

---

# Poseidon: A Source-to-Source Compiler for Optimizing Ocean Simulation Models on Modern HPC Architectures

Martin Schreiber\*<sup>1</sup>

<sup>1</sup>Université Grenoble Alpes – L’Institut National de Recherche en Informatique et en Automatique (INRIA), Laboratoire Jean Kuntzmann (LJK) – France

## Abstract

Modern high-performance computing (HPC) architectures, characterized by increasing heterogeneity and steep memory hierarchies, present significant challenges for optimizing ocean simulation models. Achieving peak performance on these architectures often requires extensive, costly code rewrites. These rewrites are not only time-consuming and error-prone but also highly architecture-specific, and require numerics experts to be proficient in parallel programming models or domain-specific languages (DSLs). To address these challenges, we introduce Poseidon, an HPC-oriented source-to-source compiler designed for Fortran-based fluid dynamics solvers used in ocean and weather models with regular grid structures. Poseidon employs a novel process called *uplifting*, which treats existing models and their coding standards as Fortran-embedded DSLs and requires minimal source code changes. This approach, which relies on co-design with model developers, allows Poseidon to robustly recover high-level information and semantics that are typically lost during the conversion of numerical algorithms to source code. By doing so, Poseidon can perform safe and holistic optimizations for specific HPC architectures using a data flow graph intermediate representation. It then generates Fortran source code augmented with parallel programming model directives, which can be further optimized by vendor or open source compilers. We detail Poseidon’s methodology and present initial results by performing architecture-specific auto-tuned kernel fusion and automatic parallelization on both CPUs and GPUs using OpenMP on a research code that implements the 2D fast barotropic solver of the CROCO 3D ocean simulation model. Our results demonstrate significant performance improvements and validate the effectiveness of Poseidon’s optimization strategies. Additionally, we discuss our ongoing research efforts for the automatic injection of communications, e.g., MPI, for latency hiding, and the implementation of automatic differentiation at the data flow graph level for data assimilation. These advancements are crucial for further improving the performance and scalability of ocean simulation models. Furthermore, we outline our current progress and future plans for integrating Poseidon with the NEMO ocean model using an elegant annotations-based *uplifter*, and leveraging its optimization techniques.

**Keywords:** HPC, regular grids, peak performance

---

\*Speaker