

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


Overview

Duration: 01.01.2017 – 31.12.2019
 Coordination: Universität Hamburg, Technische Universität München
 Partners: High-Performance Computing Center Stuttgart/Universität Stuttgart
 Technische Universität Darmstadt
 Technische Universität Berlin
 Technische Universität Kaiserslautern

Objective and Methodology


TaLPas will provide a solution to fast and robust simulation of many, inter-dependent particle systems in peta- and exascale supercomputing environments. This will be beneficial for a wide range of applications, including sampling in molecular dynamics (rare event sampling, sampling of equations of state, etc.), uncertainty quantification (sensitivity investigation of parameters on actual simulation results), or parameter identification (fitting numerical model parameters to match experiments). For this purpose, TaLPas targets

1. the development of an auto-tuning based particle simulation library *AutoPas* to leverage optimal node-level performance,
2. the development of a scalable workflow manager to optimally distribute inter-dependent particle simulation tasks on HPC compute resources,
3. the investigation of performance prediction methods for particle simulations to support auto-tuning and to feed the workflow manager with accurate runtime predictions,
4. the integration of 1-3, augmented by visualization of the sampling (parameter space exploration) and an approach to resilience. The latter will guarantee robustness at peta- and exascale.




AutoPas¹ – Auto-Tuning Library for Particle Simulations

- Foci: Node-level performance, short-range systems
 - Modular C++-template design to enable the selection of optimal SIMD/OpenMP/data structure configuration on-the-fly
 → auto-vectorization, SoA/AoS, linked cells, Verlet lists (see right), ...
- See <https://github.com/AutoPas>




Impact of OpenMP scheme and Auto-Pas performance in two transient scenarios (exploding droplet, spinodal decomposition)




Performance Prediction for Particle Simulations

- Multi-parameter performance predictions to enable efficient scheduling of heterogeneous particle simulation ensembles
- Approach 1: Automatic performance model generation with the tool *Extra-P²*
- Approach 2: Sparse grid regression³ (see right)




Workflow Manager

- Intelligent scheduling of workloads on HPC compute resources
- *Extra-P* integration: retrieves runtime models, returns runtime predictions




Resilience and Visual Parameter Space Exploration



- Analysis of ensembles by visualizing similar behavior in parameter space
- Resilience support through improvement of the checkpointing file format
- In-Memory checkpointing

Using the Workflow Manager for Equation of State Sampling



1 F. Gratl, S. Seckler, N. Tchipev, H.-J. Bungartz, P. Neumann. *AutoPas: Auto-Tuning for Particle Simulations*. In Proc. of IPDPS, workshop iWAPT, 2019 (to appear)
 2 S. Shudler, J. Vrabec, F. Wolf. *Understanding the Scalability of Molecular Simulation using Empirical Performance Modeling*. In Proc. of SC, workshop ESPT, 2018 (to appear)
 3 P. Neumann. *Sparse Grid Regression for Performance Prediction Using High-Dimensional Run Time Data*. Submitted

