TECHNISCHE UNIVERSITÄT DRESDEN

Center for Information Services and High Performance Computing (ZIH)

Scalable Dynamic Load Balancing of Detailed Cloud Physics with FD4

Minisymposium on
Advances in Numerics and Physical Modeling for Geophysical Fluid Dynamics

PASC15, June 1, 2015, Zürich, Switzerland

Matthias Lieber (matthias.lieber@tu-dresden.de)

Center for Information Services and High Performance Computing (ZIH)
Technische Universität Dresden, Germany

ZIH
Center for Information Services & High Performance Computing


"Climate models now include more cloud and aerosol processes, and their interactions, than at the time of the AR4, but there remains low confidence in the representation and quantification of these processes in models."

IPCC, 2013: Summary for Policymakers. In: Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change.


Motivation: Spectral Bin Cloud Microphysics Schemes

Widely used bulk models

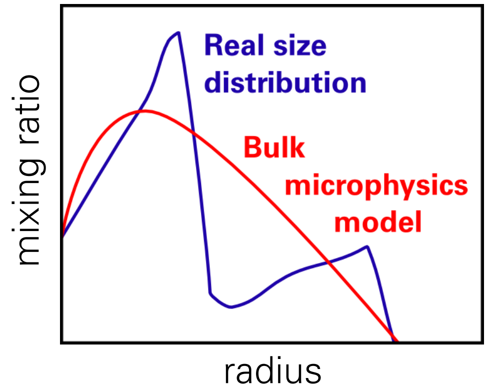

Spectral bin microphysics

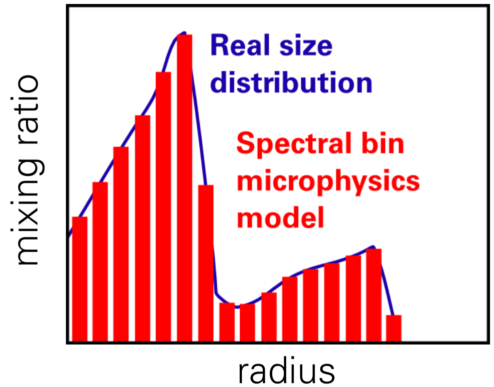

- Bin discretization of cloud particle size distribution
- Allows more detailed modeling of interaction between aerosols, clouds, and precipitation
- Computationally too expensive for forecast
- Only used for process studies up to now

Lynn et al., Mon. Weather Rev., 133:59-71, 2005

Grützun et al., Atmos. Res., 90(2-4):233-242, 2008

Khain et al., J. Atmos. Sci., 67(2):365-384, 2010

Sato et al., J. Atmos. Sci., 69:2012-2030, 2012

Planche et al., Quart. J. Roy. Meteor. Soc. Vol. 140, No. 683, 2014

Fan et al., Atmos. Chem. Phys., 14:81-101, 2014

Motivation: Tropical Cyclone Forecast with SBM?

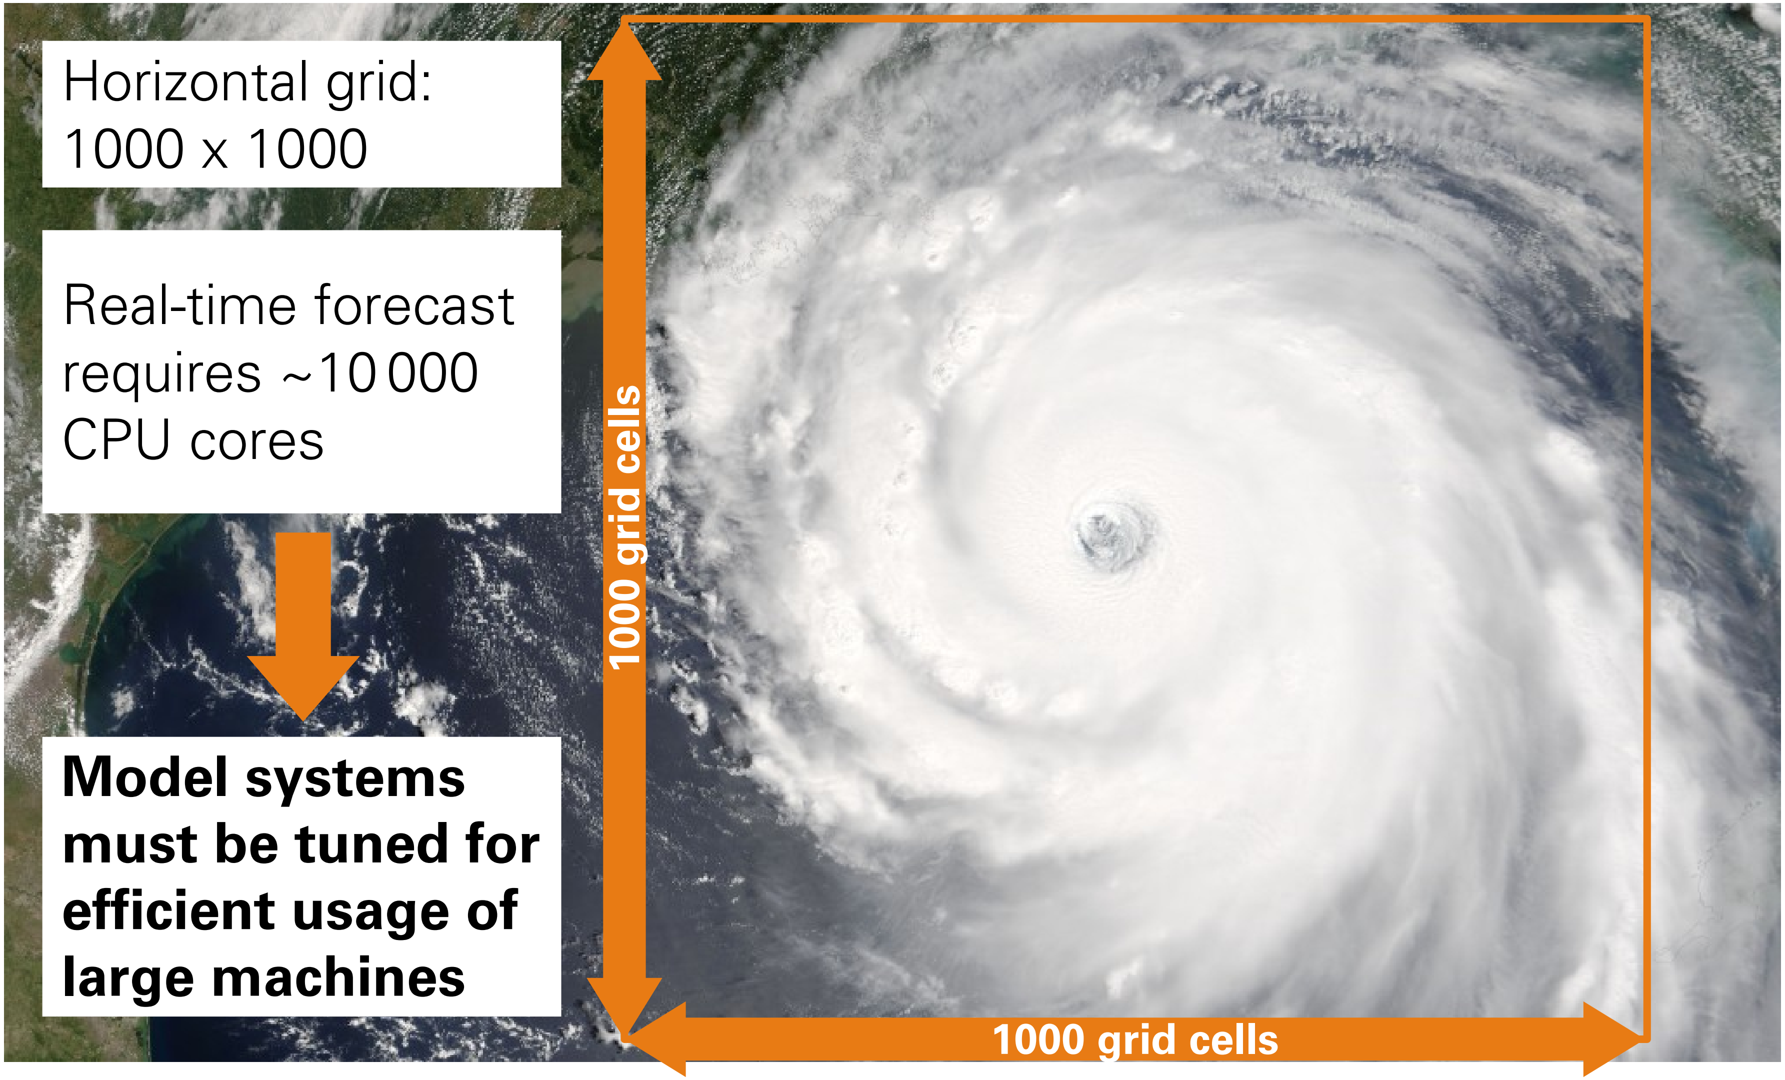

Outline

- Bottleneck Analysis
- Concept of Load-balanced Coupling
- FD4's Features
- Benchmarks
- Conclusion

- COSMO-SPECS: Atmospheric model COSMO extended with spectral bin cloud microphysics model SPECS

COSMO CONSORTIUM FOR SMALL SCALE MODELING

TROPOS Leibniz Institute for Tropospheric Research

Small 3D case with 64x64x48 grid

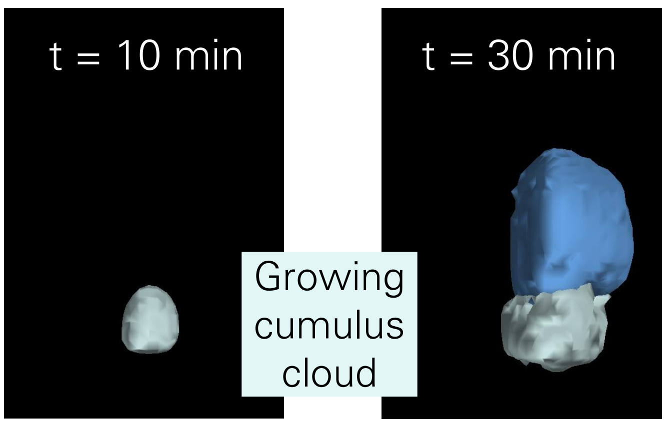

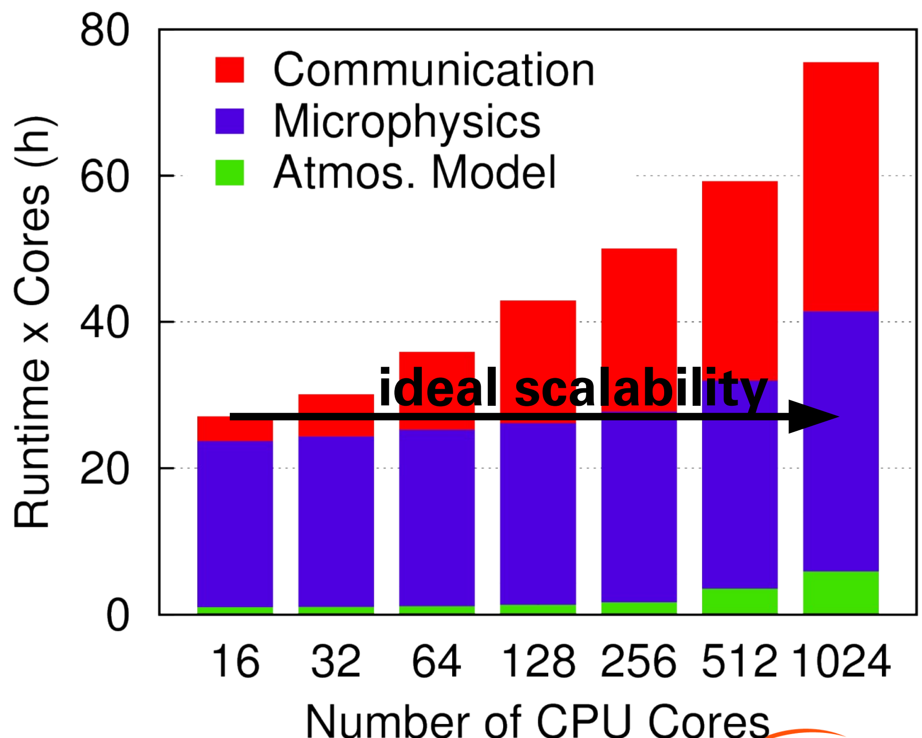

Lieber et al., Highly Scalable Dynamic Load Balancing in the Atmospheric Modeling System COSMO-SPECS+FD4, PARA 2010, 2012

Analysis: Load Imbalance due to Microphysics

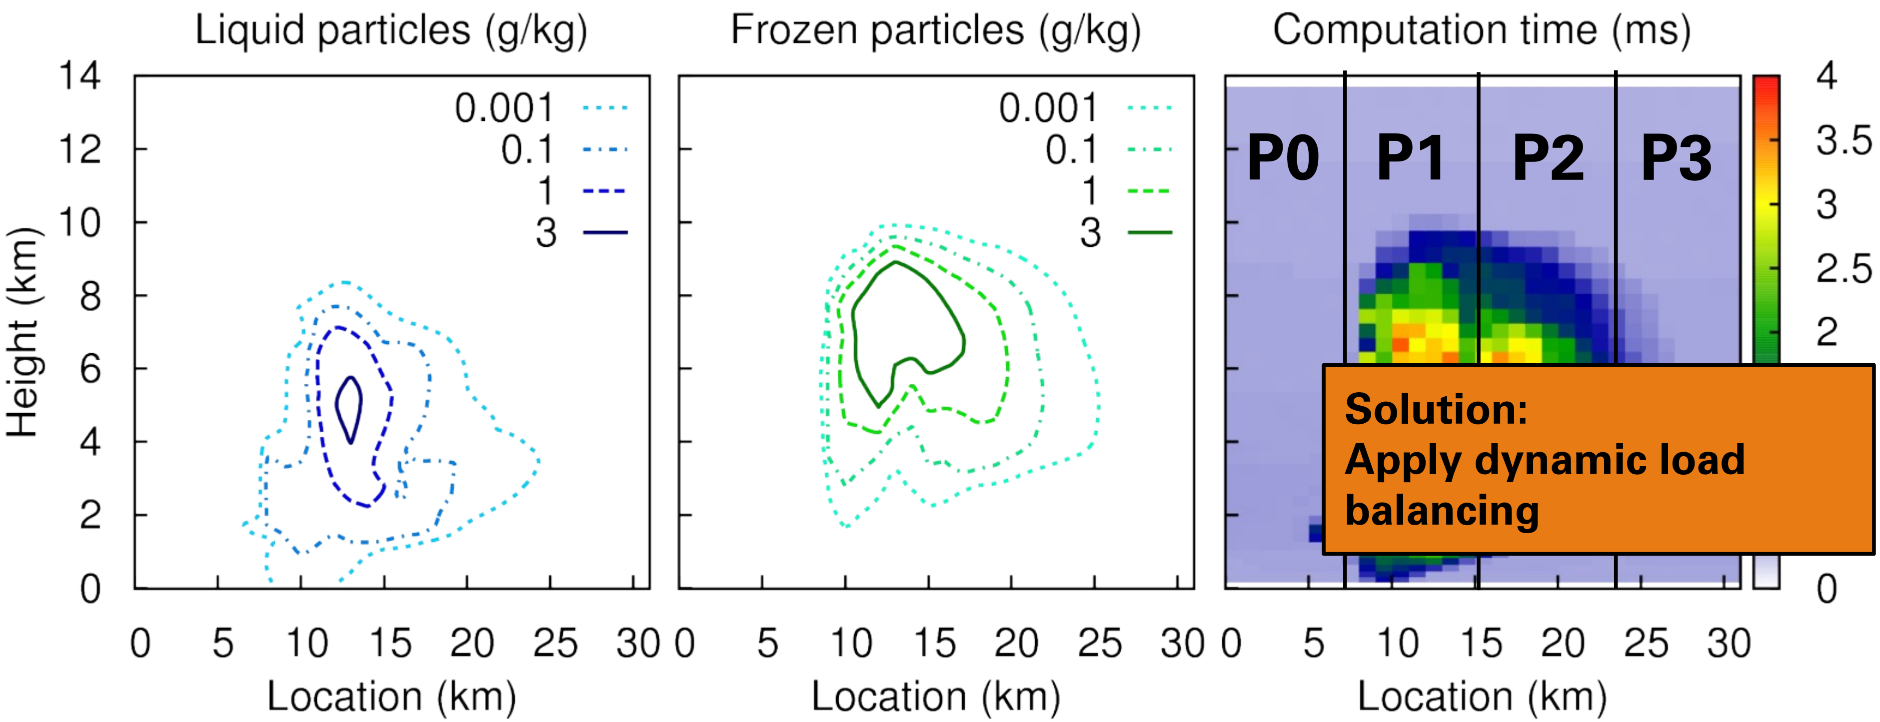

- SPECS computing time varies strongly depending on the range of the particle size distribution and presence of frozen particles
- Leads to load imbalances between partitions

TECHNISCHE UNIVERSITÄT DRESDEN

ZIH
Center for Information Services & High Performance Computing

Analysis: 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: $A^{3D}(P) = 6\ G^{2/3}\ P^{1/3} \sim P^{1/3}$

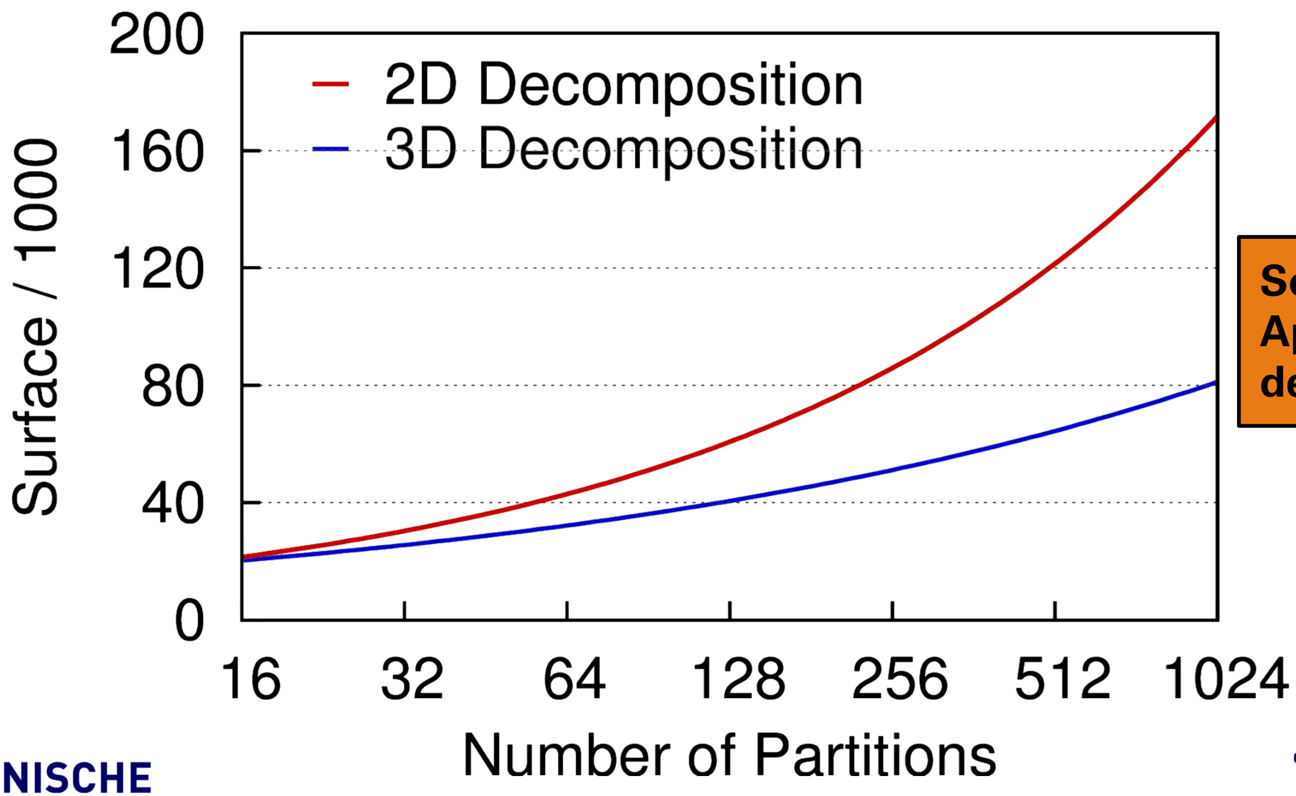

Solution: Apply 3D decomposition

Outline

- Bottleneck Analysis
- **Concept of Load-balanced Coupling**
- FD4's Features
- Benchmarks
- Conclusion

Concept of Load-Balanced Coupling

Atmospheric Model

Spectral Bin Microphysics

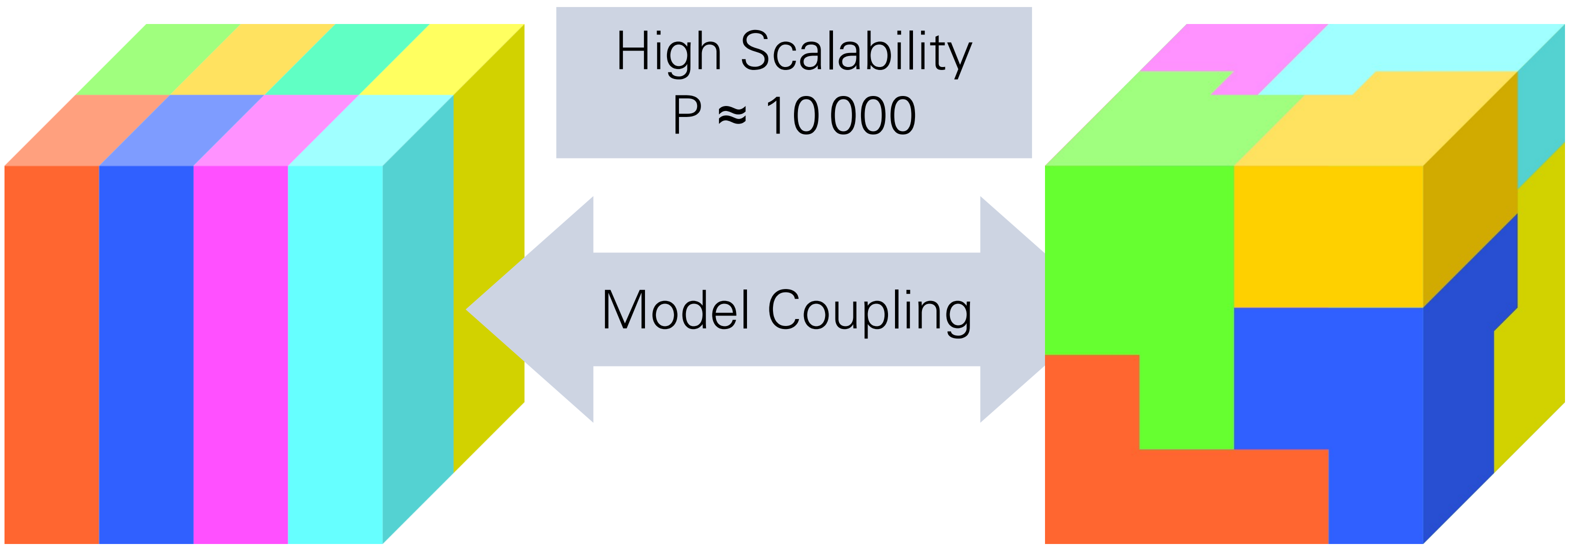

2D Decomposition

Static Partitioning

Block-based 3D Decomposition

Dynamic Load Balancing

Optimized Data Structures

Lieber et al., Highly Scalable Dynamic Load Balancing in the Atmospheric Modeling System COSMO-SPECS+FD4, PARA 2010, 2012

Concept of Load-Balanced Coupling

High Scalability
$P \approx 10\,000$

Model Coupling

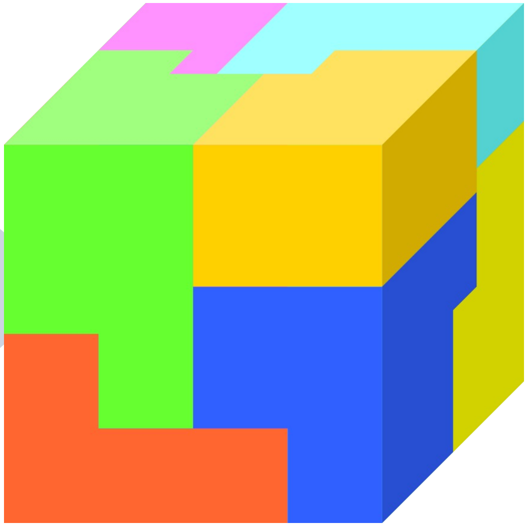

Implemented as independent framework FD4

FD4:
Four-Dimensional
Distributed
Dynamic
Data structures

Block-based 3D Decomposition

Dynamic Load Balancing

Optimized Data Structures

Outline

- Bottleneck Analysis
- Concept of Load-balanced Coupling
- **FD4's Features**
- Benchmarks
- Conclusion

FD4: Dynamic Load Balancing

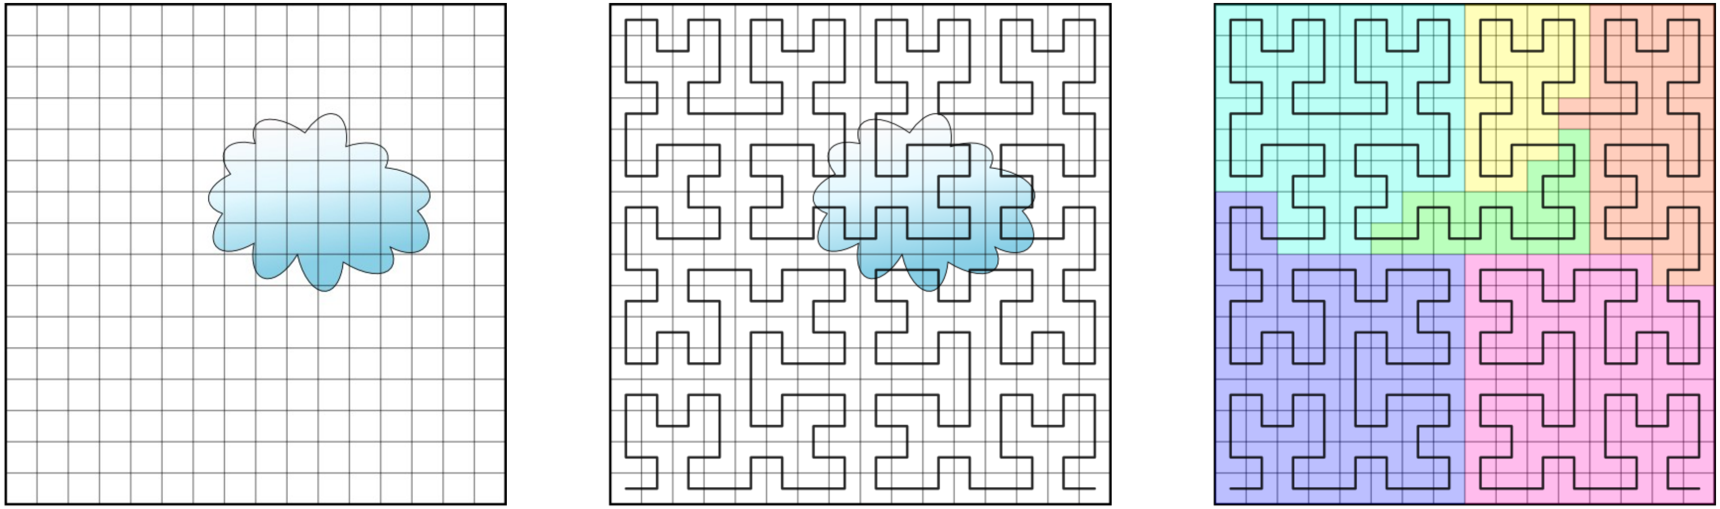

- 3D block decomposition of rectangular grid
- Space-filling curve (SFC) partitioning to assign blocks to ranks
- SFC reduces 3D partitioning problem to 1D
- High locality of SFC leads to moderate comm. costs
- Developed a highly scalable, hierarchical method for high-quality 1D partitioning of the SFC-indexed blocks

FD4: Model Coupling

- Data exchange between FD4 based model and an external model
  - E.g. weather or CFD model
  - Transfer in both directions
- FD4 computes partition overlaps after each repartitioning of FD4 grid
  - Highly scalable algorithm
- No grid transformation / interpolation
  - External model must provide data matching the FD4 grid
- "Sequential" coupling only
  - Both models run alternately on same set of MPI ranks

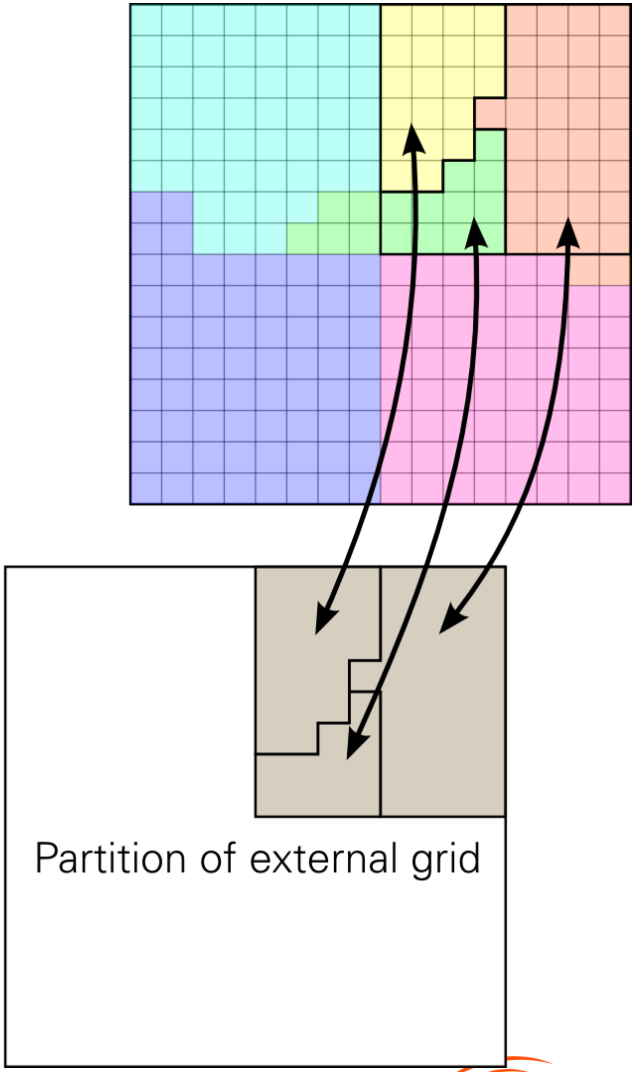

FD4: 4th Dimension

- Extra, non-spatial dimension of grid variables, e.g.
  - Size resolving models
  - Array of gas phase tracers
- FD4 is optimized for a large 4th dimension
- COSMO-SPECS requires 2 x 11 x 66 ~ 1500 values

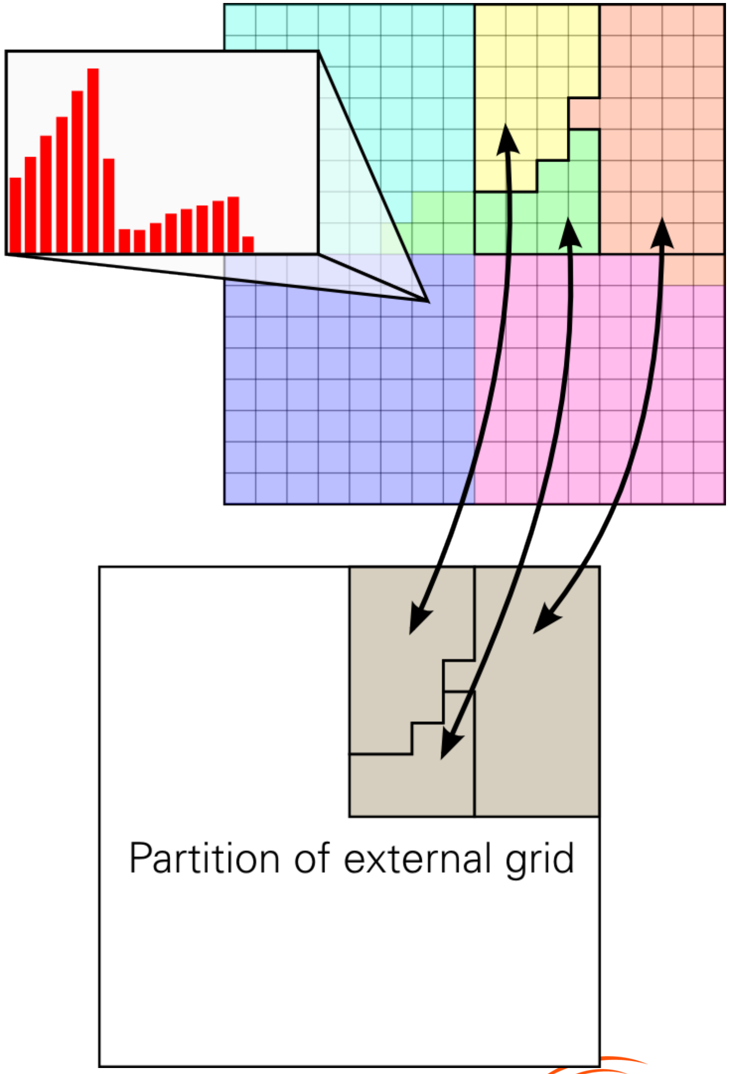

Outline

- Bottleneck Analysis
- Concept of Load-balanced Coupling
- FD4's Features
- **Benchmarks**
- Conclusion

Benchmarks: COSMO-SPECS Performance Comparison

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

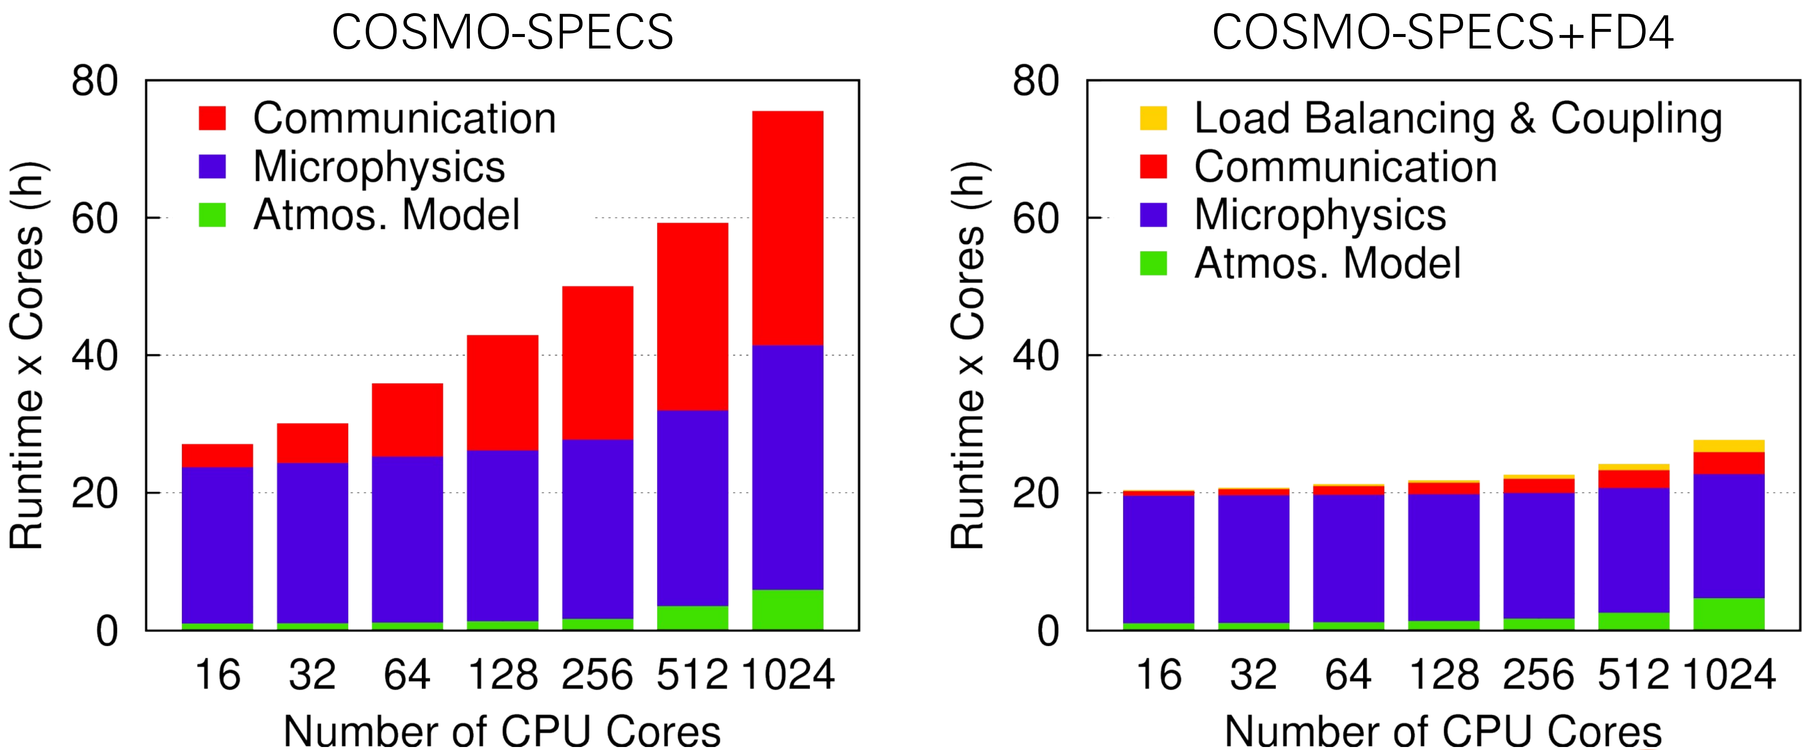

Benchmarks: Scalability on Blue Gene/Q

- Grid size: 1024 x 1024 x 48 grid cells, > 3M blocks
- 256k: 30 min forecast in <5min (w/o init and I/O)
- Runs on Blue Gene/Q with up to 262 144 MPI ranks
- 14 x speed-up from 16k to 256k

Lieber, Nagel, Mix, *Scalability Tuning of the Load Balancing and Coupling Framework FD4*, NIC Symposium 2014, pp. 363-370.

COSMO-SPECS+FD4

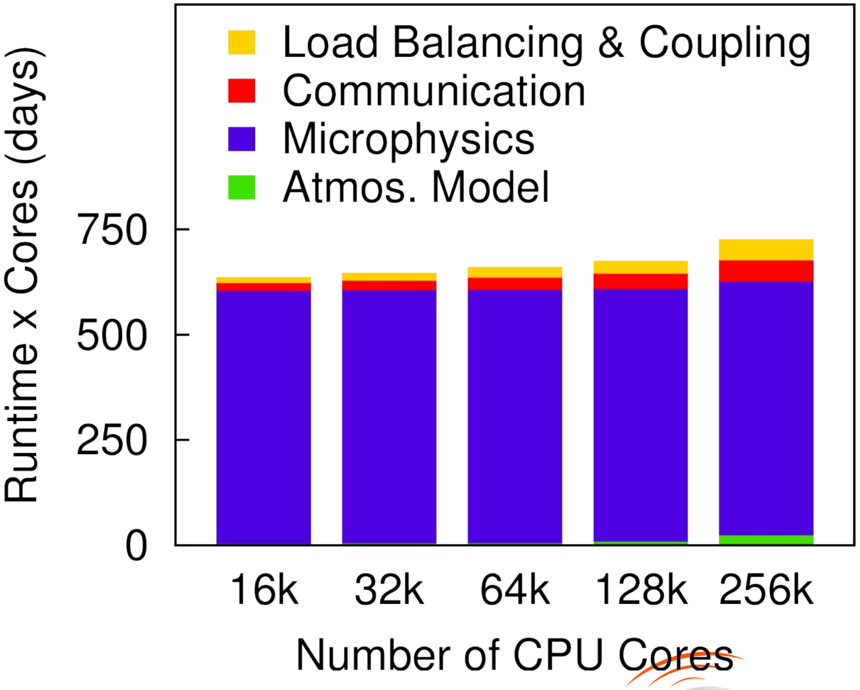

Benchmarks: Load Balancing & Coupling Scalability

- Grid size: 1024 x 1024 x 48 grid cells, > 3M blocks
- Load balancing scales comparatively very well
- Coupling scales nearly perfect

Lieber, Nagel, Mix, *Scalability Tuning of the Load Balancing and Coupling Framework FD4*, NIC Symposium 2014, pp. 363-370.

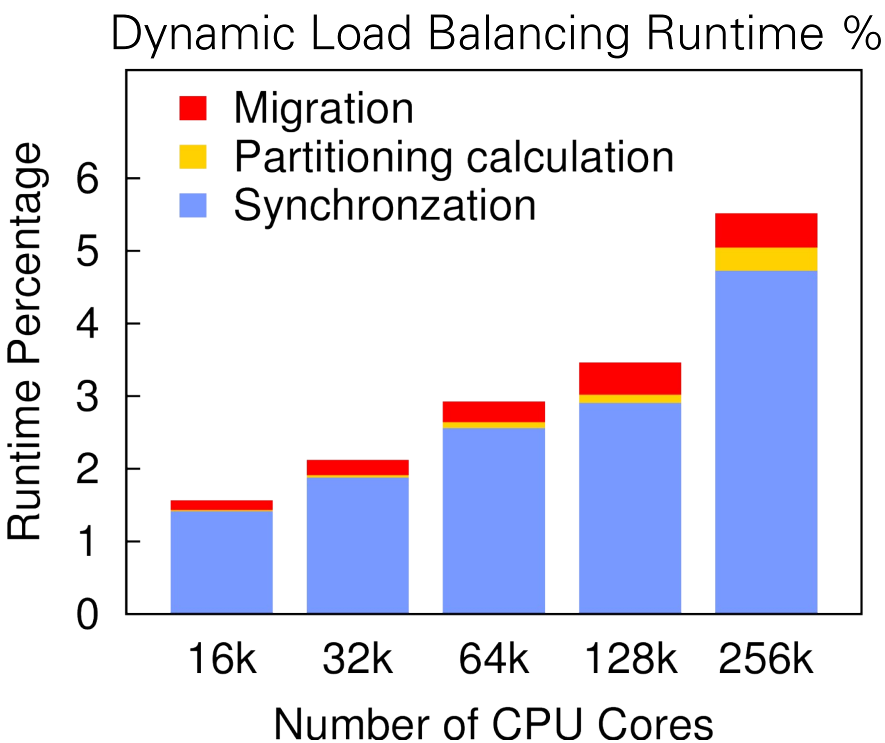

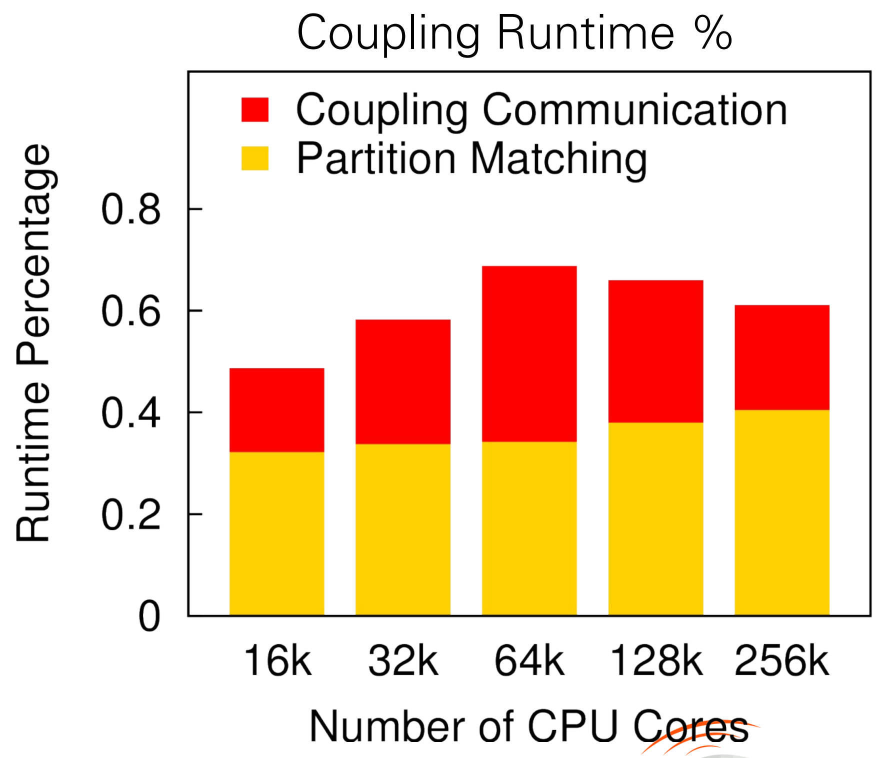

TECHNISCHE UNIVERSITÄT DRESDEN

ZIH Center for Information Services & High Performance Computing

Benchmarks: 1D Partitioning Comparison on Blue Gene/Q

- ExactBS: exact method, but slow and serial
- H2: fast heuristic, but may result in poor load balance
- HIER*: hierarchical algorithm implemented in FD4, achieves nearly optimal load balance

Lieber, Nagel, *Scalable High-Quality 1D Partitioning*, HPCS 2014, pp. 112-119, 2014

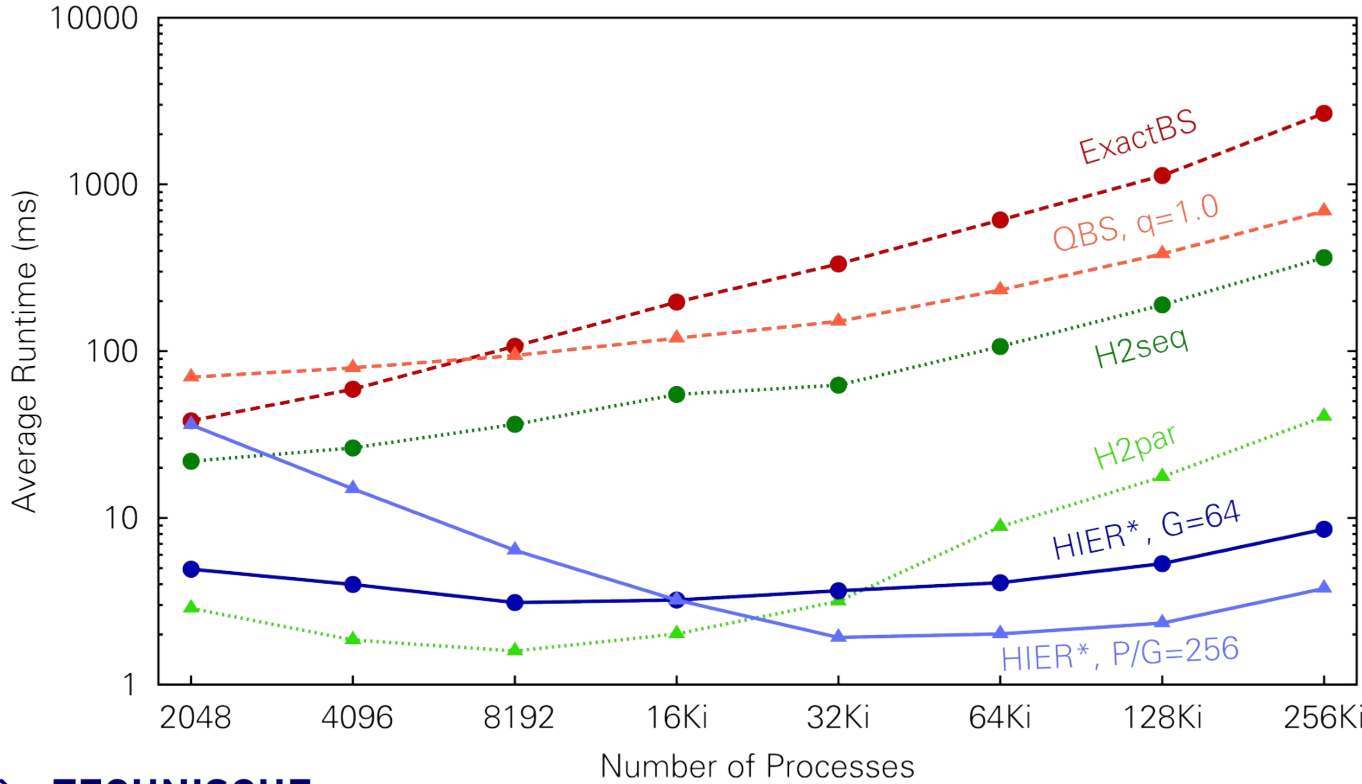

ExactBS: 2668 ms

QBS: 692 ms

H2seq: 363 ms

H2par: 40.5 ms

$\text{HIER*}_{G=64}$: 8.55 ms

$\text{HIER*}_{P/G=256}$: 3.77 ms

Heuristic H2 in Action (COSMO-SPECS+FD4)

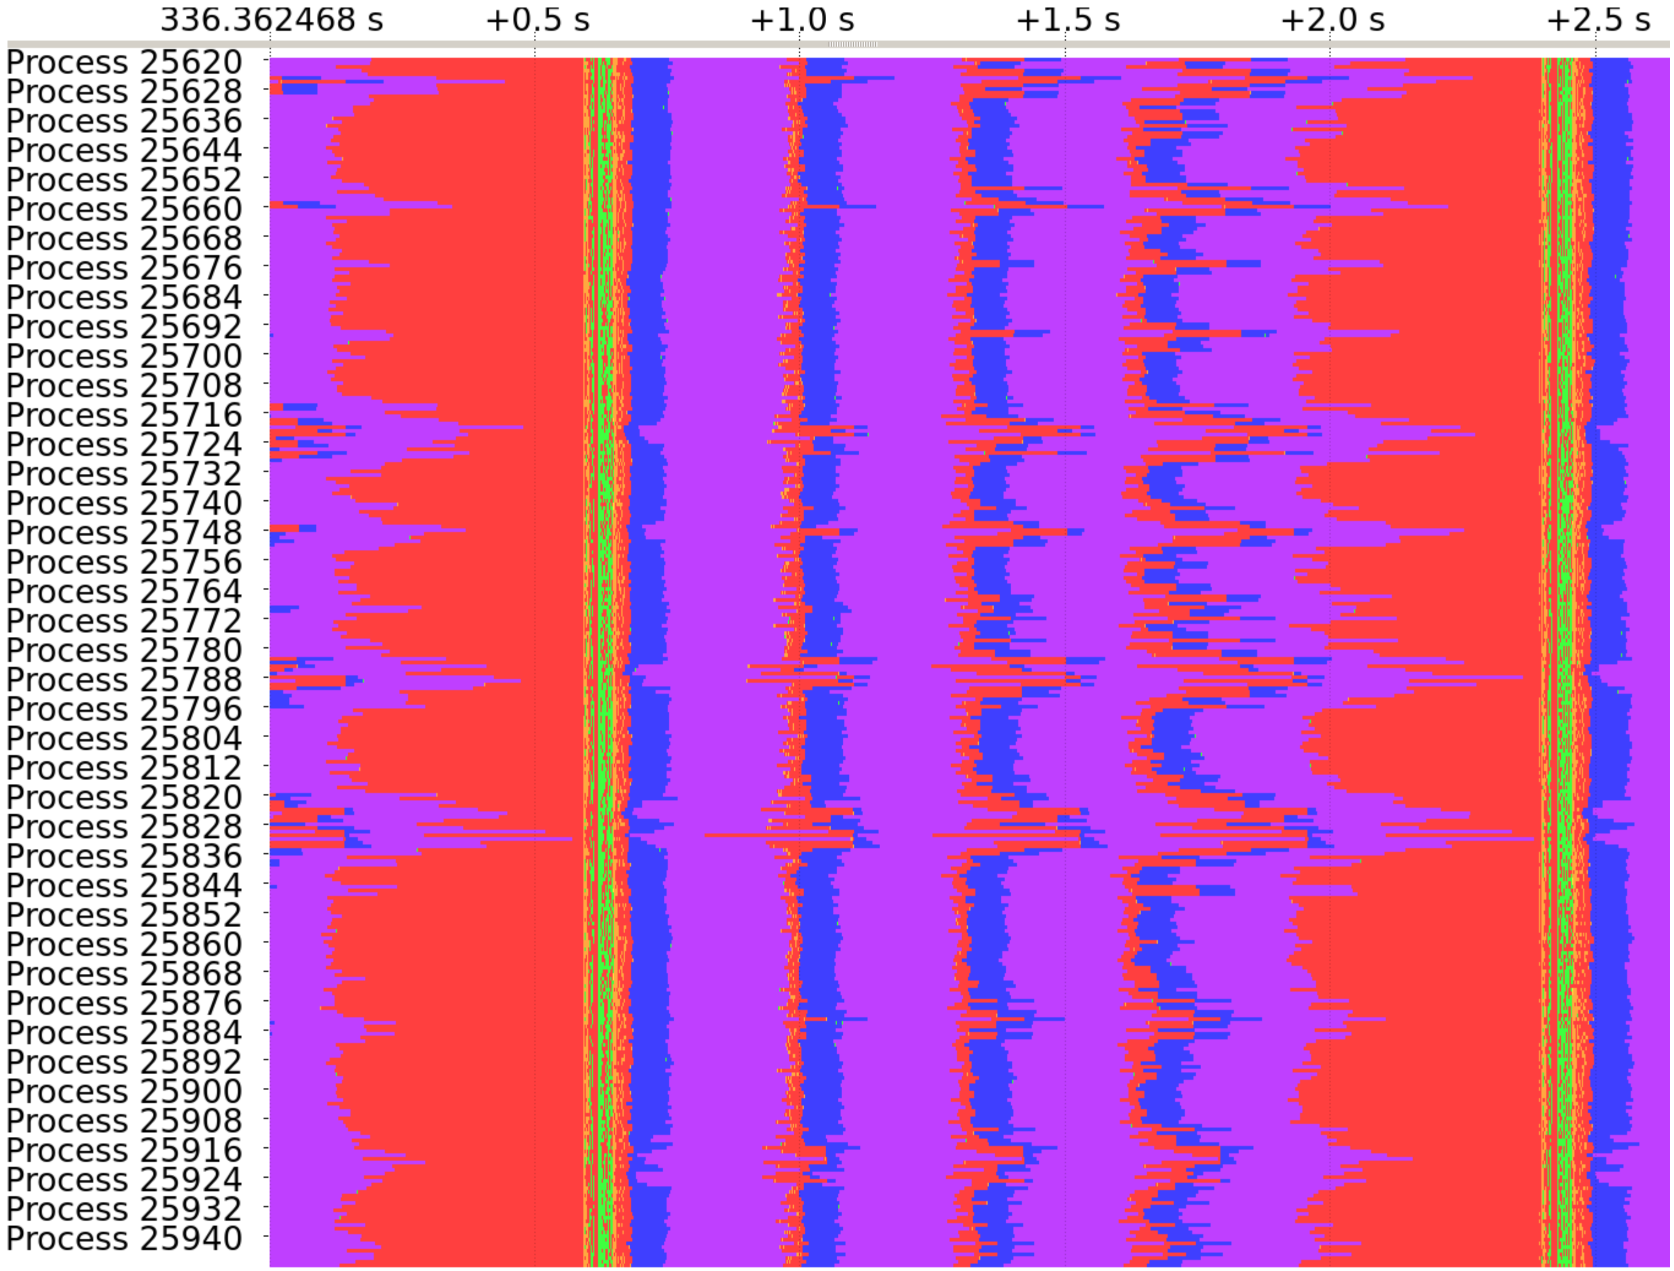

HIER* in Action (COSMO-SPECS+FD4)

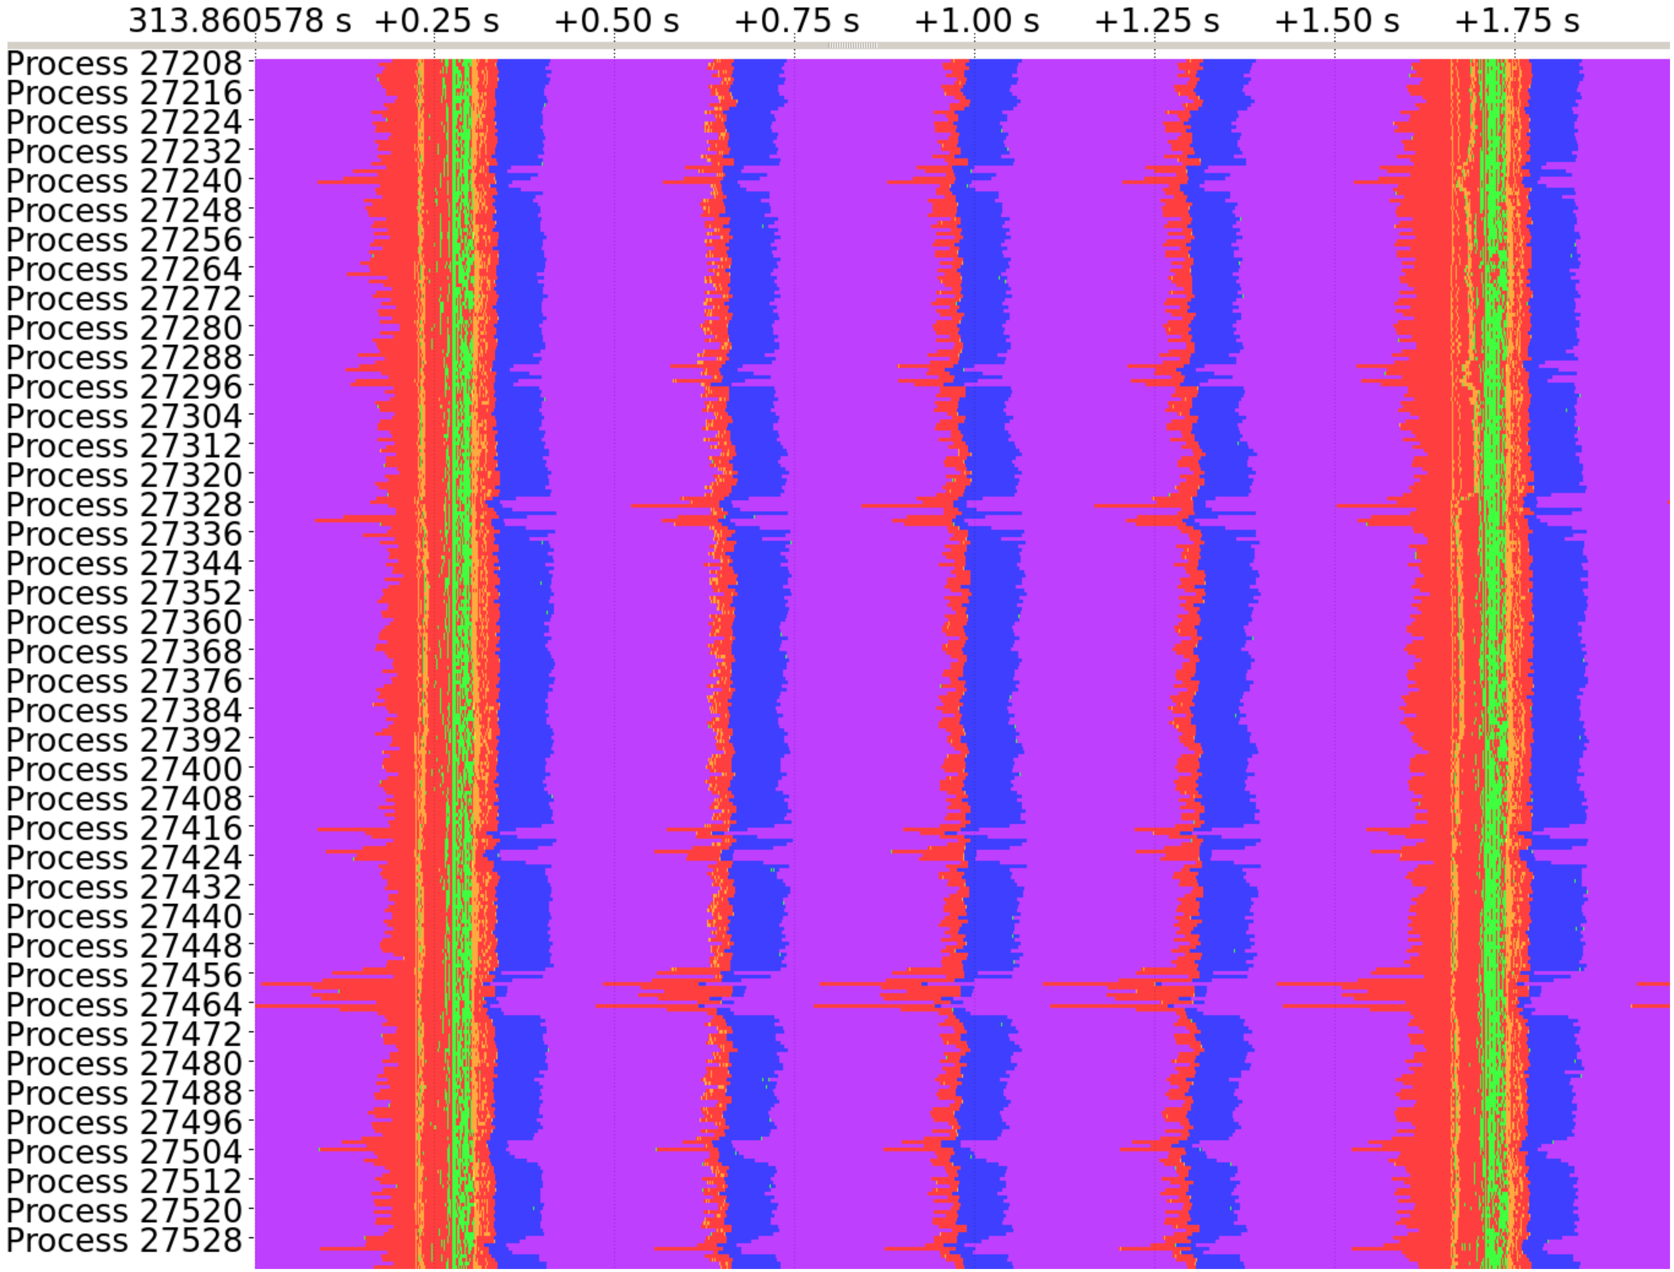

Outline

- Bottleneck Analysis
- Concept of Load-balanced Coupling
- FD4's Features
- Benchmarks
- **Conclusion**

Conclusions

- FD4 provides for simulation models
  - Dynamic load balancing
  - Model coupling
  - High scalability
- Initially developed for atmospheric modeling, but generally applicable (multiphysics, multiphase)
- FD4 is available as open source software
  - Fortran 95, MPI-2, NetCDF
  - Tested on many different HPC systems
- Future plans
  - Use FD4 for atmospheric chemistry
  - Study its potential for 2-moment microphysics

FD4 website: http://wwwpub.zih.tu-dresden.de/~mlieber/fd4

Lieber, Nagel, *Scalable High-Quality 1D Partitioning*, HPCS 2014, pp. 112-119, 2014

Lieber, Nagel, Mix, *Scalability Tuning of the Load Balancing and Coupling Framework FD4*, NIC Symposium 2014

Lieber et al., Highly Scalable Dynamic Load Balancing in the Atmospheric Modeling System COSMO-SPECS+FD4, PARA 2010, 2012

Lieber et al., *FD4: A Framework for Highly Scalable Load Balancing and Coupling of Multiphase Models*, ICNAAM 2010

Thank you very much for your attention!

Acknowledgments

Verena Grützun, Ralf Wolke,
Oswald Knoth, Martin Simmel,
René Widera, Matthias Jurenz,
Matthias Müller, Wolfgang E. Nagel

lrz Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities

VAMPIR www.vampir.eu

TROPOS Leibniz Institute for Tropospheric Research
www.tropos.de

DWD COSMO CONSORTIUM FOR SMALL SCALE MODELING
www.cosmo-model.org

PIConGPU HZDR HELMHOLTZ ZENTRUM DRESDEN ROSSENDORF
picongpu.hzdr.de

Funding

DFG Deutsche Forschungsgemeinschaft

cfaed CENTER FOR ADVANCING ELECTRONICS DRESDEN

Europa fördert Sachsen. ESF Europäischer Sozialfonds

IMData

TECHNISCHE UNIVERSITÄT DRESDEN