HELMUT SCHMIDT UNIVERSITÄT

Chair for High Performance Computing
Philipp Neumann

TaLPas: Task-Based Load Balancing and Auto-Tuning in Particle Simulations

TU München
H.-J. Bungartz
N. Tchipev
S. Seckler
F. Gratl

Helmut-Schmidt Uni
P. Neumann

Uni Hamburg
T. Ludwig

HLRS
C. Glass
J. Gracia
N. Urmersbach
C. Niethammer

VISUS/Uni Stuttgart
G. Reina
O. Fernandes
P. Gralka

TU Darmstadt
F. Wolf
S. Shudler
S. Rinke

TU Berlin
J. Vrabec
M. Heinen

TU Kaiserslautern
H. Hasse
M. Horsch
K. Langenbach
T. Vo
M. Heier

HELMUT SCHMIDT UNIVERSITÄT

Chair for High Performance Computing
Philipp Neumann

Disclaimer!

TaLPas: Task-Based Load Balancing and Auto-Tuning in Particle Simulations

TU München	~~TU München~~ ~~Uni Hamburg~~ Helmut-Schmidt Uni	HLRS	VISUS/ Uni Stuttgart	TU Darmstadt	~~Uni Paderborn~~ TU Berlin	TU Kaiserslautern
H.-J. Bungartz	P. Neumann	C. Glass	G. Reina	F. Wolf	J. Vrabec	H. Hasse
N. Tchipev	**Uni Hamburg**	J. Gracia	O. Fernandes	S. Shudler	M. Heinen	M. Horsch
S. Seckler	T. Ludwig	N. Urmersbach	P. Gralka	S. Rinke		K. Langenbach
F. Gratl		C. Niethammer				T. Vo
		X,Y				M. Heier

HELMUT SCHMIDT UNIVERSITÄT

Outline

- AutoPas
- Workflow Manager
- Performance Prediction
- In-Situ Visualization

Recap: What is TaLPas?

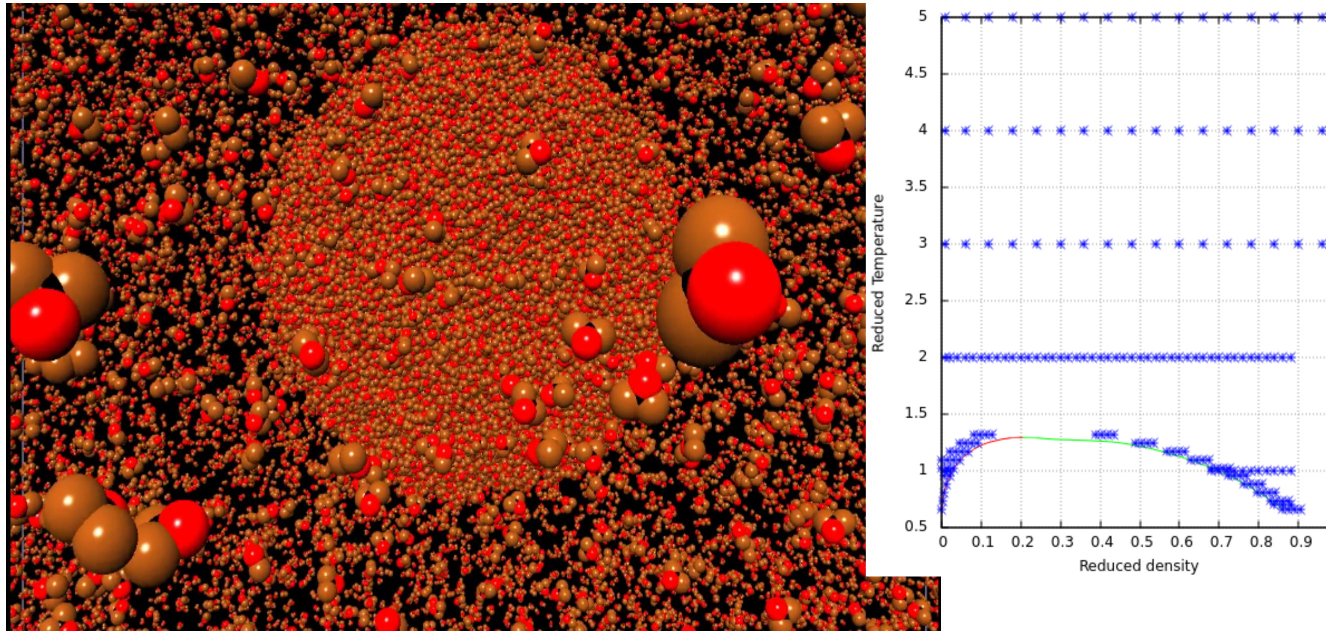

- Investigation of thermodynamic states and properties of fluids
  → vapor-liquid systems, interfacial flows, complex fluids, ...
- Particular computational challenge: problems such as equation of state sampling, rare events, ...
  → **many inter-dependent MD runs, each with different compute requirements**
  → Similar problem settings: UQ, parameter identification, ...

SPONSORED BY THE

Federal Ministry
of Education
and Research

HELMUT SCHMIDT
UNIVERSITÄT

Outline

- AutoPas
- Workflow Manager
- Performance Prediction
- In-Situ Visualization

Recap: Goals of TaLPas

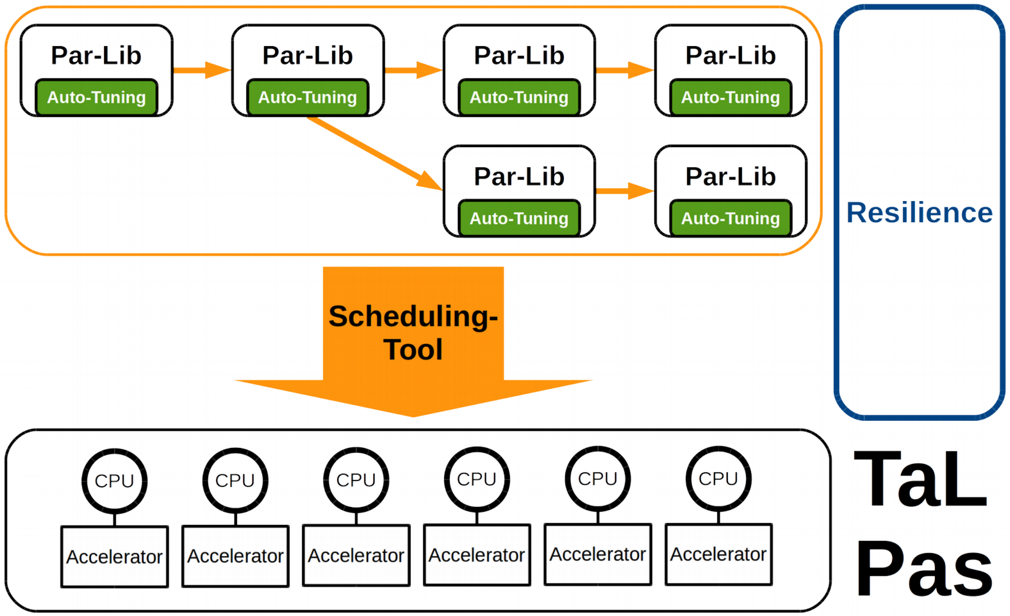

- Hardware-independent acceleration of particle simulations
  → Node level auto-tuning library AutoPas
- Self-adapting performance-optimal distribution of work load
  → workflow manager, incorporating scheduler, performance predictor, particle sampling algorithm
- Resilience for particle systems
- Integrated (in-situ) visualization

SPONSORED BY THE
Federal Ministry
of Education
and Research

HELMUT SCHMIDT UNIVERSITÄT

Outline

SPONSORED BY THE

Federal Ministry of Education and Research

HELMUT SCHMIDT UNIVERSITÄT

Outline

- AutoPas
- Workflow Manager
- Performance Prediction
- In-Situ Visualization

AutoPas[1]

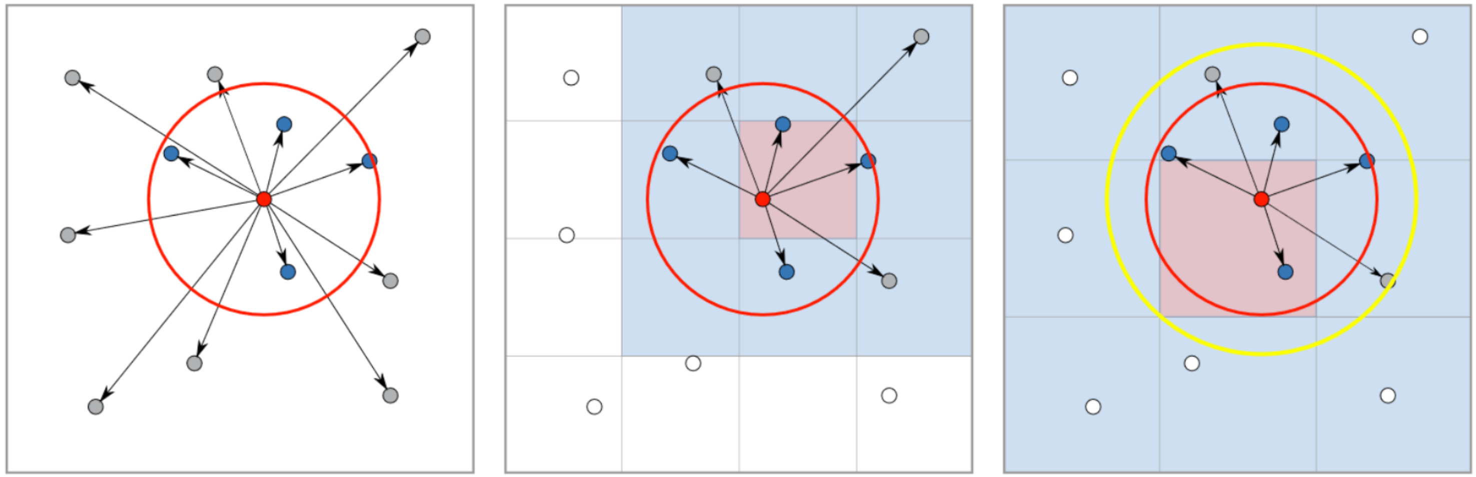

(a) DirectSum (b) LinkedCells (c) VerletLists

- Data structures: SoA vs. AoS
- 7 containers: Direct sum, linked cells, Verlet lists, cluster lists
- 17 (OpenMP) traversals: coloring, slicing, etc.
- Auto-Tuning on data structures+containers+traversals
- Integration in ls1 mardyn[2]

1 Gratl et al. AutoPas: Auto-Tuning for Particle Simulations. IPDPS proc. (iWAPT workshop), 2019
2 Tchipev et al. TweTriS: Twenty Trillion-atom Simulation. IJHPCA 33(5):838-854, 2019

SPONSORED BY THE
Federal Ministry of Education and Research

Outline

- AutoPas
- Workflow Manager
- Performance Prediction
- In-Situ Visualization

AutoPas[1]

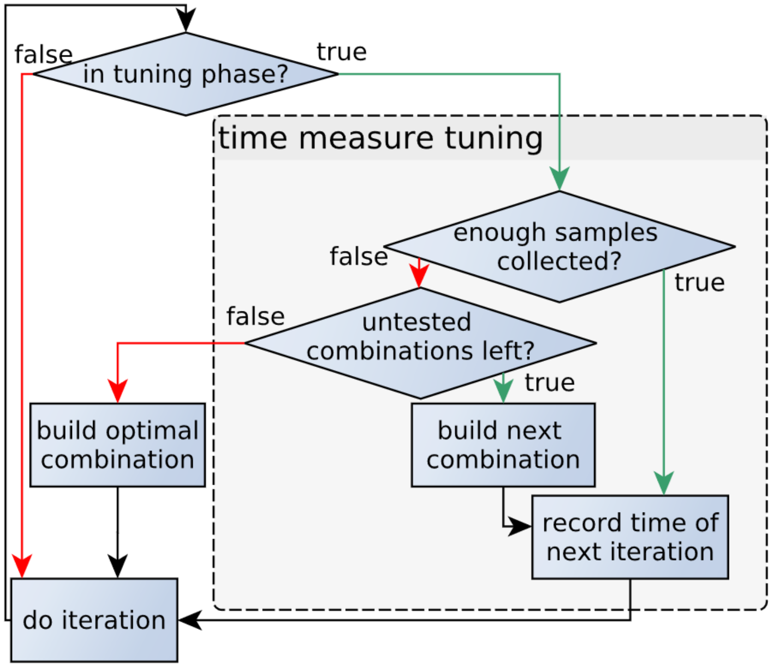

- Data structures: SoA vs. AoS
- 7 containers: Direct sum, linked cells, Verlet lists, cluster lists
- 17 (OpenMP) traversals: coloring, slicing, etc.
- Auto-Tuning on data structures+containers+traversals
- Integration in ls1 mardyn[2]

1 Gratl et al. AutoPas: Auto-Tuning for Particle Simulations. IPDPS proc. (iWAPT workshop), 2019
2 Tchipev et al. TweTriS: Twenty Trillion-atom Simulation. IJHPCA 33(5):838-854, 2019

SPONSORED BY THE
Federal Ministry of Education and Research

AutoPas[1]

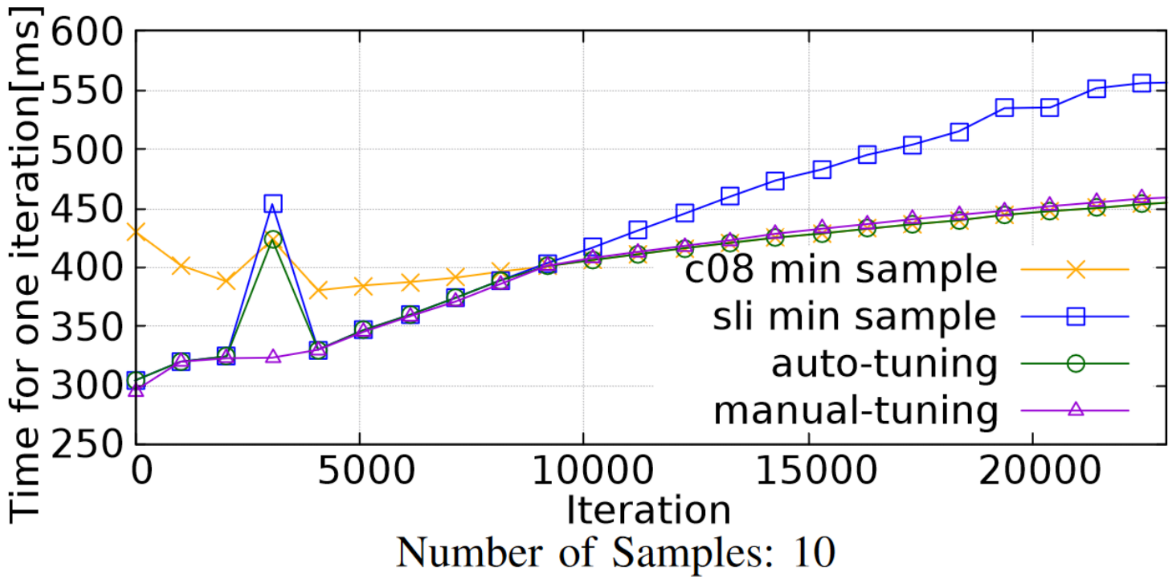

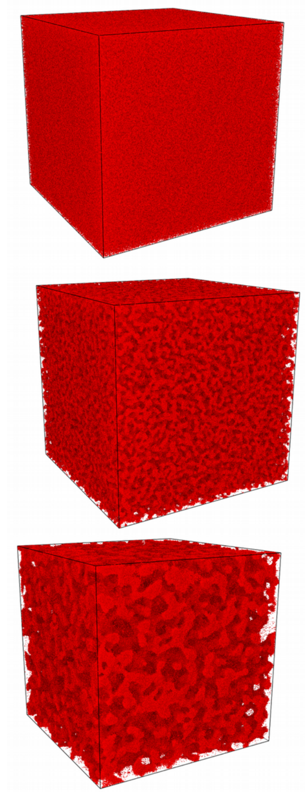

- Data structures: SoA vs. AoS
- 7 containers: Direct sum, linked cells, Verlet lists, cluster lists
- 17 (OpenMP) traversals: coloring, slicing, etc.
- Auto-Tuning on data structures+containers+traversals
- Integration in ls1 mardyn[2]

1 Gratl et al. AutoPas: Auto-Tuning for Particle Simulations. IPDPS proc. (iWAPT workshop), 2019
2 Tchipev et al. TweTriS: Twenty Trillion-atom Simulation. IJHPCA 33(5):838-854, 2019

SPONSORED BY THE
Federal Ministry of Education and Research

Outline

- AutoPas
- Workflow Manager
- Performance Prediction
- In-Situ Visualization

AutoPas[1]

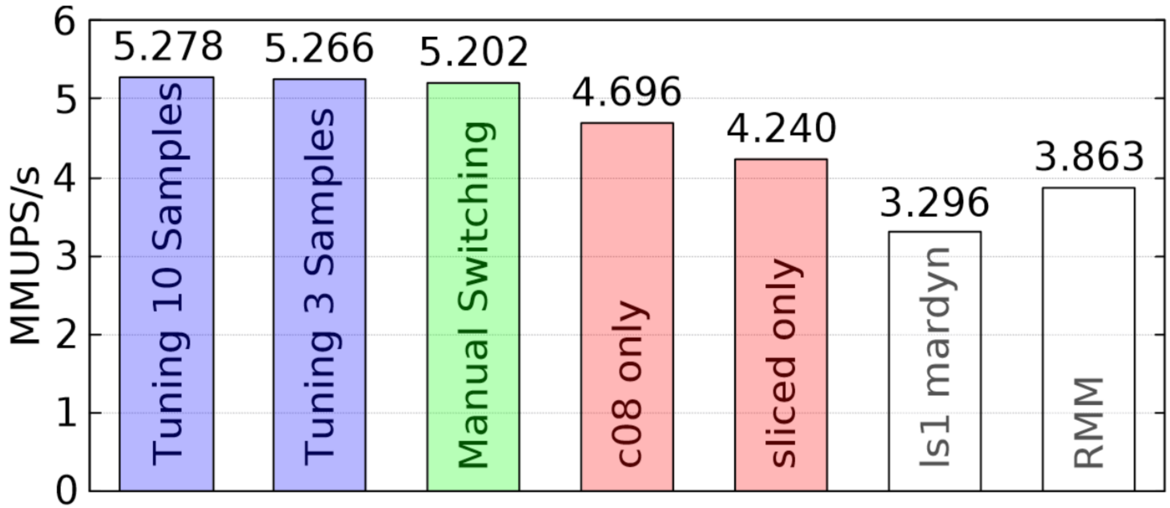

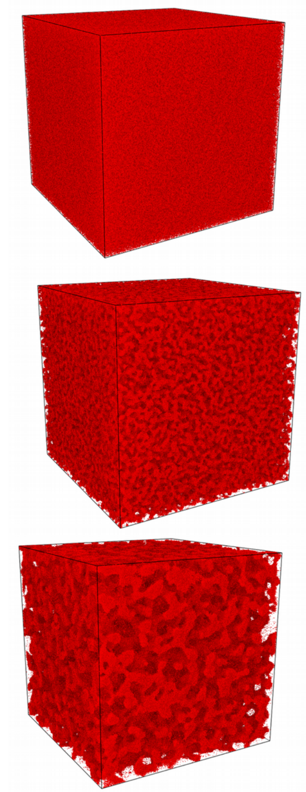

- Data structures: SoA vs. AoS
- 7 containers: Direct sum, linked cells, Verlet lists, cluster lists
- 17 (OpenMP) traversals: coloring, slicing, etc.
- Auto-Tuning on data structures+containers+traversals
- Integration in ls1 mardyn[2]

1 Gratl et al. AutoPas: Auto-Tuning for Particle Simulations. IPDPS proc. (iWAPT workshop), 2019
2 Tchipev et al. TweTriS: Twenty Trillion-atom Simulation. IJHPCA 33(5):838-854, 2019

SPONSORED BY THE
Federal Ministry
of Education
and Research

HELMUT SCHMIDT UNIVERSITÄT

Outline

- AutoPas
- Workflow Manager
- Performance Prediction
- In-Situ Visualization

AutoPas[1]

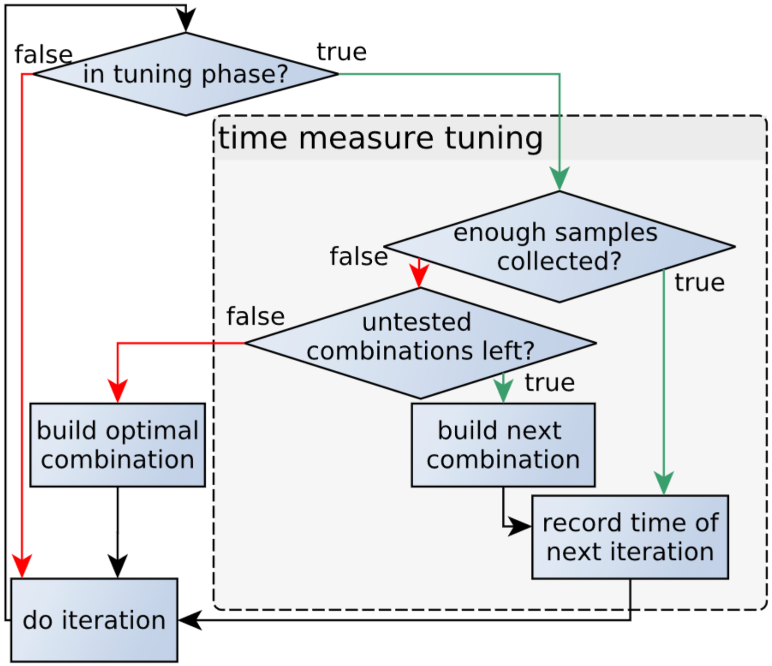

- Data structures: SoA vs. AoS
- 7 containers: Direct sum, linked cells, Verlet lists, cluster lists
- 17 (OpenMP) traversals: coloring, slicing, etc.
- Auto-Tuning on data structures+containers+traversals
- Integration in ls1 mardyn[2]

1 Gratl et al. AutoPas: Auto-Tuning for Particle Simulations. IPDPS proc. (iWAPT workshop), 2019
2 Tchipev et al. TweTriS: Twenty Trillion-atom Simulation. IJHPCA 33(5):838-854, 2019

SPONSORED BY THE
Federal Ministry of Education and Research

HELMUT SCHMIDT UNIVERSITÄT

Outline

- AutoPas
- Workflow Manager
- Performance Prediction
- In-Situ Visualization

AutoPas[1]

- Improve on "Brute-Force" testing due to outgrowing number of combinations
  → not a nice add-on, but a necessity!
  → work in progress: ML, Bayesian methods

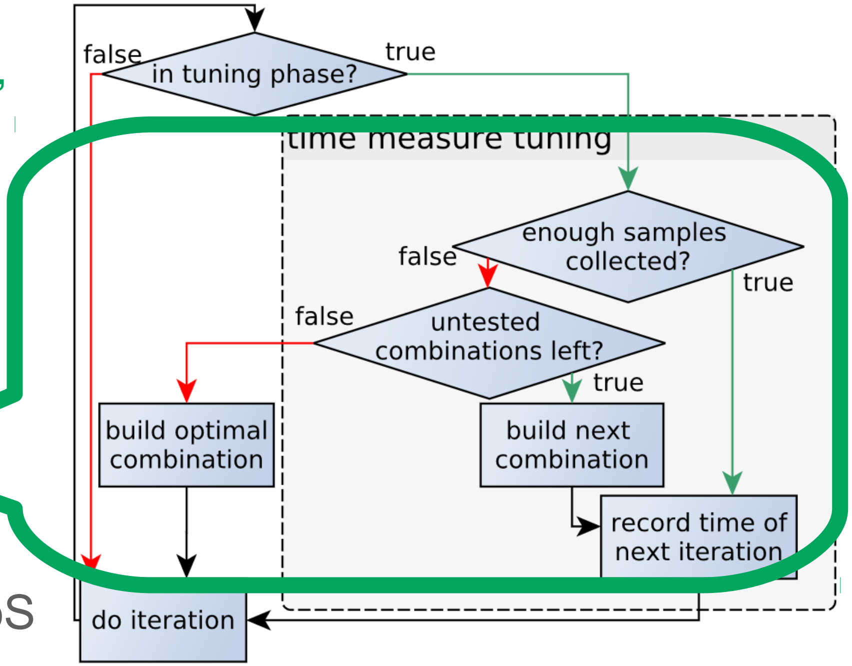

- Data structures: SoA vs. AoS
- 7 containers: Direct sum, linked cells, Verlet lists, cluster lists
- 17 (OpenMP) traversals: coloring, slicing, etc.
- Auto-Tuning on data structures+containers+traversals
- Integration in ls1 mardyn[2]

1 Gratl et al. AutoPas: Auto-Tuning for Particle Simulations. IPDPS proc. (iWAPT workshop), 2019
2 Tchipev et al. TweTriS: Twenty Trillion-atom Simulation. IJHPCA 33(5):838-854, 2019

SPONSORED BY THE
Federal Ministry of Education and Research

AutoPas[1]

- Improve on "Brute-Force" testing due to outgrowing number of combinations
  → not a nice add-on, but a necessity!
  → work in progress: ML, Bayesian methods
- CUDA & Kokkos support

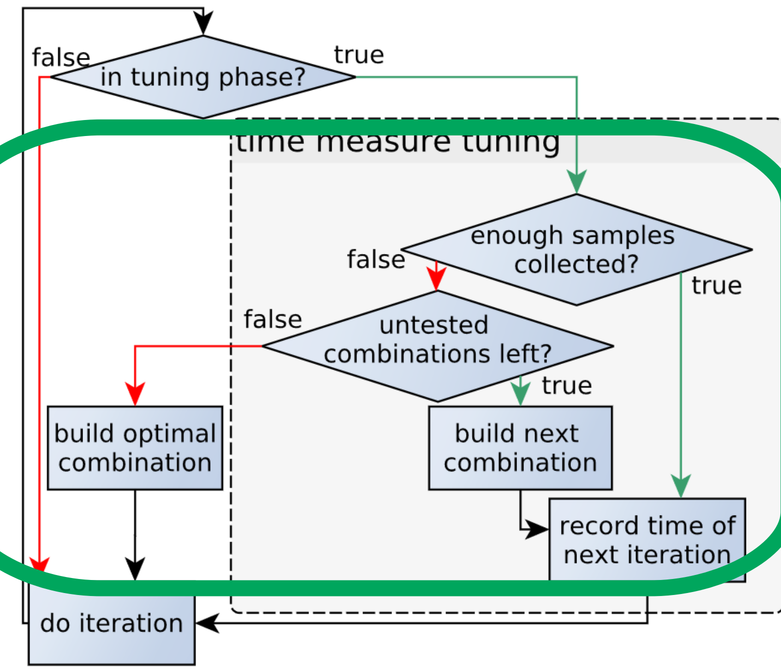

- Data structures: SoA vs. AoS
- 7 containers: Direct sum, linked cells, Verlet lists, cluster lists
- 17 (OpenMP) traversals: coloring, slicing, etc.
- Auto-Tuning on data structures+containers+traversals
- Integration in ls1 mardyn[2]

1 Gratl et al. AutoPas: Auto-Tuning for Particle Simulations. IPDPS proc. (iWAPT workshop), 2019
2 Tchipev et al. TweTriS: Twenty Trillion-atom Simulation. IJHPCA 33(5):838-854, 2019

SPONSORED BY THE
Federal Ministry of Education and Research

AutoPas: Machine Learning for Auto-Tuning

Outline

- AutoPas
- Workflow Manager
- Performance Prediction
- In-Situ Visualization

- Based on classification via neural network
- Test only five configurations that are deemed best
- Finds (equivalent to) optimal configuration with probability of 99%
- Relies on training data
  → currently, it takes days to train…
  → hardware-dependent training
- Results from a recent evaluation:
  Full search: duration: 22.79s # iterations: 27, 0.84 s/it.
  ML-based: duration: 2.05s # iterations: 5, 0.41s/it.
  → optimal configuration: 0.36s/it (detected by both approaches)
- This is work in progress :-)

SPONSORED BY THE
Federal Ministry of Education and Research

Outline

- AutoPas
- Workflow Manager
- Performance Prediction
- In-Situ Visualization

SPONSORED BY THE

Federal Ministry of Education and Research

Workflow Manager

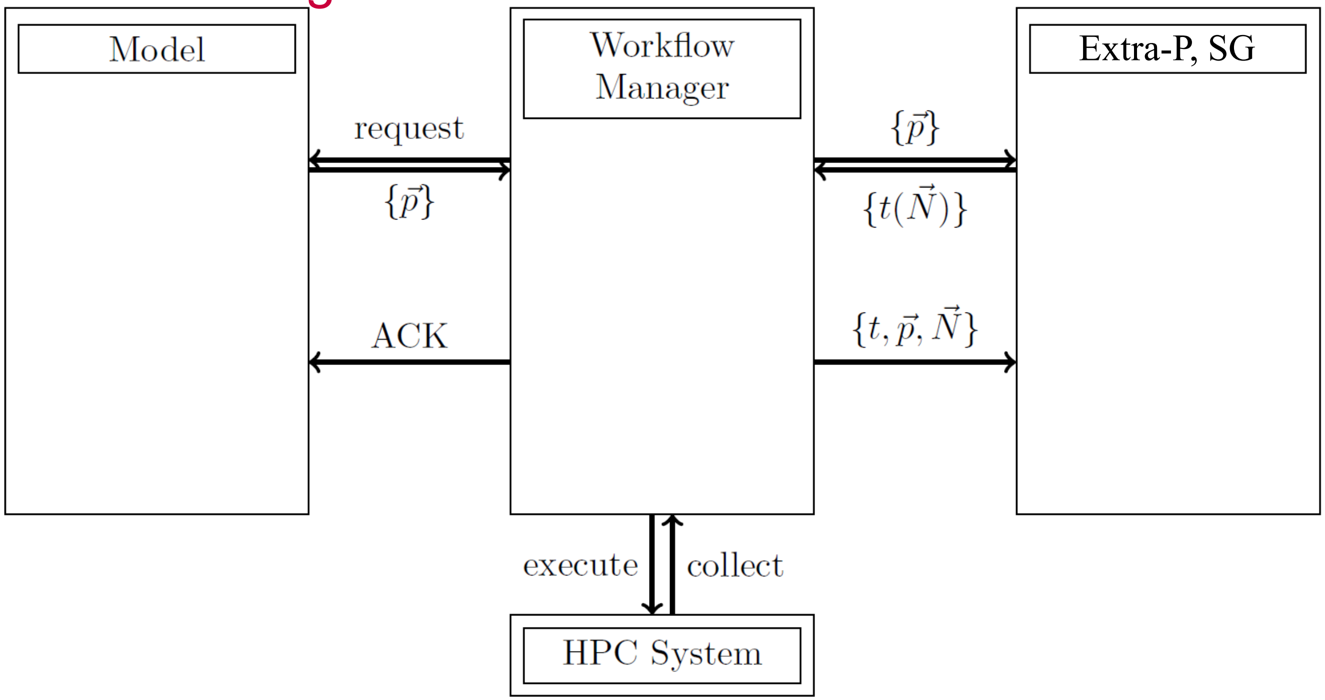

- Last time: Equation of state fitting and vapor-liquid-equilibrium envelope
- Current work: Sparse Grid-based Perf. Prediction, Adsorption processes[3]
  - → study of adsorption in dispersive systems
  - → relevant parameters: temperature, pressure, solid density, interaction fluid-solid

    → 350 simulations

3 Heier et al. J. Chem. Eng. Data 64:386-394, 2019

HELMUT SCHMIDT UNIVERSITÄT

Outline

- AutoPas
- Workflow Manager
- Performance Prediction
- In-Situ Visualization

Workflow Manager

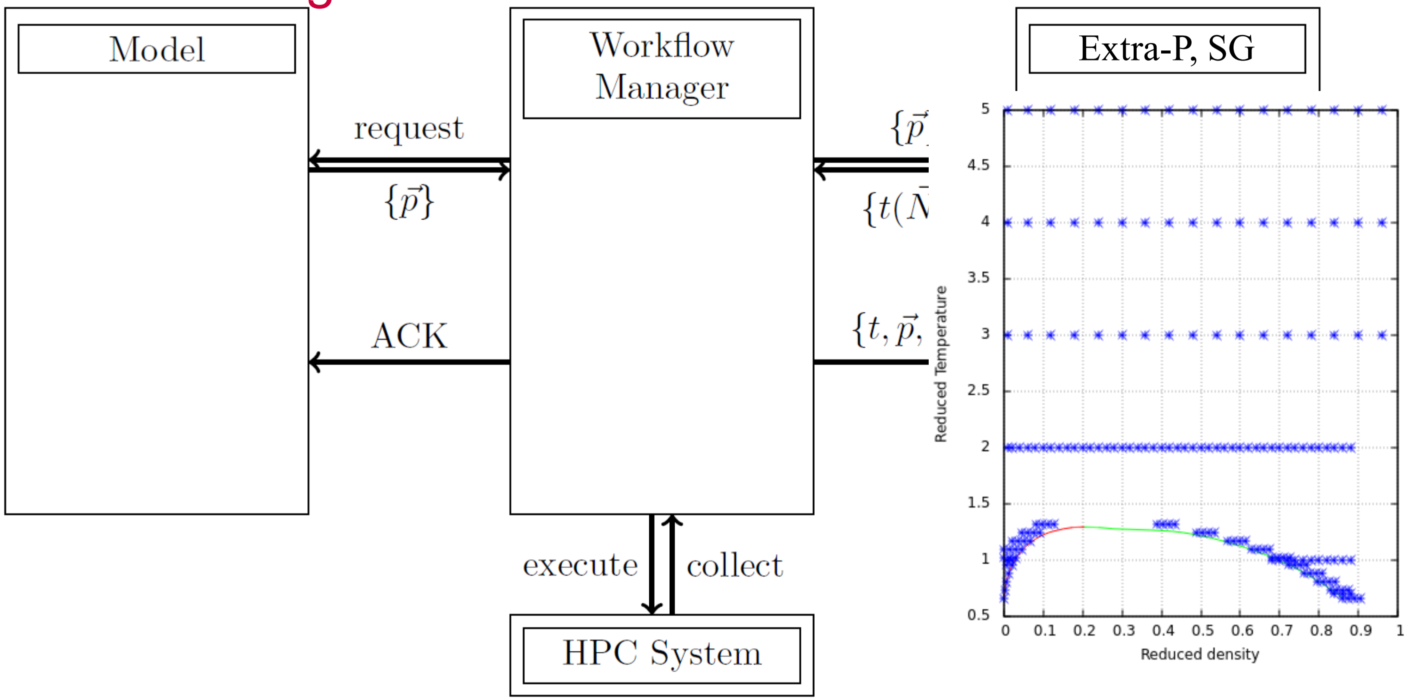

- Last time: Equation of state fitting and vapor-liquid-equilibrium envelope
- Current work: Adsorption processes[3]
  - → study of adsorption in dispersive systems
  - → relevant parameters: temperature, pressure, solid density, interaction fluid-solid
    - → 350 simulations

3 Heier et al. J. Chem. Eng. Data 64:386-394, 2019

SPONSORED BY THE Federal Ministry of Education and Research

HELMUT SCHMIDT UNIVERSITÄT

Outline

- AutoPas
- Workflow Manager
- Performance Prediction
- In-Situ Visualization

Workflow Manager

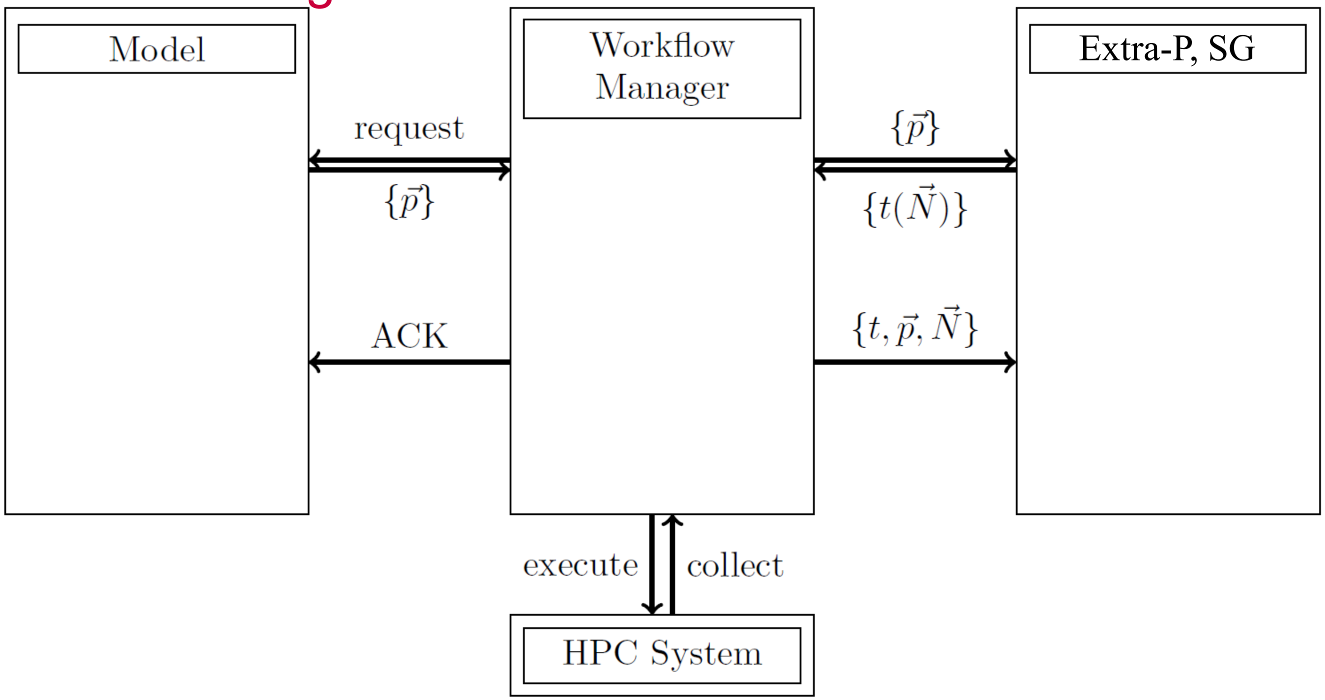

- Last time: Equation of state fitting and vapor-liquid-equilibrium envelope
- Current work: Adsorption processes[3]
  - → study of adsorption in dispersive systems
  - → relevant parameters: temperature, pressure, solid density, interaction fluid-solid

  → 350 simulations

3 Heier et al. J. Chem. Eng. Data 64:386-394, 2019

SPONSORED BY THE Federal Ministry of Education and Research

HELMUT SCHMIDT UNIVERSITÄT

Outline

- AutoPas
- Workflow Manager
- Performance Prediction
- In-Situ Visualization

Workflow Manager and Adsorption

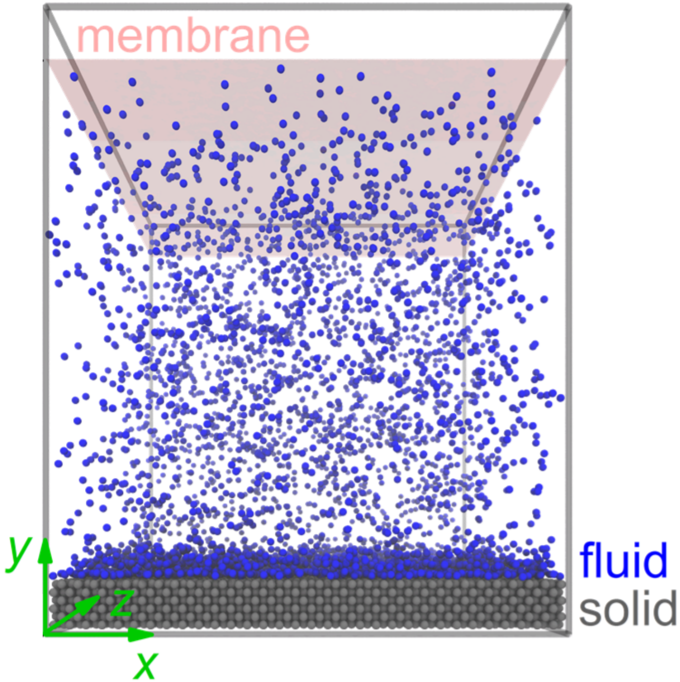

- Last time: Equation of state fitting and vapor-liquid-equilibrium envelope
- Current work: Adsorption processes[3]
  → study of adsorption in dispersive systems
  → relevant parameters: temperature, pressure, solid density, interaction fluid-solid
  → 350 simulations

3 Heier et al. J. Chem. Eng. Data 64:386-394, 2019

SPONSORED BY THE
Federal Ministry of Education and Research

HELMUT SCHMIDT UNIVERSITÄT

Workflow Manager and Adsorption

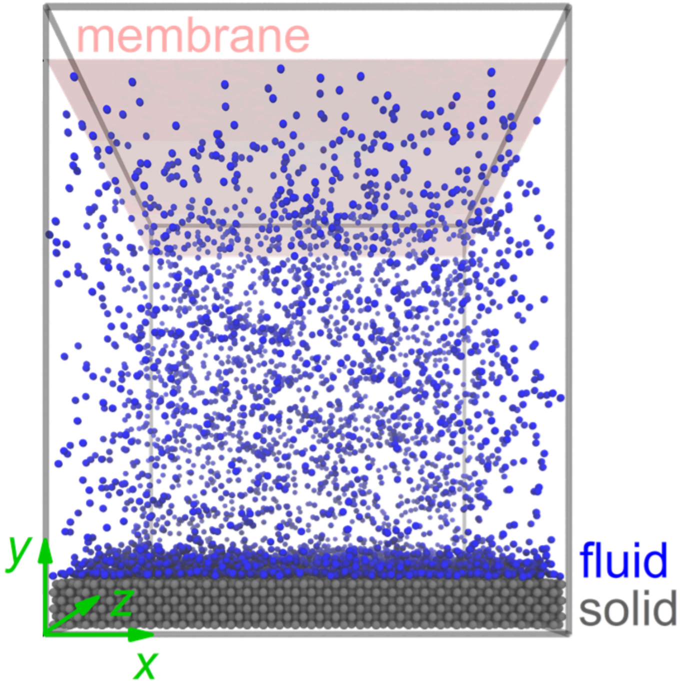

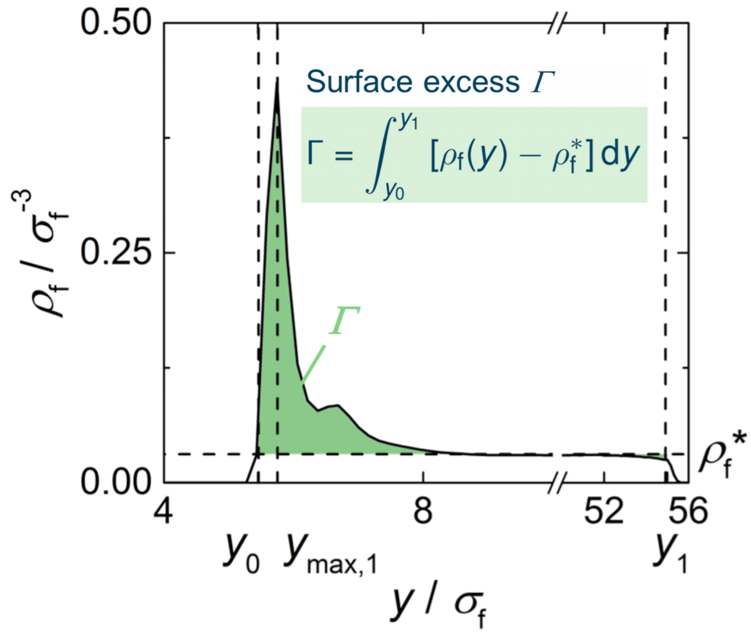

- Last time: Equation of state fitting and vapor-liquid-equilibrium envelope
- Current work: Adsorption processes[3]
  - → study of adsorption in dispersive systems
  - → relevant parameters: temperature, pressure, solid density, interaction fluid-solid

    → 350 simulations

3 Heier et al. J. Chem. Eng. Data 64:386-394, 2019

SPONSORED BY THE

Federal Ministry of Education and Research

Outline

- AutoPas
- Workflow Manager
- Performance Prediction
- In-Situ Visualization

Workflow Manager and Adsorption

Influence of fluid-solid interaction

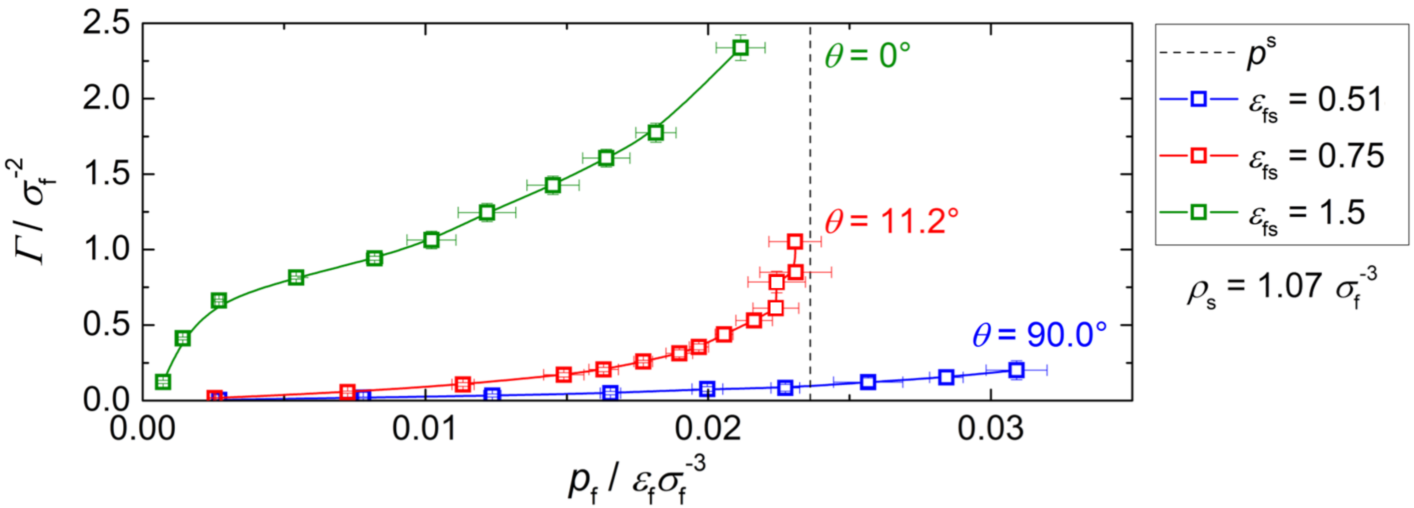

- Last time: Equation of state fitting and vapor-liquid-equilibrium envelope
- Current work: Adsorption processes[3]
  - → study of adsorption in dispersive systems
  - → relevant parameters: temperature, pressure, solid density, interaction fluid-solid
    - → 350 simulations

3 Heier et al. J. Chem. Eng. Data 64:386-394, 2019

SPONSORED BY THE
Federal Ministry of Education and Research

Outline

- AutoPas
- Workflow Manager
- Performance Prediction
- In-Situ Visualization

Workflow Manager and Adsorption

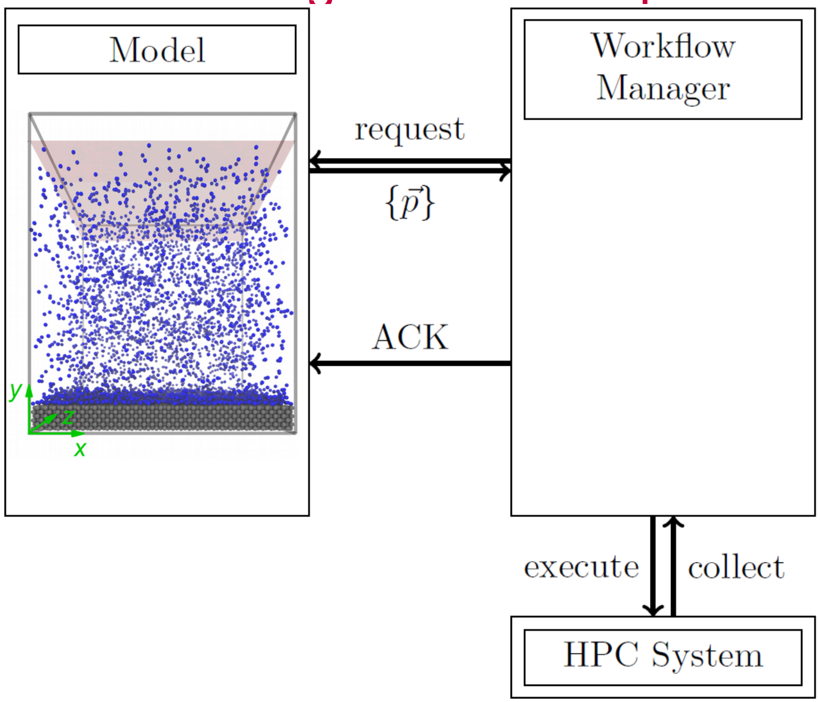

- Still some user interaction required → work in progress
- Debugging: some SLURM scalability issues…

- Last time: Equation of state fitting and vapor-liquid-equilibrium envelope
- Current work: Adsorption processes[3]
  - → study of adsorption in dispersive systems
  - → relevant parameters: temperature, pressure, solid density, interaction fluid-solid
    - → 350 simulations

3 Heier et al. J. Chem. Eng. Data 64:386-394, 2019

SPONSORED BY THE
Federal Ministry of Education and Research

HELMUT SCHMIDT UNIVERSITÄT

Performance Prediction with Extra-P: Higher-Dimensional Parameter Spaces

Outline

- AutoPas
- Workflow Manager
- Performance Prediction
- In-Situ Visualization

HELMUT SCHMIDT UNIVERSITÄT

- Extra-P
  - Originally:
    N parameters → 5^N data points to generate performance model
    normal form

$$f(r_1, r_2, \ldots, r_q) = \sum_{k=1}^{n} c_k \cdot \prod_{l=1}^{q} r_l^{i_{k_l}} \cdot \log^{j_{k_l}}(r_l)$$

  - Idea: Select subset of data points without degrading accuracy
    → reinforcement learning approach

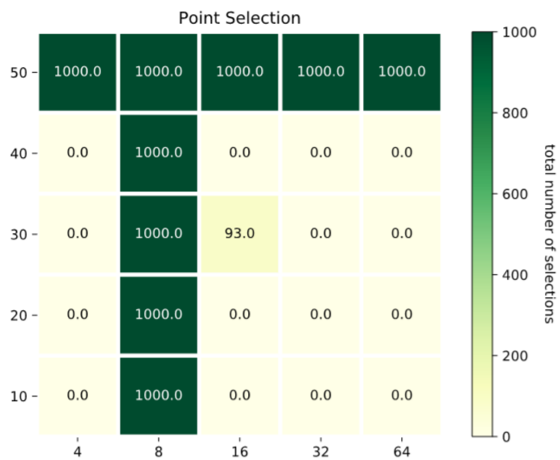

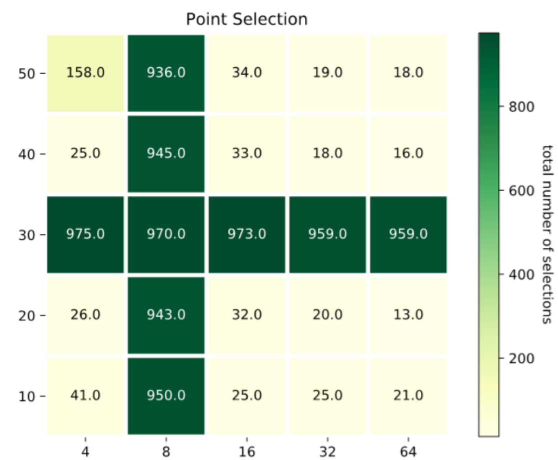

SPONSORED BY THE
Federal Ministry of Education and Research

Performance Prediction with Extra-P: Higher-Dimensional Parameter Spaces

HELMUT SCHMIDT UNIVERSITÄT

- Extra-P
  - Originally:
    N parameters → 5^N data points to generate performance model normal form

$$f(r_1, r_2, ..., r_q) = \sum_{k=1}^{n} c_k \cdot \prod_{l=1}^{q} r_l^{i_{k_l}} \cdot \log^{j_{k_l}}(r_l)$$

  - Idea: Select subset of data points without degrading accuracy
    → reinforcement learning approach
    → example: 3 params, 5% noise

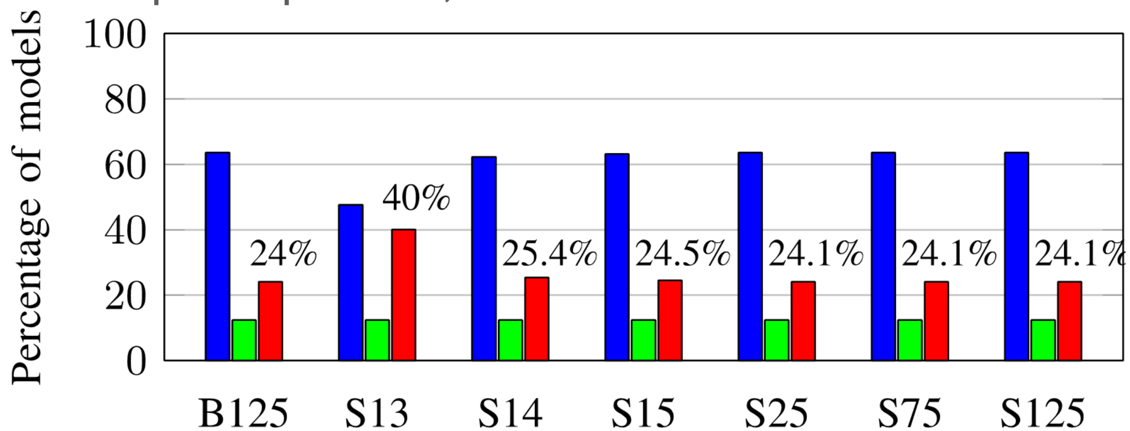

SPONSORED BY THE
Federal Ministry of Education and Research

Outline

- AutoPas
- Workflow Manager
- Performance Prediction
- In-Situ Visualization

Performance Prediction with Sparse Grids[4]

- SG=Efficient hierarchical discretization for high-dimensional problems
- Method: Sparse grid regression with local mesh refinement
- Example: Molecular dynamics
  - Particle density $\rho \in [0.3;0.9]$
  - # particles $N \in [10^3;10^5]$
  - Cut-off radius $r_c \in [1.2;4.5]$
  - Blocksize $b \in [10;10^3]$
  - # MPI procs $p \in \{1,2,4,8\}$

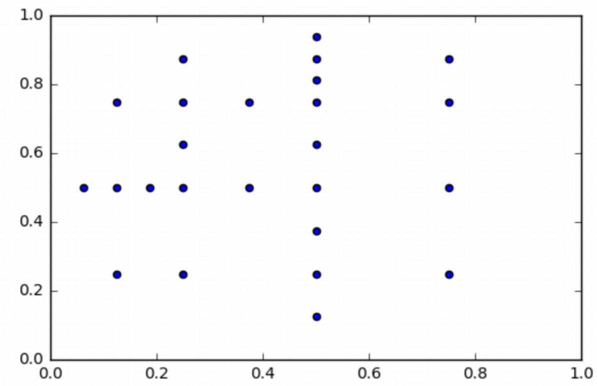

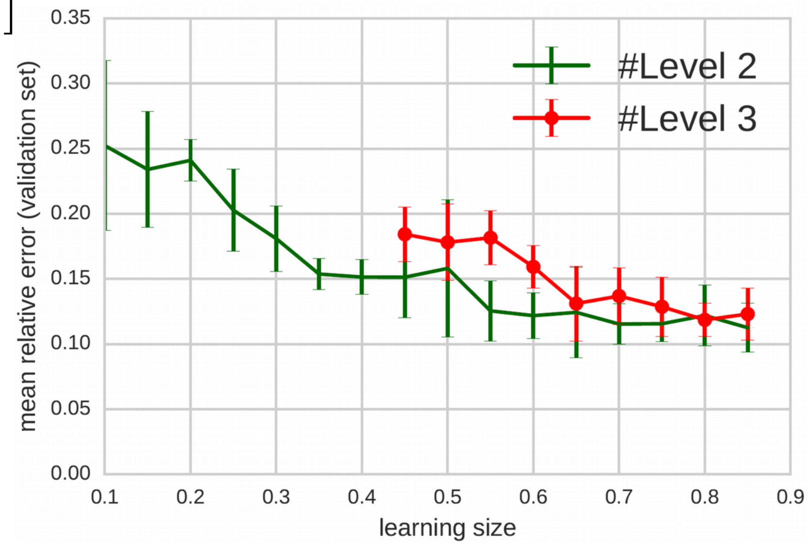

4 P. Neumann. Sparse Grid Regression for Performance Prediction Using High-Dimensional Run Time Data. Euro-Par Workshops (PMACS) 2019. Accepted

SPONSORED BY THE
Federal Ministry of Education and Research

HELMUT SCHMIDT UNIVERSITÄT

Outline

- AutoPas
- Workflow Manager
- Performance Prediction
- **In-Situ Visualization**

In-Situ Visualization

- MegaMol: Integration in ls1
  - Data exchange via shared memory
  - Different MPI worlds
    → if visualization fails, simulation continues

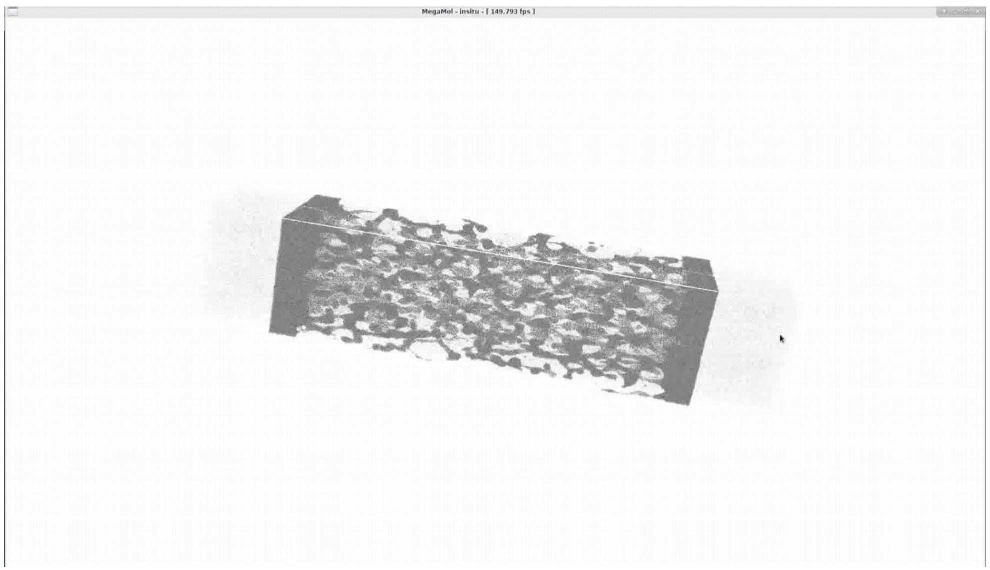

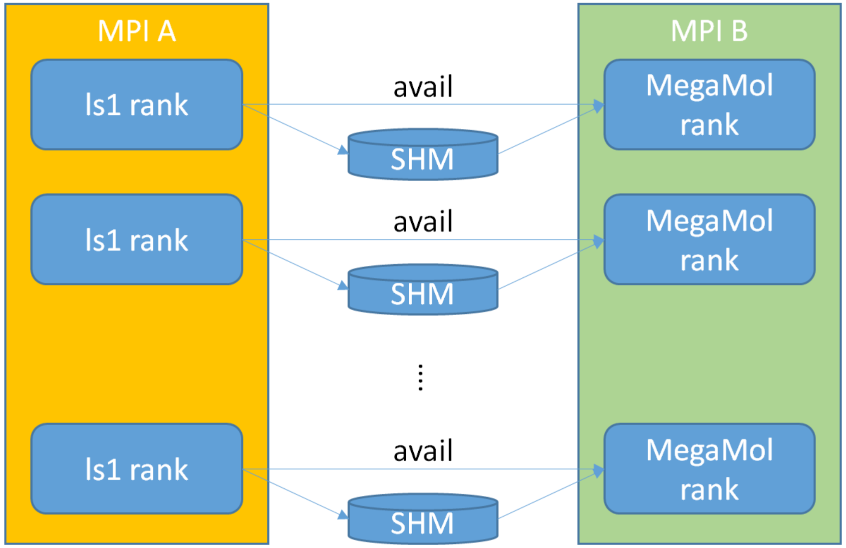

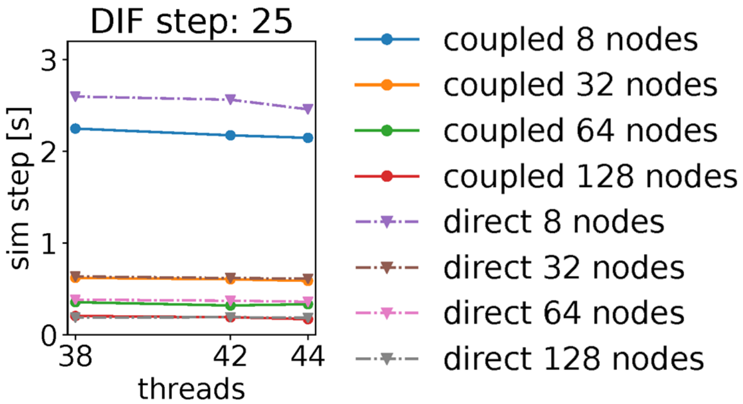

HELMUT SCHMIDT UNIVERSITÄT

Outline

- AutoPas
- Workflow Manager
- Performance Prediction
- In-Situ Visualization

SPONSORED BY THE

Federal Ministry of Education and Research

Resilience

- Checkpointing-based
- Future work: evaluation of quality for molecular films

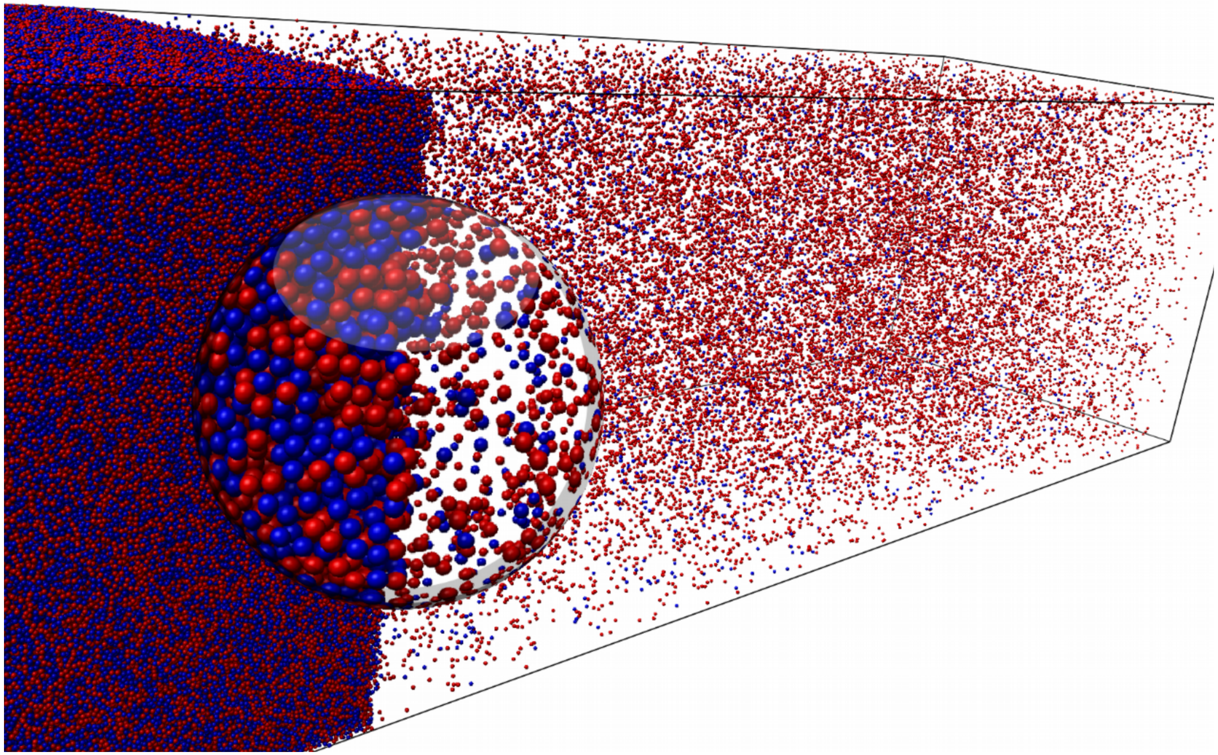

Outline

- AutoPas
- Workflow Manager
- Performance Prediction
- In-Situ Visualization

Summary

- AutoPas
  - Integration in ls1, first application-relevant tests
  - Work in progress: improved tuning procedure, CUDA support
- Workflow Manager
  - New scenario: adsorption processes
    → first (hand-steered) workflow solution using ls1
  - Work in progress: Scheduling with performance prediction, integration of components, scalability
- Performance Prediction
  - New methods for high-dimensional parameter spaces (Extra-P, Sparse Grids)
  - Work in progress: Integration with Workflow Manager
- In-Situ Visualization: Integration of MegaMol and ls1

SPONSORED BY THE

Federal Ministry
of Education
and Research

P. Neumann and colleagues acknowledge funding by the Federal Ministry of Education and Research, grants 01IH16008, project TaLPas.

Outline

- AutoPas
- Workflow Manager
- Performance Prediction
- In-Situ Visualization

Summary

- AutoPas
  - Integration in ls1, first application-relevant tests
  - Work in progress: improved tuning procedure, CUDA support
- Workflow Manager
  - New scenario: adsorption processes
    → first (hand-steered) workflow solution using ls1
  - Work in progress: Scheduling with performance prediction, integration of components, scalability
- Performance Prediction
  - New methods for high-dimensional parameter spaces (Extra-P, Sparse Grids)
  - Work in progress: Integration with Workflow Manager
- In-Situ Visualization: Integration of MegaMol and ls1

SPONSORED BY THE

Federal Ministry of Education and Research

P. Neumann and colleagues acknowledge funding by the Federal Ministry of Education and Research, grants 01IH16008, project TaLPas.

Advertisement

HELMUT SCHMIDT UNIVERSITÄT

Outline

- AutoPas
- Workflow Manager
- Performance Prediction
- In-Situ Visualization

- ISC HPC 2020: Project Posters, Submission Deadline: 12 Feb 2020
  → https://www.isc-hpc.com/project-posters-2020.html
- (If you know someone who's) Interested in a PhD/Postdoc on HPC, multiscale flow simulation, molecular dynamics, performance prediction, etc.
  → contact me, philipp.neumann@hsu-hh.de :-)

SPONSORED BY THE
Federal Ministry of Education and Research