# In Situ Flow Analysis for MPAS-Ocean Simulations

Mukund Raj (ANL), Hanqi Guo (ANL), Orcun Yildiz(ANL), Phillip Wolfram (LANL), Tom Peterka (ANL)

## Overview

Next generation earth system models that scale to emergent HPC hardware systems will need in situ analysis coupled to climate simulation components. In situ analysis can enable quicker and higher fidelity results than postprocessing.

In this work, we enable in situ analysis for MPAS-Ocean simulation using Decaf, a data flow system for managing dataflow among tasks in a workflow. We also develop methods to improve load balancing in decoupled analysis tasks.

## Decaf: Decoupled dataflow for in-situ workflows

A **workflow** is a directed graph of tasks and communication between them. Graph nodes are tasks and graph links are communication; both of which can be parallel processes.

A **dataflow** is the communication over links in a workflow. While workflows consists of high level tasks, dataflows consists of communication between processes.

The **Decaf** software stack consists of a workflow definition layer, a dataflow definition layer, and a transport layer.

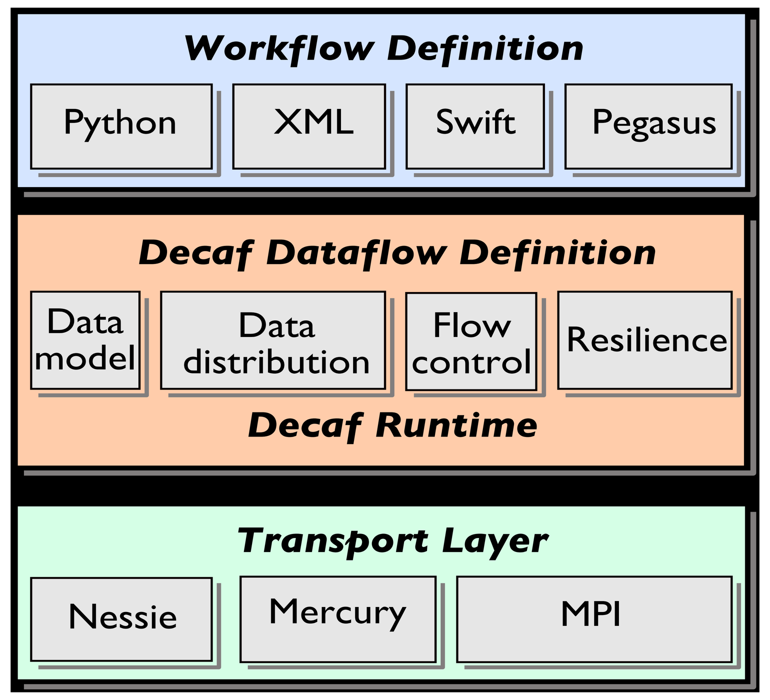


Decaf software organization

## Coupling MPAS-O with stand alone particle tracer

We move the inbuild MPAS-Ocean particle tracer LIGHT to a stand alone tracer using Decaf. This allows for better control over sharing of resources between simulation and analysis (tracer).

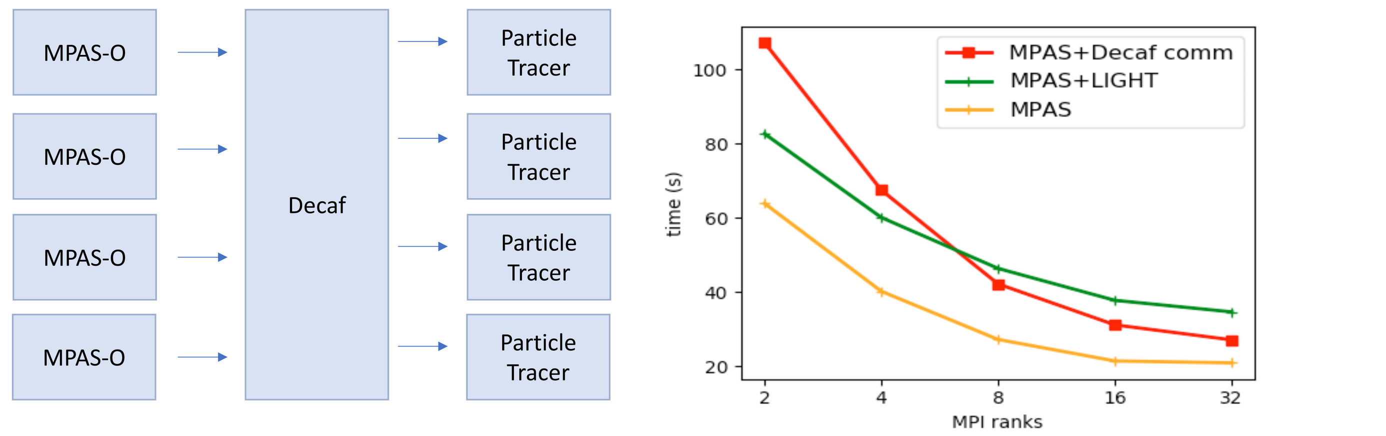


Coupling MPAS-O with particle tracer using Decaf (left) and preliminary timing results for MPAS-Ocean built-in particle tracer (LIGHT), and data movement using Decaf

## Dynamic load balancing for unstructured data using constrained graph partitioning

State-of-the art methods to load balance based on constrained k-d tree do not directly extend to unstructured data. We introduce a constrained graph partitioning based load balancing method that naturally handles data on both structured and unstructured grids.

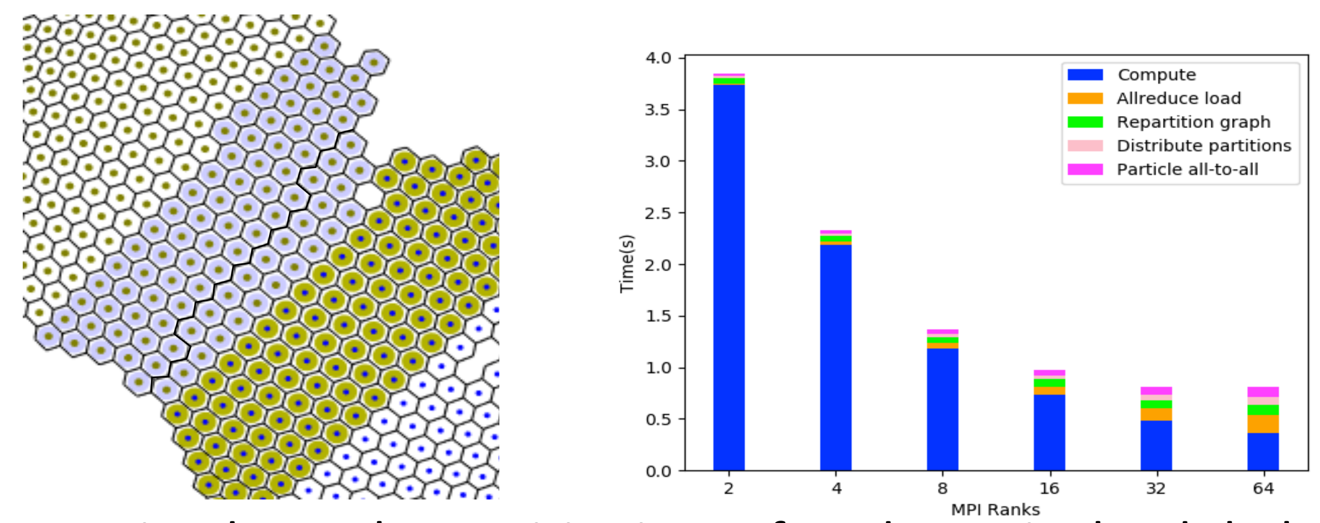


Constrained graph partitioning for dynamic load balancing particle tracing with unstructured data (left) and preliminary timing results for constrained graph partitioning using BFS expansion (right)

## Dynamic load balancing for unstructured data using graph distance based embedding

The use of graph partitioning to balance workload is challenged by computational cost of obtaining balanced partitions. We propose a method to address this issue by adapting the constrained k-d tree method for graph distance based embeddings.

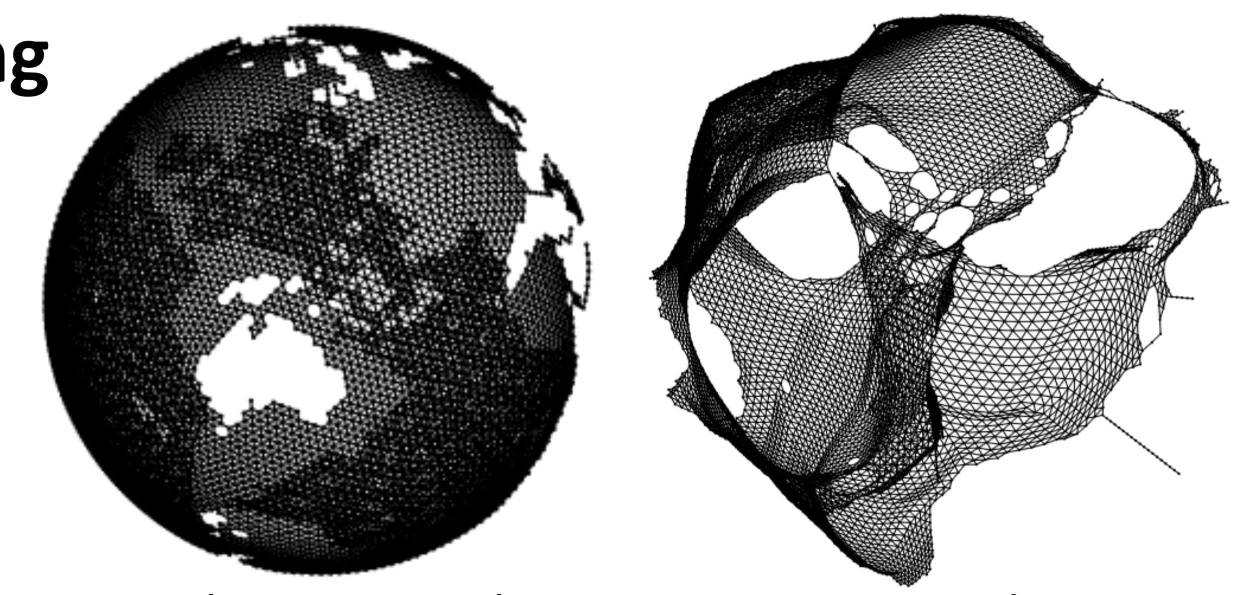

Vertex layout in the MPAS-Ocean grid using intrinsic spatial coordinates (left) and graph distance based embedding via Pivot MDS (right)

## Dynamic load balancing for parallel particle tracing using workload prediction

Existing methods to dynamically load balance parallel particle tracing assumes workload distribution at upcoming epoch to be same as the current workload when partitioning data for upcoming epoch. We propose a method to achieve better load balance by partitioning the data based on *predicted* workload for upcoming epoch, rather than *current* workload distribution.

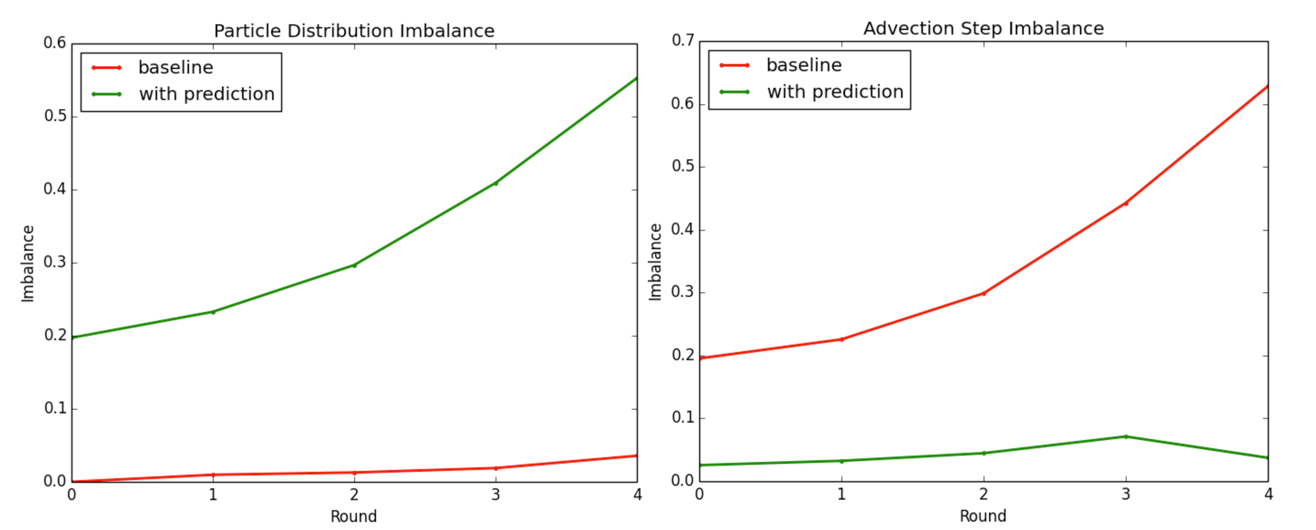


Imbalance in particle distribution (left) and advection workload (right) for synthetic data in prediction and baseline cases