

Scaling JAX-SPH

Interdisciplinary Project (IDP), Semester Thesis

In recent years, JAX-based differentiable fluid mechanics solvers like [PhiFlow](#), [JAX-CFD](#), and [JAX-Fluids](#) have been developed to aid the exploration of Machine Learning (ML) approaches to Computational Fluid Dynamics (CFD). However, these three CFD solvers implement grid-based approaches, e.g. Finite Differences or Finite Volumes schemes. We are the first to develop a particle-based CFD solver in JAX termed [JAX-SPH](#) and just published it in [Toshev et al. \(2024\)](#).

In this work, we want to extend the existing codebase in one major way – scalability to larger particle systems. Currently, we can simulate up to roughly 30k particles on an 8GB consumer GPU, while C++ codes can easily go beyond 1M particles. Ultimately, we aim at reaching 1M particles with our code by serializing the two most memory-intensive operations: neighbor search and core SPH algorithm. Some prior work on reducing the memory footprint of the [JAX-MD](#) neighbor search has already been done as part of LagrangeBench ([lagrangebench/neighbors_search](#)), on which we want to extend in this project.




Figure 1: 2D and 3D lid-driven cavity simulation.

Milestones

- Performance (memory / runtime / scaling) comparison of the neighbor search from (a) JAX-MD, (b) serialized JAX-MD from LagrangeBench, (c) [Matscipy](#), and (d) a naive $\mathcal{O}(N^2)$ implementation.
- Block-serialize the SPH code and identify the memory bottlenecks.
- Block-serialize the neighbor search and SPH code together.

Requirements

- Experience with Python, specifically JAX.
- Ability to work independently.

Contact

Artur Toshev artur.toshev@tum.de

References

Toshev, A., Erbesdobler, J. A., Galletti, G., Ramachandran, H., Brandstetter, J., and Adams, N. A. (2024). JAX-SPH: A differentiable smoothed particle hydrodynamics framework. In *ICLR 2024 Workshop on AI4DifferentialEquations In Science*.