HELMUT SCHMIDT
UNIVERSITÄT

Chair for High-Performance Computing
Philipp Neumann

Sparse Grid Regression for Performance Prediction Using High-Dimensional Run Time Data

HELMUT SCHMIDT UNIVERSITÄT

Outline

- Performance Analysis and Higher Dimensions
- Sparse Grids in a Nutshell
- Regression on Sparse Grids
- Results: Molecular Dynamics, Climate, Weather
- Summary

Performance Analysis and Higher Dimensions: Parameters Affecting Performance

HELMUT SCHMIDT UNIVERSITÄT

- Algorithmic parameters
  → convergence criteria, mesh size, time step, …
- Hardware-aware optimization
  → params for cache blocking, data alignment, vector widths, …
- Parallelization settings
  → number of MPI processes, OMP threads, …
- Scenario-dependent parameters
  → domain size/shape, number of cells/particles, …

→ High-Dimensional Parameter Space

Performance Analysis and Higher Dimensions: Exploring High-Dimensional Spaces

HELMUT SCHMIDT UNIVERSITÄT

Outline

- (Semi-)Analytical models
  → Only available for small subset of params
- Neural networks/ deep learning
  → Effective approach
  → Interesting for hard (e.g., combinatorial) problems
  → Decisions/results not necessarily transparent
- Regression and related methods
  → Effective approach
  → Application in higher dimensions?

Outline

- Performance Analysis and Higher Dimensions
- **Sparse Grids in a Nutshell**
- Regression on Sparse Grids
- Results:
  Molecular Dynamics
  Climate
  Weather
- Summary

HELMUT SCHMIDT UNIVERSITÄT

Sparse Grids in a Nutshell

Theorem 1 *For the interpolation error of a function $f \in H^2_{0,mix}$ in the sparse grid space $V^s_{0,n}$ holds*

$$||f - f^s_n||_2 = \mathcal{O}(h^2_n \log(h^{-1}_n)^{d-1}).$$

J. Garcke.

Sparse grids in a nutshell

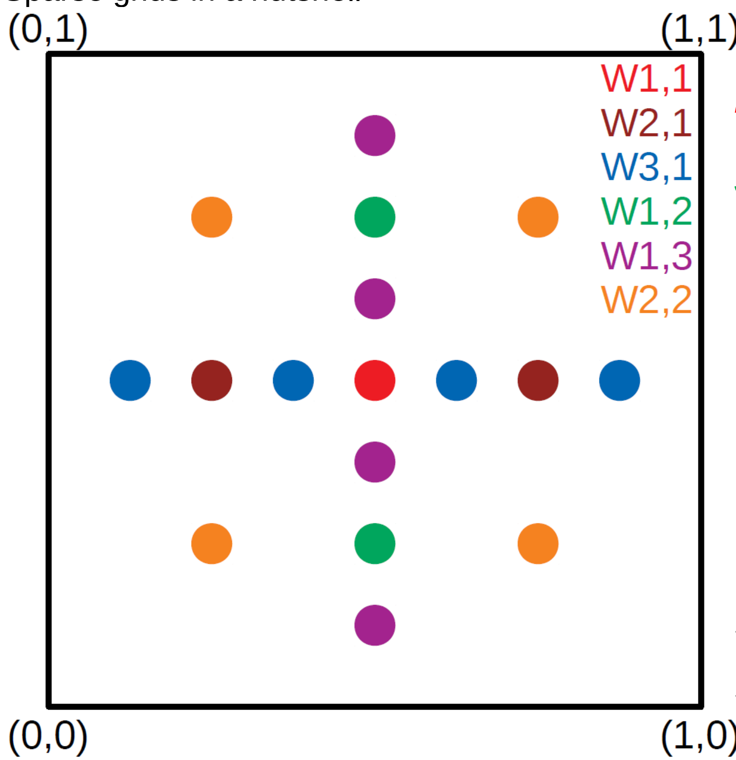

Full Cart. grid: O(N^d) points

SG: O(N(log N)$^{d-1}$) points

→hierarchical representation

→prerequisite for “good” approximations: sufficiently smooth settings/params

Outline

- Performance Analysis and Higher Dimensions
- Sparse Grids in a Nutshell
- Regression on Sparse Grids
- Results:
  Molecular Dynamics
  Climate
  Weather
- Summary

Sparse Grids: Local Mesh Refinement

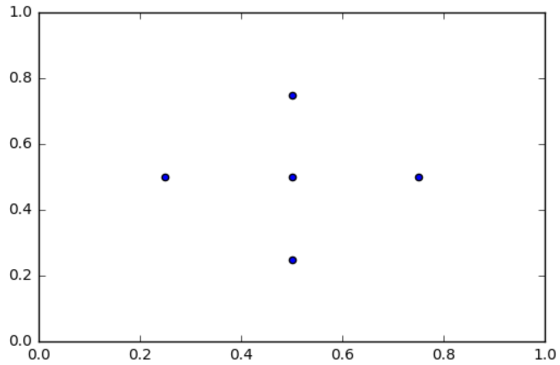

- No. refinement iterations: 3
- No. adaptable grid points: 3
- Example:
  2 refinement iterations,
  3 adaptable grid points,
  start from level-2 grid
- Software in use: SG++

HELMUT SCHMIDT UNIVERSITÄT

Sparse Grids: Local Mesh Refinement

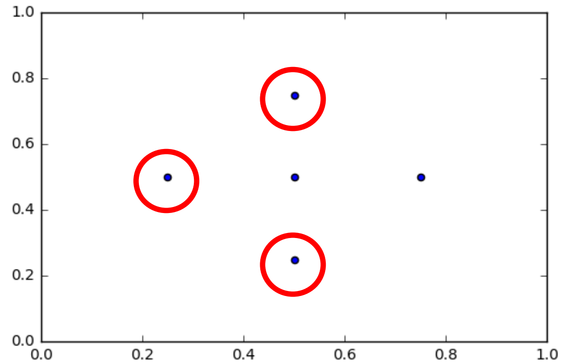

- No. refinement iterations: 3
- No. adaptable grid points: 3
- Example:
  2 refinement iterations,
  3 adaptable grid points,
  start from level-2 grid
- Software in use: SG++

Outline

- Performance Analysis and Higher Dimensions
- Sparse Grids in a Nutshell
- Regression on Sparse Grids
- Results:
  Molecular Dynamics
  Climate
  Weather
- Summary

Sparse Grids: Local Mesh Refinement

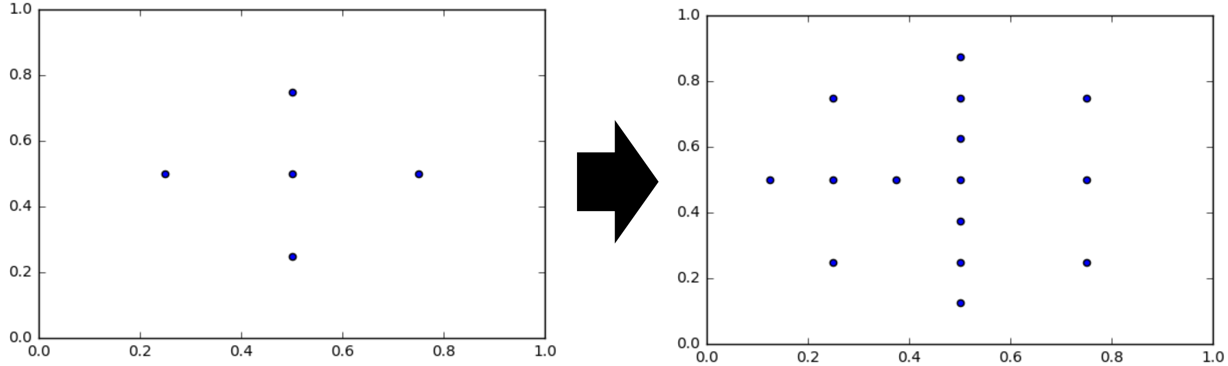

- No. refinement iterations: 3
- No. adaptable grid points: 3
- Example:
  2 refinement iterations,
  3 adaptable grid points,
  start from level-2 grid
- Software in use: SG++

HELMUT SCHMIDT UNIVERSITÄT

Outline

- Performance Analysis and Higher Dimensions
- Sparse Grids in a Nutshell
- Regression on Sparse Grids
- Results:
  Molecular Dynamics
  Climate
  Weather
- Summary

Sparse Grids: Local Mesh Refinement

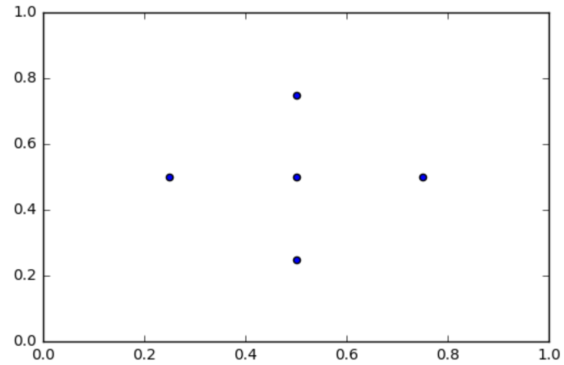

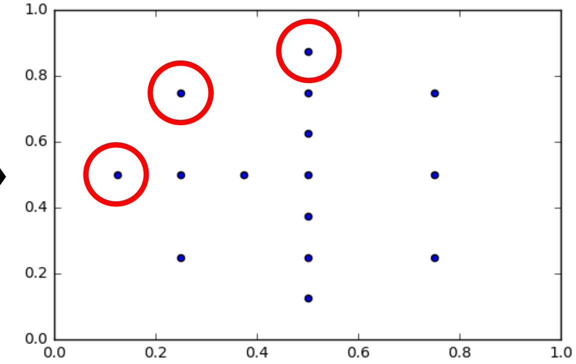

- No. refinement iterations: 3
- No. adaptable grid points: 3
- Example:
  2 refinement iterations,
  3 adaptable grid points,
  start from level-2 grid
- Software in use: SG++

HELMUT SCHMIDT UNIVERSITÄT

Outline

- Performance Analysis and Higher Dimensions
- Sparse Grids in a Nutshell
- Regression on Sparse Grids
- Results:
  Molecular Dynamics
  Climate
  Weather
- Summary

Sparse Grids: Local Mesh Refinement

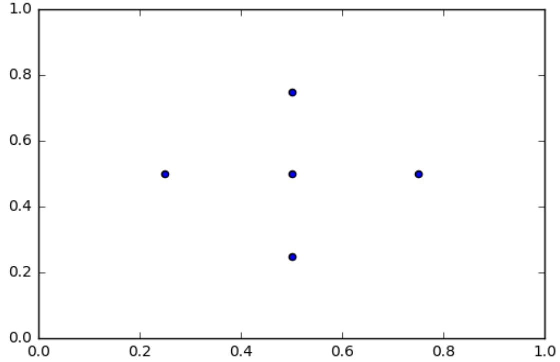

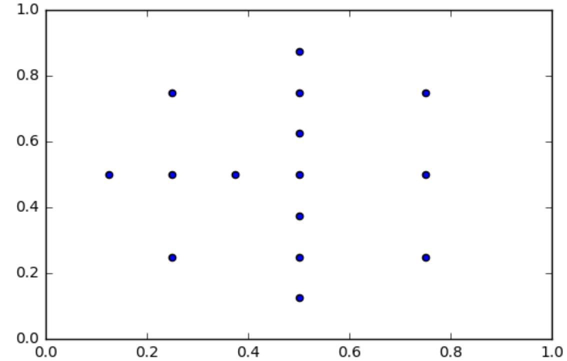

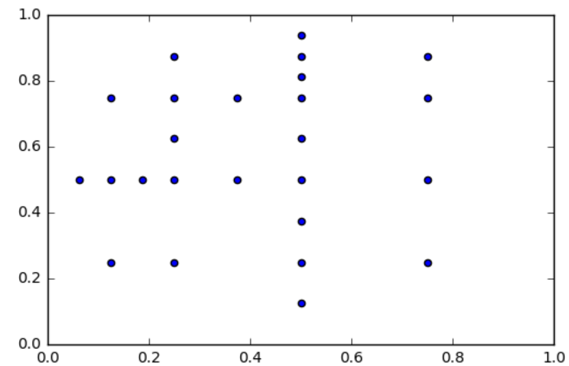

- No. refinement iterations: 3
- No. adaptable grid points: 3
- Example:
  2 refinement iterations,
  3 adaptable grid points,
  start from level-2 grid
- Software in use: SG++

HELMUT SCHMIDT UNIVERSITÄT

Outline

Regression on Sparse Grids

- Define linear hat function φ_i per sparse grid point
  → defines function space V_n
- Solve regression problem on run time data y_j, given parameter combinations x_j:

$$u = \arg\min_{v \in V_n} \left(\frac{1}{M} \sum_{j=1}^{M} (y_j - v(\vec{x}_j))^2 + \lambda C(v) \right)$$

with $v(\vec{x}) := \sum_i \alpha_i \varphi_i(\vec{x})$

- Results in linear system: $\left(\frac{1}{M} BB^\top + \lambda \mathbb{I} \right) \vec{\alpha} = \frac{1}{M} B\vec{y}$

Outline

- Performance Analysis and Higher Dimensions
- Sparse Grids in a Nutshell
- Regression on Sparse Grids
- Results:
  Molecular Dynamics
  Climate
  Weather
- Summary

Results: Evaluation Procedure

HELMUT SCHMIDT UNIVERSITÄT

- Data splitting:
  Use s % of data for learning and 1-s % for validation
- Mean relative error:
  - Start from one data split
  - Compute and average relative errors for this data split
  - Repeat this procedure for 10 data splits and average errors
- Consider different initial sparse grid level refinements (level-2 and level-3 grids)

HELMUT SCHMIDT UNIVERSITÄT

Outline

- Performance Analysis and Higher Dimensions
- Sparse Grids in a Nutshell
- Regression on Sparse Grids
- **Results: Molecular Dynamics** Climate Weather
- Summary

Results: Molecular Dynamics (1)

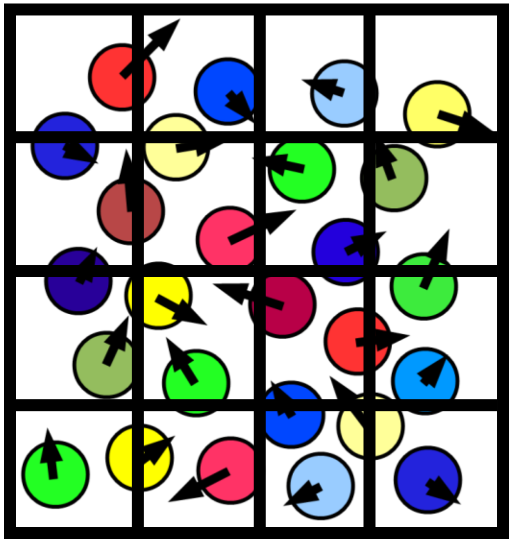

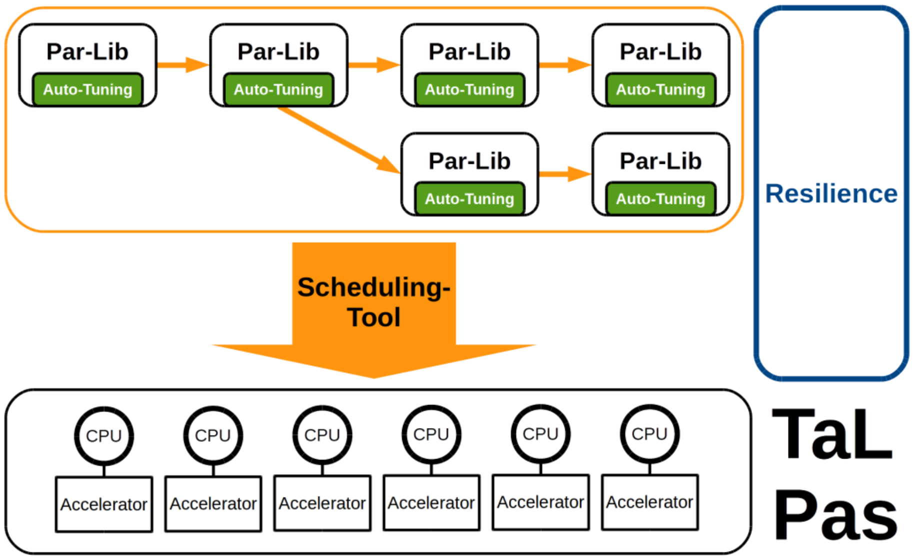

- particle density $\rho \in [0.3; 0.9]$, number of particles $N \in [1e3; 1e5]$, cut-off radius $r_c \in [1.2; 4.5]$, blocksize $\in [1e1; 1e3]$, no MPI processes $P \in \{1,2,4,8\}$
- SimpleMD: Single-Site Lennard-Jones, Linked Cells

SPONSORED BY THE

Federal Ministry of Education and Research

HELMUT SCHMIDT UNIVERSITÄT

Outline

- Performance Analysis and Higher Dimensions
- Sparse Grids in a Nutshell
- Regression on Sparse Grids
- Results: Molecular Dynamics
  Climate
  Weather
- Summary

SPONSORED BY THE

Federal Ministry of Education and Research

Results: Molecular Dynamics (2)

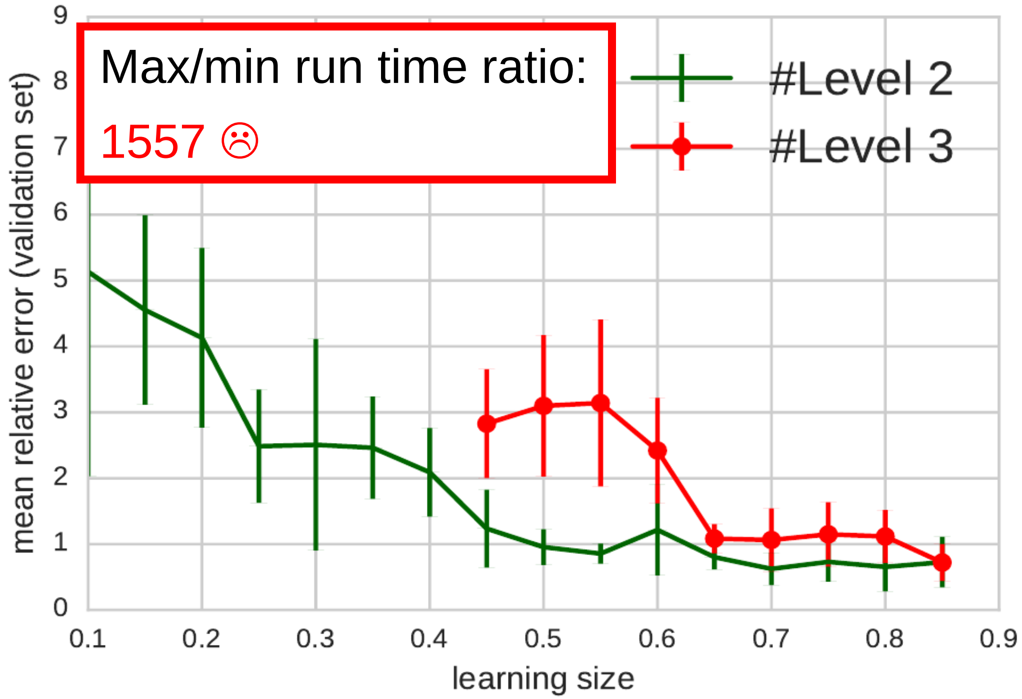

- particle density $\rho \in [0.3; 0.9]$, number of particles $N \in [1e3; 1e5]$, cut-off radius $r_c \in [1.2; 4.5]$, blocksize $\in [1e1; 1e3]$, no MPI processes $P \in \{1,2,4,8\}$
- Random sampling of run time space → 357 samples

HELMUT SCHMIDT UNIVERSITÄT

Outline

- Performance Analysis and Higher Dimensions
- Sparse Grids in a Nutshell
- Regression on Sparse Grids
- **Results: Molecular Dynamics**
  Climate
  Weather
- Summary

SPONSORED BY THE
Federal Ministry of Education and Research

Results: Molecular Dynamics (2)

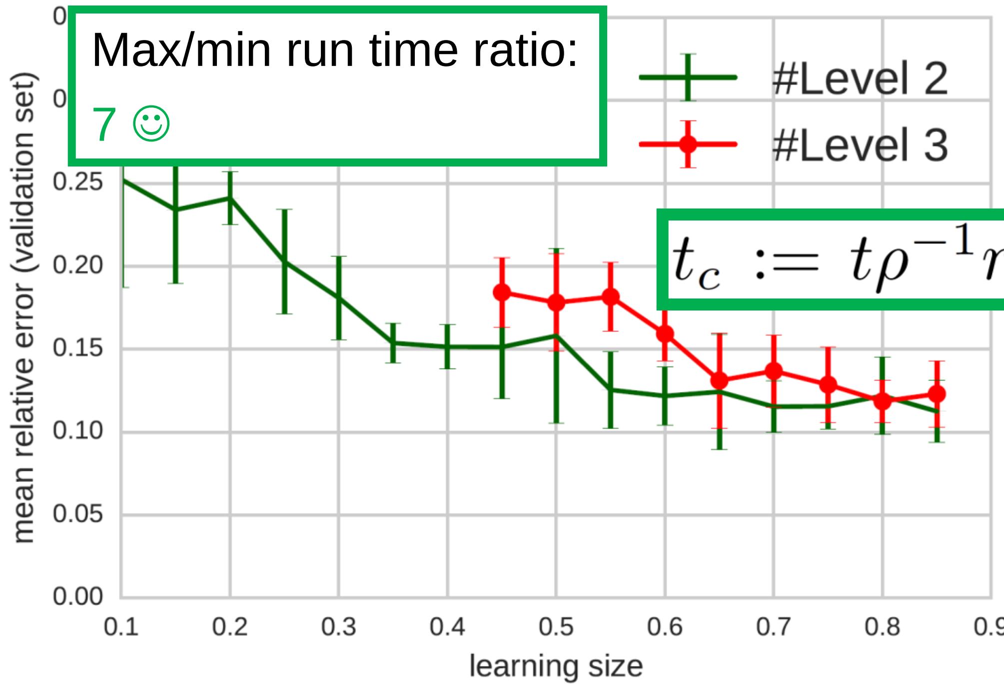

Max/min run time ratio: 7 ☺

$$t_c := t\rho^{-1} r_c^{-3} N^{-1} P$$

- particle density $\rho \in [0.3; 0.9]$, number of particles $N \in [1e3; 1e5]$, cut-off radius $r_c \in [1.2; 4.5]$, blocksize $\in [1e1; 1e3]$, no MPI processes $P \in \{1,2,4,8\}$
- Random sampling of run time space → 357 samples

HELMUT SCHMIDT
UNIVERSITÄT

Results: Molecular Dynamics (3)

Outline

- Performance Analysis and Higher Dimensions
- Sparse Grids in a Nutshell
- Regression on Sparse Grids
- Results:
  Molecular Dynamics
  Climate
  Weather
- Summary

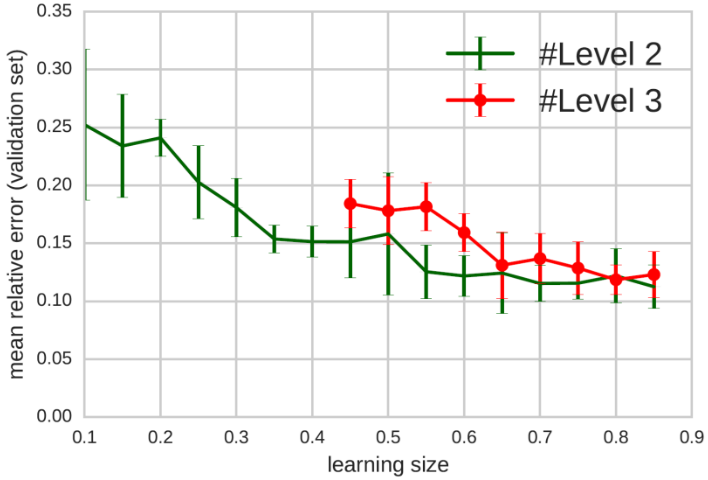

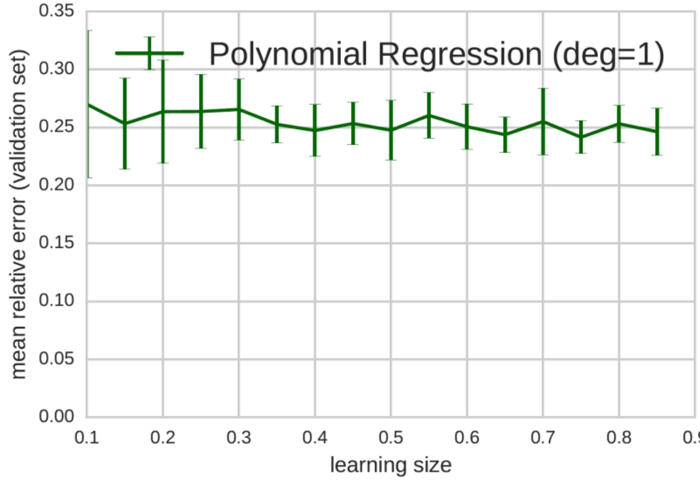

- Upper left: SG
- Upper right: 1st order reg.
- Lower right: 2nd order reg.

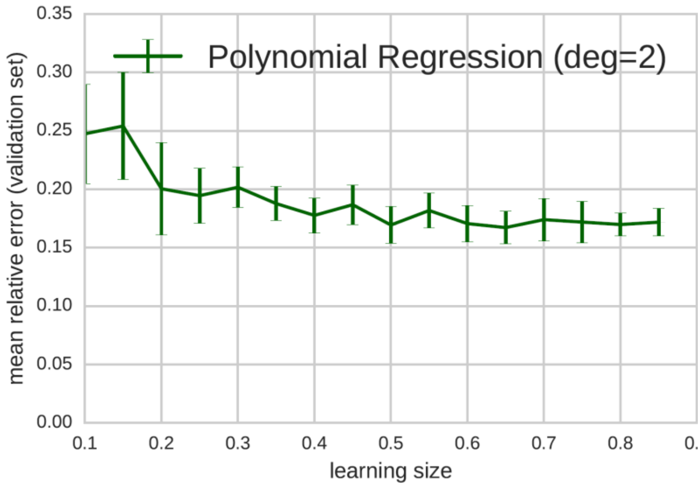

SPONSORED BY THE
Federal Ministry of Education and Research

Results: Weather and Climate – ICON Model

Outline

- Performance Analysis and Higher Dimensions
- Sparse Grids in a Nutshell
- Regression on Sparse Grids
- Results:
  Molecular Dynamics
  Climate
  Weather
- Summary

- ICON=ICOsahedral Non-hydrostatic model
- Developed by Deutscher Wetterdienst/
  Max-Planck-Institut für Meteorologie
- Triangular grids on the sphere + vertical columns
- Multiscale, multiphysics: dynamical core, climate/weather physics, radiation, land surface interaction, …

Outline

- Performance Analysis and Higher Dimensions
- Sparse Grids in a Nutshell
- Regression on Sparse Grids
- **Results:**
  Molecular Dynamics
  Climate
  Weather
- Summary

Results: Climate – ICON V16.0 Benchmark[1]

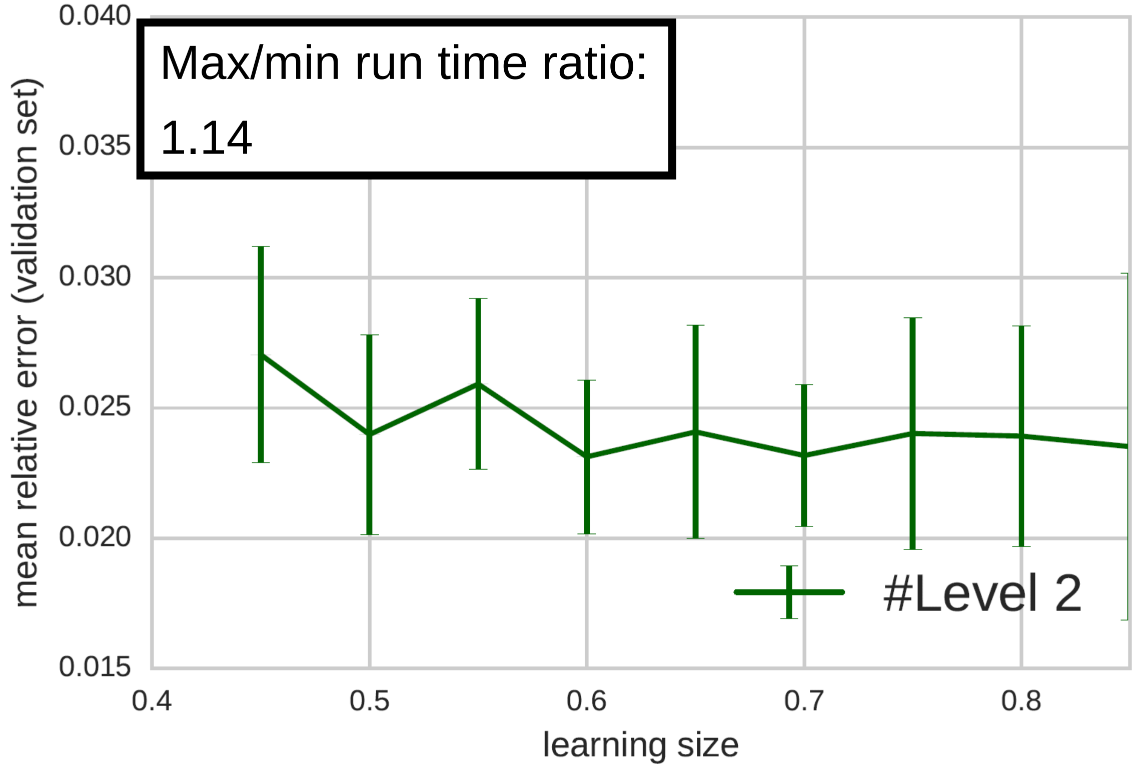

- Params: # OpenMP threads (1,2,4,6,8,12,18,36), nproma (col. blocking; 2,8,16,24,32)

1 https://redmine.dkrz.de/projects/icon-benchmark/wiki/Instructions_on_download_execution_and_analysis_ICON_Benchmark_v160

esiwace CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

HELMUT SCHMIDT UNIVERSITÄT

Helmut Schmidt Universität

Outline

- Performance Analysis and Higher Dimensions
- Sparse Grids in a Nutshell
- Regression on Sparse Grids
- Results:
  Molecular Dynamics
  Climate
  Weather
- Summary

Results: Weather

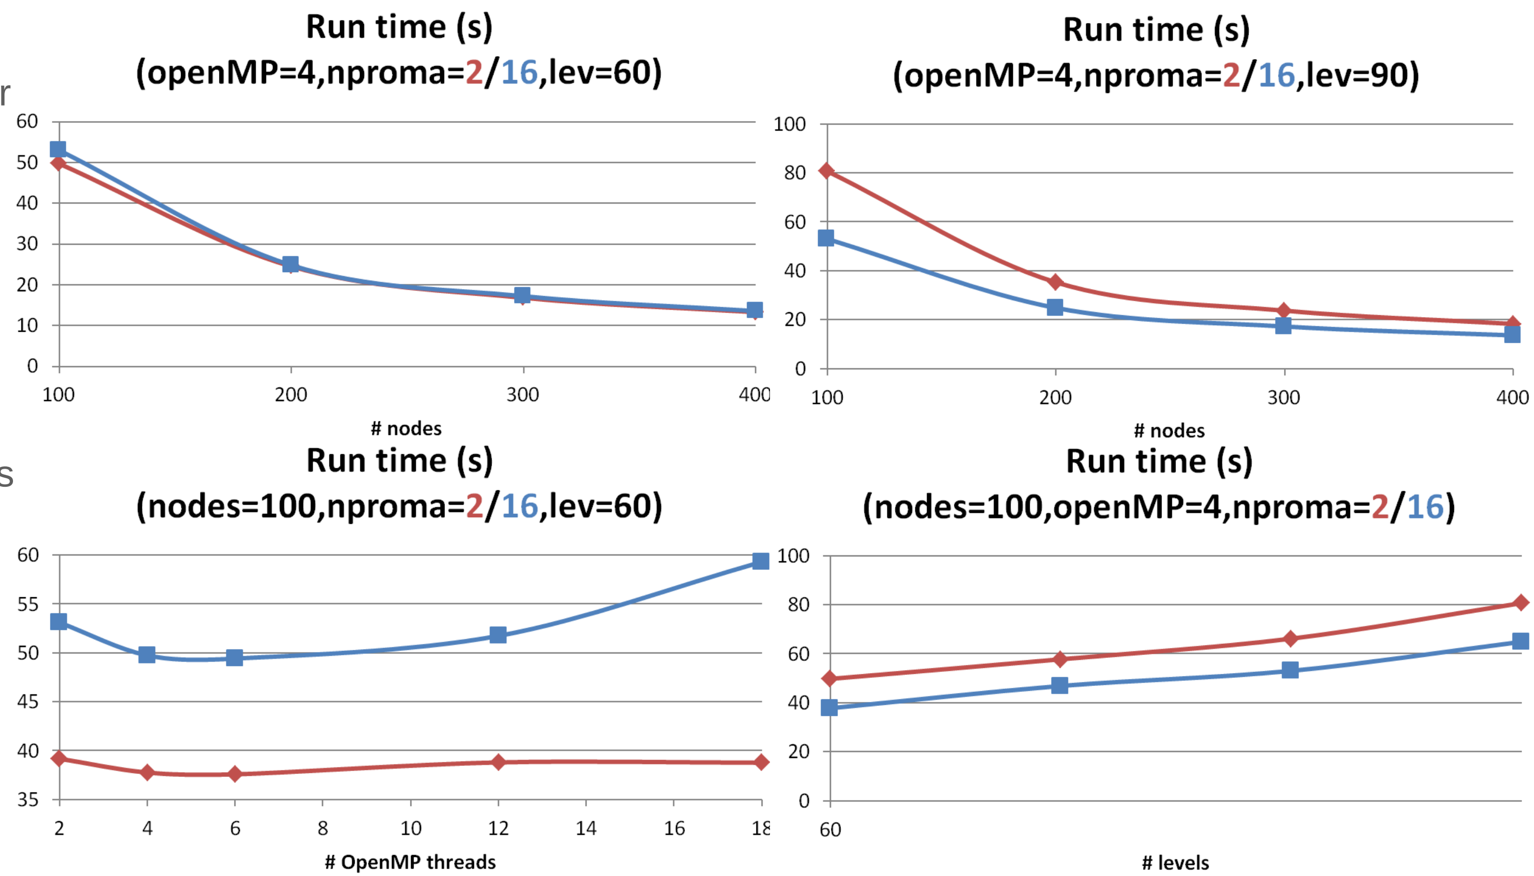

- Params: # OpenMP threads (2,4,6,12,18), # nodes (100,200,300,400), nproma (col. blocking; 2,4,8,16,32), # vert. levels (60,70,80,90)

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

Outline

- Performance Analysis and Higher Dimensions
- Sparse Grids in a Nutshell
- Regression on Sparse Grids
- Results: Molecular Dynamics Climate Weather
- Summary

Results: Weather

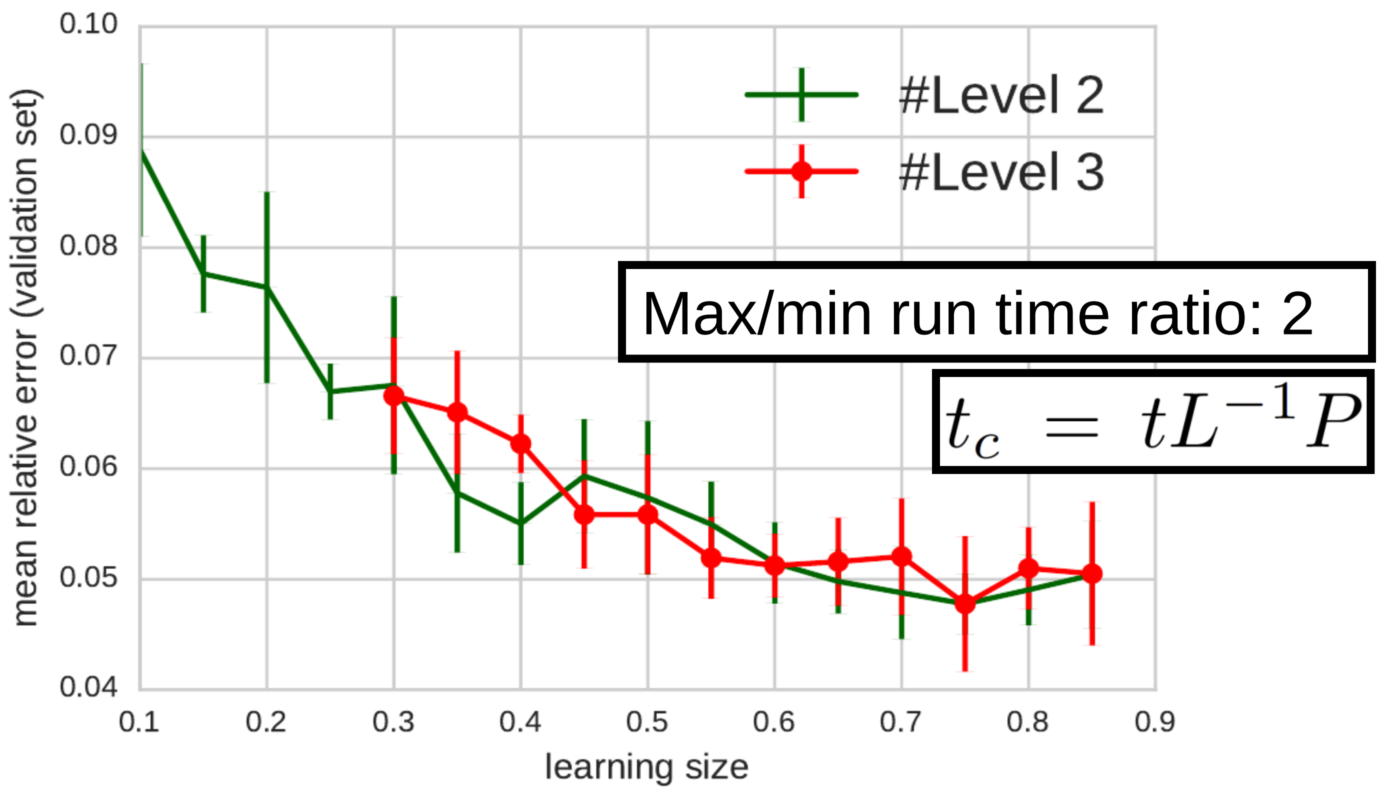

- Params: # OpenMP threads (2,4,6,12,18), # nodes (100,200,300,400), nproma (col. blocking; 2,4,8,16,32), # vert. levels (60,70,80,90)

HELMUT SCHMIDT UNIVERSITÄT

Outline

- Performance Analysis and Higher Dimensions
- Sparse Grids in a Nutshell
- Regression on Sparse Grids
- Results:
  Molecular Dynamics
  Climate
  Weather
- Summary

Summary

- Application of the sparse grid regression
  → Training of SG with performance data
  → Prediction of run times via SG basis functions
- Molecular dynamics: Accurate prediction (ca 15% dev.) using >= 180 samples to describe nonlinear 5D parameter space
- Climate: ca 2.5% deviation for small-deviation case (max/min run time ratio: 1.14)
- Future work:
  - Comparison with other methods
    → Neural networks, Gaussian process regression
  - On-the-fly data collection and prediction

P. Neumann acknowledges ESiWACE. ESiWACE has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 675191. This material reflects only the author's view and the European Commission is not responsible for any use that may be made of the information it contains.
P. Neumann acknowledges funding by the Federal Ministry of Education and Research, grant No 01IH16008B, project TaLPas.

HELMUT SCHMIDT UNIVERSITÄT

Outline

- Performance Analysis and Higher Dimensions
- Sparse Grids in a Nutshell
- Regression on Sparse Grids
- Results: Molecular Dynamics Climate Weather
- **Summary**

Interested in PhD or Postdoc?

HPC and …

- Multiscale flow simulation
- Particle simulations
- Computational fluid dynamics/ Lattice Boltzmann
- Data analytics
- Auto-tuning
- Load balancing
- Performance analysis and profiling

Contact: philipp.neumann@hsu-hh.de