Online Machine Learning for Exascale CFD

Riccardo Balin

Postdoctoral Appointee

ALCF

rbalin@anl.gov

ALCF AI for Science Training Series

10/18/2022

Overview

- Introduction
- Online ML Workflows
- Performance on Polaris
- Beyond Turbulence Closure Models
- Conclusions

Introduction

- There are 4 main modeling approaches to computations of turbulent flow

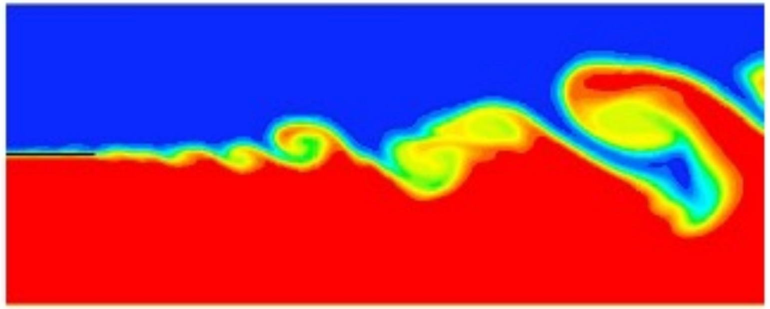

Direct numerical simulation (DNS)
- Solve unsteady Navier-Stokes (NS) equations directly
- Resolve all spatial and temporal turbulent scales, no modeling
- Most accurate
- Most computational expense

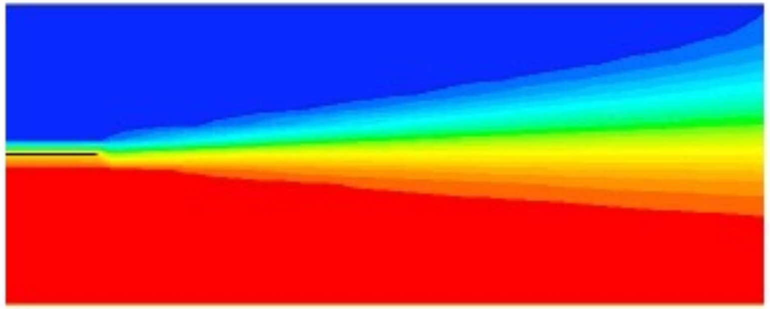

Reynolds-averaged NS (RANS)
- Solve for the steady mean flow directly
- Model all spatial and temporal turbulent scales
- Inaccurate for complex flows
- Least computational expense

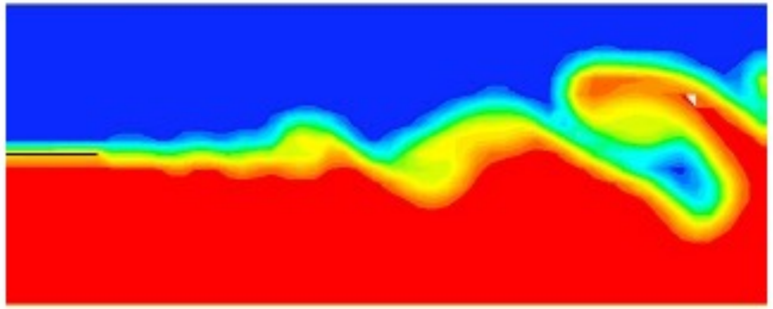

Large eddy simulation (LES)
- Solve unsteady filtered NS equations
- Resolve largest spatial scales and model smallest (sub-grid) scales
- Modest accuracy
- Modest computational expense

Hybrid RANS/LES and Wall-Modeled LES (WMLES)
- Solve unsteady RANS and/or LES equations
- Model all turbulent scales of the near-wall flow

Developing closure models for LES using ML approaches

Introduction

- NN models for LES and WMLES closures must be trained on instantaneous flow data
  - —For petascale and exascale simulations, expensive multi-terabyte databases are needed to store training data
  - —Online (in situ) learning offers attractive solution to avoid I/O and storage bottleneck
  - —Data production and training performed concurrently

DNS of a Turbulent Boundary Layer Over a Bump at $Re = 2\text{M}$

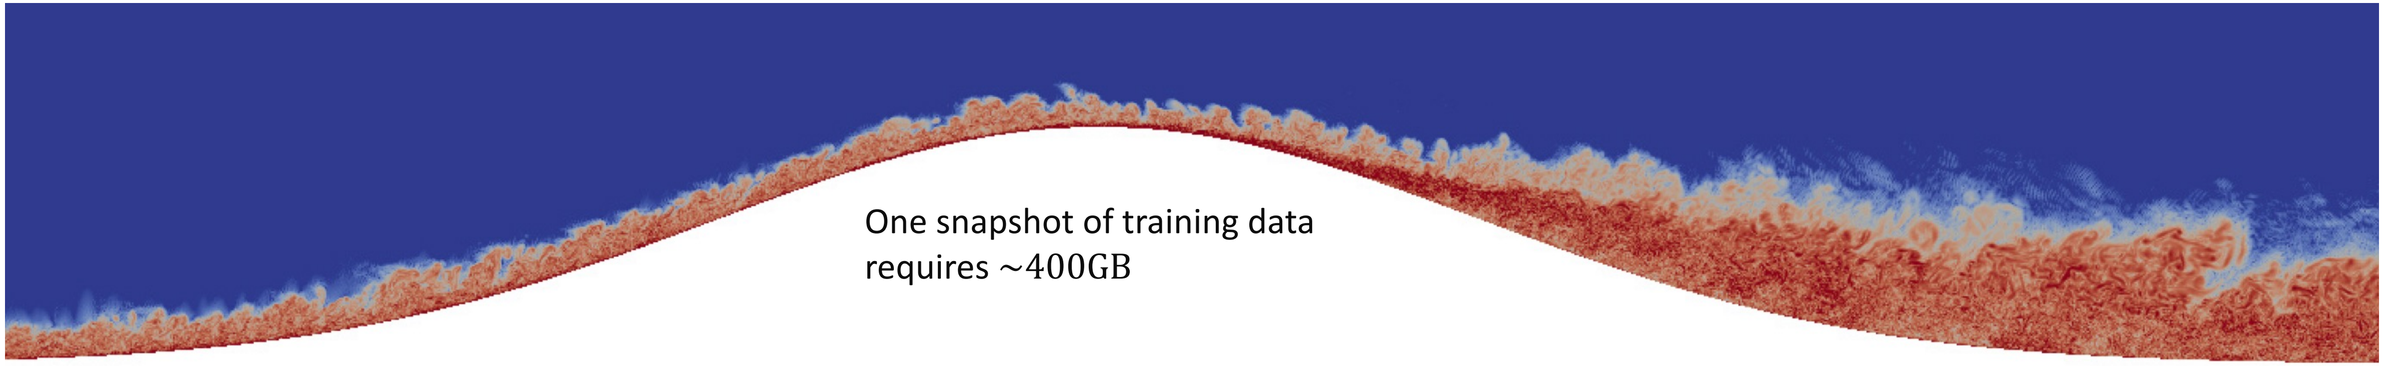

Introduction

- Why do we need to train on the entire flow domain?
  - —A simple bump geometry introduces many physical disturbances that change the turbulence physics
  - —A predictive model must be accurate for all these physical effects
- Why do we need to train on multiple snapshots?
  - —Single snapshot likely does not contain all time-dependent phenomena that model should handle during inference
  - —Some flows of interest are not statistically stationary and evolve over time (e.g., internal combustion engines, active flow control applications, wind turbine under variable wind conditions, etc.)

DNS of a Turbulent Boundary Layer Over a Bump

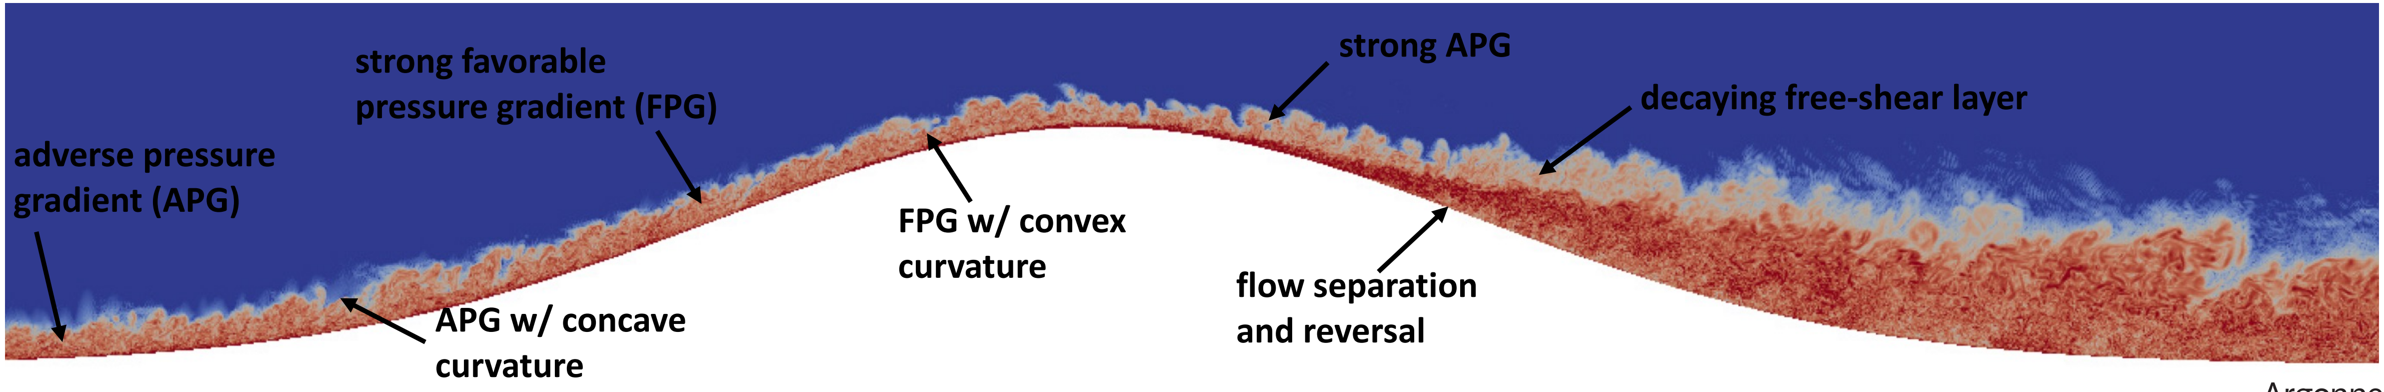

Introduction

- Flows over bumps are insightful, but what are we really after?
  - — Full CFD simulations of complex aeronautical and aerospace systems
  - — NASA vision 2030: towards aircraft certification by simulation

Hybrid RANS/LES of Flow Around Aircraft Vertical Tail with Active Flow Control

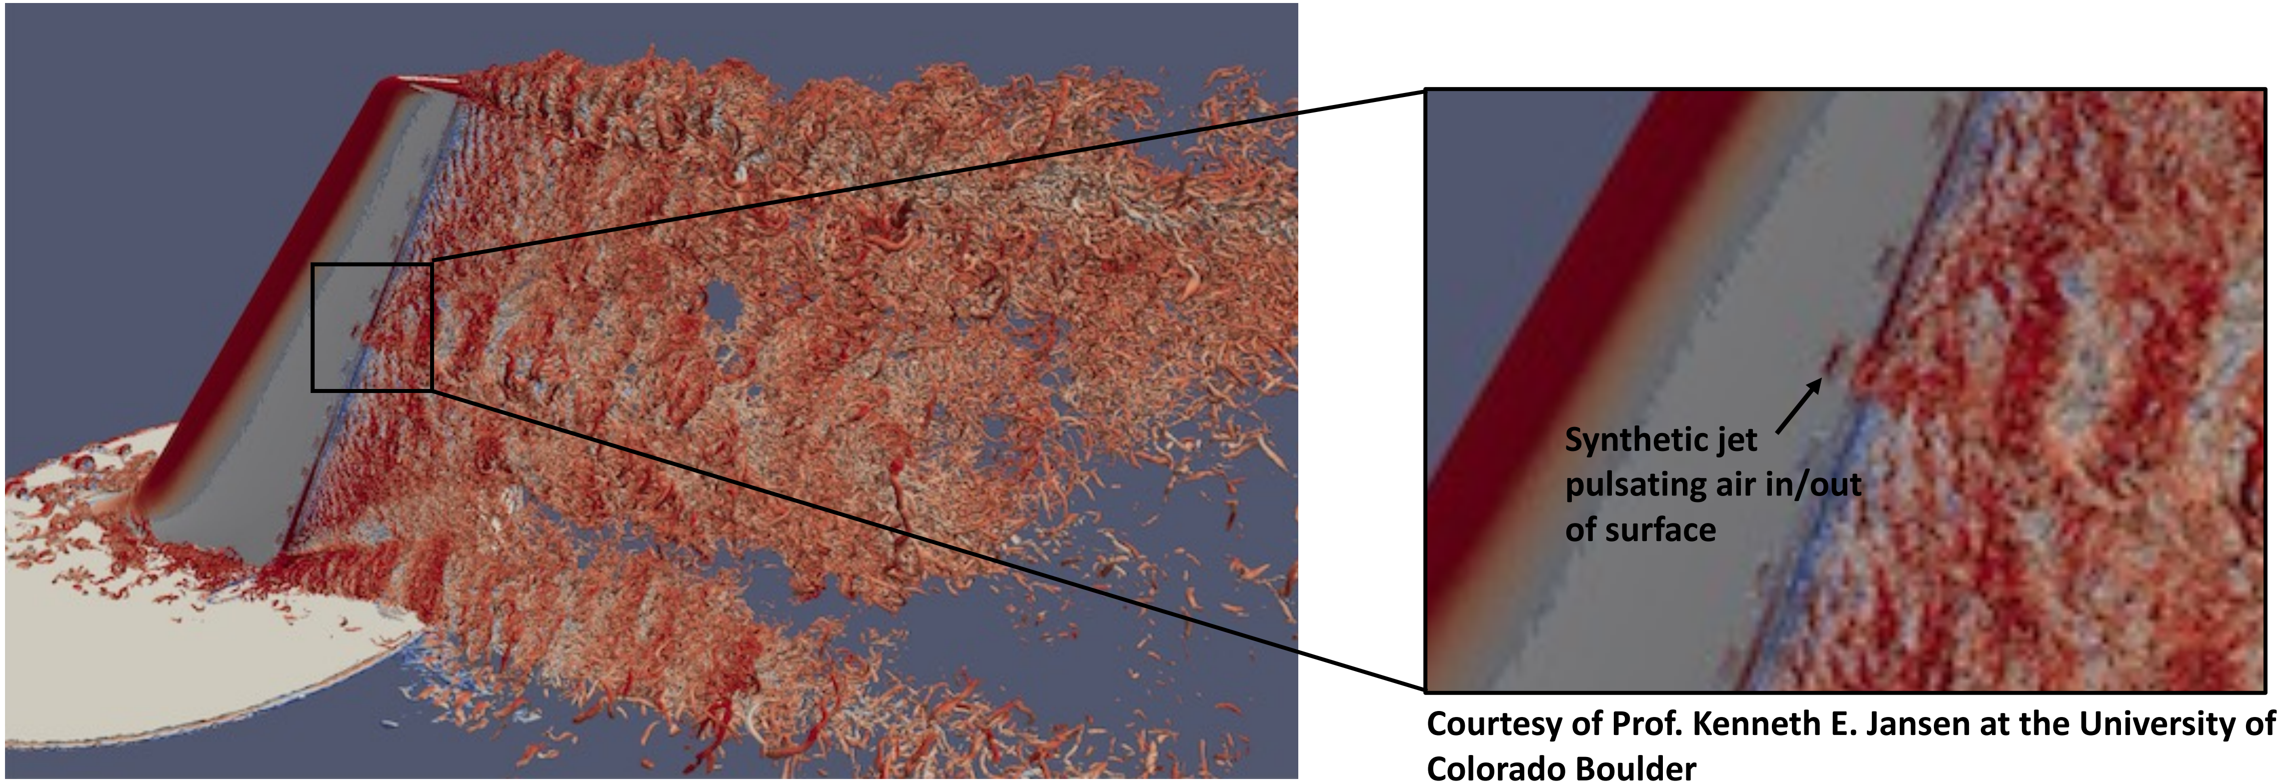

Courtesy of Prof. Kenneth E. Jansen at the University of Colorado Boulder

Online Machine Learning Workflows

- Four components to the workflow: simulation, ML, in-memory database and data transporter

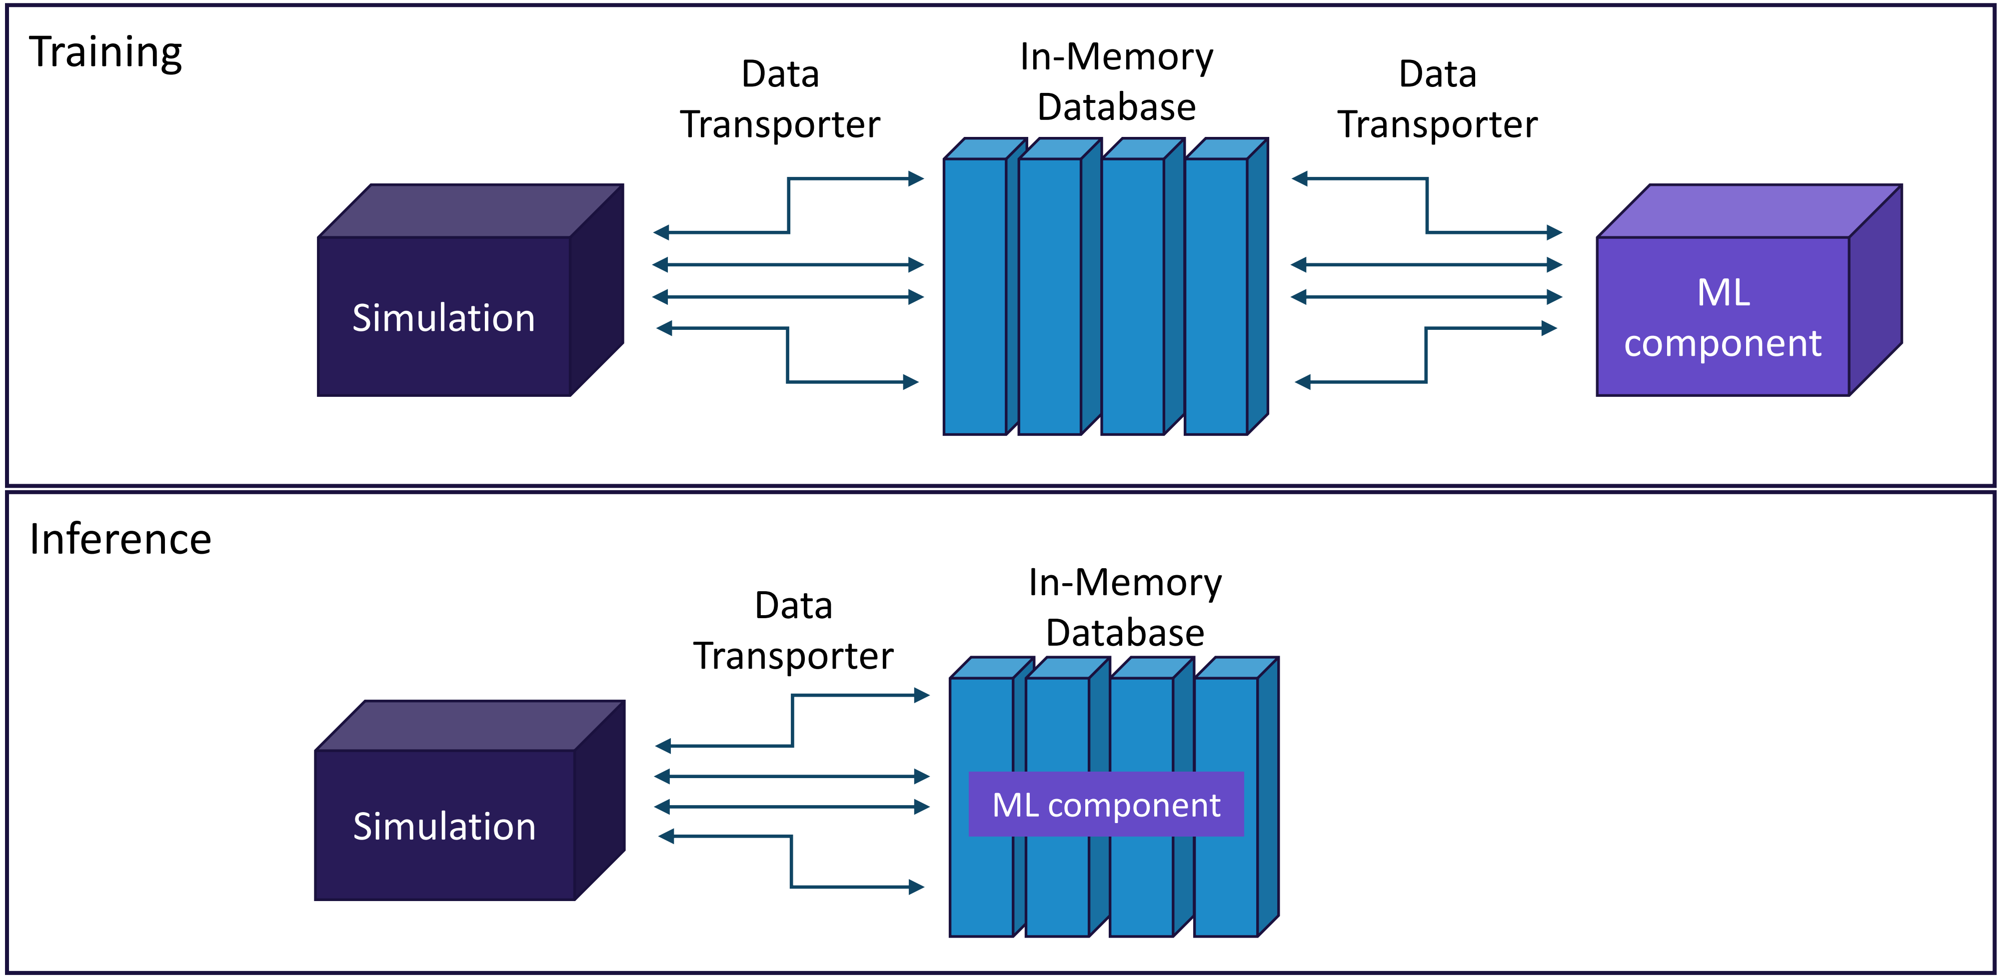

Online Machine Learning Workflows

- We leverage SmartSim and SmartRedis for the database and data transporter
- SmartSim and SmartRedis are open source tools (https://github.com/CrayLabs/SmartSim)

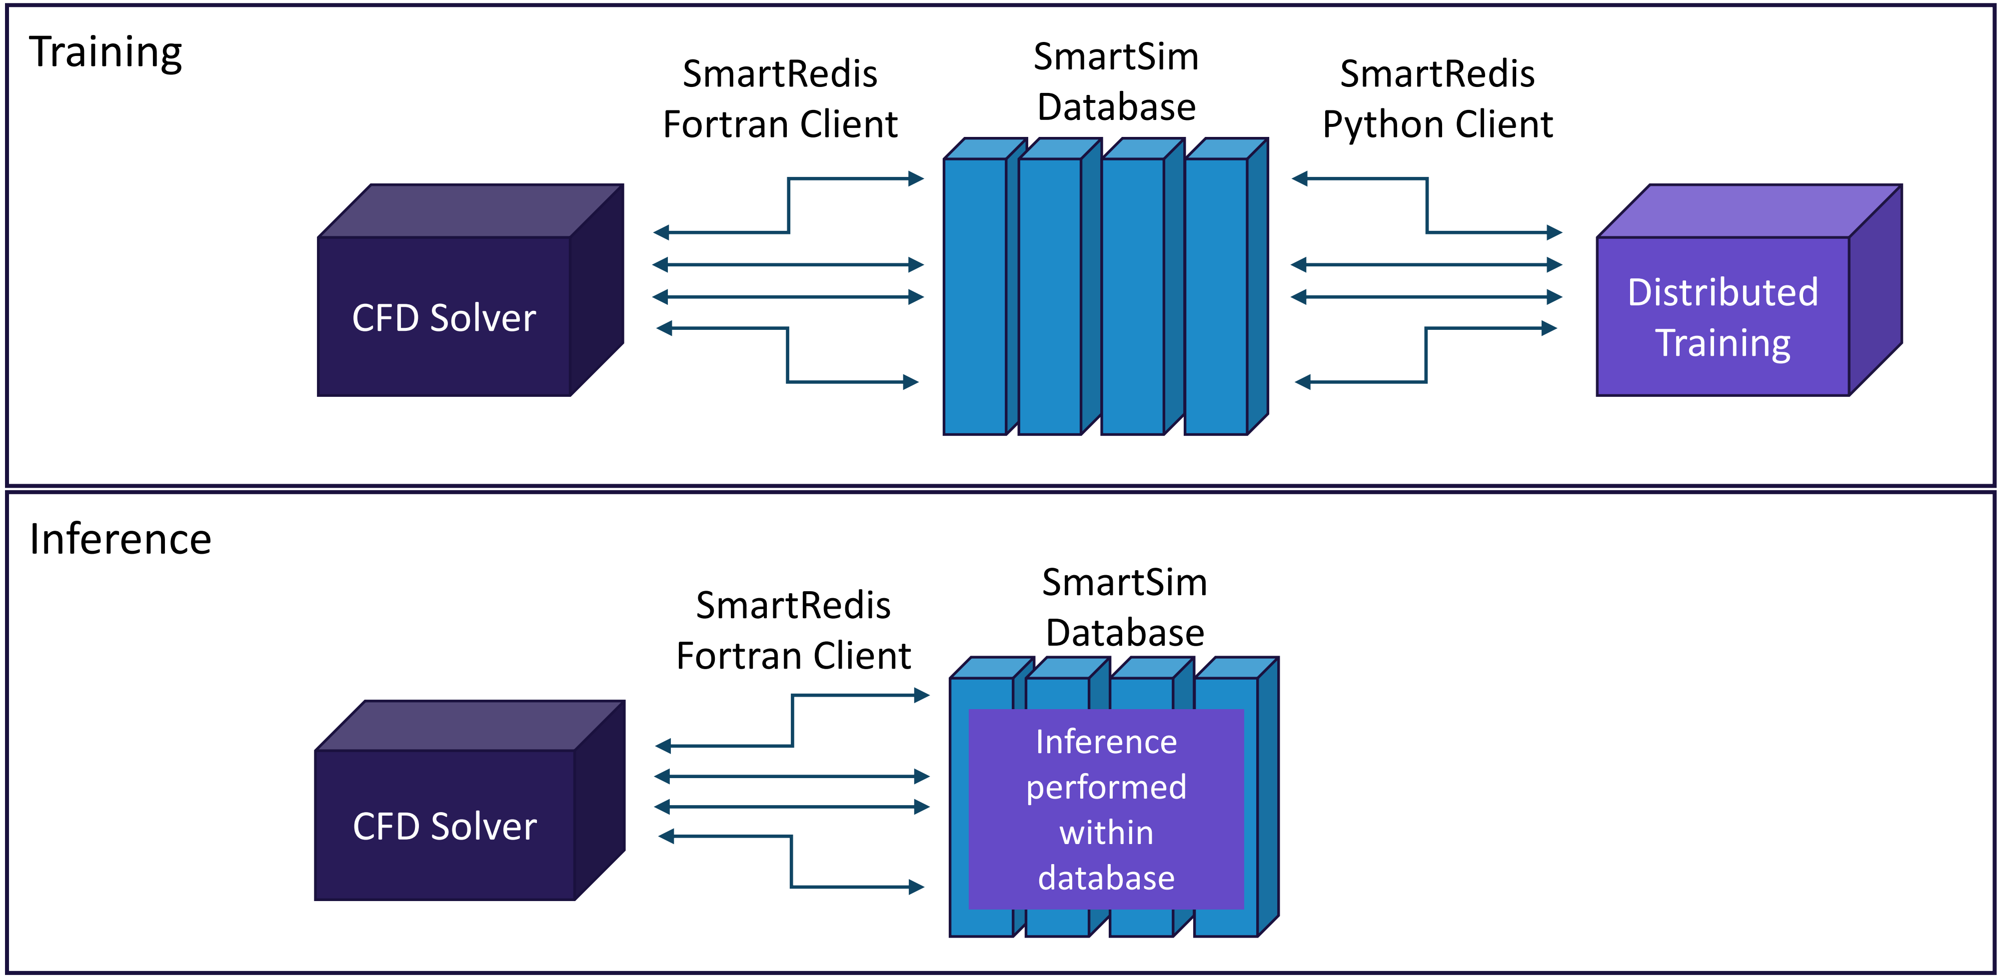

Scaling Performance on Polaris

- Average data transfer overhead during online inference
- Clustered and co-located refer to separate implementations of the workflow

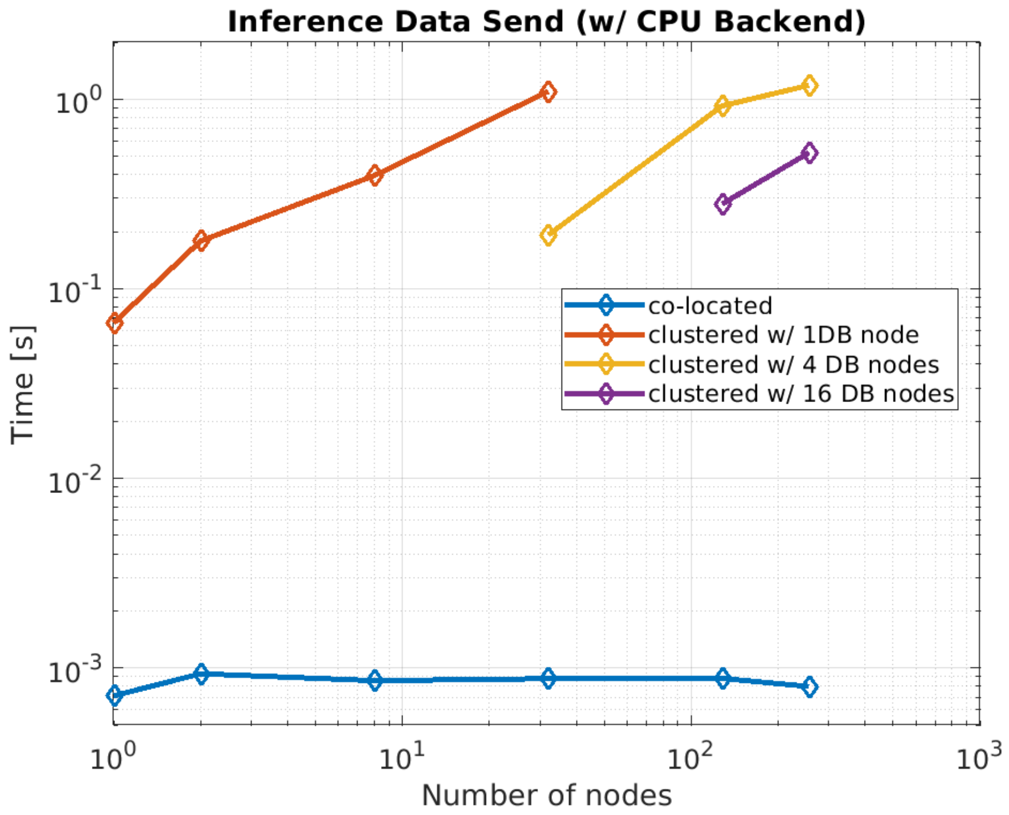

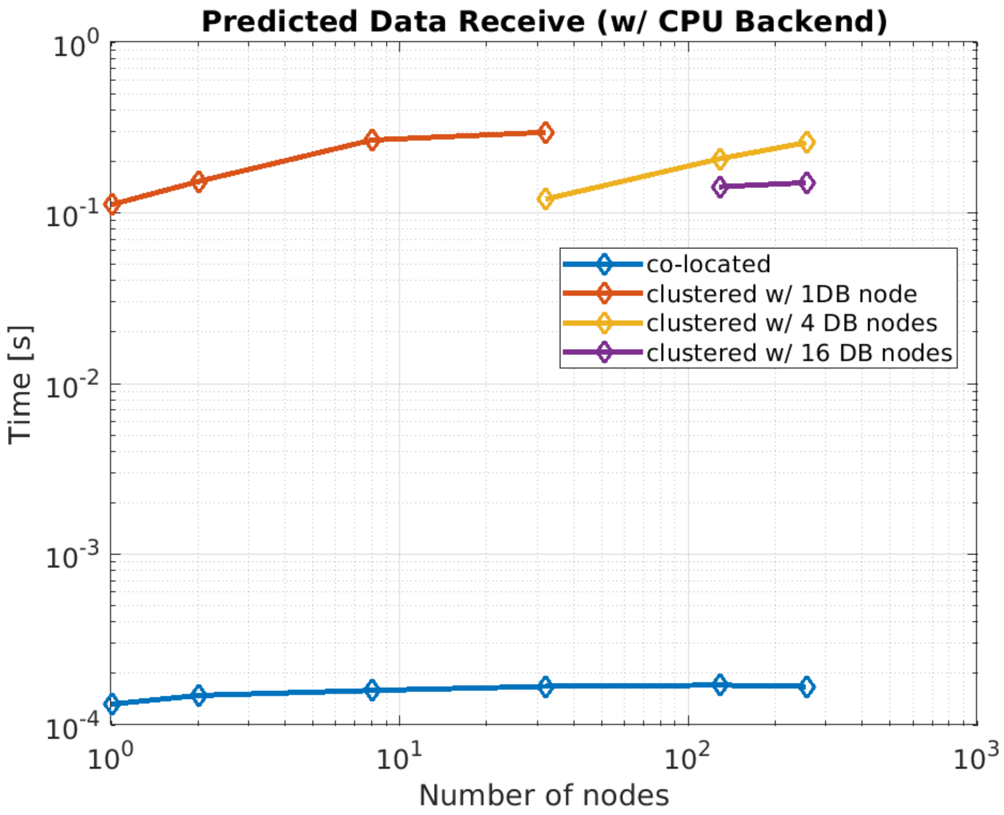

Beyond Turbulence Closure Modeling

- Train multiple models simultaneously on data from multiple simulations
- Perform hyperparameter sweeps for a single model

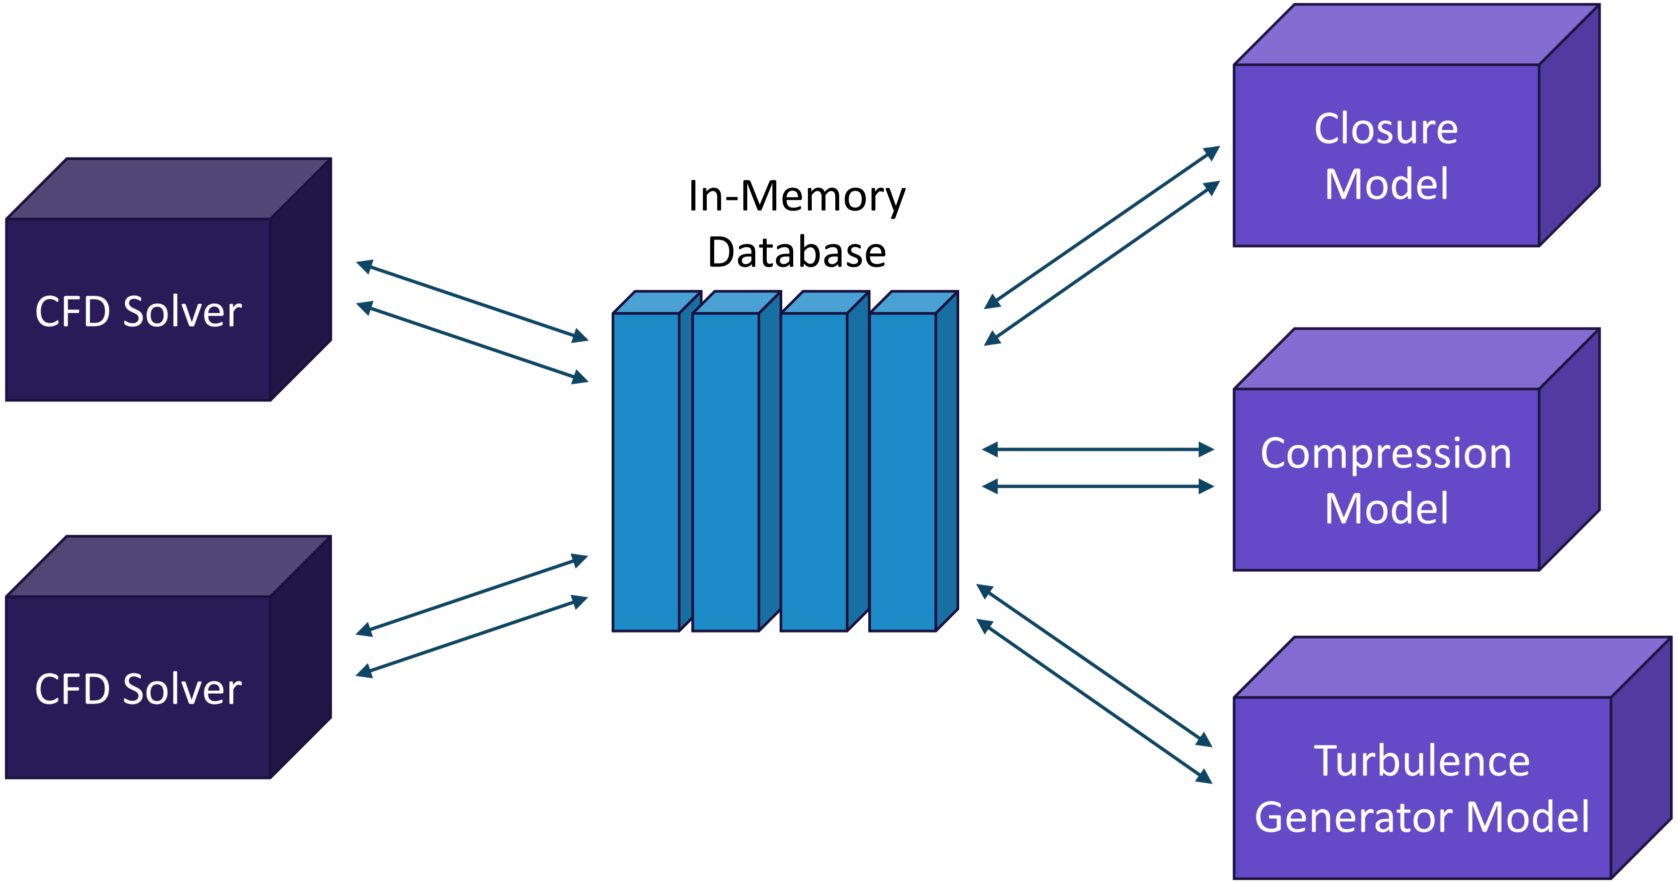

Beyond Turbulence Closure Modeling

- Adding user interactivity to online workflow
  - User-dedicated node to run scripts and notebooks interacting with database
  - Visualize training curves, validate model checkpoints, change hyperparameters on the fly
  - Adding online visualization with Paraview

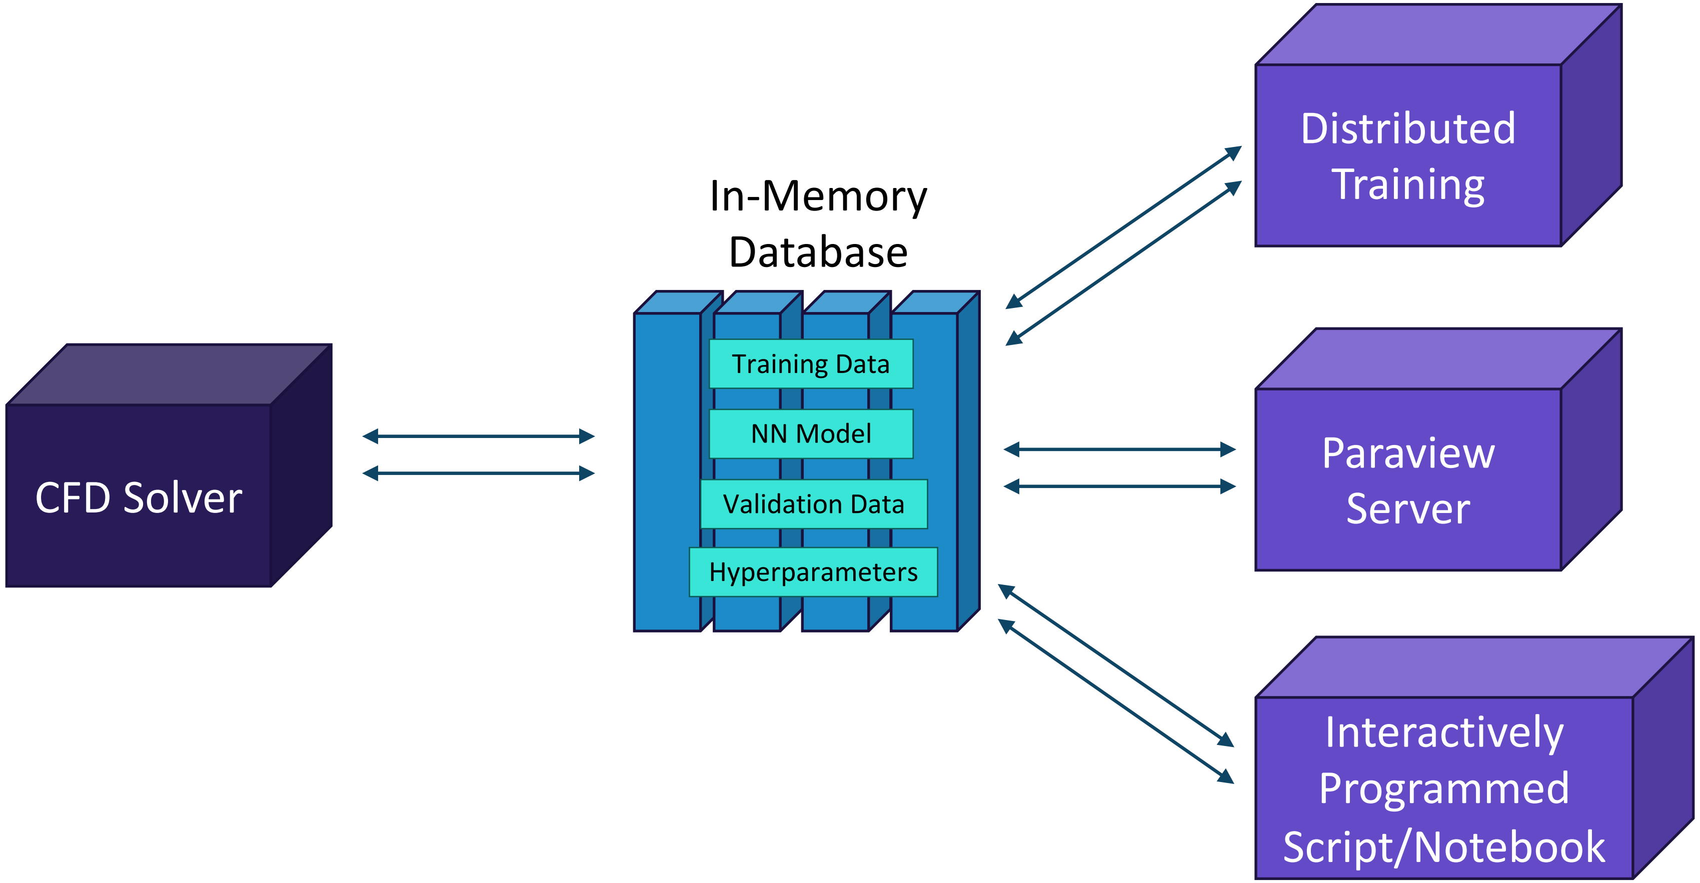

Conclusions

- Developed a software infrastructure to combine CFD with online (in situ) ML at scale
- Capable of performing online training and inference of NN models efficiently
- Currently developing turbulence closure models for aerospace flows, but easily extendable to other applications

Questions?