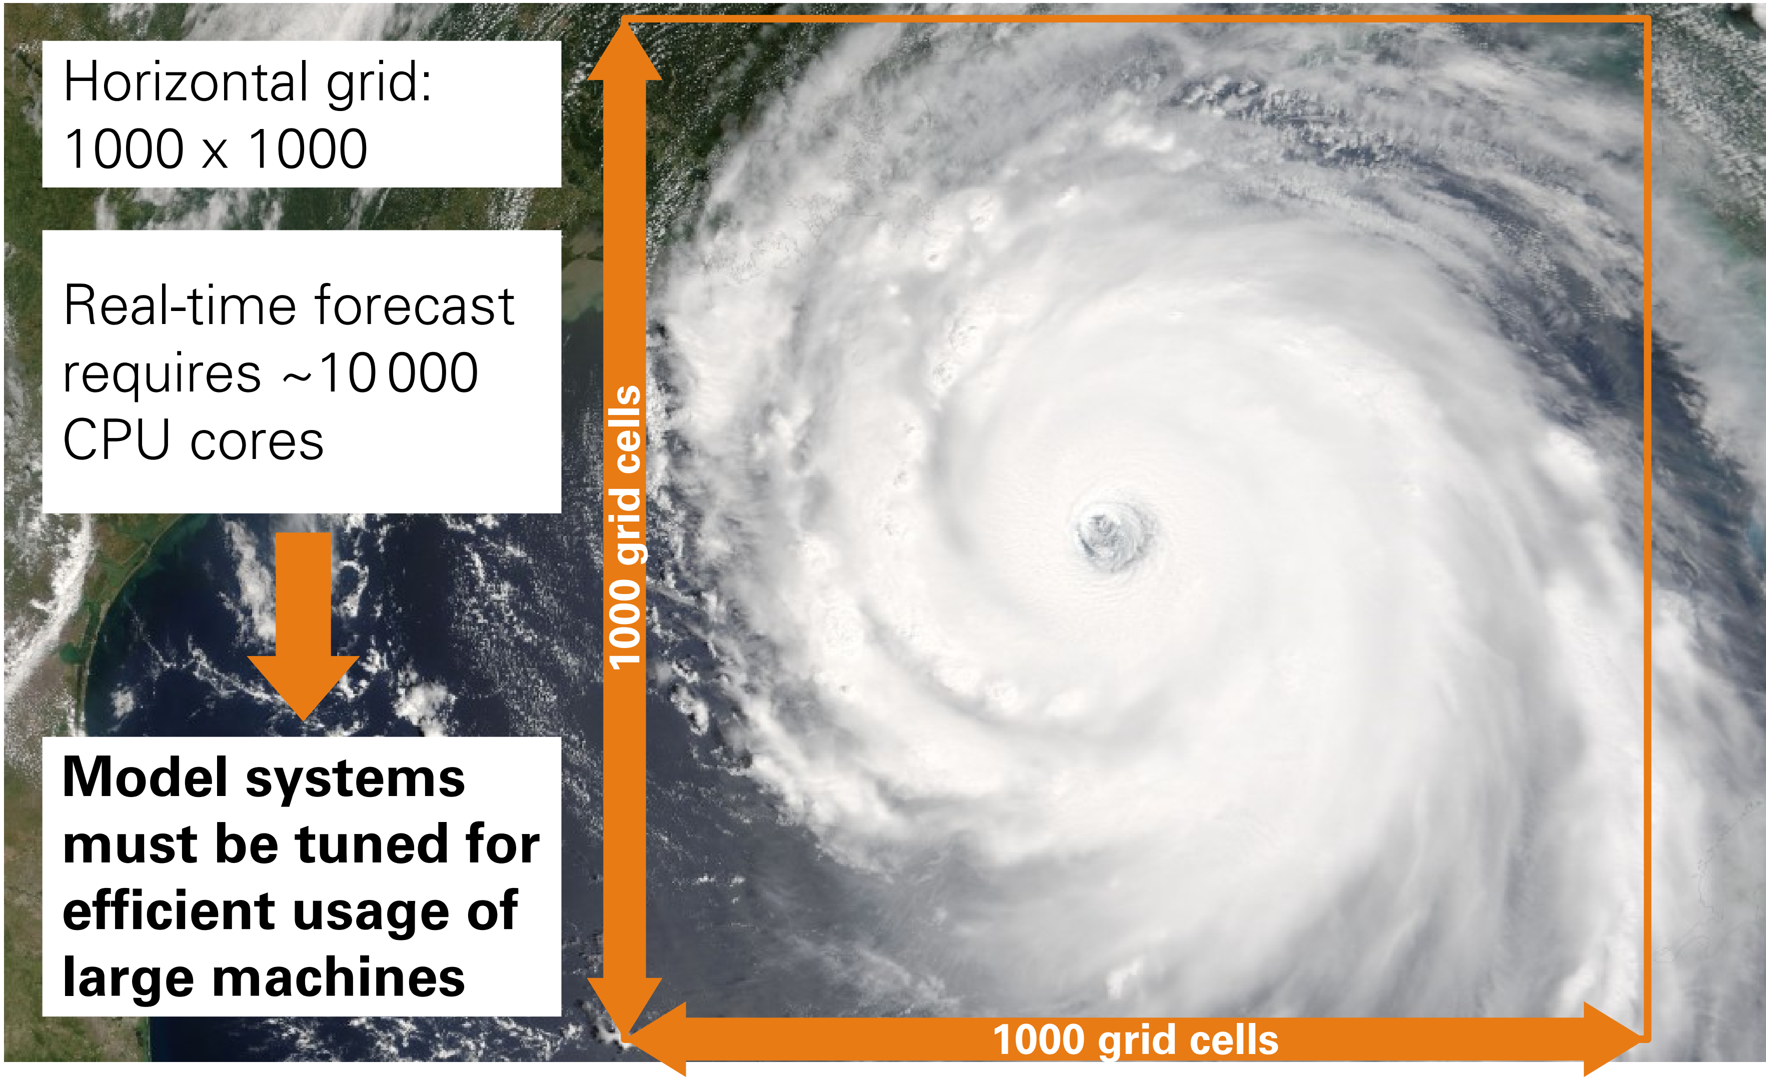

Project Overview

- DFG-funded project *"Parallel coupling framework and advanced time integration methods for detailed cloud processes in atmospheric models"* (IfT & ZIH)
  - Run-time tuning of atmospheric models with spectral bin microphysics
- Model system COSMO-SPECS
  - COSMO: forecast model of the German Weather Service (www.cosmo-model.org)
  - SPECS: spectral bin microphysics model developed at the Leibniz Institute for Tropospheric Research [Simmel06, Grützun08]
  - 3 particle classes, 66 size bins, 726 variables

COSMO-SPECS Analysis: Scalability Benchmark

- Growth of idealized cumulus cloud
- 64 x 64 x 48 grid
- Scalability up to max. possible number of cores

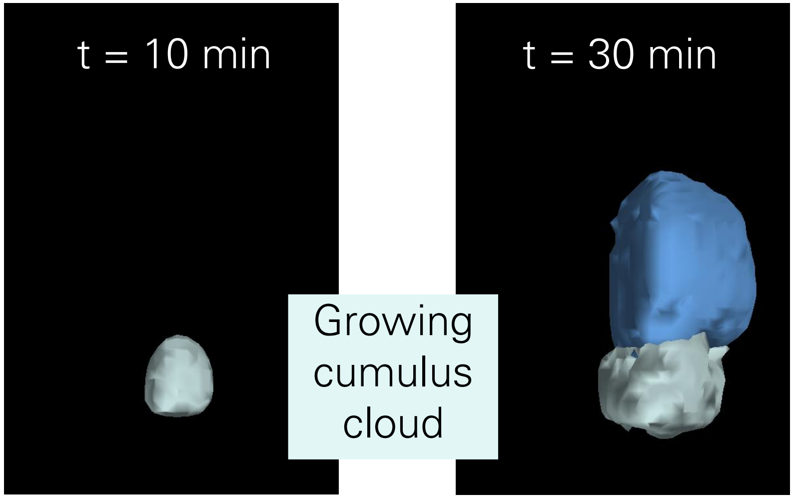

TECHNISCHE UNIVERSITÄT DRESDEN

ZIH Center for Information Services & High Performance Computing

COSMO-SPECS Analysis: Scalability Benchmark

- Growth of idealized cumulus cloud
- 64 x 64 x 48 grid
- Scalability up to max. possible number of cores
- Strongly increasing comm. costs
- **Could do more science with more efficient HPC usage!**

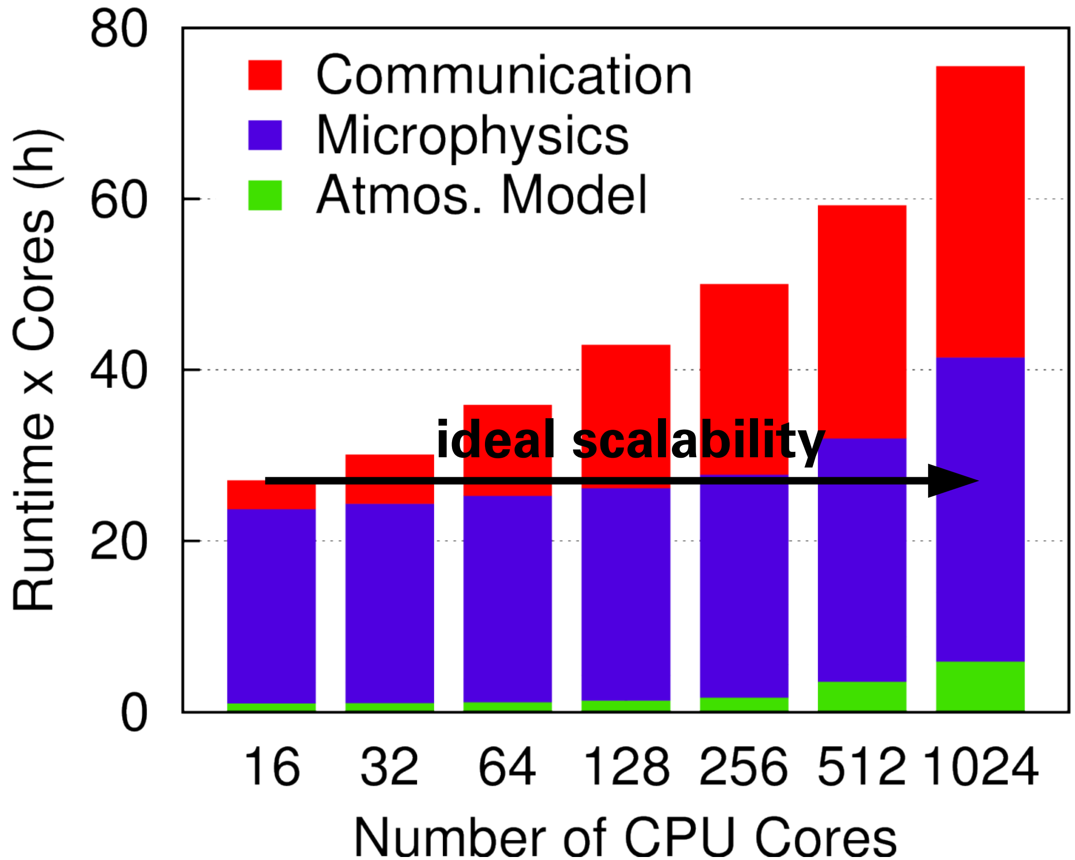

COSMO-SPECS Analysis: Parallelization Scheme

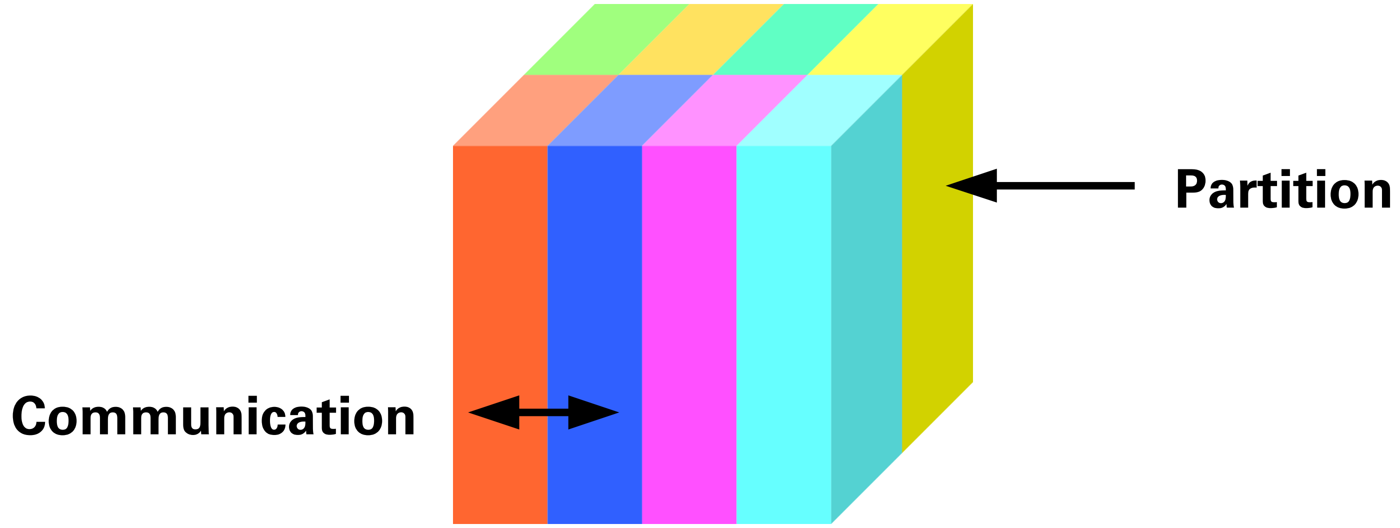

- 3D domain partitioned into rectangular boxes
- 2D decomposition (horizontal dimensions)
- Regular communication with 4 direct neighbors required (periodic boundary conditions)
- Based on MPI (Message Passing Interface)

Vampir Global Timeline

Vampir Global Timeline: One Time Step

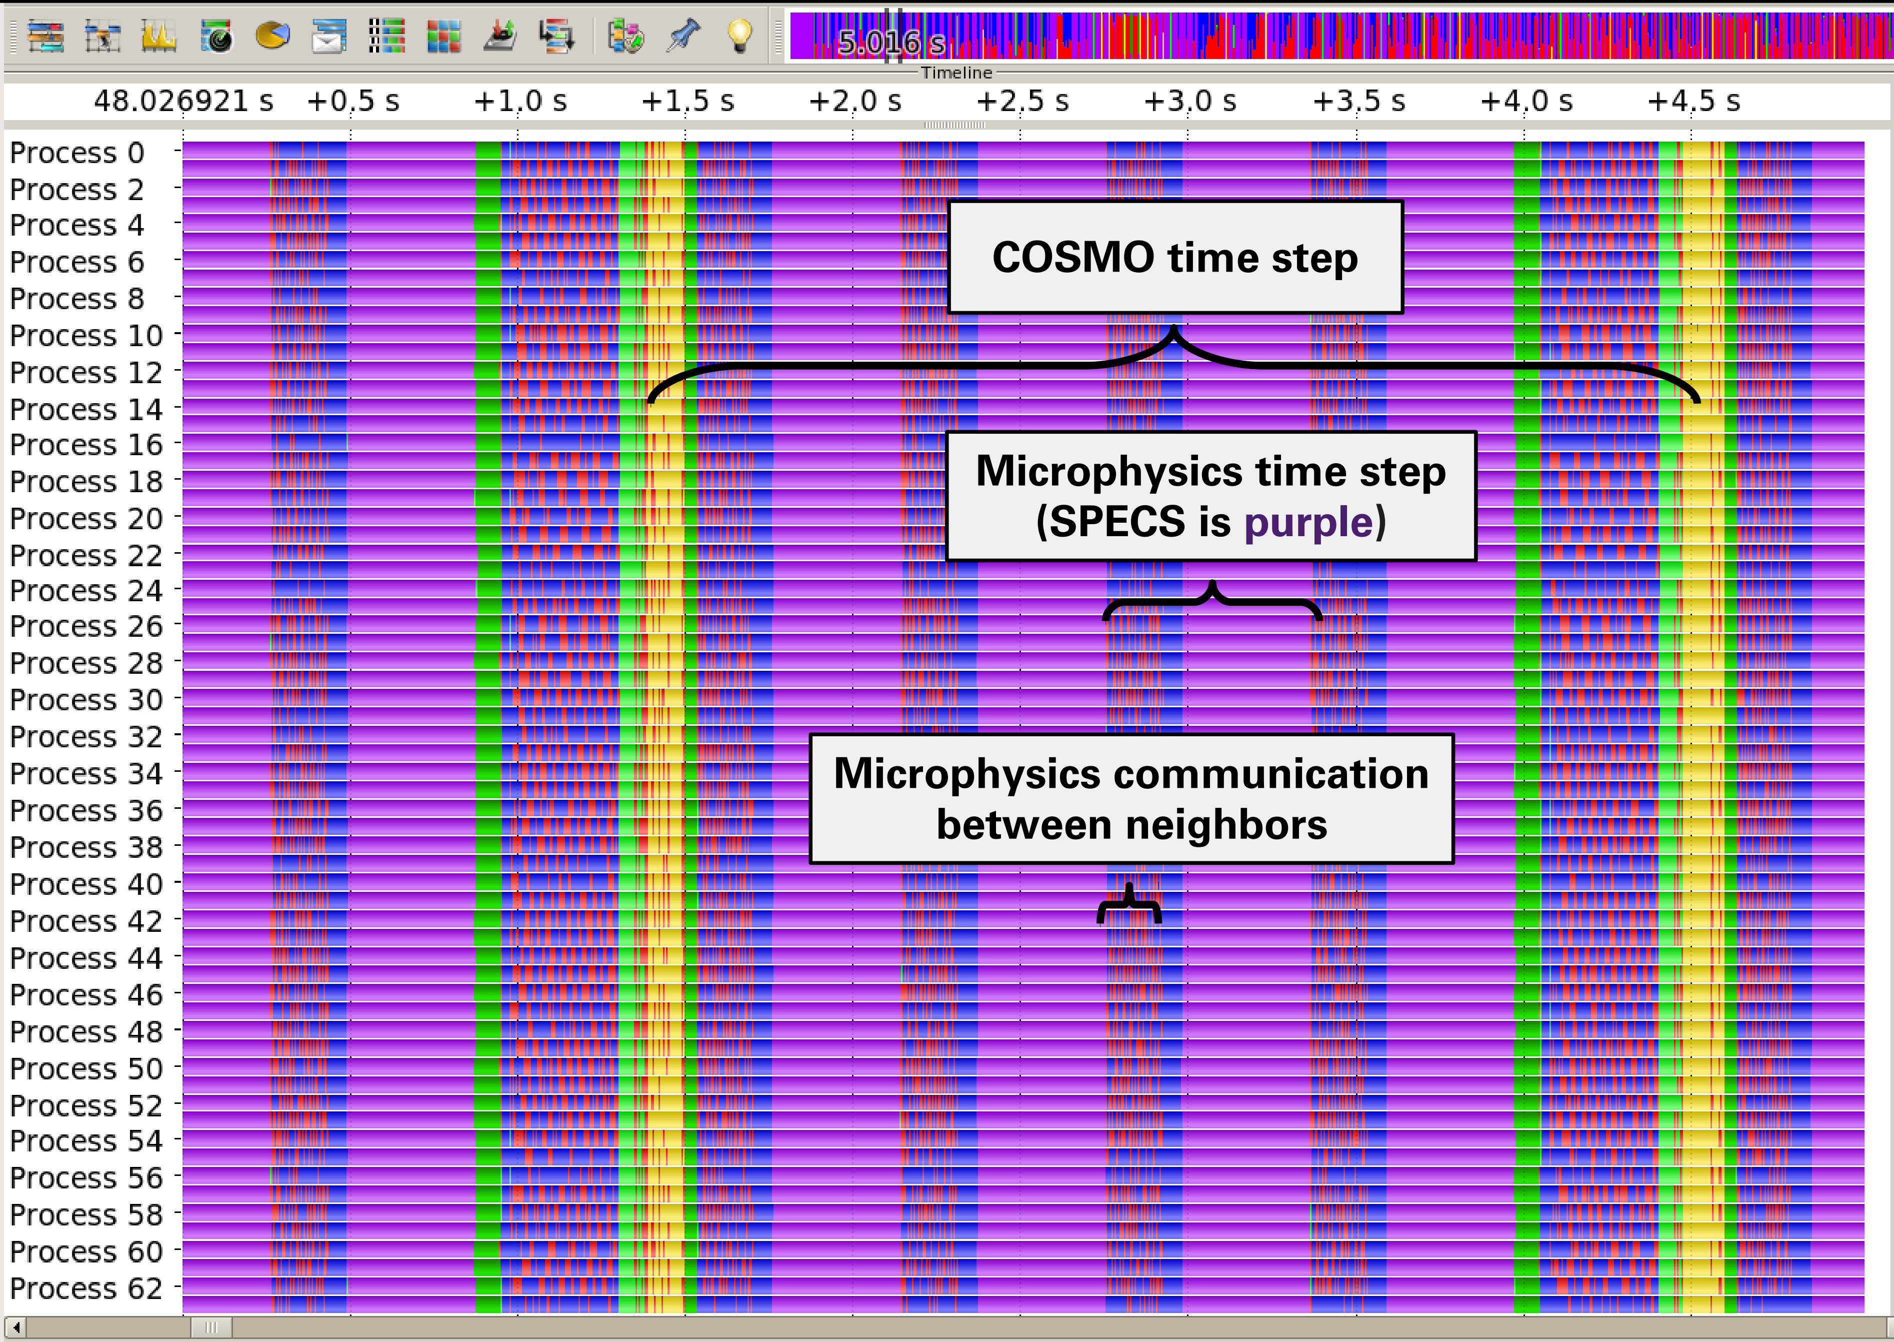

Vampir Global Timeline: Communication between Neighbors

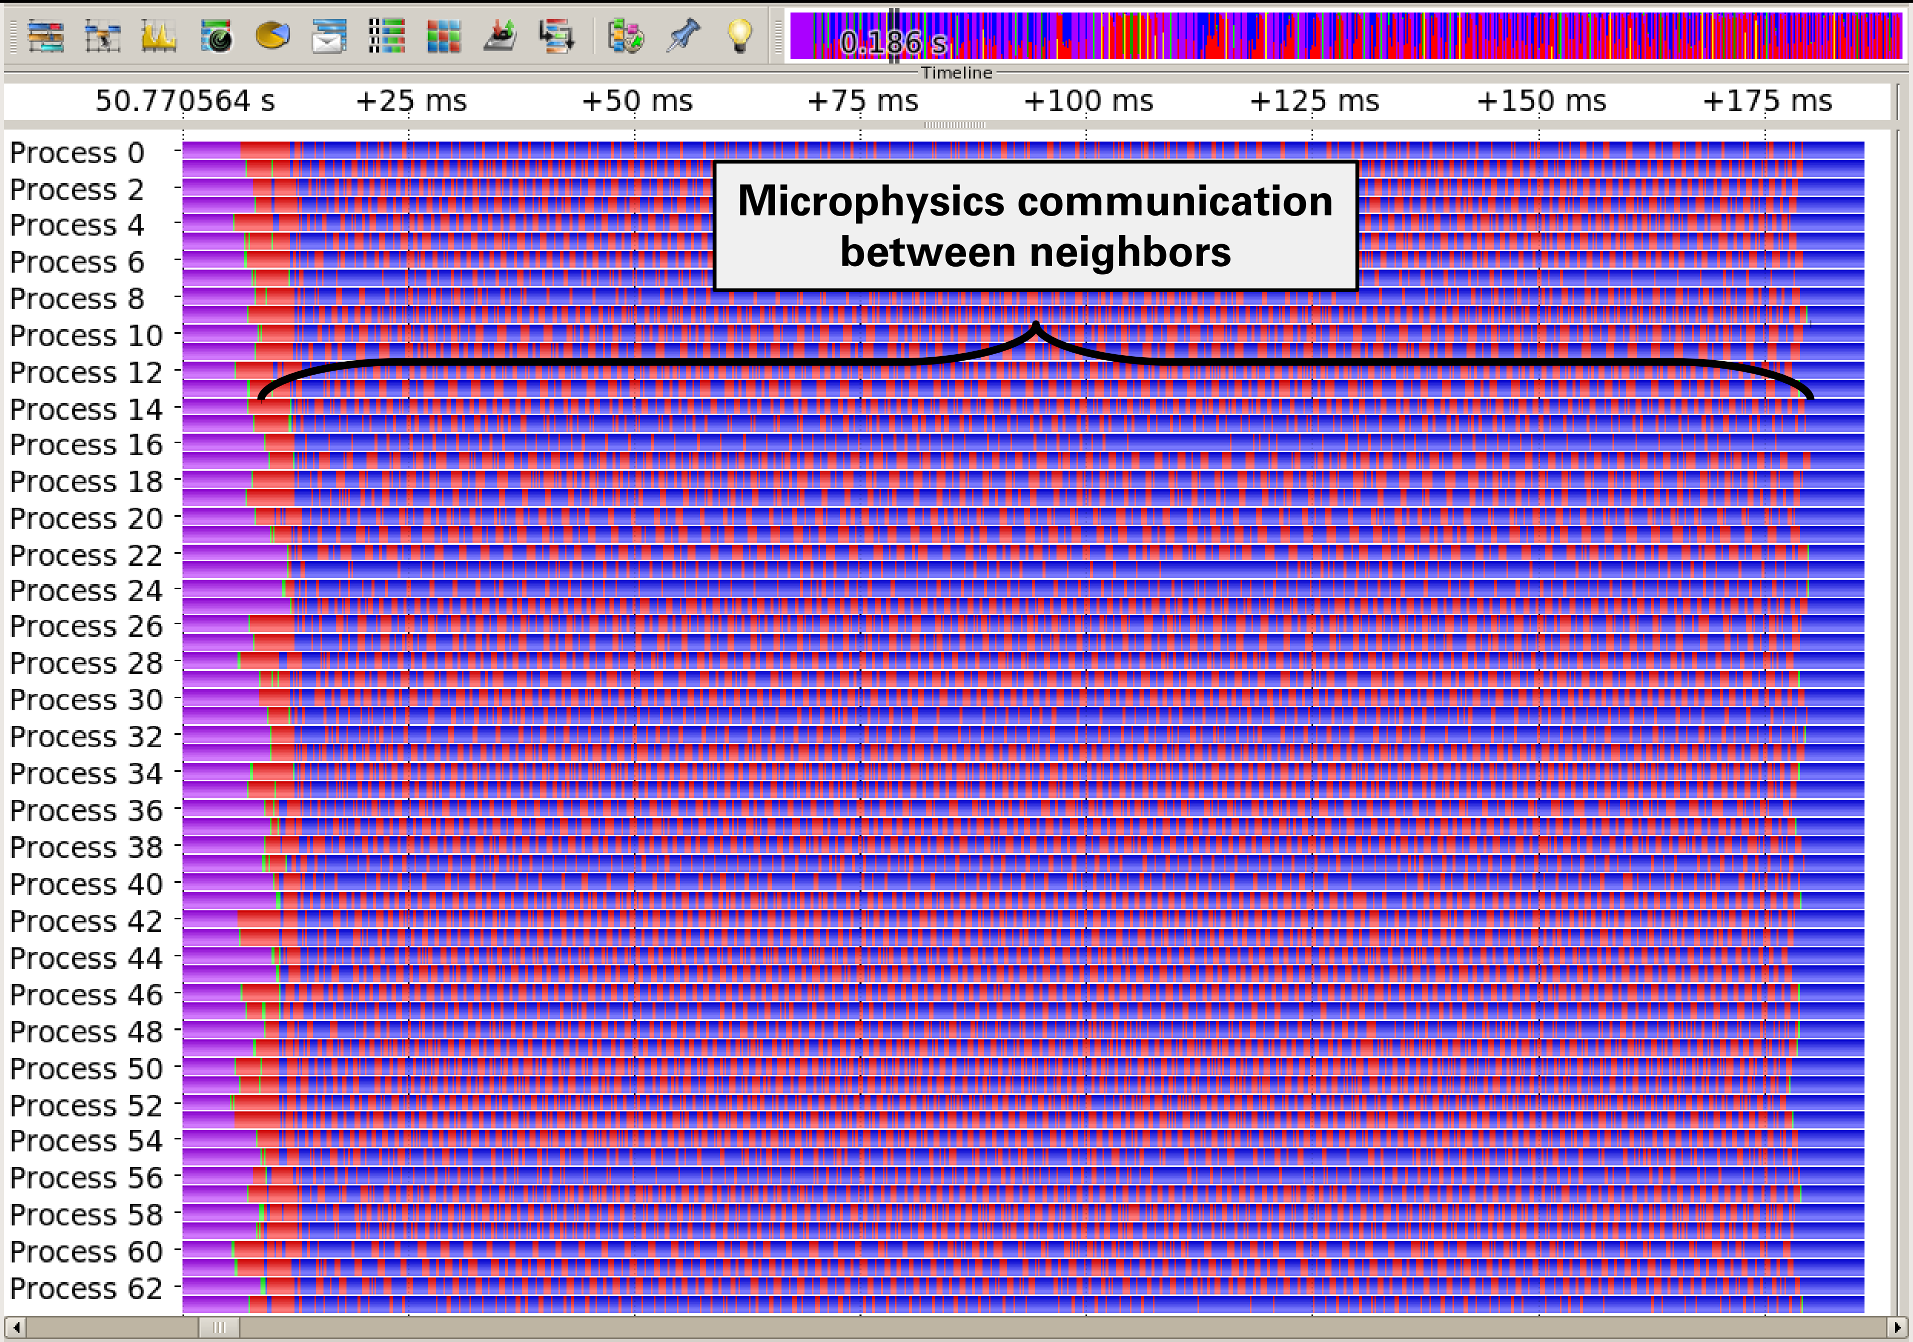

Vampir Global Timeline: Communication between Neighbors

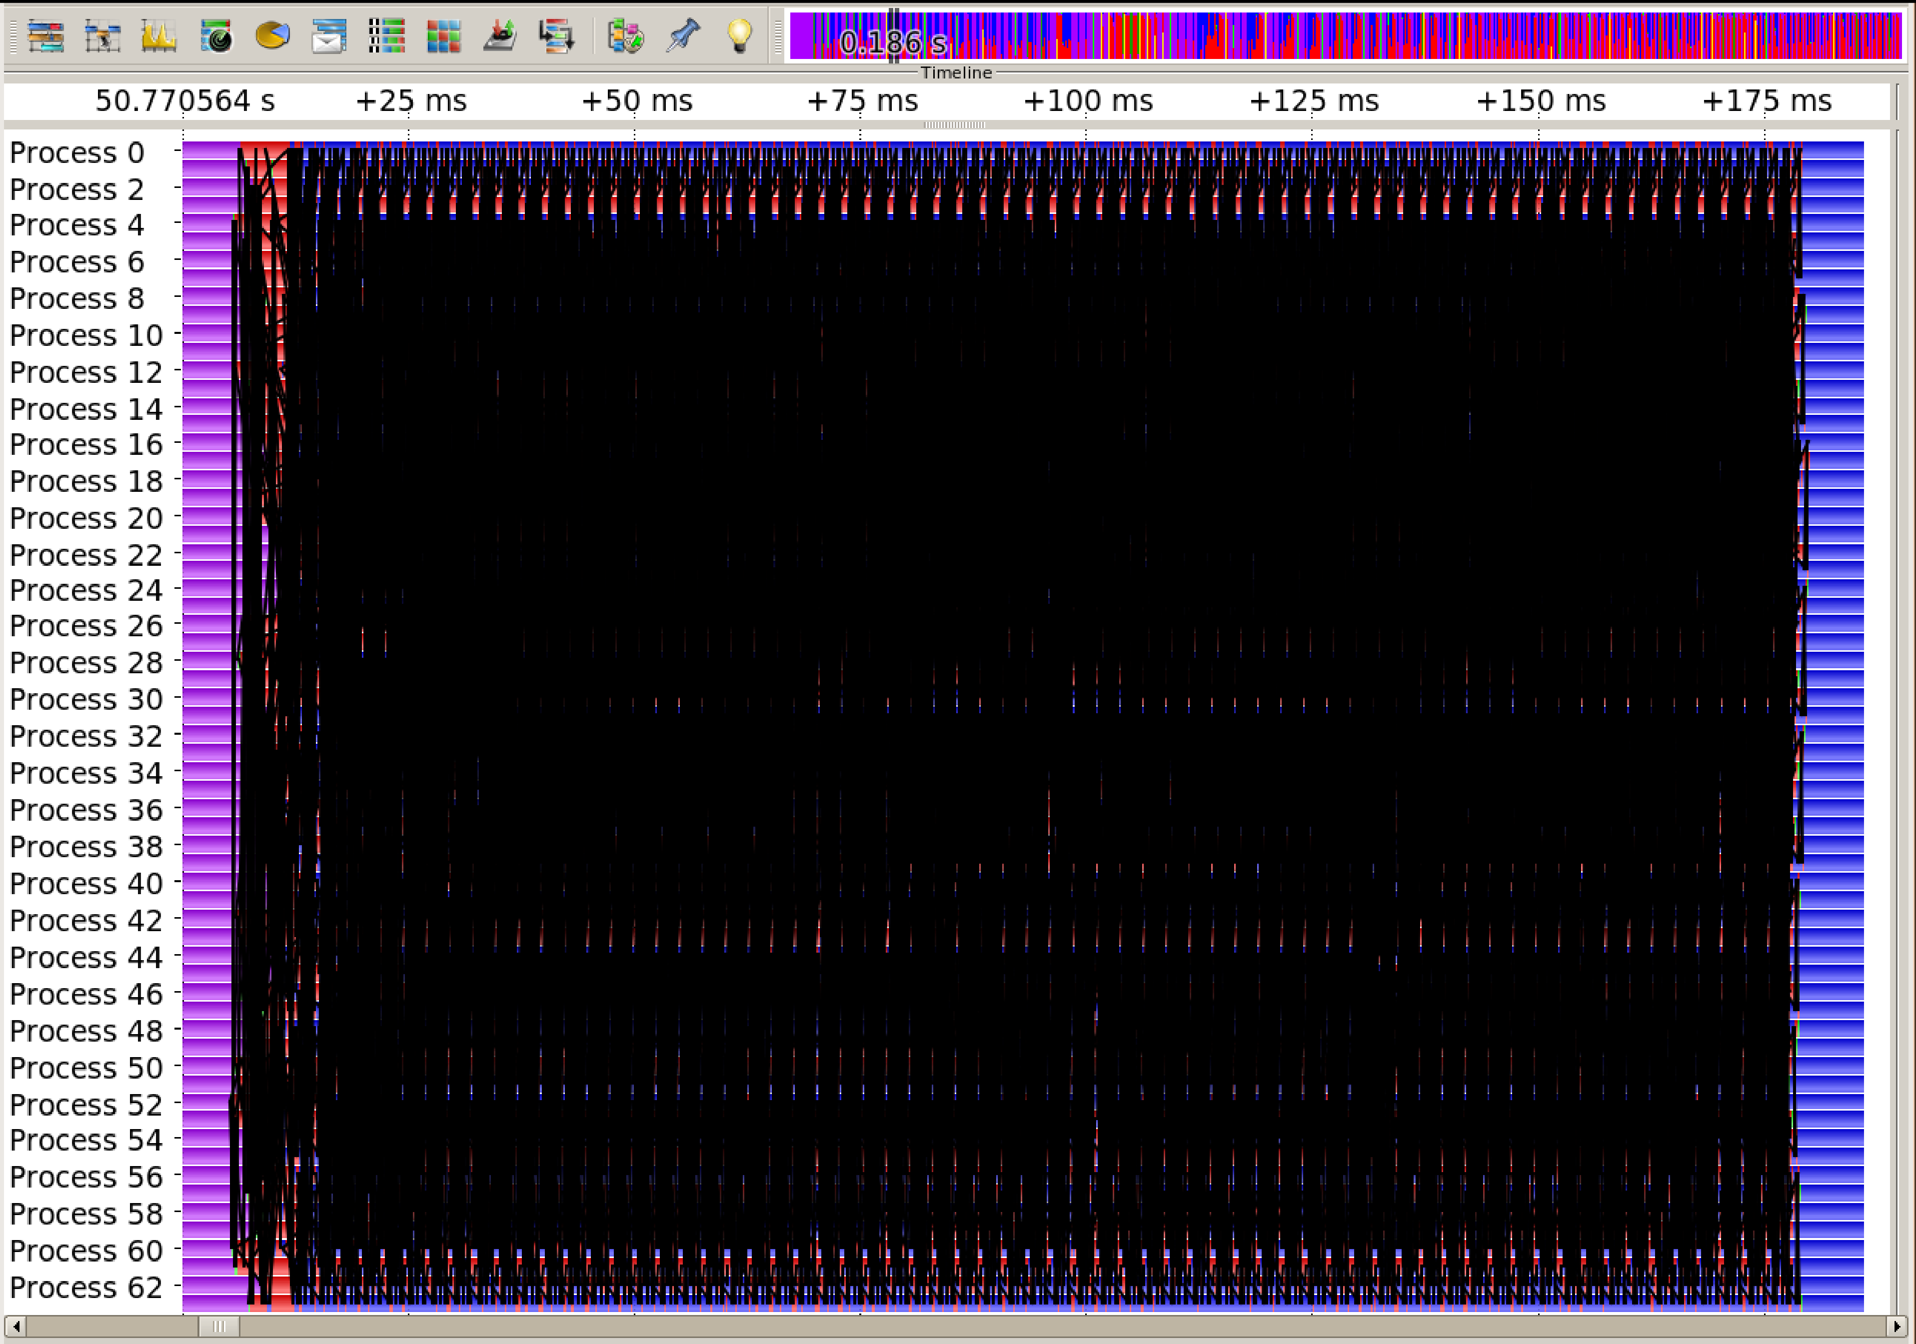

Problem 1: Many small Messages

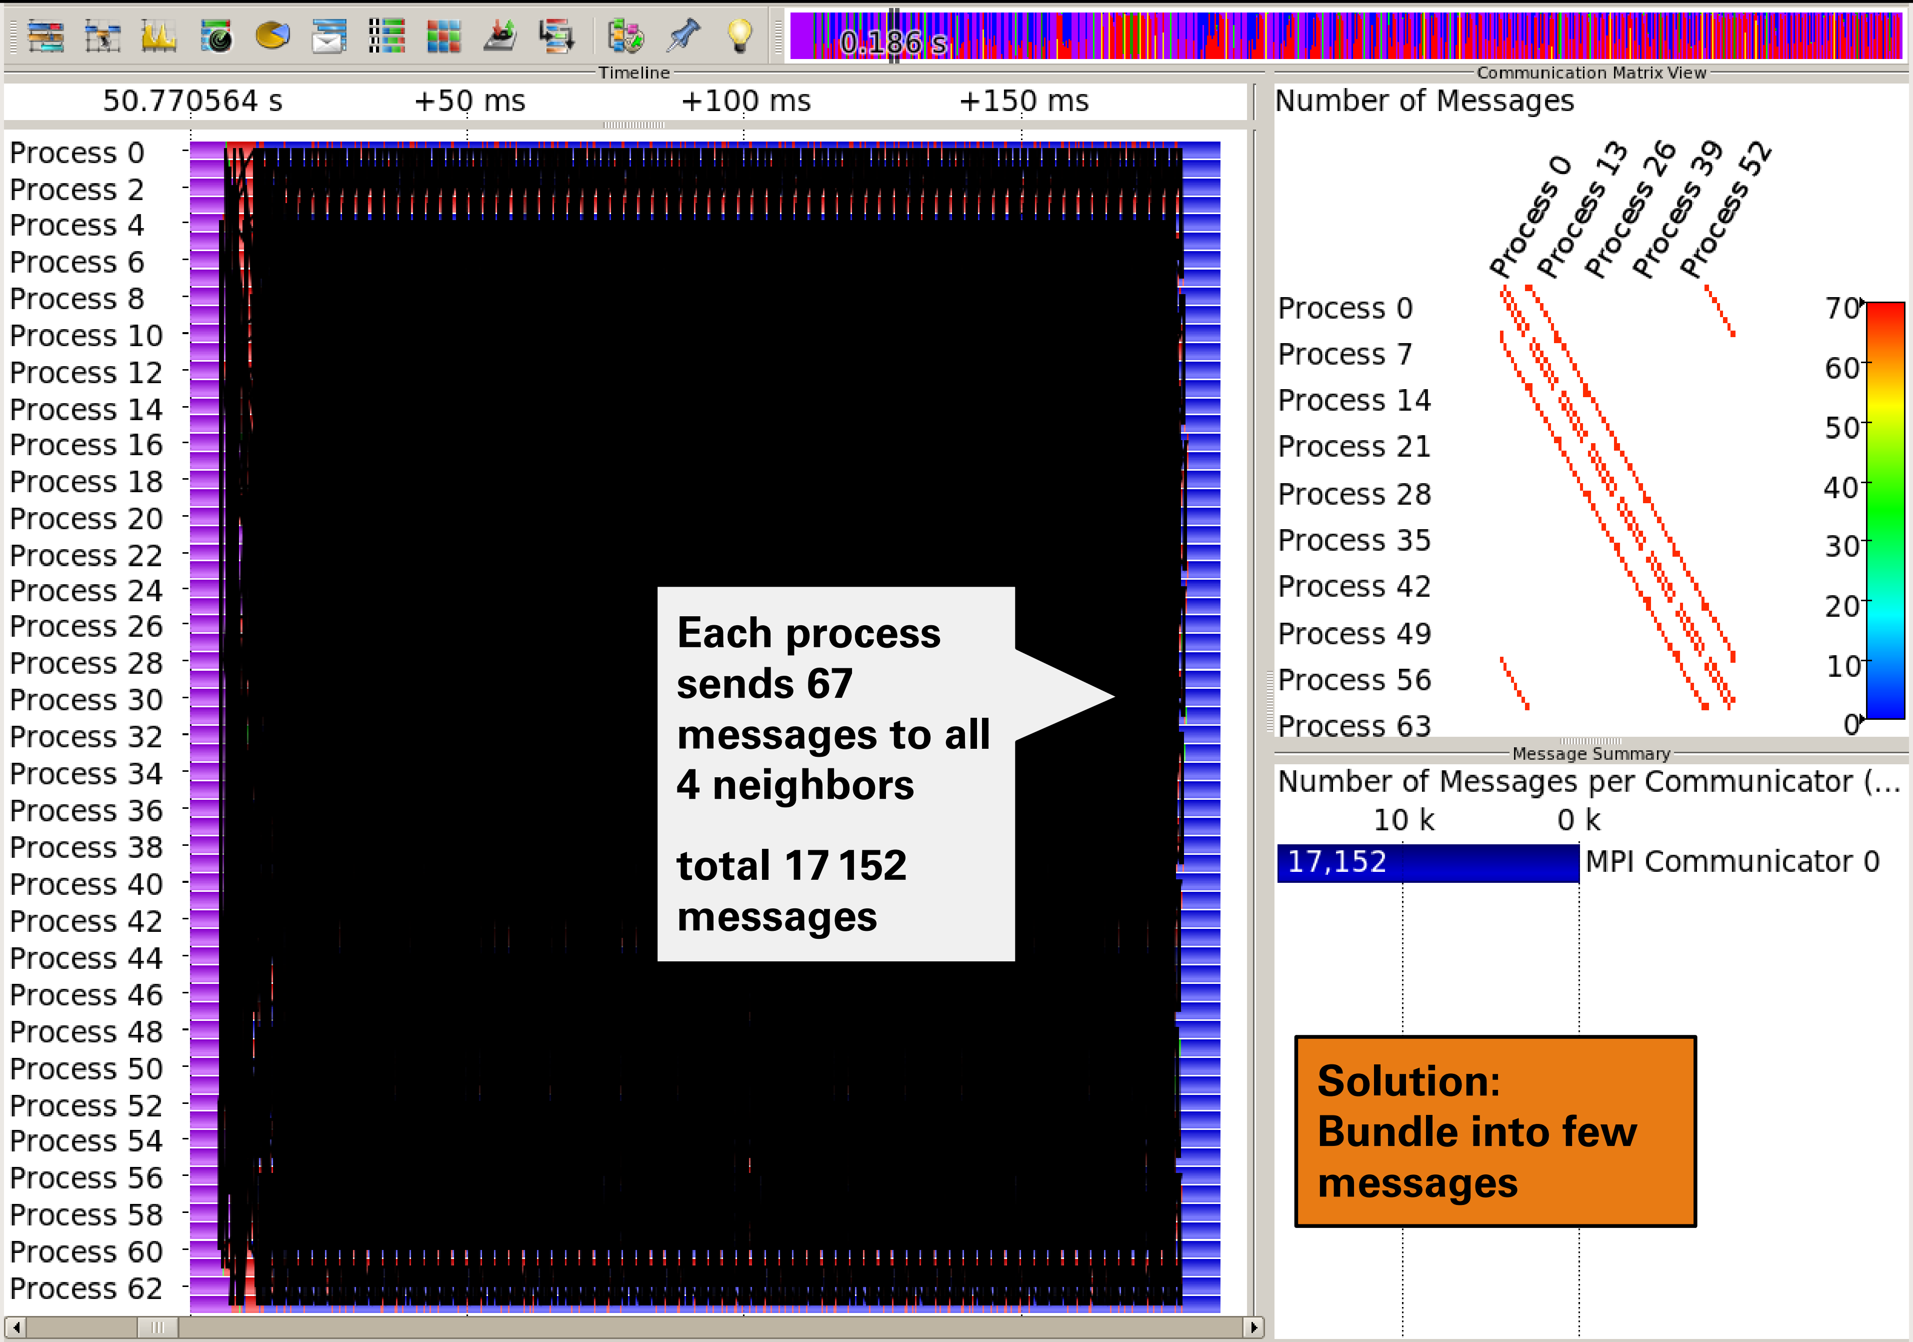

Vampir Global Timeline

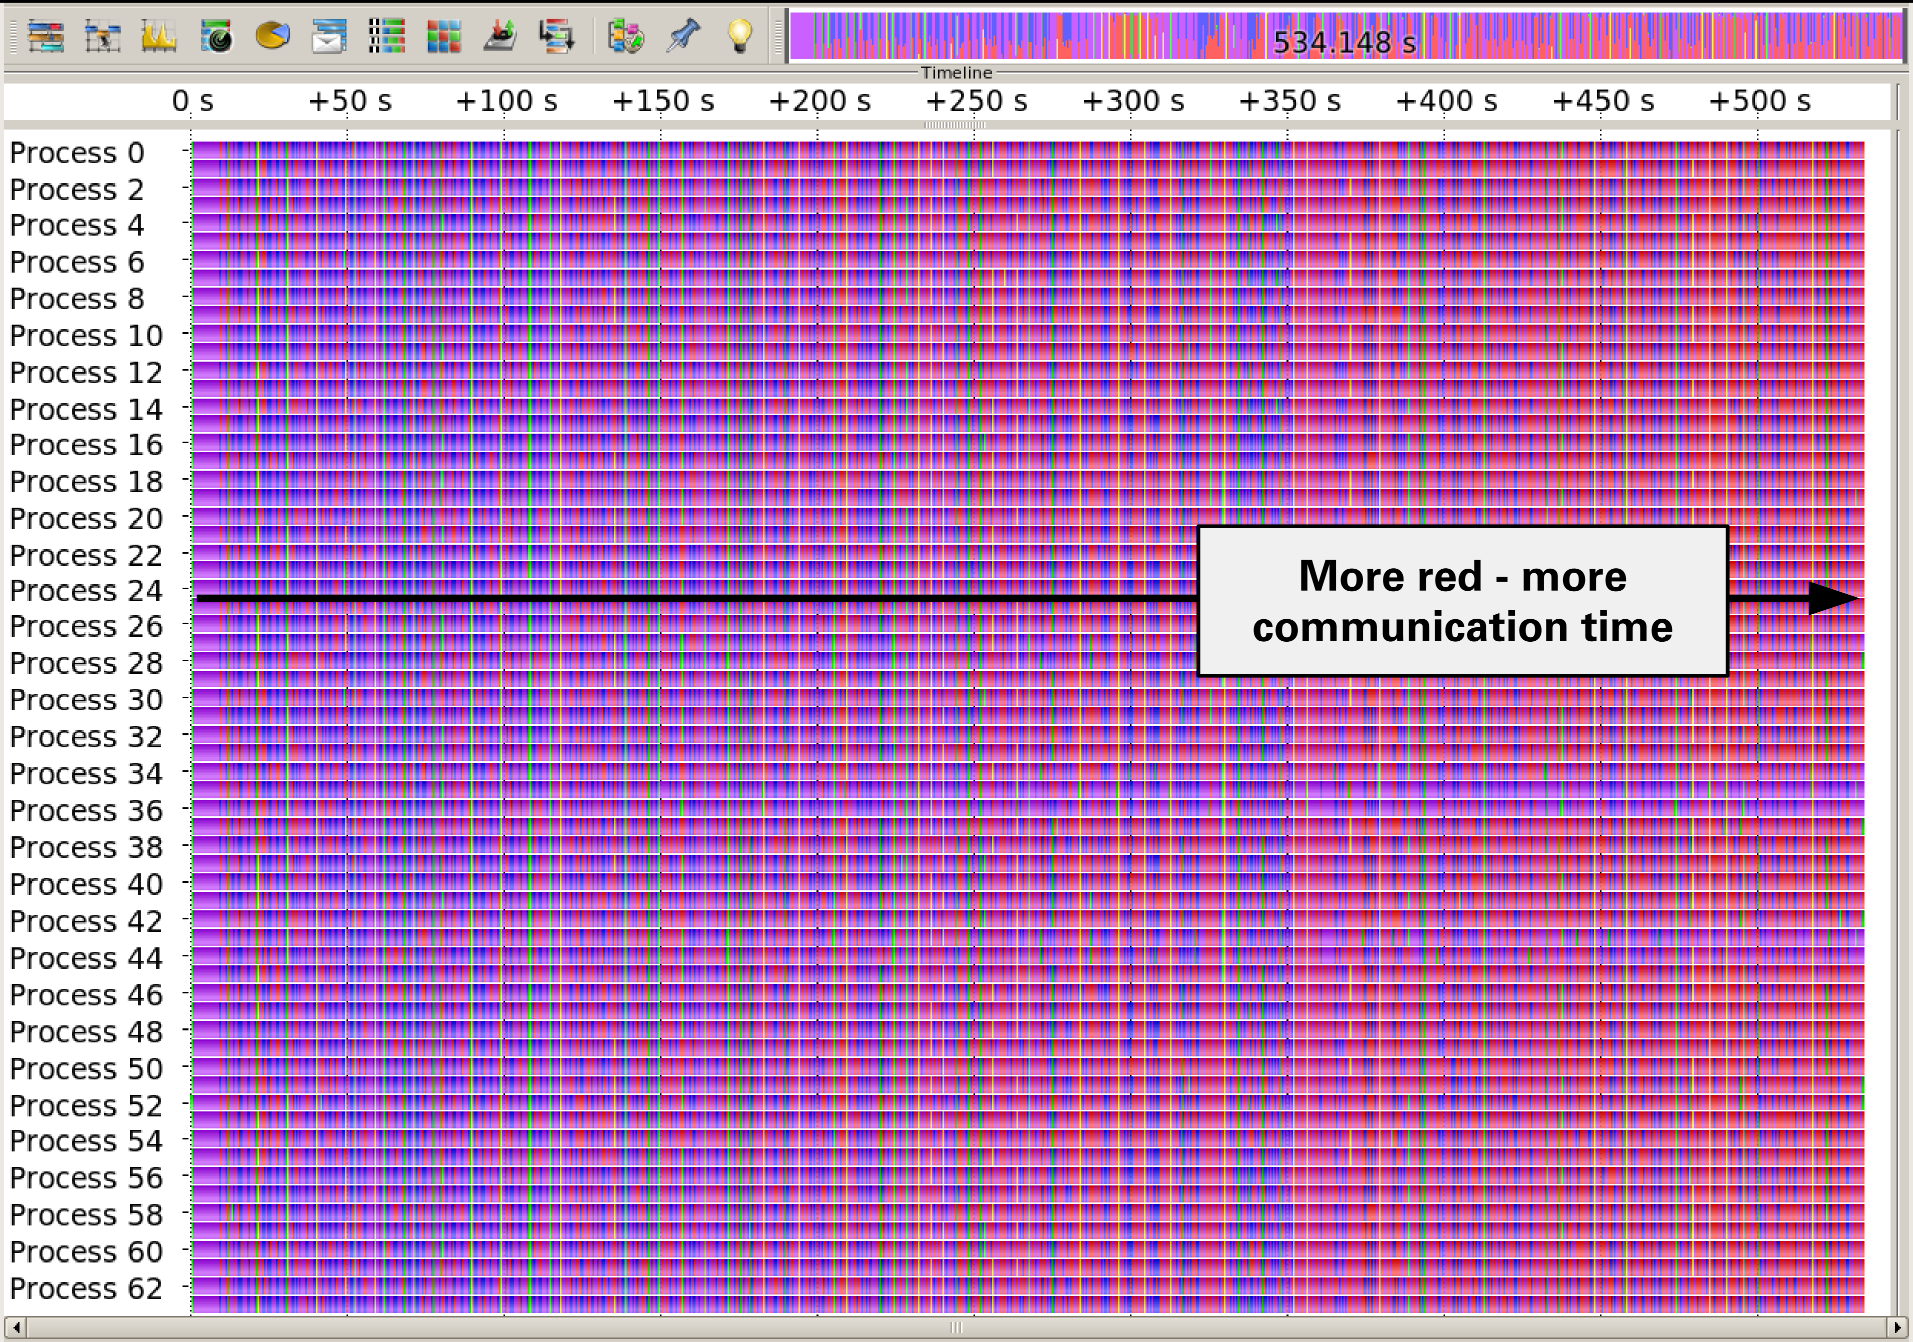

Problem 2: Load Imbalance due to Microphysics

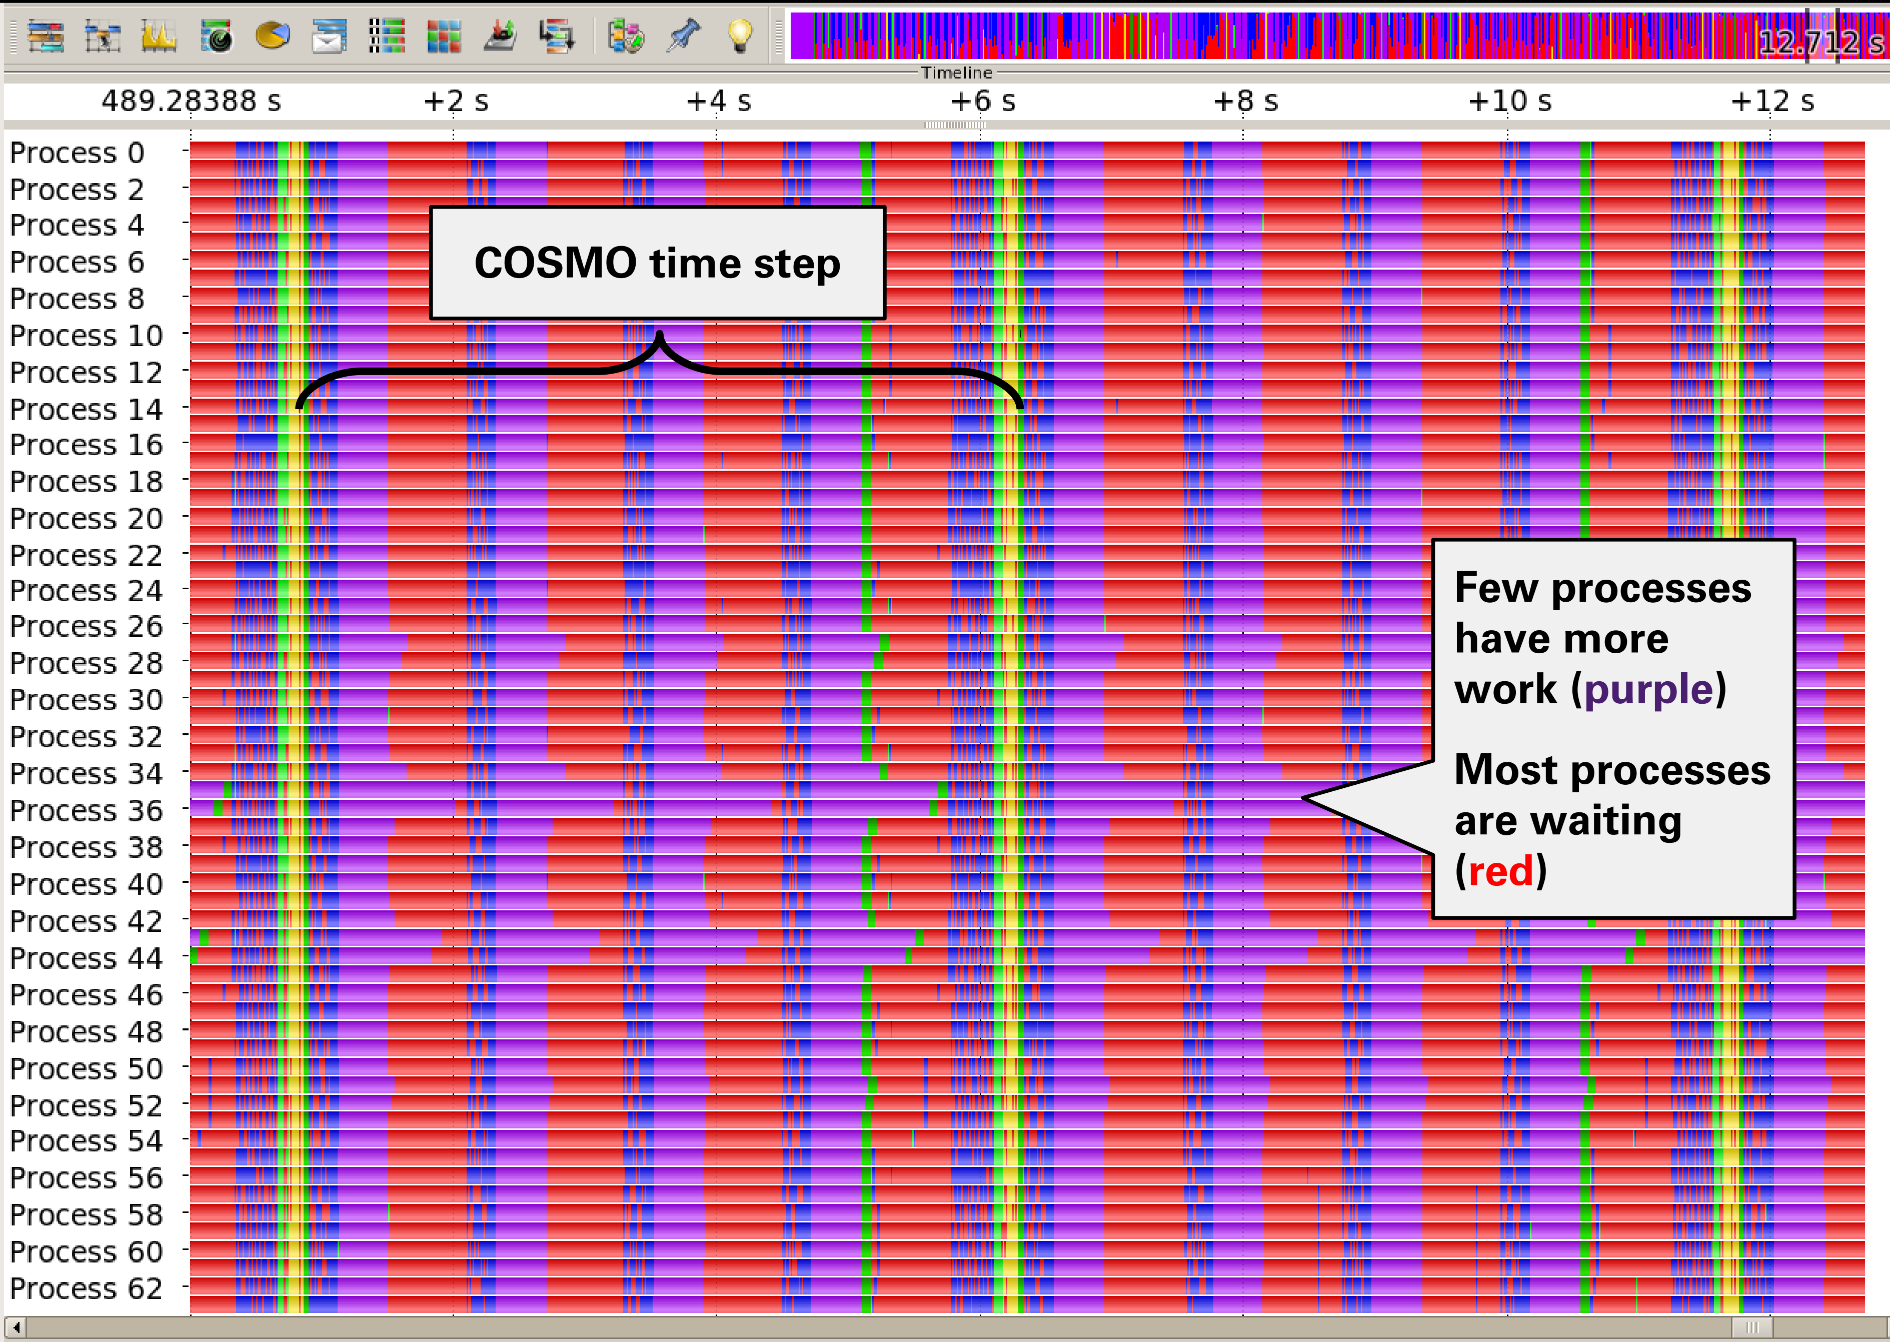

Problem 2: Load Imbalance due to Microphysics

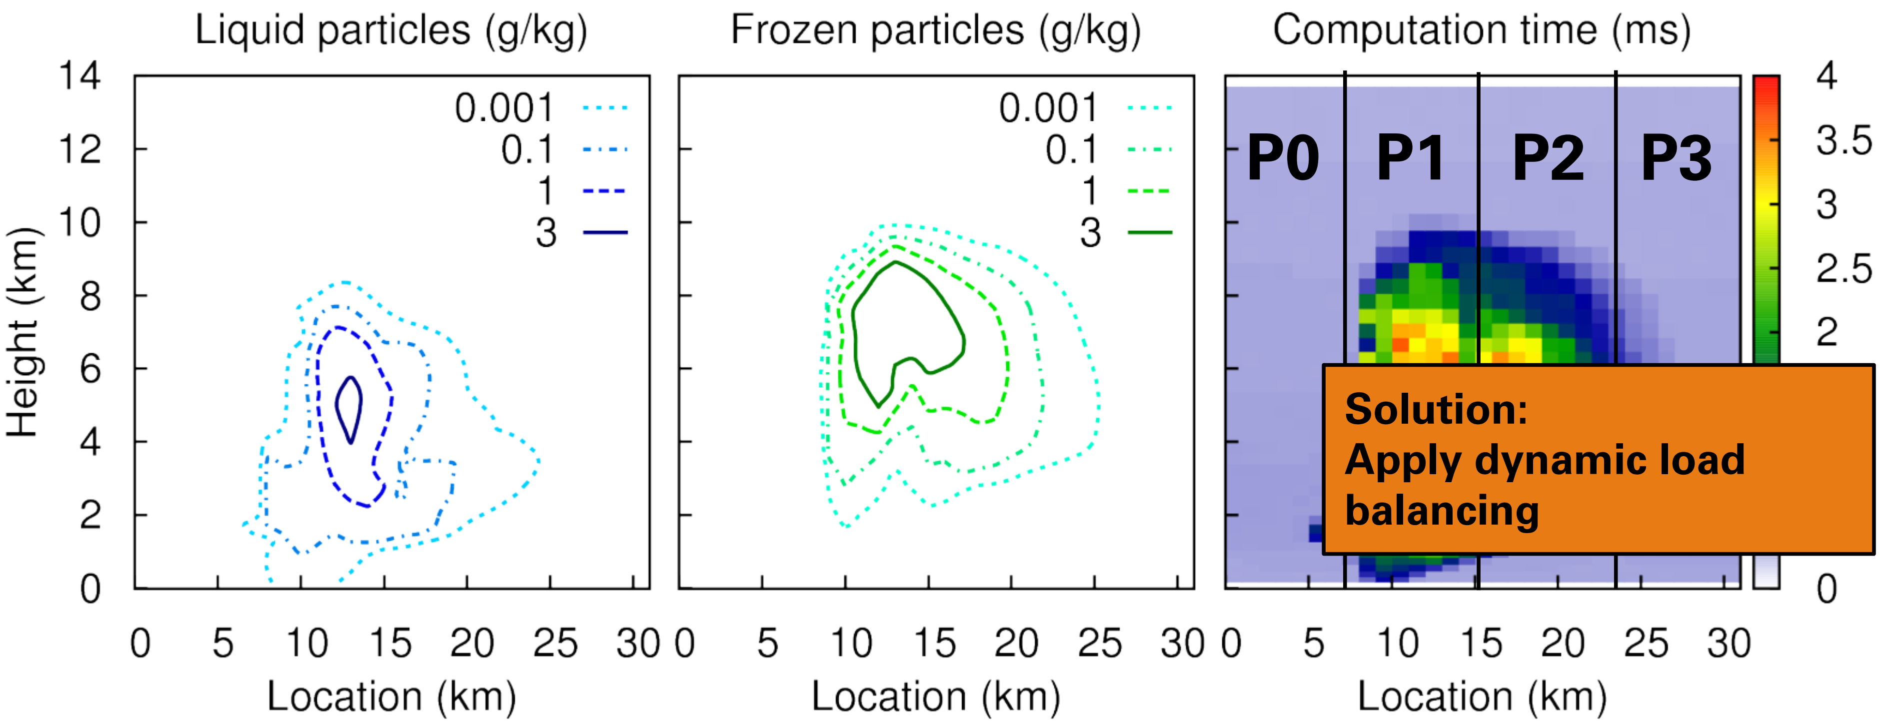

Solution:
Apply dynamic load balancing

- SPECS computing time varies strongly depending on the range of the particle size distribution and presence of frozen particles
- Leads to load imbalances between partitions

Problem 3: Increasing Communication Volume

- Surface-to-volume-ratio of partitions grows with number of partitions, in theory (best case):
  - 2D decomposition: $A^{2D}(P) = 4\ G^{2/3}\ P^{1/2} \sim P^{1/2}$
  - 3D decomposition: $A^{3D}(P) = 6\ G^{2/3}\ P^{1/3} \sim P^{1/3}$

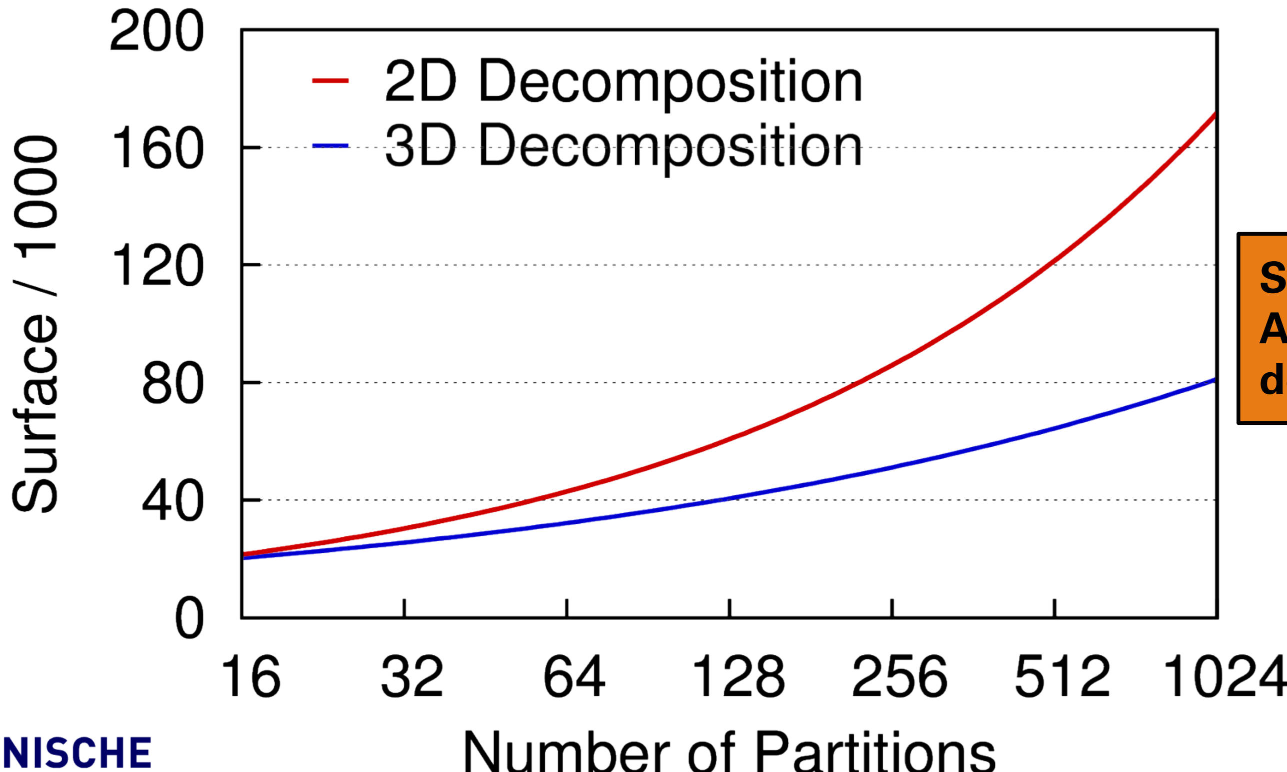

Solution: Apply 3D decomposition

Concept of Load-Balanced Coupling

Atmospheric Model & Spectral Bin Microphysics

Spectral Bin Microphysics

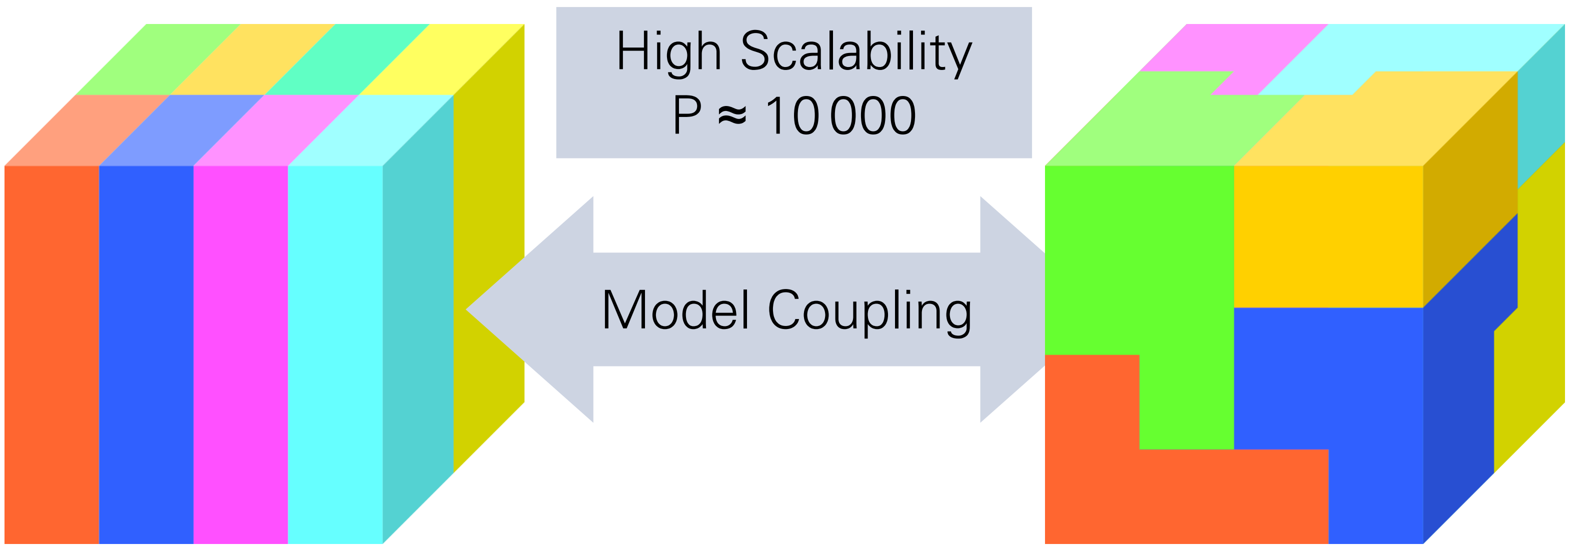

2D Decomposition

Static Partitioning

Block-based 3D Decomposition

Dynamic Load Balancing

Optimized Data Structures

Concept of Load-Balanced Coupling

Implemented as independent framework FD4

High Scalability
P ≈ 10 000

Model Coupling

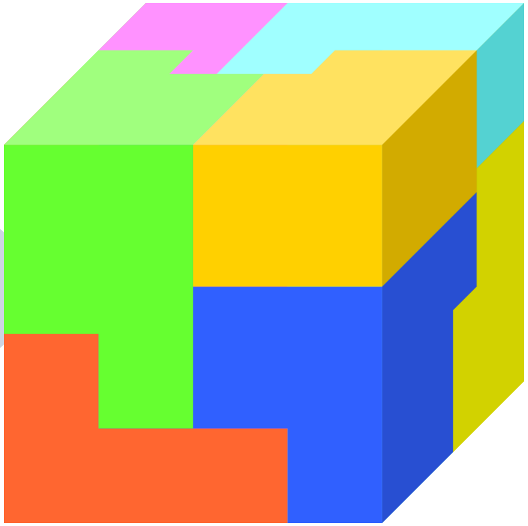

Block-based 3D Decomposition

Dynamic Load Balancing

Optimized Data Structures

Framework FD4: Optimized Data Structure

- A large number of small blocks are good for performance:
  - Size-resolved approach / ~1000 variables per grid cell: Only small blocks do not exceed processor cache
  - Load balancing: #blocks > #partitions to enable fine-grained balancing
- Additional memory costs for a boundary of ghost cells
  - Too high for small blocks!
- Add ghost blocks at the partition borders only

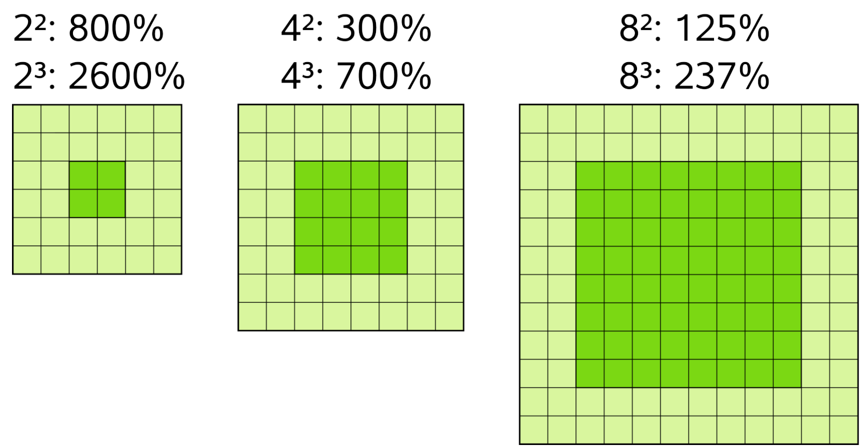

Framework FD4: Dynamic Load Balancing

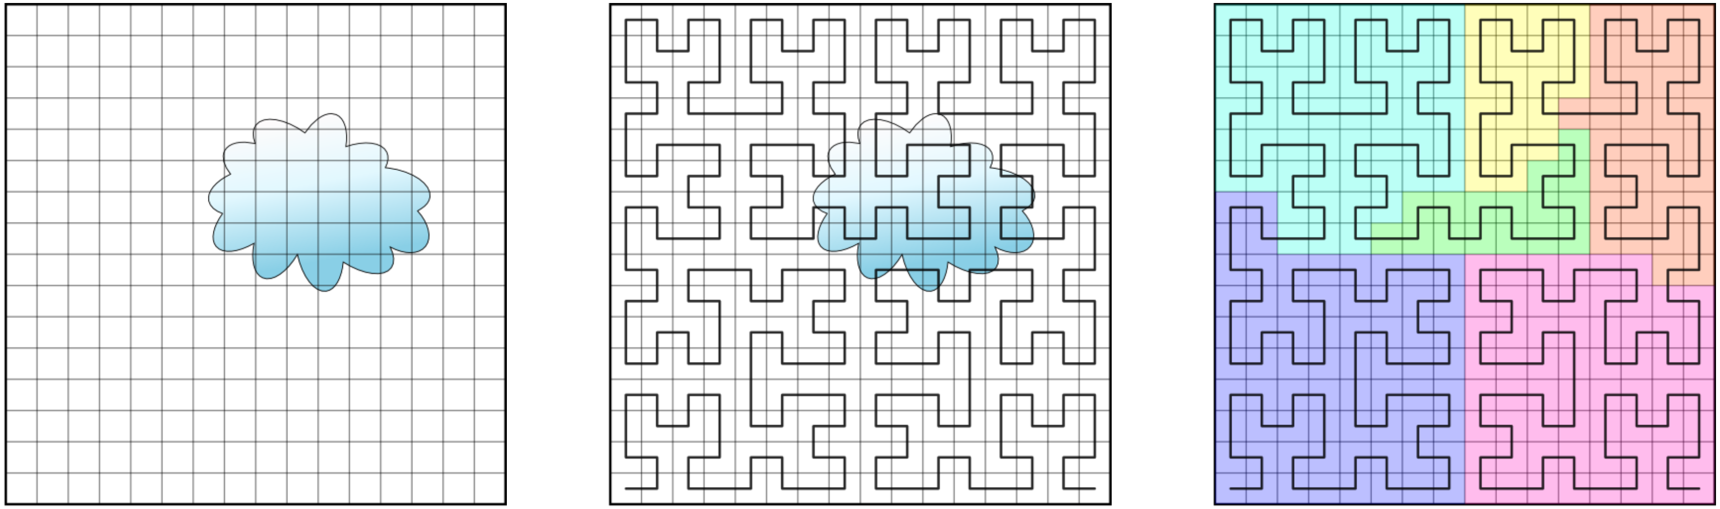

- Based on 3D block decomposition of grid
- Task: Assign blocks to processes in a balanced way which also minimizes communication costs. And do this quickly.
- Space-filling curve (SFC) partitioning [Sagan94, Teresco06]
- SFC reduces 3D partitioning problem to 1D

COSMO-SPECS+FD4: Visualization of Load Balancing

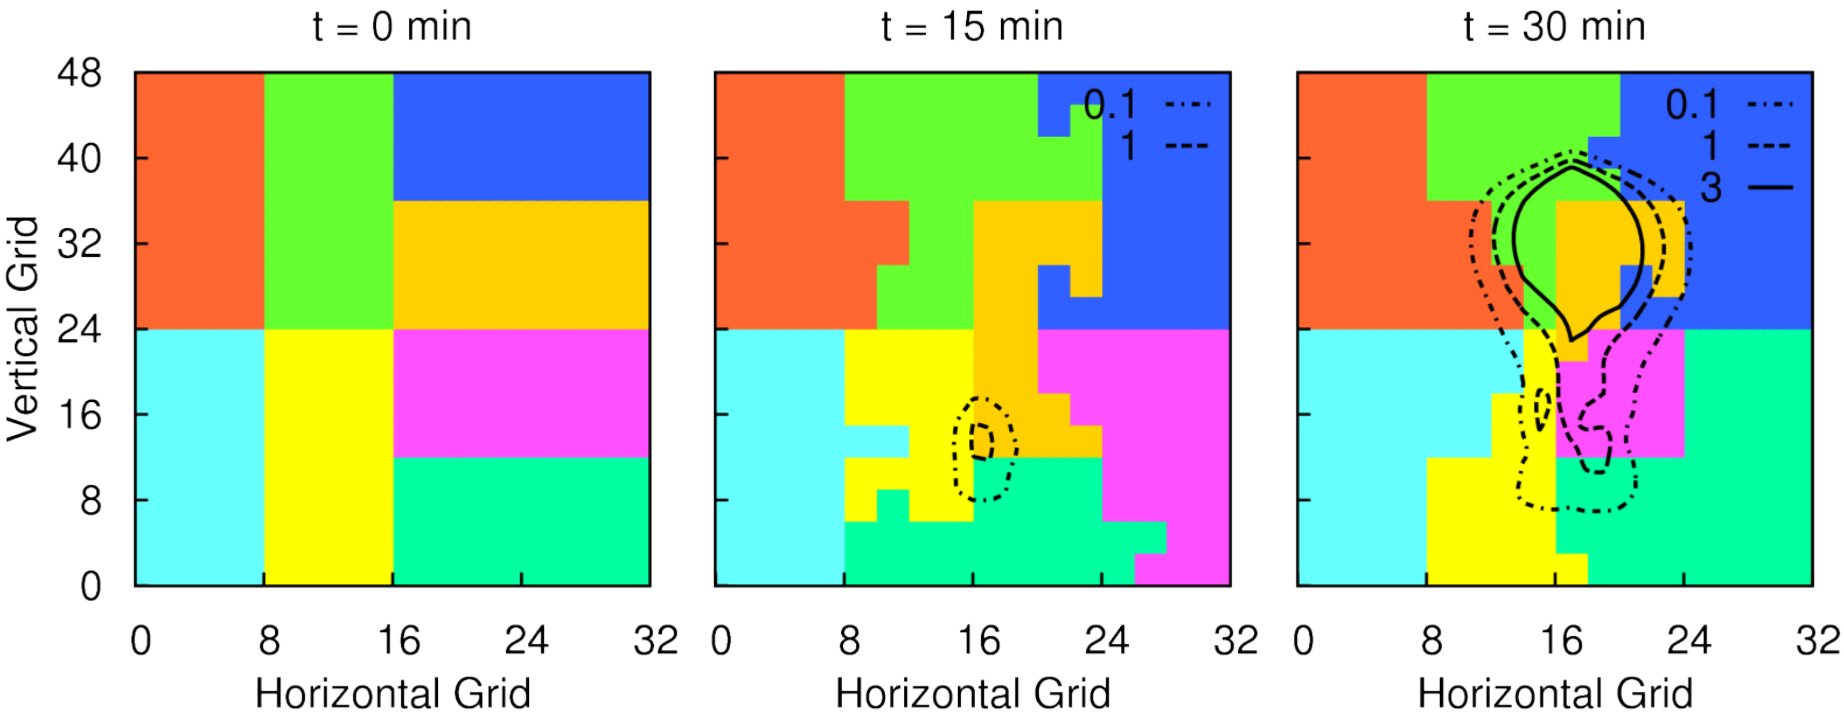

- Small 2D case, 32 x 48 grid, 16 x 16 FD4 blocks
- Colors: 8 partitions for SPECS balanced by FD4
- Contours: liquid and frozen particle concentration (g/kg)

Load Balance Comparison

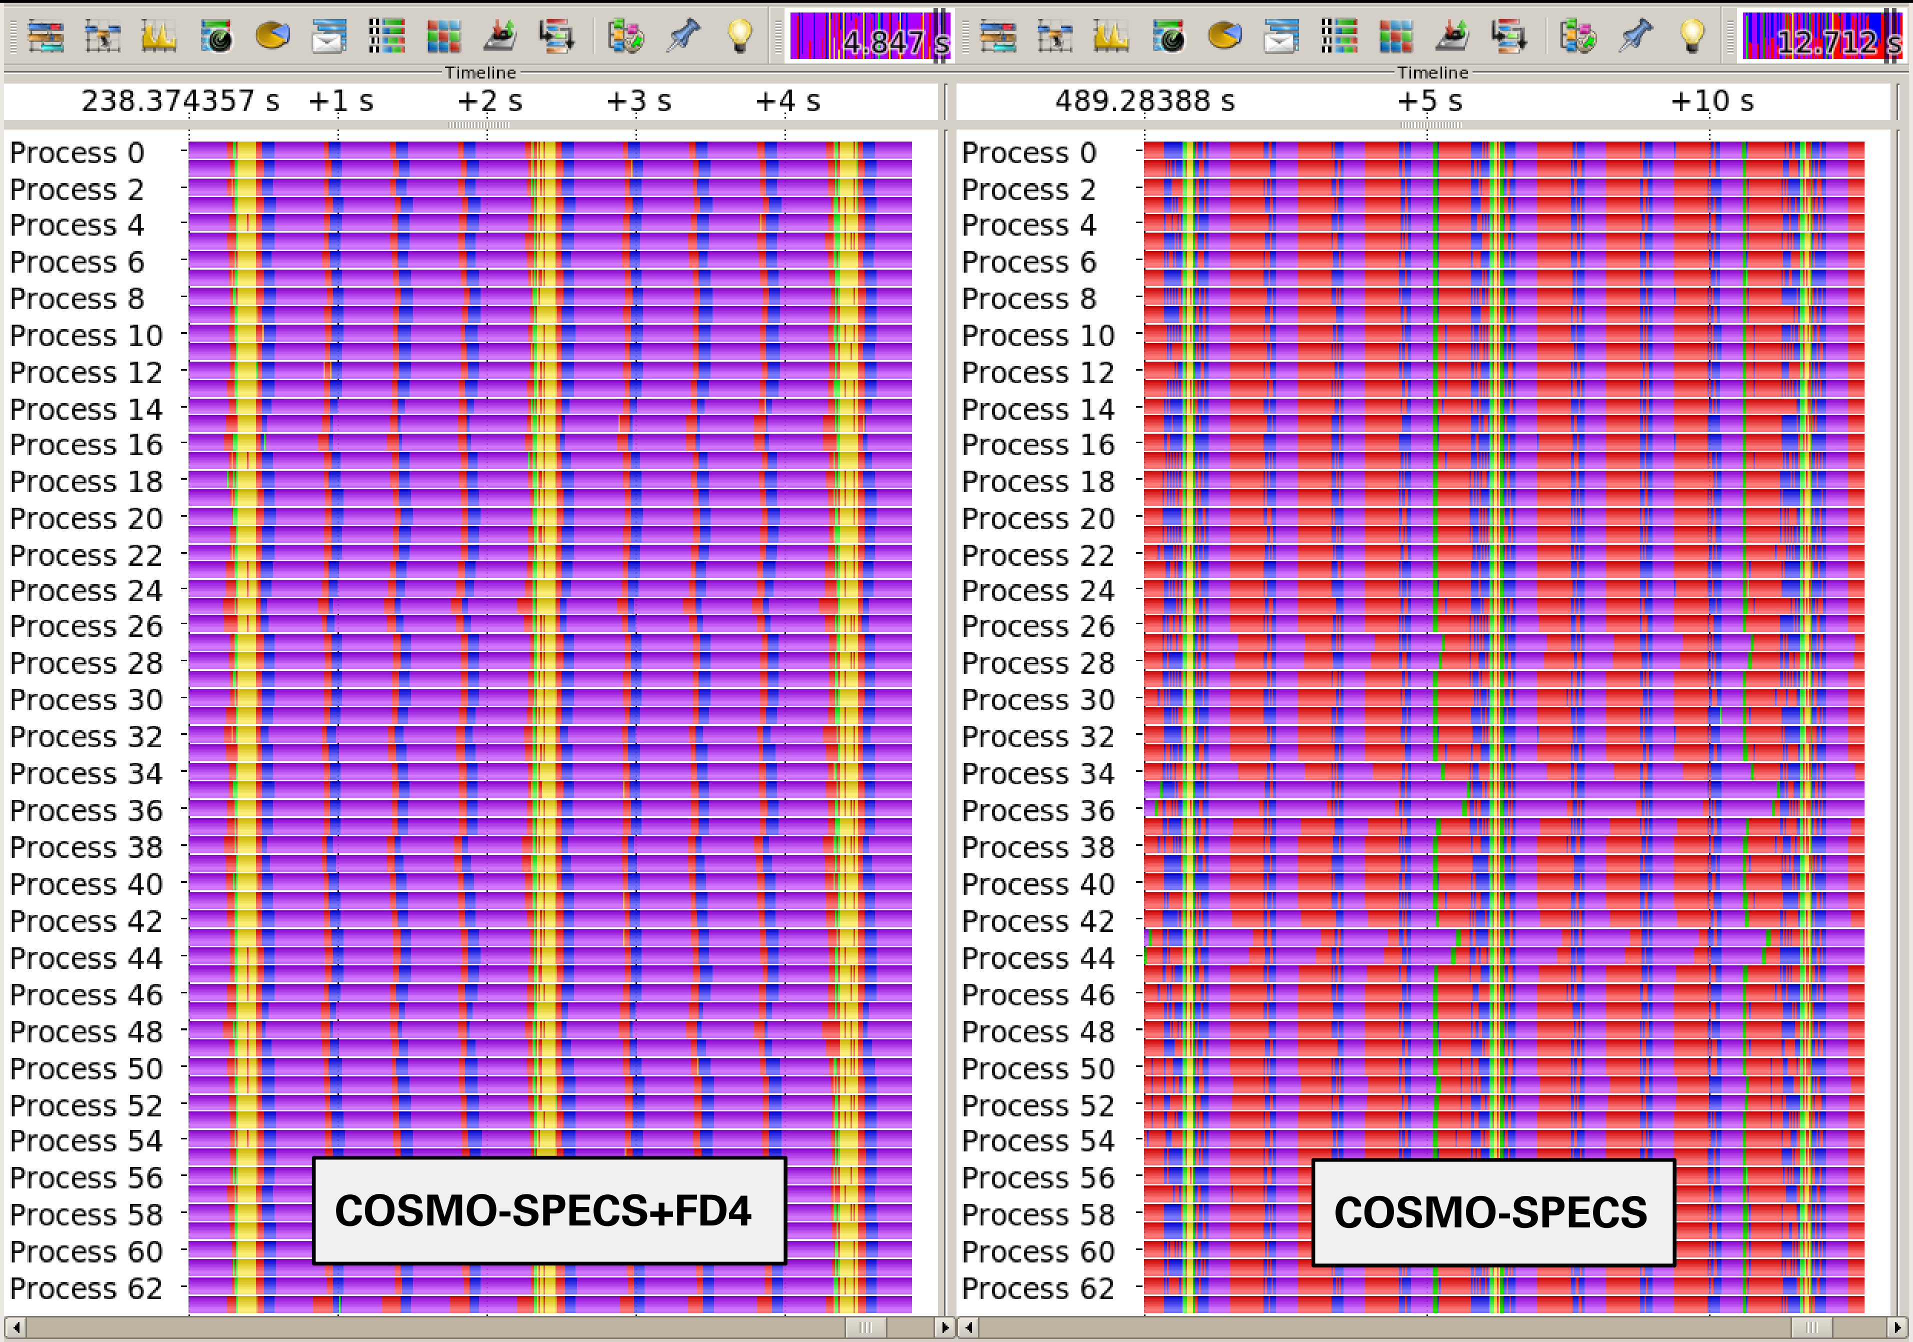

COSMO-SPECS+FD4: Performance Comparison

- Almost 3 times faster at 1024 CPU cores
- Load balancing & coupling scale well, but can we reach > 10 000 processes?

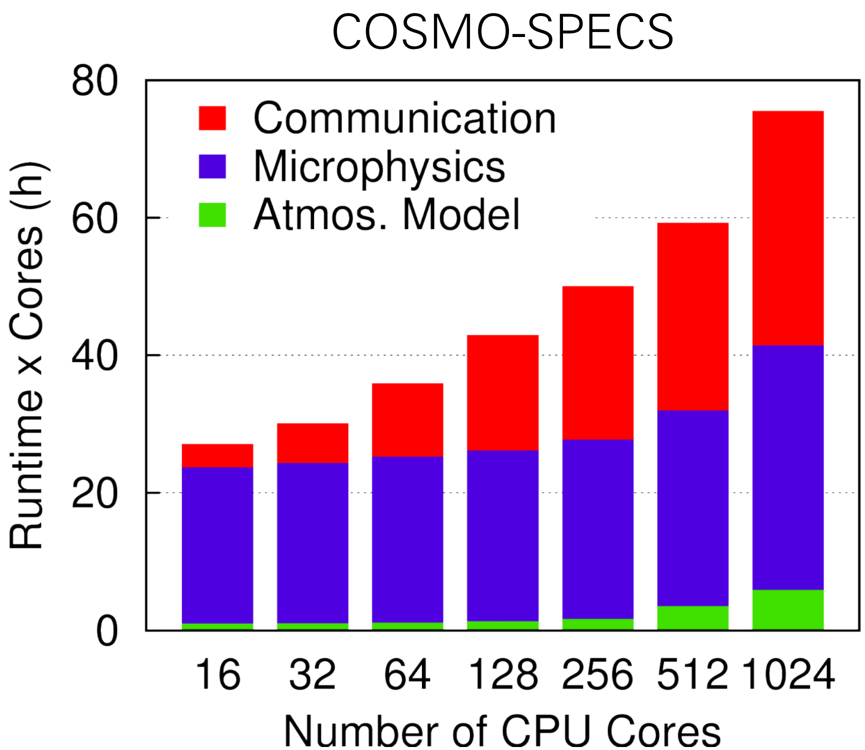

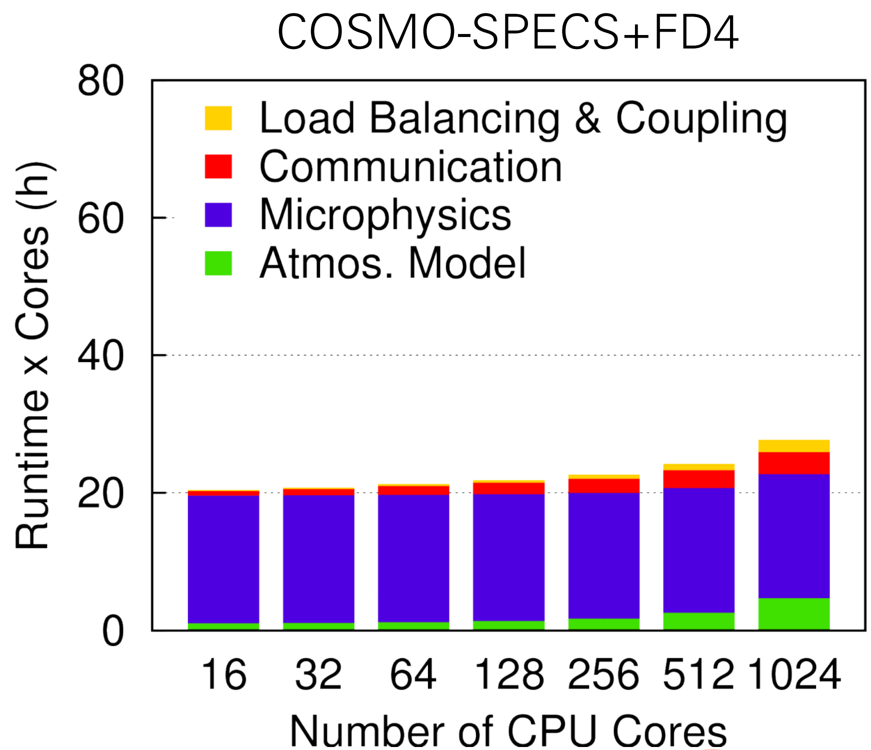

COSMO-SPECS+FD4: High Scalability Benchmark

- IBM BlueGene/P: 294 912 CPU cores
- Resume previous benchmark at 1024 cores as weak scaling
- i.e. grid size grows with number of CPU cores

CPU cores	Grid size	FD4 blocks
1 024	64 x 64 x 48	12 288
2 048	128 x 64 x 48	24 576
...	...	...
131 072	1024 x 512 x 48	1 572 864

COSMO-SPECS+FD4: High Scalability Benchmark

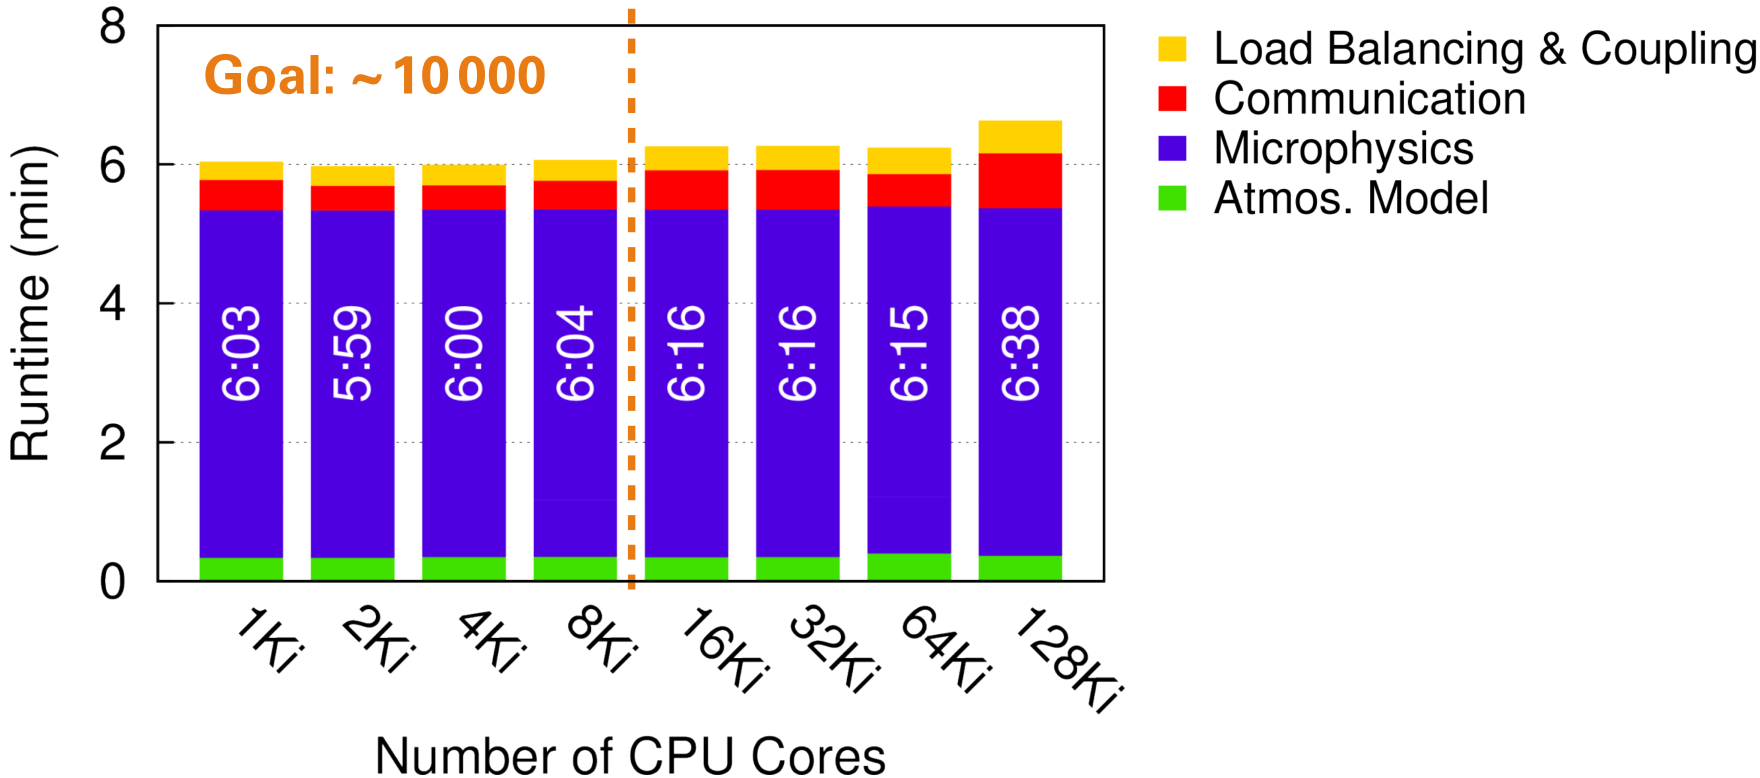

- 30 min forecast on 1024 x 512 x 48 grid in 6:38 min
- Enables large-scale applications of spectral bin microphysics with efficient HPC usage

COSMO-SPECS+FD4: Scalability Challenges

- SFC-based load balancing: Highly scalable and high-quality 1D partitioning algorithm for >1 million blocks
- Model coupling: Identify overlaps between 10 000s of (dynamically changing) partitions
- Avoid global (meta)data as much as possible

Conclusion & Outlook

- FD4 provides an efficient and scalable way of coupling spectral bin microphysics with atmospheric models
- Large-scale studies with spectral bin microphysics feasible
- FD4 is not limited to spectral bin microphysics!
- FD4 is open source: http://www.tu-dresden.de/zih/clouds

- Outlook:
  - Investigate adaptive time stepping for SPECS
  - Coupling with MUSCAT
  - Parallel I/O in FD4

Thank you for your attention!

Refences

[Grützun08] V. Grützun, O. Knoth, and M. Simmel. *Simulation of the influence of aerosol particle characteristics on clouds and precipitation with LM–SPECS: Model description and first results.* Atmos. Res., 90:233–242, 2008.

[CCSP09] *Atmospheric Aerosol Properties and Climate Impacts*, U.S. Climate Change Science Program and the Subcommittee on Global Change Research, 2009.

[IPCC07] S. Solomon et al. *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Forth Assessment Report of the Intergovernmental Panel on Climate Change (IPCC)*, 2007.

[Lieber12] M. Lieber, V. Grützun, R. Wolke, M.S. Müller, and W.E. Nagel. *Highly Scalable Dynamic Load Balancing in the Atmospheric Modeling System COSMO-SPECS+FD4*, in Proc. of PARA 2010, LNCS Vol. 7133, pp. 131-141, 2012

[LieberPhD] M. Lieber, Dynamische Lastbalancierung und Modellkopplung zur hochskalierbaren Simulation von Wolkenprozessen, PhD Thesis, Technische Universität Dresden, 2012

[Pinar04] A. Pinar and C. Aykanat. *Fast optimal load balancing algorithms for 1D partitioning.* J. Parallel Distrib. Comput., 64(8):974-996, 2004.

[Sagan94] H. Sagan. *Space-filling curves*, Springer, 1994

[Simmel06] M. Simmel and S. Wurzler. *Condensation and activation in sectional cloud microphysical models*, Atmos. Res., 80:218-236, 2006.

[Teresco06] J.D. Teresco, K.D. Devine, and J.E. Flaherty. *Partitioning and Dynamic Load Balancing for the Numerical Solution of Partial Differential Equations,* in Numerical Solution of Partial Differential Equations on Parallel Computers, pages 55-88, Springer, 2006.