SAND2022-5239C

NREL NATIONAL RENEWABLE ENERGY LABORATORY | Sandia National Laboratories | Oak Ridge National Laboratory | TEXAS The University of Texas at Austin

ExaWind: Exascale Predictive Wind Plant Flow Physics Modeling

NREL: M.A. Sprague (PI), M. Brazell, N. Brunhart-Lupo, M. Henry de Frahan, P. Mullowney, J. Rood, A. Sharma, S. Thomas, G. Vijayakumar
SNL: P. Crozier, L. Berger-Vergiat, L. Cheung, N. de Velder, J. Hu, R. Knaus, D.H. Lee, N. Matula, J. Overfelt, P. Sakievich, T. Smith, A. Williams, I. Yamazaki
ORNL: J. Turner, A. Prokopenko, R. Wilson
UTA: R. Moser, J. Melvin

EXASCALE COMPUTING PROJECT

Project Overview

Objective: Create a predictive physics-based simulation capability that will provide a validated "ground truth" foundation for wind plant siting, operational controls, and reliably integrating wind energy into the grid

Motivation: Validated, predictive wind plant simulations will reduce the cost of energy by providing
- a path to better understanding of wind plant flow physics, which will lead to
  - new plant layout design in complex terrain
  - new turbine technologies to optimize plant performance
- a foundation for improved computer-aided engineering models, which will enable better design optimization

Primary Application Codes:
- **Nalu-Wind**
  - https://github.com/exawind/nalu-wind
  - Unstructured-grid computational fluid dynamics (CFD) code
  - C/C++
  - Built on Trilinos/STK/hypre/TIOGA
- **AMR-Wind**
  - https://github.com/exawind/amr-wind
  - Structured-grid adaptive-mesh-refinement (AMR) computational fluid dynamics (CFD) code
  - C/C++
  - Built on the AMReX library
- **OpenFAST**
  - https://github.com/openfast/openfast
  - Whole-turbine simulation code; blades, control system, tower, etc.
  - Fortran 90; dedicated Intel Parallel Computing Center (IPCC) for parallelization

ExaWind Software/Library Partnerships

- ***Trilinos***, https://trilinos.org/
  - ***MueLu:*** provides aggregation-based multigrid preconditioners
  - ***Ifpack2:*** provides SOR-based, polynomial and incomplete factorization preconditioners
  - ***Kokkos-Kernels:*** provides shared memory algorithms: graph-coloring, SpMV, SPMM, iterative and incomplete factorization preconditioners
  - ***Tpetra:*** provides distributed memory, sparse linear algebra objects
  - ***Belos:*** provides templated Krylov and recycling solvers
  - ***Amesos2:*** provides sparse direct solvers
  - ***Sierra Toolkit (STK):*** provides an unstructured-mesh in-memory, parallel-distributed database as well as I/O, load-balancing, proximity search, transfers, etc.
- ***hypre***, https://github.com/LLNL/hypre
  - Multigrid solvers and preconditioners based on classic Ruge-Stüben AMG algorithm
- ***Kokkos***, https://github.com/kokkos
  - Programming model in C++ for writing performance portable applications targeting all major HPC platforms
- ***TIOGA***, https://github.com/jsitaraman/tioga
  - Library for overset-grid assembly on parallel distributed systems
- **ALExa / ArborX**, https://github.com/arborx/ArborX
  - Performance portable algorithms for geometric search
- ***VTK-m***, https://gitlab.kitware.com/vtk/vtk-m
  - New *in situ* visualization and analysis capabilities
- **Spack**, https://github.com/spack/spack
  - Package manager for exascale software
- **AMReX**, https://github.com/AMReX-Codes/amrex
  - Software Framework for Block Structured AMR

ECP Key Performance Parameter (KPP-2)

Challenge Problem:
Predictive simulation of a wind farm with tens of megawatt-scale wind turbines dispersed over an area of 50 square kilometers

Minimum Requirements:
- 2x2 array of megawatt-scale turbines operating at rated speed
- 3 km x 3 km domain with height of 1 km
- Hybrid-RANS/LES model
- At least 30-billion grid points

2021-2022 Highlights

ExaWind hybrid solver approach established in FY21

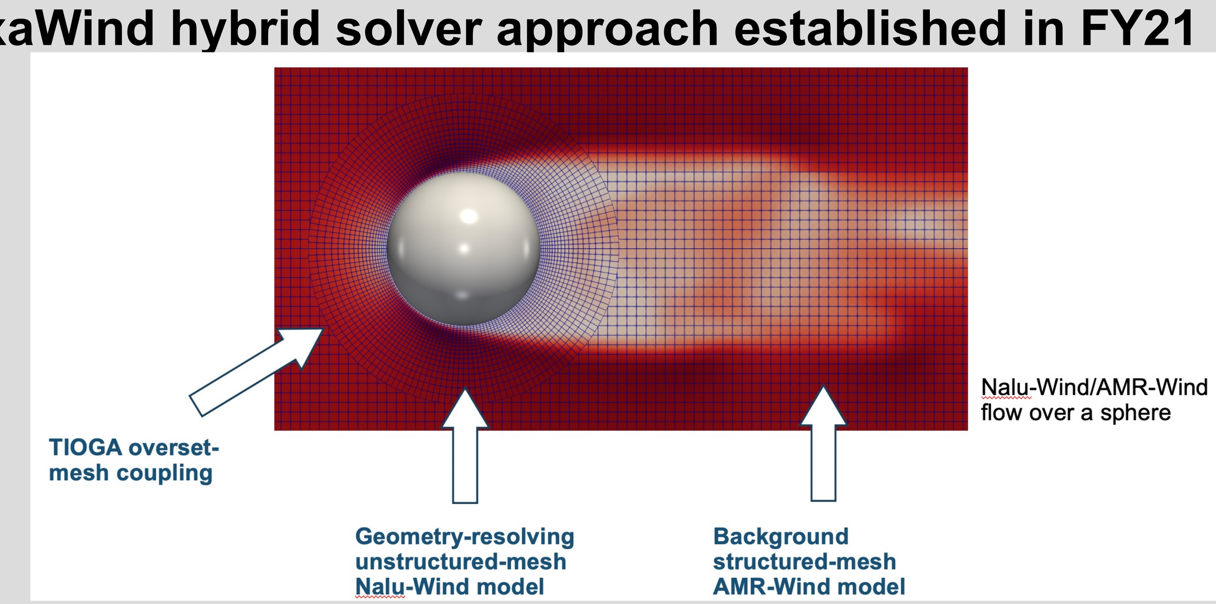

Flow over a sphere proof-of-concept simulation.

ExaWind: Hybrid CFD on hybrid HPC

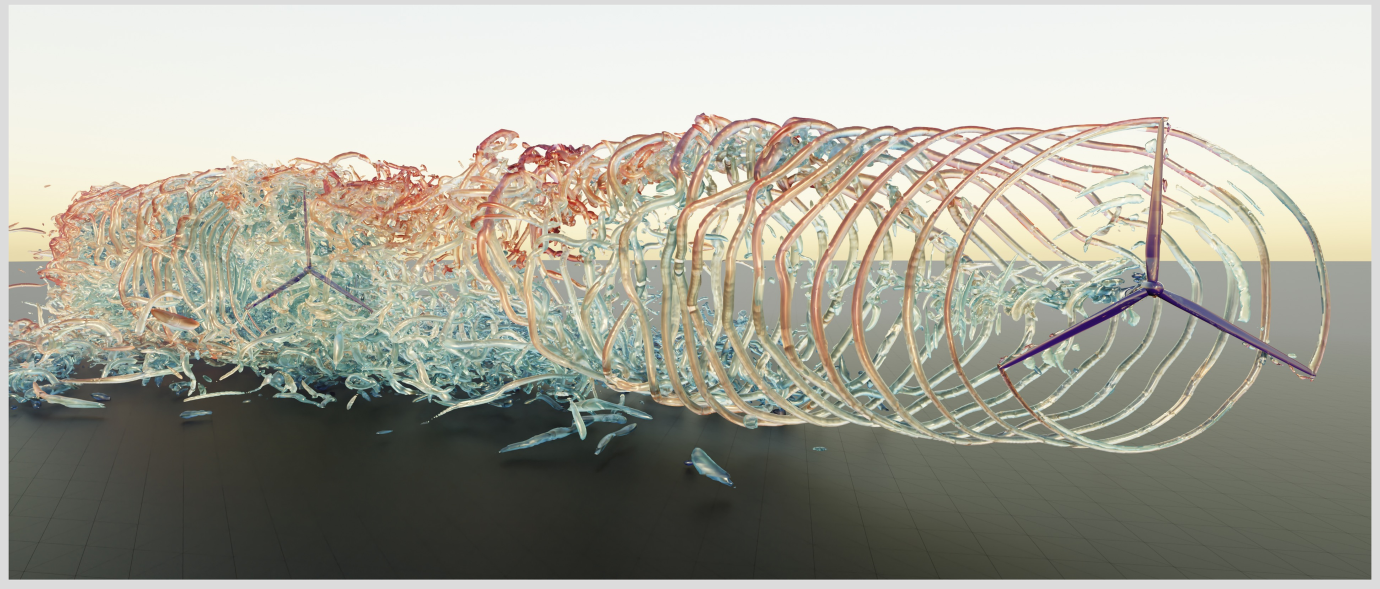

Representative two-turbine blade-resolved simulation with the ExaWind hybrid solver on NREL's Eagle computer. The near-body flow is resolved with an unstructured-grid Nalu-Wind model and the background flow is resolved with a structured-grid AMR-Wind model. The models are coupled through overset meshes handled by the Topology Independent Overset Grid Assembler (TIOGA).

The ExaWind team successfully demonstrated execution of the hybrid-AMR-Wind/Nalu-Wind solver on 32 Crusher nodes with a single-wind-turbine simulation. AMR-Wind ran on 256 Crusher GPU Tiles while Nalu-Wind ran on 256 CPU cores, which was made possible through the generalized MPI coupling between the codes. This approach allows exploration of using the full software stack on this early system while work continues to enable Nalu-Wind on AMD GPUs.

Crusher Update

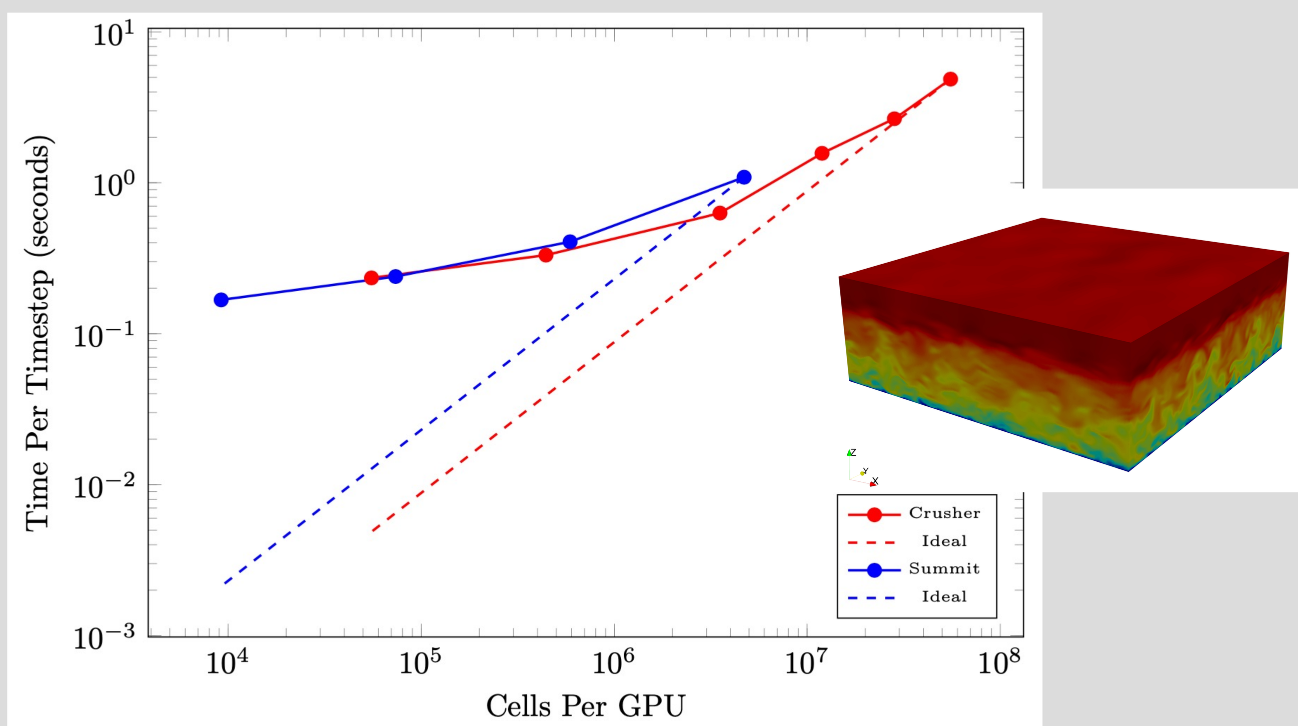

Scaling plot comparing AMR-Wind performance for atmospheric boundary layer simulations on Summit and Crusher GPUs. Simulations of different model sizes were run on a fixed resource of 16 Crusher nodes (128 GCDs) or 16 Summit Nodes (96 GPUs). The largest models were 1920^3 and 768^3 on Crusher and Summit, respectively.

- Crusher scales well down to about 2M cells per GCD.
- Crusher shows similar time per timestep for these resources
- One can fit a model on a Crusher node that is five-times larger that what can fit on a Summit Node.

ExaWind software at core of collaboration with GE

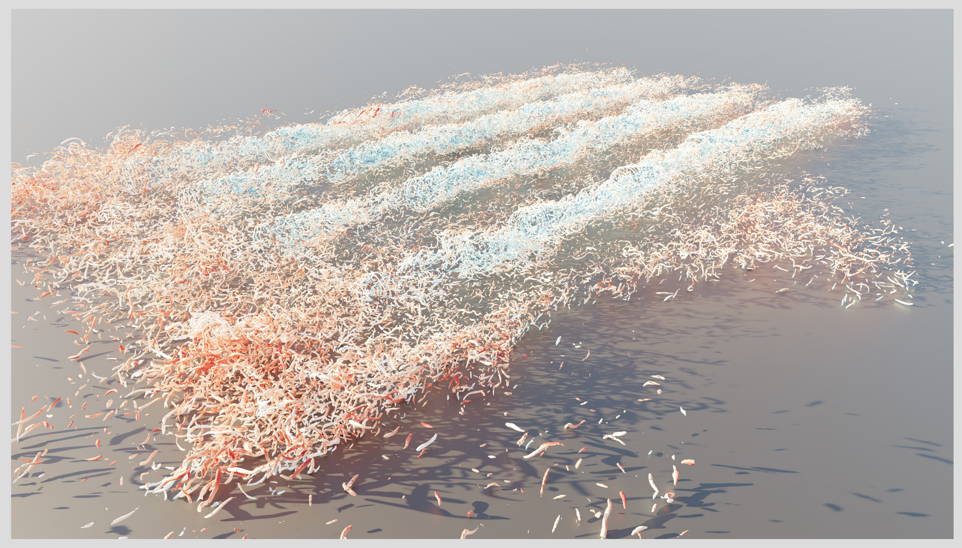

AMR-Wind Summit simulation results for a 20-turbine wind farm off the coast of NY, where wind turbines were represented as "actuator lines". Shown are iso-surfaces of q-criterion colored by velocity magnitude. (Figure: S. Yellapantula)

- NOWRDC project on Impact of Low-Level-Jets (LLJs) on wind farm performance being jointly run by GE Research (prime) and NREL (sub) has been using AMR-Wind to simulate LLJs off the coast of New York.
- As a part of this project, the project team implemented mesoscale-microscale coupling strategy in AMR-Wind and successfully simulated LLJs along the NY coastline and validated against LIDAR data.
- The simulation was one of our largest production runs using AMR-Wind on Summit; the model had 1.3 billion grid points, used 552 GPUs, and was enabled through an ALCC allocation.

AMR-Wind speed up (on Summit GPUs)

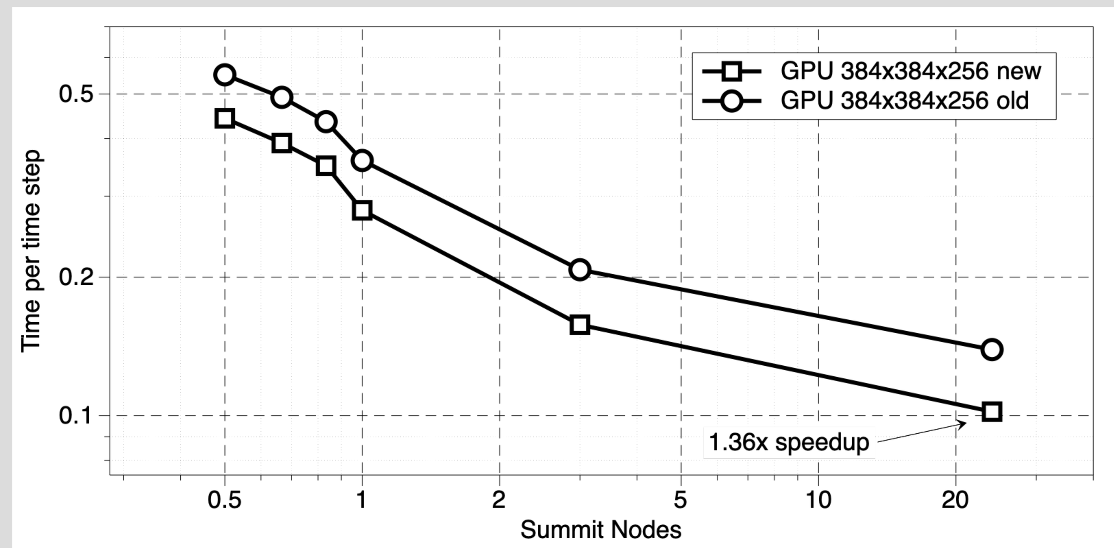

Strong scaling results on Summit GPUs for an atmospheric boundary layer simulation; "old" is from AMR-Wind in April 2021, whereas "new" is AMR-Wind in April 2022.

Through, in part, a collaboration with the ECP CEED and AMReX teams, the ExaWind team has significantly reduced the time per timestep in AMR-Wind by using the switching to the BiCG solver for momentum, solving segregated momentum systems, and using the latest version of AMReX.

ExaWind is going offshore

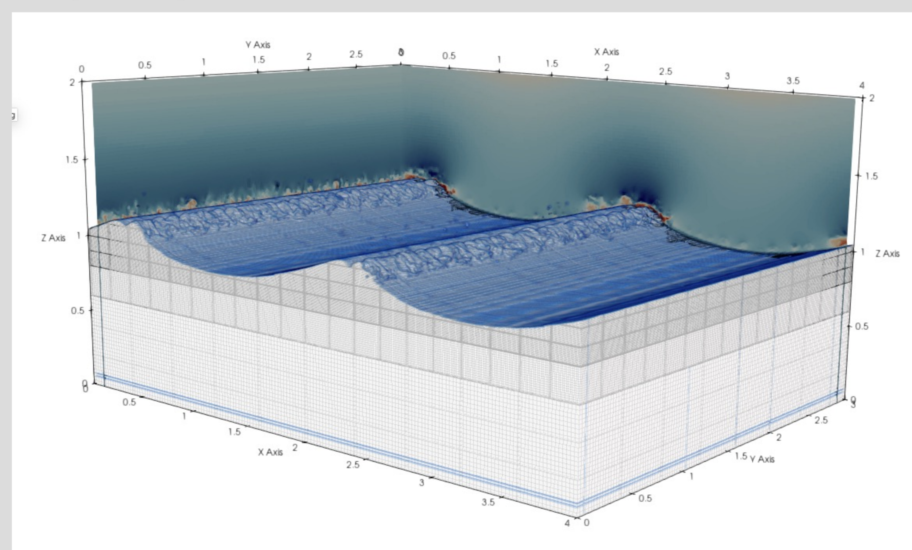

Multi-phase flow simulation of waves in AMR-Wind (Figure: Georgios Deskos)

Through funding for the DOE EERE Wind Energy Technologies Office, the team is equipping ExaWind for multi-phase fluid dynamics which is the first step to simulating floating offshore wind turbines.

This work was authored in part by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.This report followed the Sandia National Laboratories formal review and approval process (SAND2022-TBD). As such, the technical report is suitable for unlimited release.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of two U.S. Department of Energy organizations (Office of Science and the National Nuclear Security Administration) responsible for the planning and preparation of a capable exascale ecosystem, including software, applications, hardware, advanced system engineering, and early testbed platforms, in support of the nation's exascale computing imperative.

Exascale Computing Project Annual Meeting
Virtual
May 2-6, 2022