TaLPas

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

Overview

Duration: 01.01.2017 – 31.12.2019
Coordination: Universität Hamburg (*UHH*), Technische Universität München (*SCCS*)
Partners: High-Performance Computing Center Stuttgart (*HLRS*)
Technische Universität Darmstadt (*PARA*)
Universität Stuttgart (*VISUS*)
Universität Paderborn (*THET*)
Technische Universität Kaiserslautern (*LTD*)

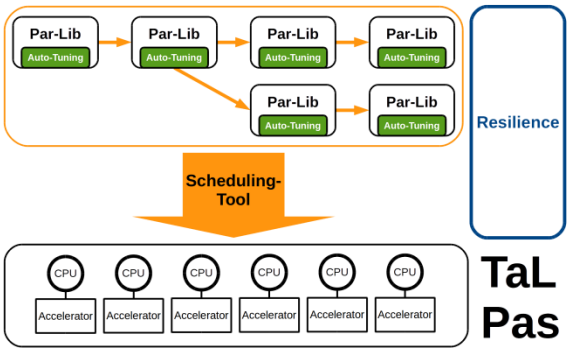

Objective and Methodology

***TaLPas* will provide a solution to fast and robust simulation of many, inter-dependent particle systems in peta- and exascale supercomputing environments.**
Therefore, ***TaLPas* bundles interdisciplinary expert knowledge** on high-performance computing (*SCCS,UHH,HLRS*), visualisation and resilience (*VISUS*), performance modeling (*PARA*), and particle applications (*THET,LTD*).

Application Fields

TaLPas findings and software which is developed in *TaLPas* will be beneficial for a wide range of applications, including

- sampling in molecular dynamics: so-called “rare events”, e.g. droplet formation, require a multitude of molecular dynamics simulations to investigate the actual conditions of phase transition. Another example is given by sampling of (highly nonlinear) equations of state for complex fluids. The latter necessitates the investigation of the phase space (spanned by temperature, pressure, etc.),
- uncertainty quantification: various simulations are performed using different parametrisations to investigate parameter sensitivities,
- parameter identification: given, e.g., a set of experimental data and a molecular model, an optimal set of model parameters needs to be found to fit the model to the experiment.

Development

For this purpose, *TaLPas* targets

1. the **development of innovative auto-tuning based particle simulation software in form of an open-source library to leverage optimal node-level performance and enable performance portability** for small- to mid-sized particle systems (MD, SPH, DPD, applications in visualisation, etc.),
2. the **development of a scalable task scheduler to optimally distribute inter-dependent particle simulation tasks** on available HPC resources, extended by **interfaces for performance prediction** for individual simulations **and application-specific scheduling algorithms**,
3. the **investigation of performance prediction methods** for particle simulations to support auto-tuning and to feed the scheduler with accurate runtime predictions, with one focus on having only a moderate number of runtimes in a multi-dimensional parameter space available,
4. the **integration of auto-tuning based particle simulation, task scheduler and performance prediction, augmented by visualisation of the sampling (parameter space exploration) and an approach to resilience**. The latter will guarantee robustness at peta- and exascale.

Current work focuses on steps 1-3. The integration of all components (step 4) is anticipated for the year 2019.

Towards an Auto-Tuning Particle Library

- Base for N-body interaction computations
- Manages node-level performance internally via auto-tuning (checking parameter and particle states such as the number of particles)
- Modular C++-template design to select optimal SIMD/OpenMP/data structure configuration on-the-fly

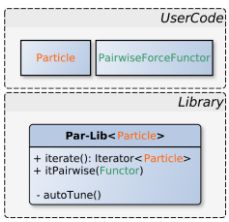

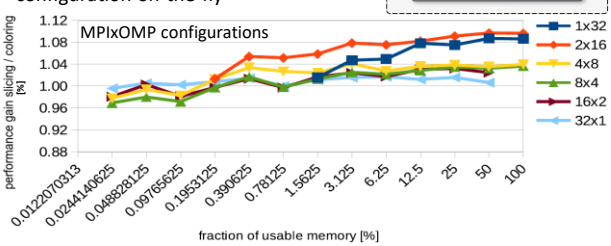

Resilience and Visual Parameter Space Exploration

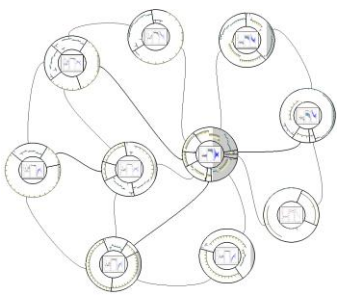

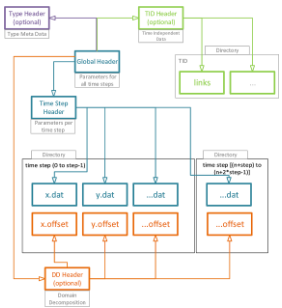

- Analysis of ensembles by visualizing similar behavior in parameter space
- Supporting resilience by improving the checkpointing file format

The Scheduler

- Gets tasks from the simulation model
- Intelligently manages the assignment of available HPC compute resources
- Interacts with *Extra-P*: retrieves task runtime models and returns execution times

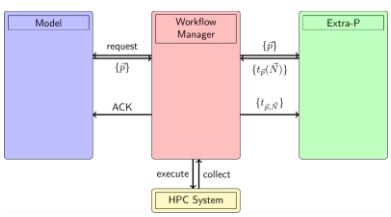

Performance Prediction with the Automatic Performance-Modeling Tool *Extra-P*

- Automated performance model generation
- Supports performance models with multiple parameters
- GUI and library interface available

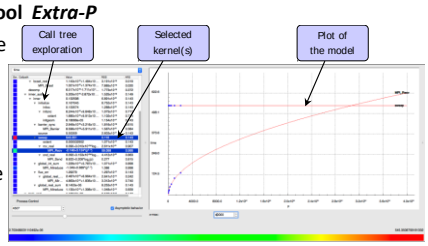

Contact: Dr. Philipp Neumann, Universität Hamburg, Scientific Computing, philipp.neumann@uni-hamburg.de
Website: www.talpas.de
This project has received funding from the Federal Ministry of Education and Research, grant numbers 01IH16008A-F.
SPONSORED BY THE Federal Ministry of Education and Research