Wind Energy Science Leadership Series

Discussions on the leading edge of wind energy science.

NREL
Transforming ENERGY

The Future of High-Performance Computing for Wind Energy

Thursday, July 30
9:00 to 10:15 a.m. MT

The future of high-performance computing for wind energy

Michael A Sprague[1], michael.a.sprague@nrel.gov
Shreyas Ananthan[2], shreyas.ananthan@nrel.gov
Ganesh Vijayakumar[1], ganesh.vijayakumar@nrel.gov
Luis 'Tony' Martínez Tossas[1], luis.martinez@nrel.gov

[1]National Wind Technology Center
[2]Computational Science Center

Overview: Speakers and topics

- **Mike Sprague**, PI for the DOE ExaWind and High-Fidelity Modeling Projects
    - Motivation for predictive high-fidelity modeling for wind energy
    - Overview of our team efforts and open-source software stack
- **Shreyas Ananthan,** Senior Research Scientist, ExaWind Chief Software Architect
    - *Next-generation high-performance computing*
- **Ganesh Vijayakumar,** Research Engineer
    - *Model validation: Building confidence in our simulated reality*
- **Luis 'Tony' Martínez Tossas,** Research Engineer
    - *HPC and HFM as the foundation for next-generation engineering models*

Motivation for next-generation high-fidelity models (which require HPC)

- More wind energy at low cost is a good thing
- High penetration of wind energy requires large wind farms composed of megawatt-scale turbines
  - *Both land-based and offshore*
- Wind farm flow dynamics and coupled turbine structural dynamics are extremely complex and models are lacking
  - *Relevant dynamics span many orders of magnitude*
- Only when we can **model well** the wind system can we **optimize** that system
  - *Maximize energy extraction*
  - *Maximize turbine life, minimize downtime*

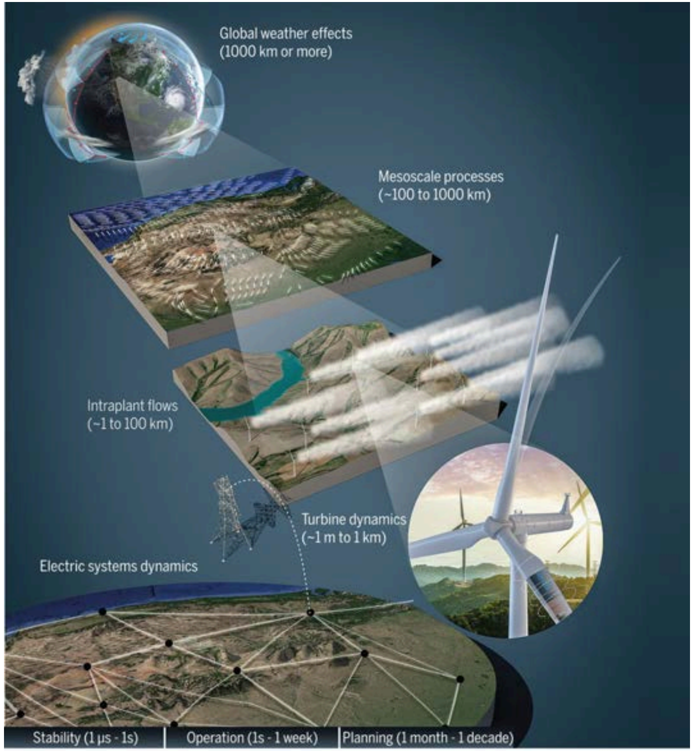

Grand challenges in the science of wind energy

https://science.sciencemag.org/content/366/6464/eaau2027

HFM & HPC can illuminate the path to reducing the cost of wind energy

Photo by Gitte Nyhus Lundorff, Bel Air Aviation Denmark – Helicopter Services

Can we <u>predict</u> and <u>understand:</u>

- Impact of wakes on downstream turbines?
- Evolution of the wakes?
- Formation of the wakes?

… and all in a highly complex, dynamic metocean environment

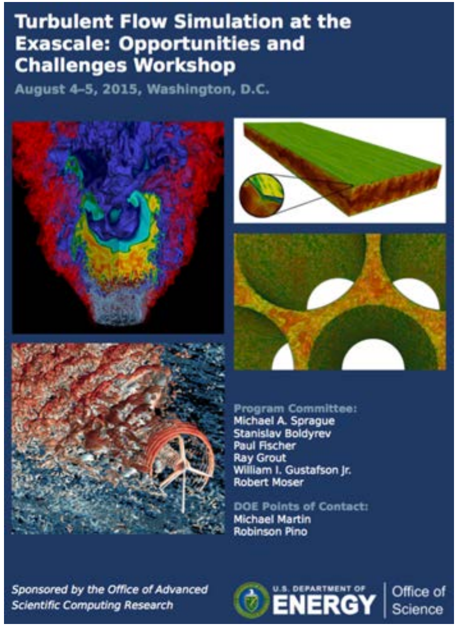

Turbulent Flow Simulation at the Exascale: Opportunities and Challenges Workshop

August 4-5, 2015, Washington, D.C.

Program Committee:
Michael A. Sprague
Stanislav Boldyrev
Paul Fischer
Ray Grout
William I. Gustafson Jr.
Robert Moser

DOE Points of Contact:
Michael Martin
Robinson Pino

Sponsored by the Office of Advanced Scientific Computing Research

U.S. DEPARTMENT OF ENERGY | Office of Science

http://www.nrel.gov/docs/fy17osti/67648.pdf

DOE Advanced Scientific Computing Research (ASCR) workshop highlighted predictive wind farm simulation as a grand-challenge requiring next-gen **exascale-class** supercomputers

Primer: What is high-fidelity modeling (HFM)?

- ***Mathematical model: a description of a system using mathematical concepts***
  - ***i.e., an equation or a bunch of equations***
- For HFM of wind farms, we strive to adhere to **first-principles** to increase **predictivity**
  - Conservation of mass
  - Newton’s second law of motion
- The accepted model for many fluid motions is embodied in the **Navier-Stokes equations**, e.g.,

$$\rho\left(\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u}\cdot\nabla\mathbf{u}\right) = -\nabla p + \mu\nabla^2\mathbf{u} + \mathbf{g}\alpha\,(T \quad T_0)$$

$$\nabla\cdot\mathbf{u} = 0$$

$$\frac{\partial T}{\partial t} + \mathbf{u}\cdot\nabla T = \kappa\nabla^2 T$$

$\mathbf{u}(\mathbf{x}, t)$: velocity
$p(\mathbf{x}, t)$: pressure
$T(\mathbf{x}, t)$: temperature
$\mathbf{g}$: gravity
ρ: density
μ: viscosity
α: coefficient of thermal expansion
κ: coefficient of thermal diffusivity

Set of non-linear, partial-differential equations governing fluid velocity, pressure, and temperature

Primer: Can we solve Navier-Stokes for wind energy systems?

- Analytical solutions only exist for the **most simple laminar** problems
- **Turbulence** brings in orders more complexity (many scales to be captured)

“Big whirls have little whirls that feed on their velocity, and little whirls have lesser whirls and so on to viscosity.” -- Lewis Richardson

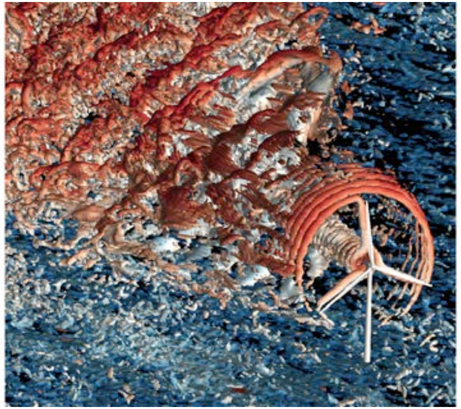

The wind energy system is extremely **NOT SIMPLE**, and **HIGHLY TURBULENT**

The only way to solve the Navier-Stokes equations for wind energy flows is:

- **Introduce some approximations**
- **“Discretize” the equations (millions to billions of equations)**
- **Solve on high-performance computing (HPC, e.g., supercomputers)**

This can be very difficult.

Development efforts are focused on the open-source ExaWind software stack

ExaWind: An open-source **multi-fidelity** modeling & simulation software stack designed to run on **laptops** and **next-gen supercomputers**

Nalu-Wind

- https://github.com/exawind/nalu-wind
- Incompressible-flow computational fluid dynamics (CFD) code
- Unstructured-grid finite-volume discretization
- Closely tied to Trilinos & *hypre* libraries
- Blade-resolved and actuator-type simulations

AMR-Wind

- https://github.com/exawind/amr-wind
- Incompressible-flow CFD code
- Structured-grid finite-volume atmospheric-boundary-layer “background solver”
- Built on AMReX libraries
- Coupled to Nalu-Wind through overset meshes

OpenFAST

- https://github.com/openfast
- Whole-turbine simulation code

Software described in NAWEA 2019 paper

https://iopscience.iop.org/article/10.1088/1742-6596/1452/1/012071/pdf

ExaWind modeling approach defined in 2015 meeting of experts

DOE Strategic Planning meetings established the modeling and simulation requirements for **predictive** wind farm simulations

- Compressible- **or** incompressible-flow model
- Geometry-resolving meshes
- Fluid-structure interaction
- Hybrid RANS/LES modeling
- Nonlinear structural dynamics

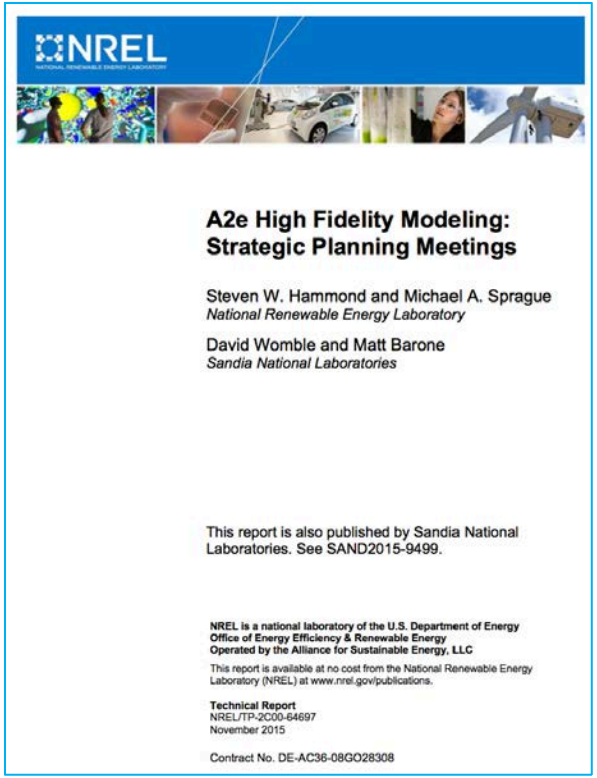

A2e High Fidelity Modeling: Strategic Planning Meetings

Steven W. Hammond and Michael A. Sprague
National Renewable Energy Laboratory

David Womble and Matt Barone
Sandia National Laboratories

This report is also published by Sandia National Laboratories. See SAND2015-9499.

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Technical Report
NREL/TP-2C00-64697
November 2015

Contract No. DE-AC36-08GO28308

https://www.nrel.gov/docs/fy16osti/64697.pdf

HPC performance portability is central to ExaWind development

Supercomputer architecture is evolving rapidly

- U.S. Department of Energy (DOE) supercomputers are increasingly relying on Graphical Processing Units (GPUs) for computational speed at low power
- The the first **exascale** supercomputers will all have hybrid CPU-GPU architectures
  - Coming online in 2021-2022
  - Aiming for power requirements below 30 MW
- Hybrid CPU-GPU architecture is expected to become more common amongst all clusters
- **CFD codes will need to be able to utilize GPUs!**

https://www.flickr.com/photos/olcf/41941941904/in/album-72157683655708262/

OLCF Summit:

- 200 x 10^{15} floating-point operations per sec.
  - 200 PetaFLOPS
- #2 fastest supercomputer
- 4608 Nodes:
  - 2 IBM Power9 CPUs + **6 NVIDIA Volta GPUs**
- 10 MW system

Exascale systems will be at least 5 times faster, but require no more than 3 times the power

ExaWind development is funded by two DOE offices

DOE Wind Energy Technologies Office:

- “High-Fidelity Modeling” project
- Period of performance: 2016-2023
- Partnership between the National Renewable Energy Laboratory (NREL) and Sandia National Laboratories (SNL)

ATMOSPHERE
TO ELECTRONS
U.S. DEPARTMENT OF ENERGY

DOE NNSA/SC Exascale Computing Project:

- “ExaWind” project; https://www.exawind.org/
- Period of performance: 2016-2023
- Partnership between NREL, SNL, Oak Ridge National Laboratory, University of Texas at Austin, & Parallel Geometric Algorithms, Inc.

ECP
EXASCALE COMPUTING PROJECT

The ExaWind-HFM team: 40+ researchers

NREL
NATIONAL RENEWABLE ENERGY LABORATORY

- **M. Sprague, HFM & ExaWind PI**
- S. Ananthan
- R. Binyahib
- M. Brazell
- M. Churchfield
- G. Deskos
- A. Glaws
- K. Gruchalla
- M. Henry de Frahan
- R. King
- J. Jonkman
- T. Martinez
- P. Mullowney
- M. Natarajan
- R. Mudafort
- J. Rood
- A. Sharma
- K. Swirydowicz
- S. Thomas
- G. Vijayakumar
- S. Yellapantula

Sandia National Laboratories

- **P. Crozier, ExaWind & HFM co-PI**
- L. Berger-Vergiat
- M. Blaylock
- L. Cheung
- D. Glaze
- A. Hsieh
- J. Hu
- R. Knaus
- D.H. Lee
- D. Maniaci
- T. Okusanya
- J. Overfelt
- S. Rajamanickam
- P. Sakievich
- T. Smith
- J. Vo
- A. Williams
- I. Yamazaki

OAK RIDGE
National Laboratory

- **J. Turner**
- A. Prokopenko
- R. Wilson

TEXAS
The University of Texas at Austin

- **R. Moser**
- J. Melvin

Parallel Geometric Algorithms, Inc.

- J. Sitaraman

Next-generation high-performance computing

Shreyas Ananthan

DOE high-performance computing (HPC) systems

Petascale systems

EERE

Eagle

NREL Eagle
8 PF; ~1 MW
51 in HPC Top500

Peregrine

NREL Peregrine
2.24 PF; ~700 kW
Retired 2019

DOE Leadership Computing Facilities (LCF)

ANL Theta
11.6 PF
34 in HPC Top500

ORNL Titan
27 PF
Retired 2019

ORNL Summit
200 PF; ~10 MW
2 in HPC Top500

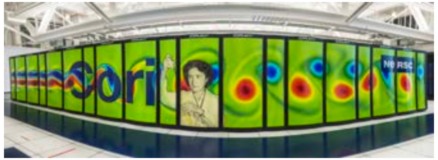

NERSC Cori
27.8 PF; ~4 MW
16 in HPC Top500

Exascale systems
2021 – 2023

ORNL Frontier
~ 1.5 Exaflops

ANL Aurora
> 1 Exaflops

1 Petaflop = 10^{15} calculations per second

DOE HPC systems present an unprecedented opportunity for unlocking novel insights into wind farm physics

However, wind research has not fully harnessed all the available computing power

Existing wind-energy codes are not suited for running on state-of-the-art HPC systems

Supercomputer hardware is very different

Summit node layout

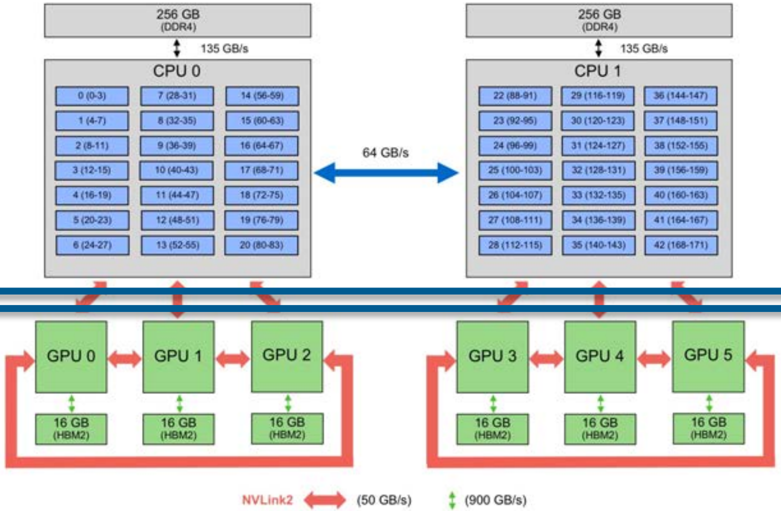

Most wind-energy codes are designed to run on these chips

Almost all computing power in modern supercomputers on these chips

Supercomputer hardware is very different

Future supercomputers will have more exotic architectures

NVIDIA, AMD, Intel are all making their own GPUs

Tomorrow's NREL system might look a lot like today's Summit

Exascale system (2021)
Frontier node layout

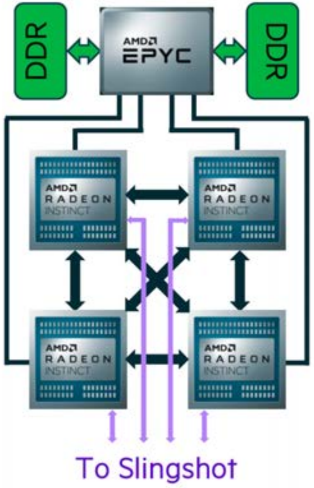

To run efficiently on current & exascale systems, codes must be able target different types of hardware

Running on different kinds of hardware is a key priority for ExaWind codes

Laptops, workstations

NREL Eagle

ORNL Summit

ORNL Frontier (2021)

Cloud computing

Amazon EC2

Microsoft Azure

Containerization

Docker

Singularity

Measuring performance on supercomputers

Strong-scaling **– a way to measure how well we are utilizing the supercomputer**

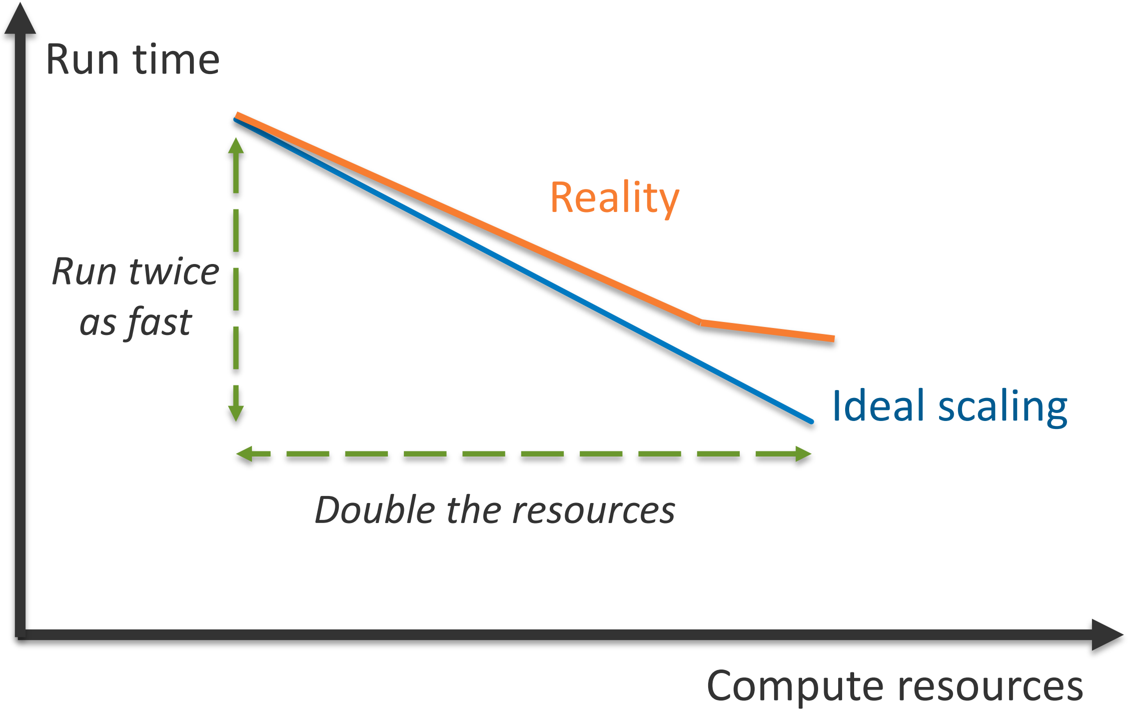

Closing the gap between ideal and reality is a big focus of ExaWind project

Performance-measurement example

- Atmospheric-boundary-layer simulation
- Compute the complex flow in which turbines operate
- Simulated using millions of *grid points*
- Double compute resources (cores) and measure run time

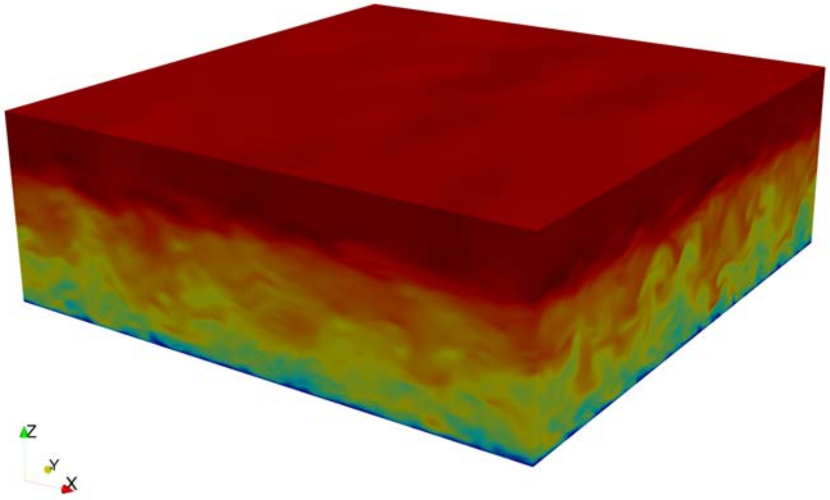

Complex atmospheric flow over a 3 sq. km area

ExaWind strong-scaling performance on Eagle

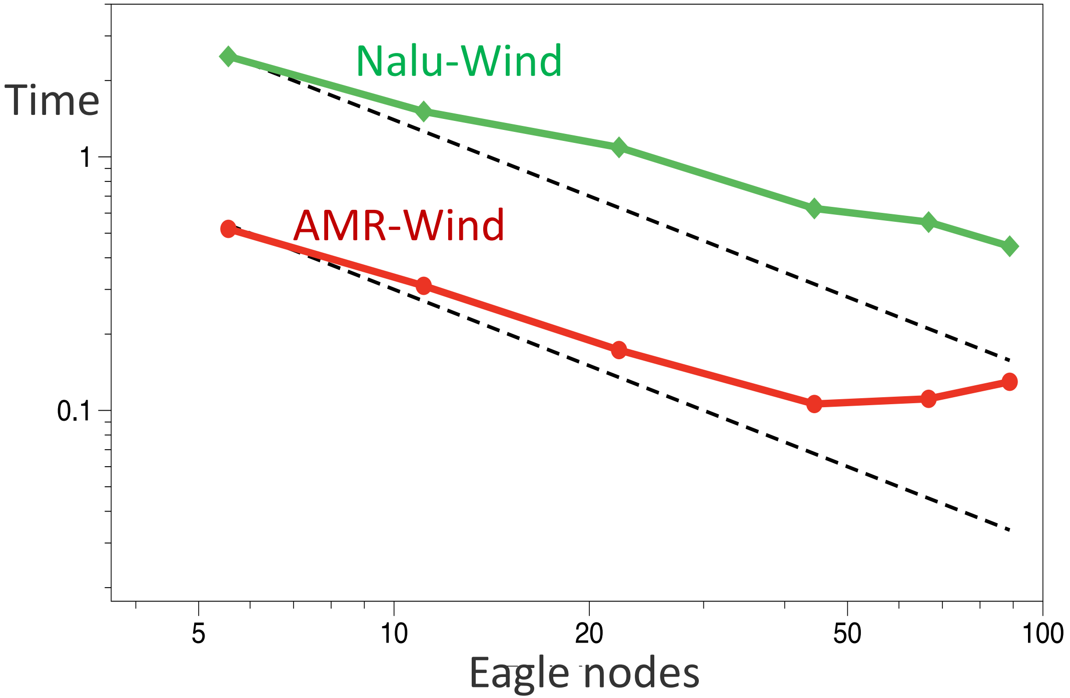

- Measure performance of Nalu-Wind and AMR-Wind on Eagle
- 25 million grid points where equations are solved
- Starts trailing-off when we use large portions of the machine
- AMR-Wind is *5x faster* than Nalu-Wind on the same problem

Comparing performance for CPUs and GPUs on Summit

3D flow field from ABL simulations

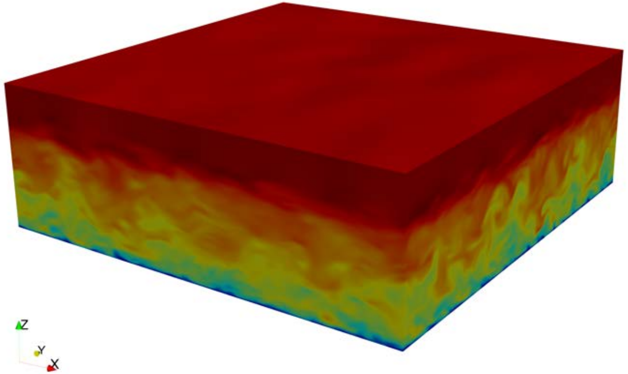

Strong-scaling study for the ABL LES *precursor* simulation on a 3 km x 3 km x 2 km domain with uniform mesh resolution on ORNL Summit

AMR-Wind Strong-scaling performance

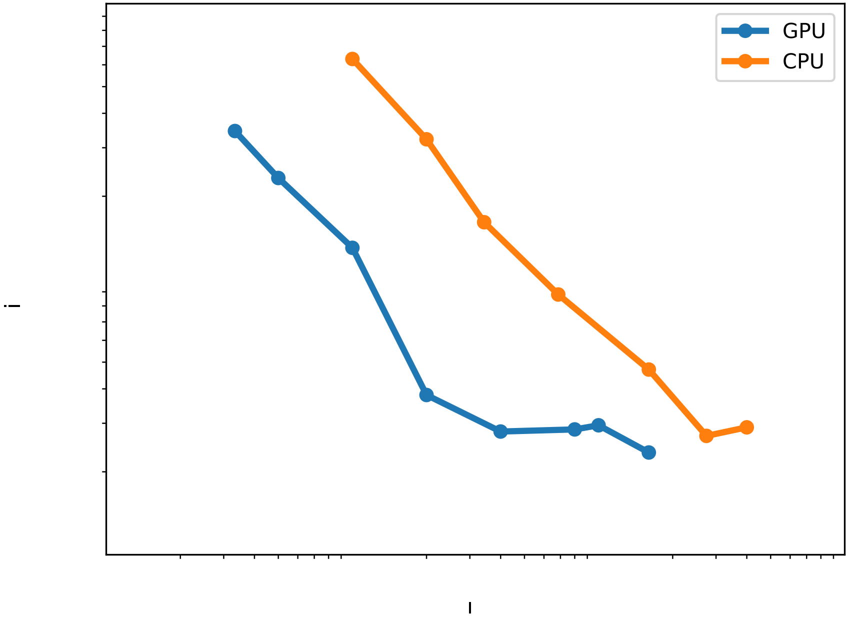

Model validation: Building confidence in our simulated reality

Ganesh Vijayakumar

Validation is key to establishing predictivity of wind-energy simulations

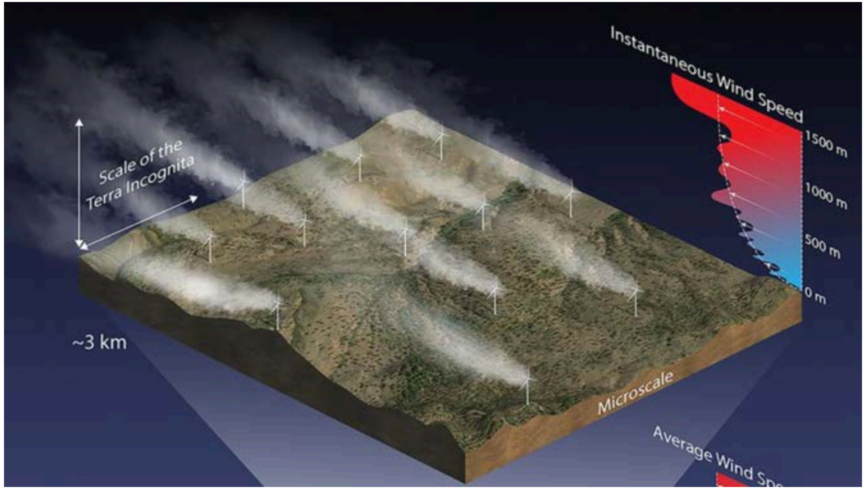

How can we trust Exawind framework results?

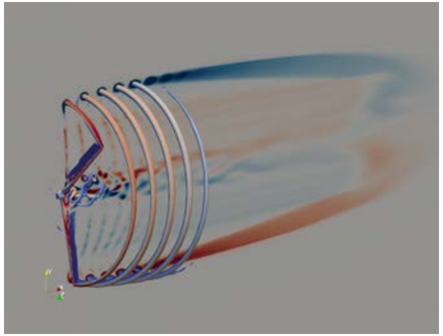

Q-criterion and contours of velocity through NREL-5MW turbine simulated using Exawind framework

Exawind framework – solve model equations on HPC for wind-energy problems

Validation – Are the equations a true representation of the physics?

Comparison to experimental results

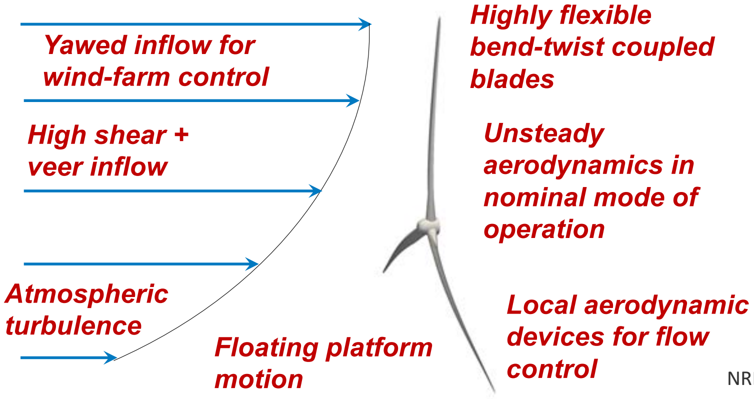

Yawed inflow for wind-farm control

High shear + veer inflow

Atmospheric turbulence

Floating platform motion

Highly flexible bend-twist coupled blades

Unsteady aerodynamics in nominal mode of operation

Local aerodynamic devices for flow control

Hierarchy of validation studies for increasing complexity

International Benchmark Working Group from industry and universities to frame benchmark validation cases.

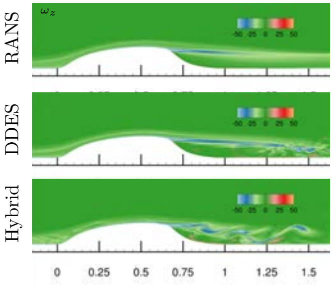

Comparison of turbulence models for the NASA validation case of a 2D wall mounted hump (Moser)

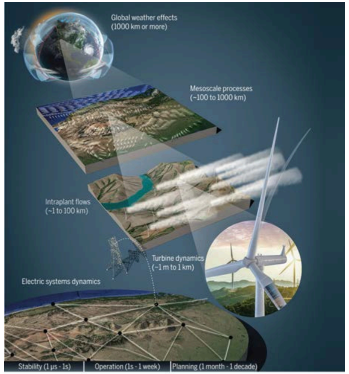

Complexity

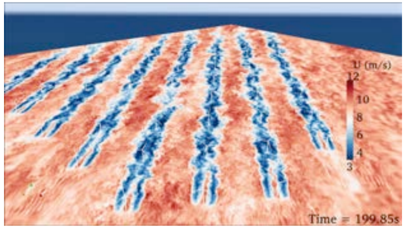

Full-scale wind farm simulations, e.g., Lillgrund wind farm (Churchfield)

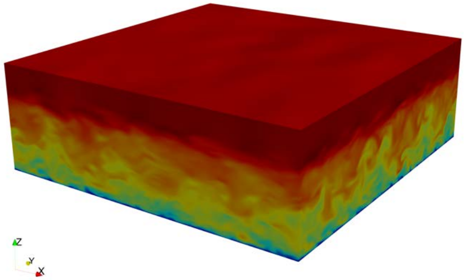

ABL simulation (Brazell)

Single turbine simulation in uniform inflow (Ananthan)

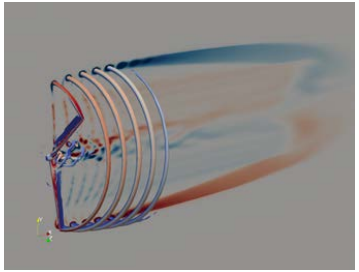

Single turbine simulation in ABL (Churchfield)

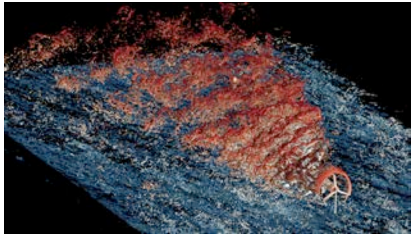

McAlister-Takahashi fixed-wing wind-tunnel validation

NACA0012 airfoil benchmark problem from NASA Langley Turbulence modeling resource[1]

k-w-SST – RANS turbulence model

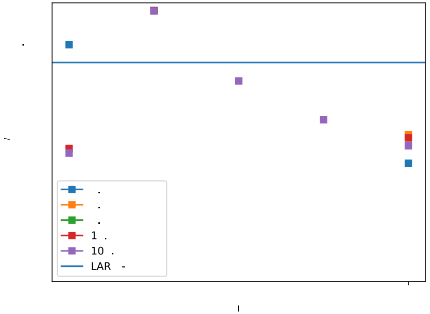

Aerodynamic performance: Lift and drag coefficients with grid refinement

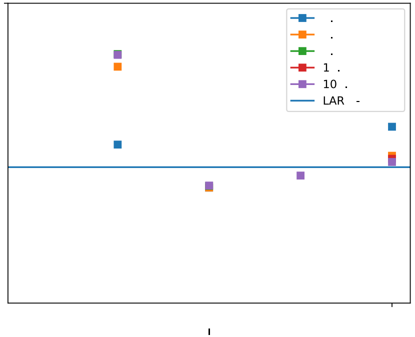

McAlister-Takahashi fixed-wing[2]

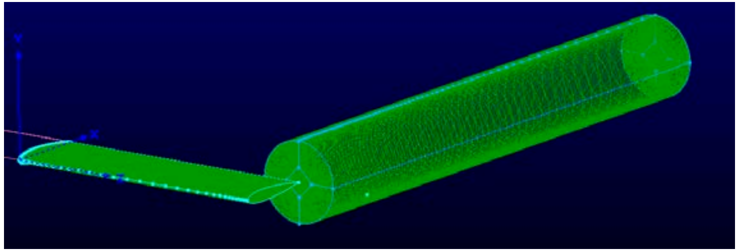

Pressure profiles along wing

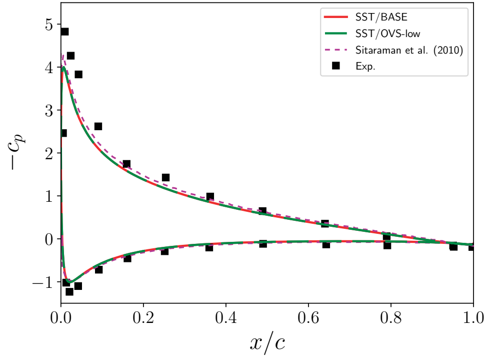

Tip vortex capture

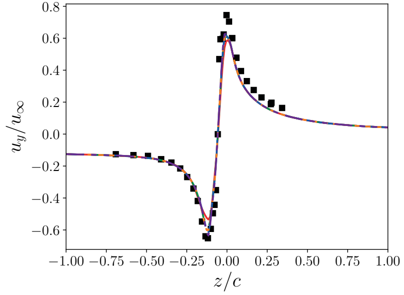

1 – https://turbmodels.larc.nasa.gov/naca0012_val_sst.html

2 – McAlister and Takahashi, NACA 0015 wing pressure and trailing vortex measurements, Tech. Rep. NASA-A-91056, NASA Ames

NREL Phase-VI: Wind-tunnel full-turbine experiment

Unsteady Aerodynamic Experiments in the 1990s in NASA Ames wind tunnel[1]

Two-bladed extremely stiff teetering turbine with fixed rpm/pitch

10m diameter

Compare with detailed measurements on blade

k-w-SST RANS turbulence model

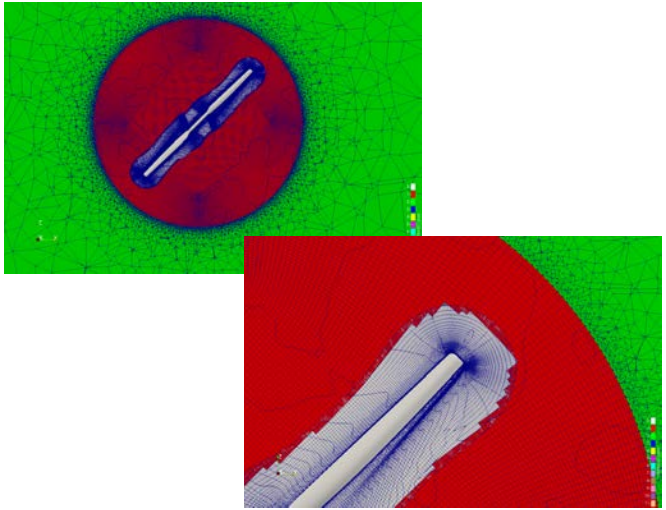

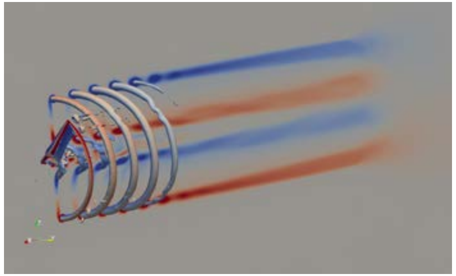

Integrated rotor torque

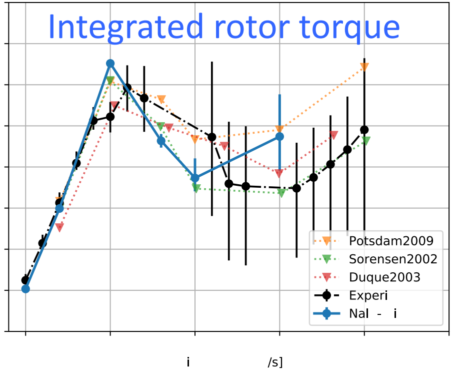

Pressure profiles at different points along the blade

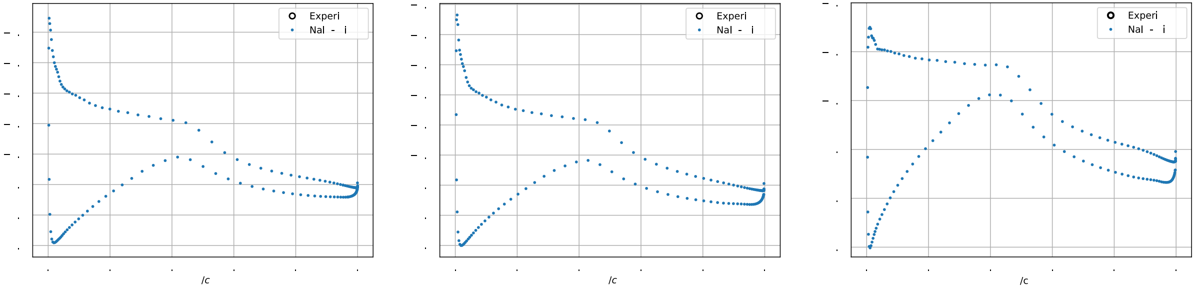

r/R = 0.3

r/R = 0.63

r/R = 0.9

---

1 — Hand et. al., Unsteady Aerodynamics Experiment Phase VI: Wind tunnel test configurations and available data campaigns, Tech. Rep. NREL/TP-500-29955, 2001.

IEA Task 29: Dan-Aero 2-MW NM-80 turbine field experiment

Three-bladed upwind turbine with full rpm and pitch control

International group focused on validation and code-to-code comparison

Collaboration with DTU: Use very similar grids, algorithm and turbulence model

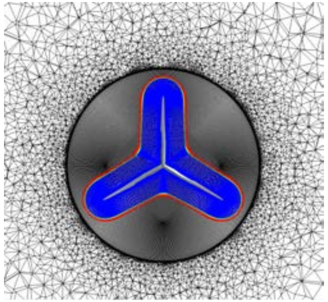

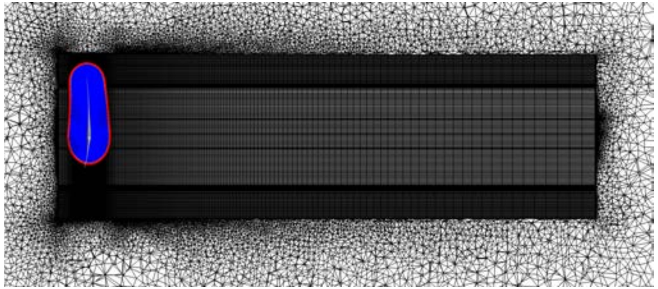

Pressure profiles

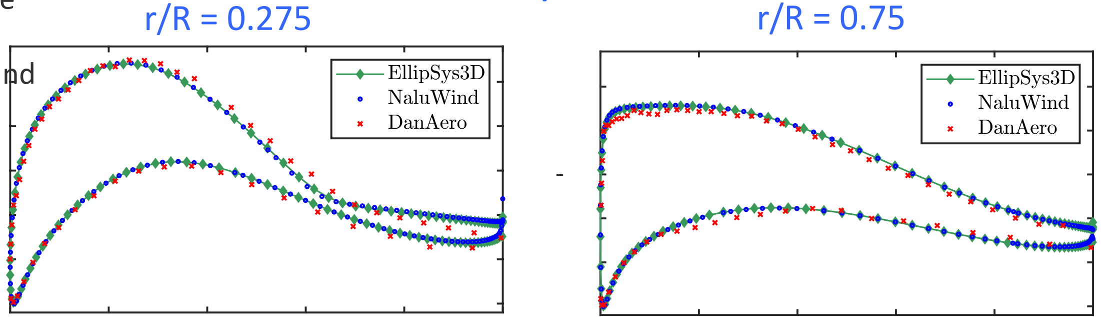

Normal force profile

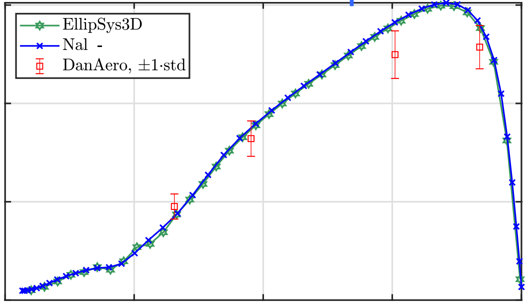

Tangential force profile

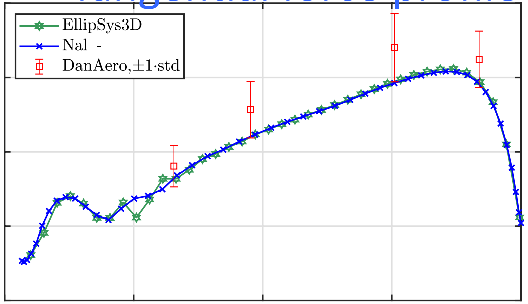

---

1 – Grinderslev et. al., Validation of blade-resolved computational fluid dynamics for a MW-scale turbine rotor in atmospheric flow, To be presented at Torque 2020.

NREL 5-MW: Offshore-relevant demonstration

Reference turbine established by NREL in 2009[1]

Three-bladed upwind turbine with full rpm and pitch control

No validation data: Compare to other codes in literature

k-w-SST RANS turbulence model

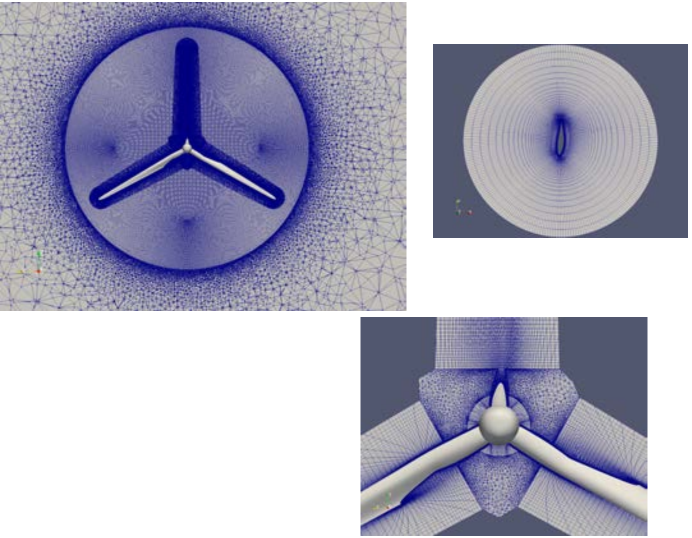

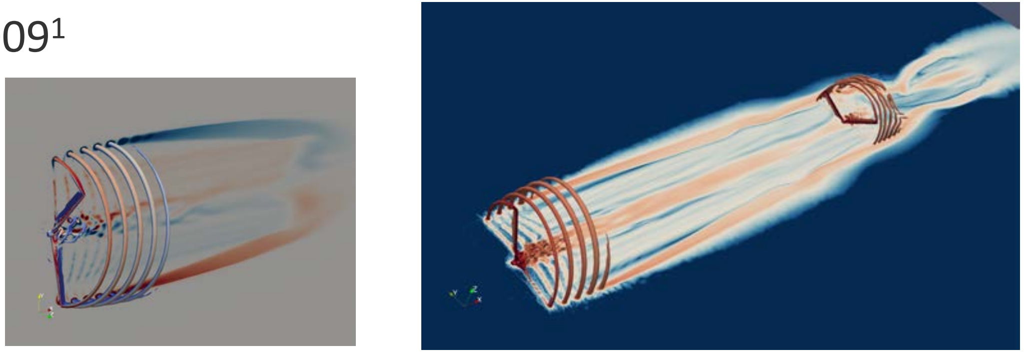

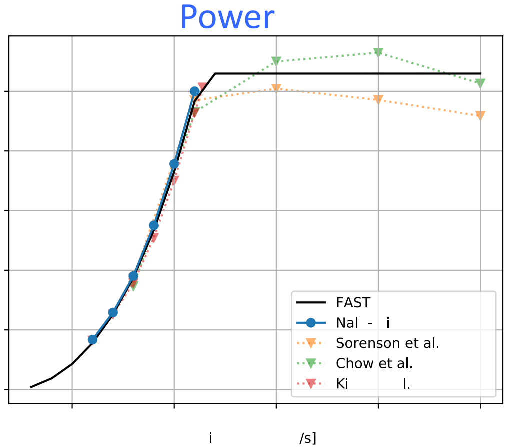
Power

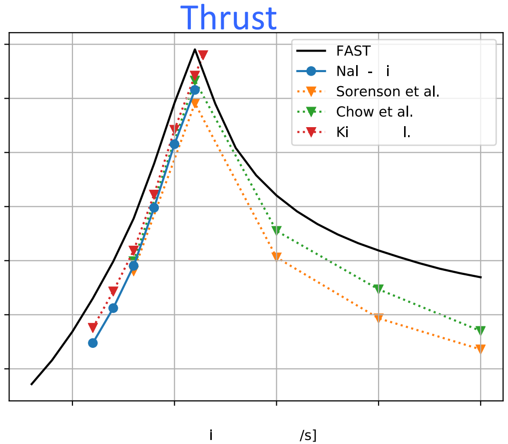
Thrust

---

1 – Jonkman et. al., Definition of a 5-MW reference wind turbine for offshore system development, Tech. Rep. NREL/TP-500-38060, 2009.

Current work

Fluid-Structure Interaction

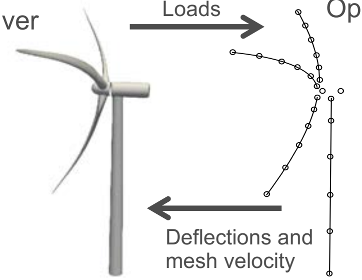

Nalu-wind
CFD solver

Loads

OpenFAST

Deflections and
mesh velocity

Demonstration using NREL 5MW turbine with k-w-SST RANS turbulence model[1]

Generator power

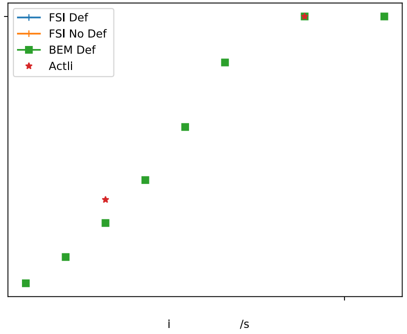

Rotor thrust

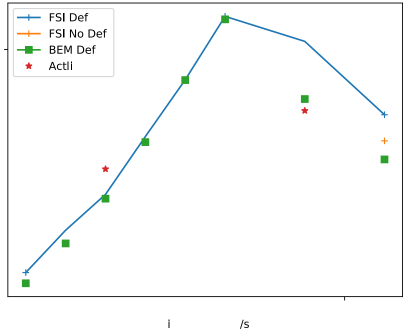

Bridging length scales using hybrid-RANS/LES turbulence modeling

Blade-resolved simulation in ABL turbulent inflow

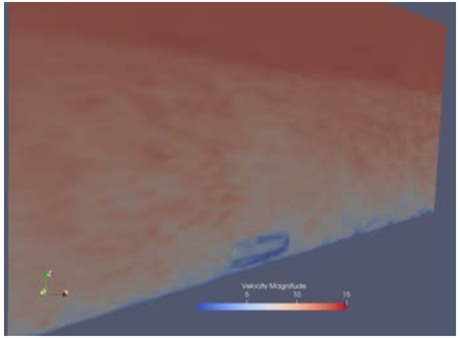

Ongoing validation with IEA Task 29

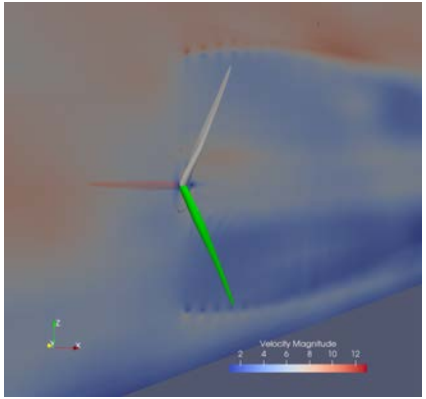

---

1 – Vijayakumar et. al., Effect of Fluid-Structure Interaction on wind turbine loads, Wind Energy Science Conference, Cork, Ireland.

Conclusion

Exawind framework is built on a strong validation base with increasing levels of complexity

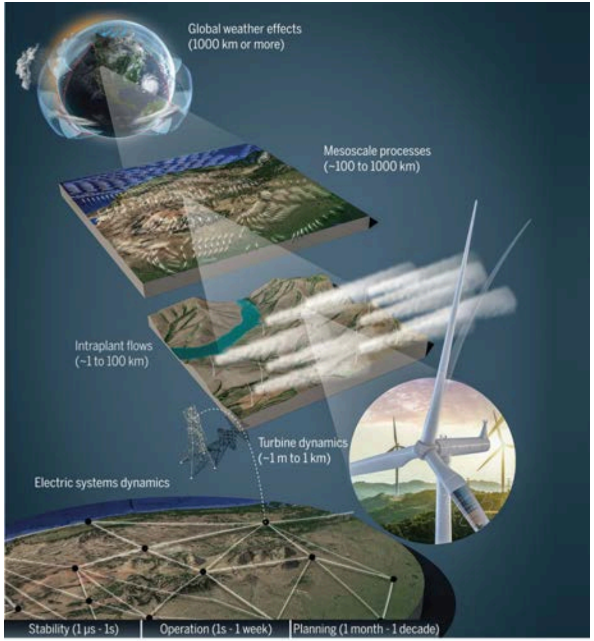

Goal: Wind farm with flexible turbines in ABL with complex terrain/offshore

Single turbine in ABL with FSI

Single turbine in ABL

Single turbine in uniform inflow

Fixed wing

Airfoil

HPC and HFM as the foundation for next-generation engineering models

Luis 'Tony' Martínez Tossas

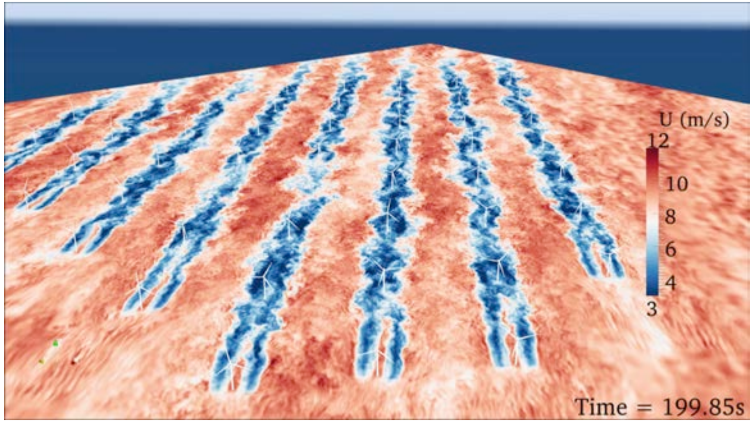

Image of SOWFA simulation of the Lillgrund offshore wind farm (Churchfield, NREL)

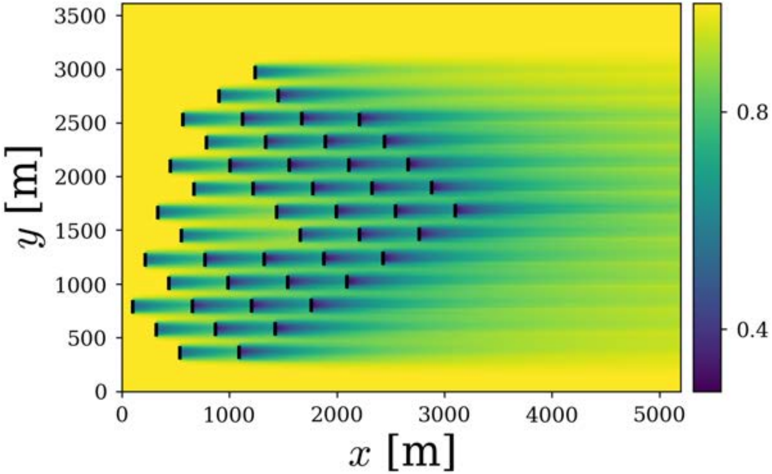

Image of a simplified wake model simulation of the Lillgrund offshore wind farm

The challenge

How can we use high-fidelity modeling to design better engineering models?

Modeling-fidelity spectrum

Higher fidelity
More predictive (less tuning)
More computational expense

Best for:

- Untangling complex dynamics
- Exploring/demonstrating technology innovations
- Validating/creating lower-fidelity engineering models

Lower fidelity
Less predictive (more tuning)
Less computational expense

Best for:

- Optimization studies
- Sensitivity/uncertainty-quantification studies
- Certification studies

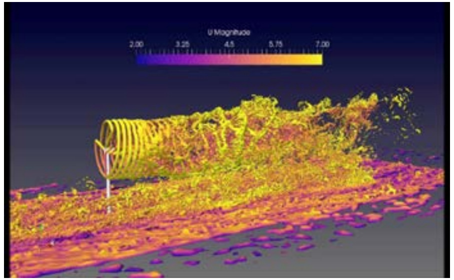

Need a modeling suite that spans the fidelity spectrum

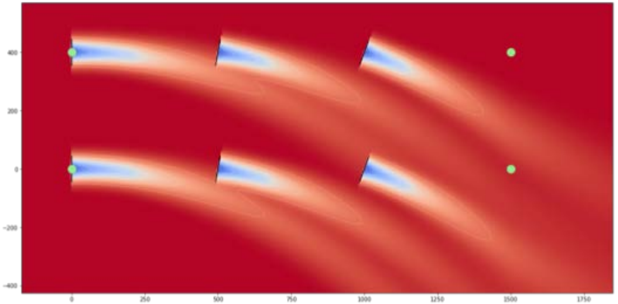

Newton’s second law

$$\boldsymbol{F} = m\boldsymbol{a}$$

$$\rho\left(\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u}\cdot\nabla\mathbf{u}\right) = -\nabla p + \mu\nabla^2\mathbf{u} + \mathbf{g}\alpha\left(T \quad T_0\right)$$

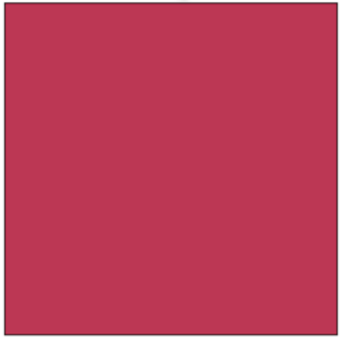

Many approximations
(lower fidelity)
inexpensive

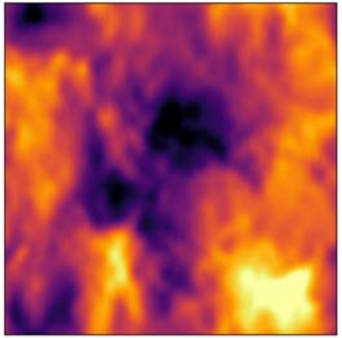

Some approximations
(medium fidelity)
expensive

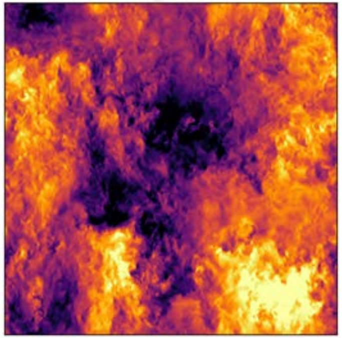

Zero approximations
(highest fidelity)
most expensive

Turbulence data obtained from the JHU database http://turbulence.pha.jhu.edu/

Computational cost vs model fidelity

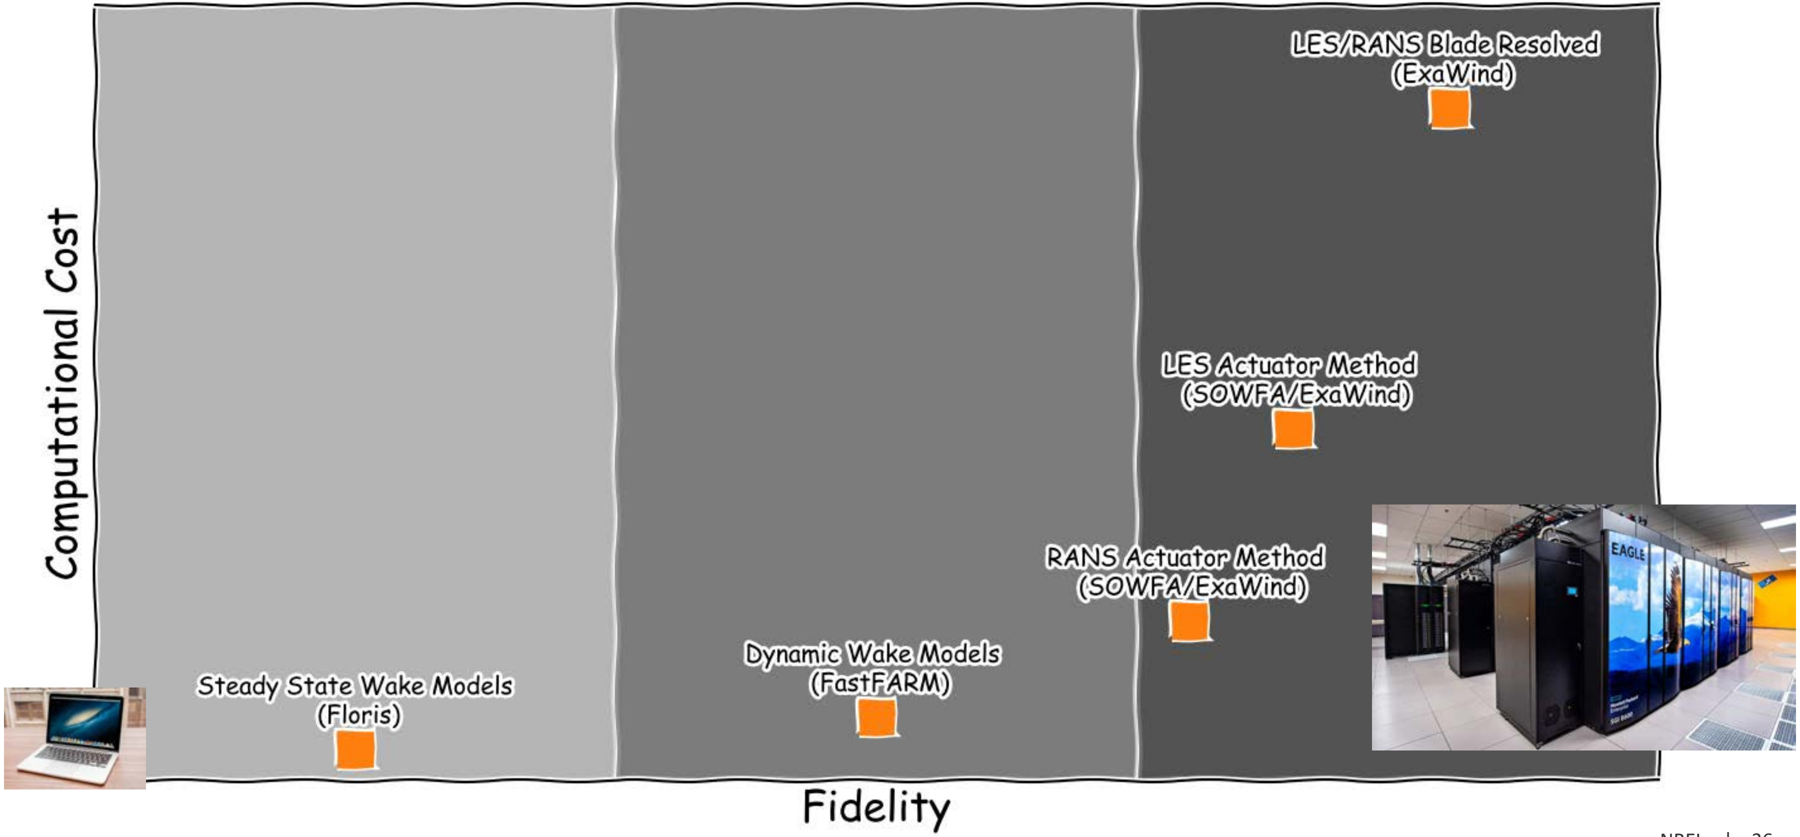

Example 1: The curled wake

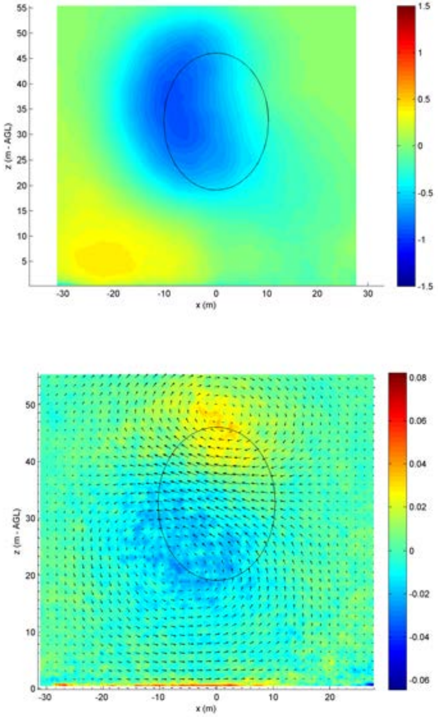

Images from brainstorming with Matt J Churchfield about wakes in yaw (2016)

APS Gallery of Fluid Motion (2016)

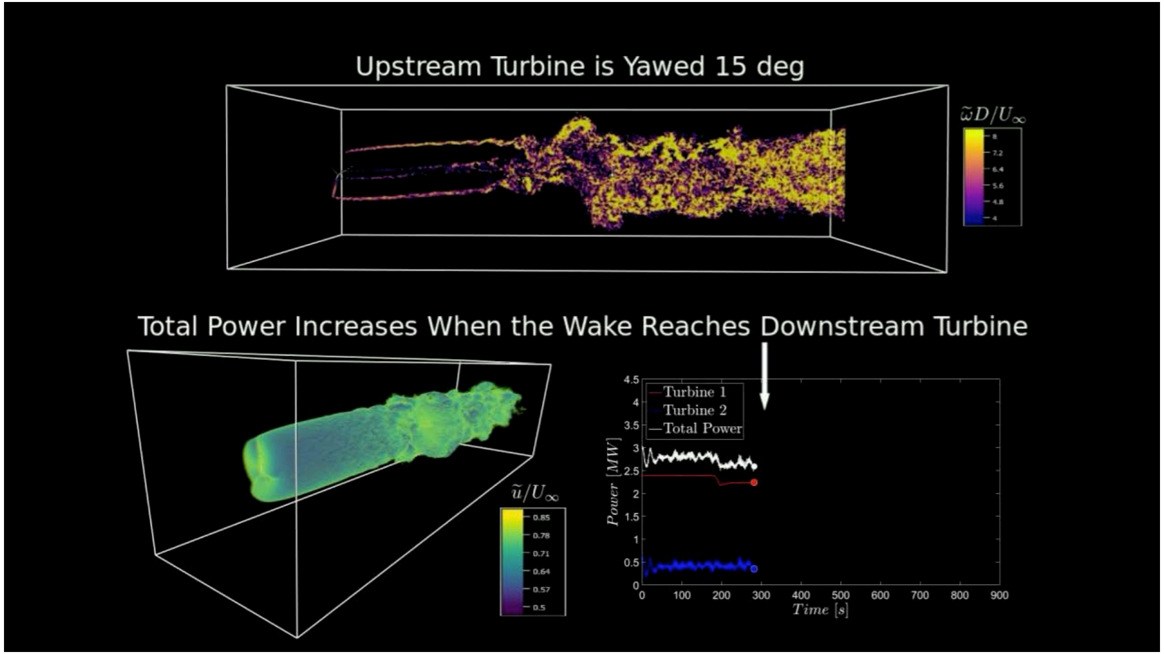

Simulation comparison of wake mitigation control strategies for a two-turbine case. Paul Fleming, Pieter M.O. Gebraad , Sang Lee, Jan-Willem van Wingerden, Kathryn Johnson, Matt Churchfield, John Michalakes, Philippe Spalart, Patrick Moriarty. Wind Energy 2016

Example 1: The curled wake

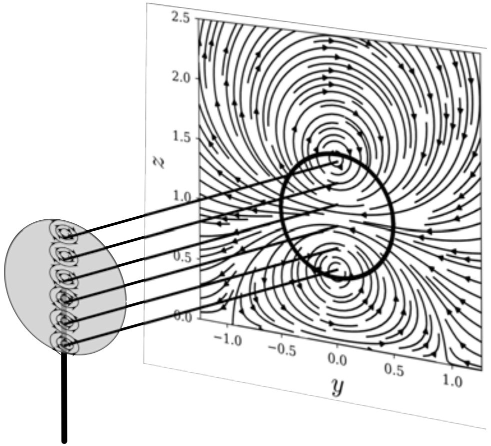

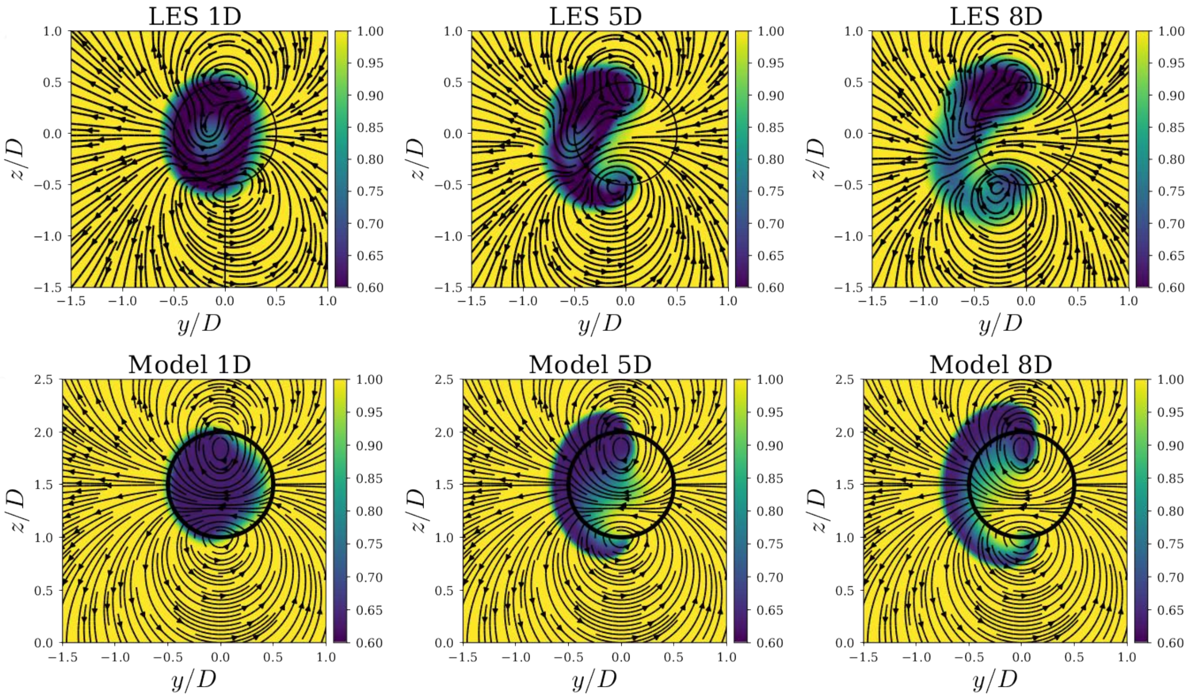

LA Martínez-Tossas, J Annoni, PA Fleming, and MJ Churchfield. **The aerodynamics of the curled wake: a simplified model in view of flow control.** Wind Energy Science 2019.

CJ Bay, J King, LA Martinez-Tossas, R Mudafort, P Hulsman, M Kuhn, and P Fleming. **Towards flow control: an assessment of the curled wake model in the FLORIS framework,** Torque 2020

LA Martinez-Tossas, J King, E Quon, CJ Bay, R Mudafort, N Hamilton, and P Fleming. **The curled wake model: A three-dimensional and extremely fast steady-state wake solver for wind plant flows**, in review, 2020.

J King, P Fleming, R King, LA Martínez-Tossas, CJ Bay, R Mudafort, and E Simley: **Controls-Oriented Model for Secondary Effects of Wake Steering**, Wind Energ. Sci. Discuss, in review, 2020.

Example 2: Flow over a fixed wing

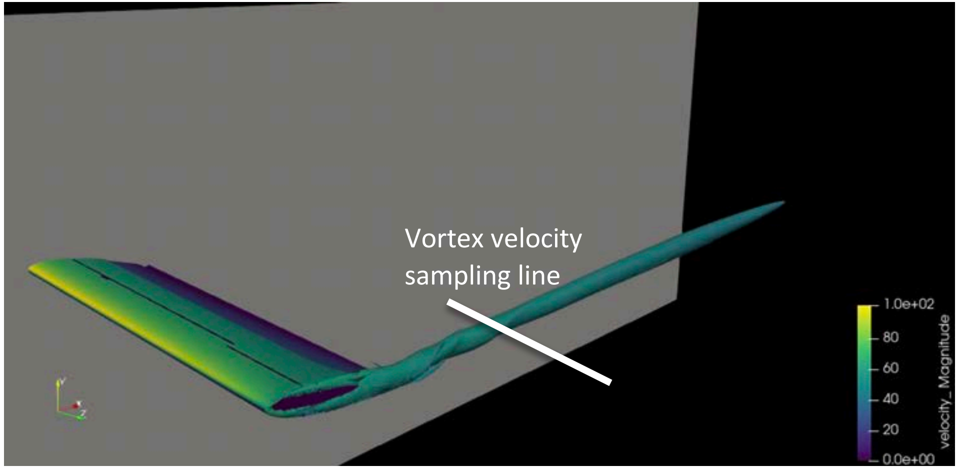

Mcalister fixed wing simulation velocity magnitude highlighting the tip vortex

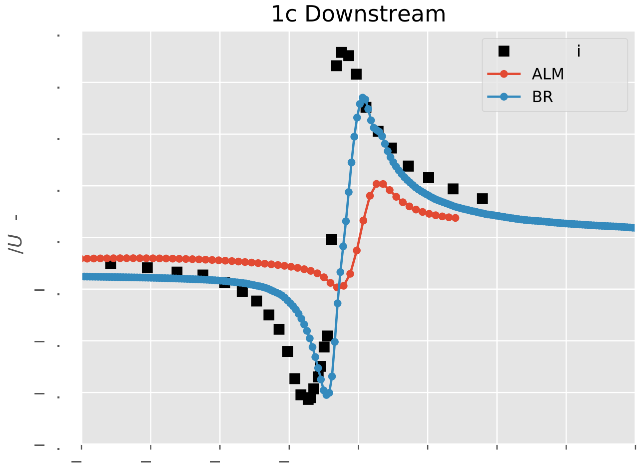

Cross-stream tip vortex velocity sampled along white line shown on the left figure

Comparison with experiments, high-fidelity simulations, and reduced-order models can expose pathways to next-generation design models & codes

Results presented at 2019 APS Division of Fluid Dynamics conference
C09.00010: Comparison of theory and large-eddy simulations with experiments of flow over a wing
Luis Martinez, Marc Henry de Frahan, Ganesh Vijayakumar, Shreyas Ananthan

Example 3: HFM results can be used to train neural-network models

New unsteady aerodynamics model using machine learning performs better than state-of-the-art unsteady aerodynamics models

Can be trained using Exawind framework CFD data.

Will improve fatigue load estimations for bigger wind turbines of the future on floating platforms

Ananthan, S., Vijayakumar, G., Yellapantula, S., *A DNN surrogate unsteady aerodynamic model for wind turbine loads calculations*, To be presented at TORQUE 2020.

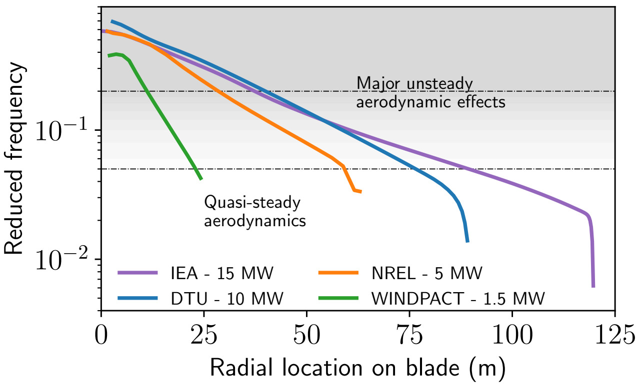

Measure of unsteadiness (reduced frequency) over the blade length for 4 commercially relevant wind turbines.

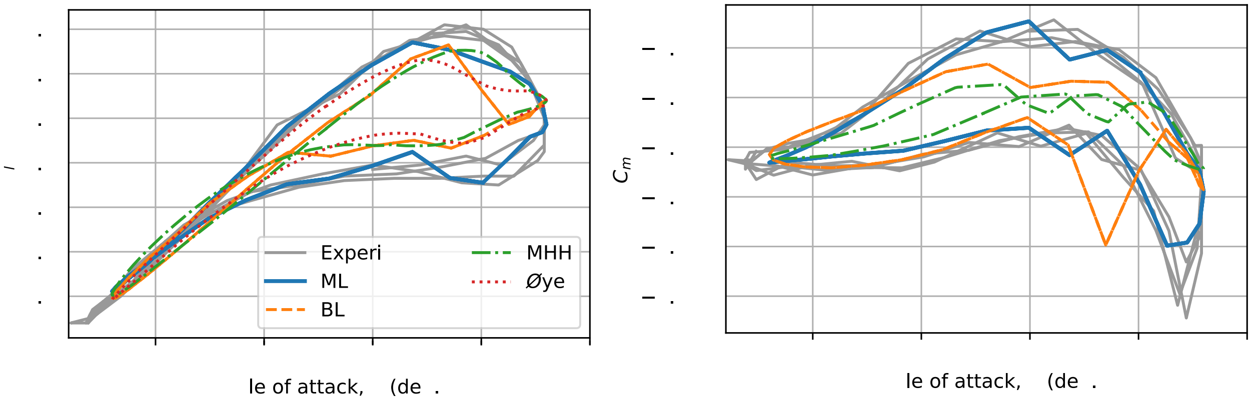

Comparison of the new ML unsteady aerodynamics model and state-of-the-art models against experimental data for a pitching N4415 airfoil.

Current work: High-thrust coefficient

- Wake models have a hard time predicting wakes for high-thrust conditions
- Can we simulate these conditions using HFM tools?
- How can we improve the current wake models to account for high thrust?
- Simulations using Nalu-Wind and OpenFAST (ExaWind)

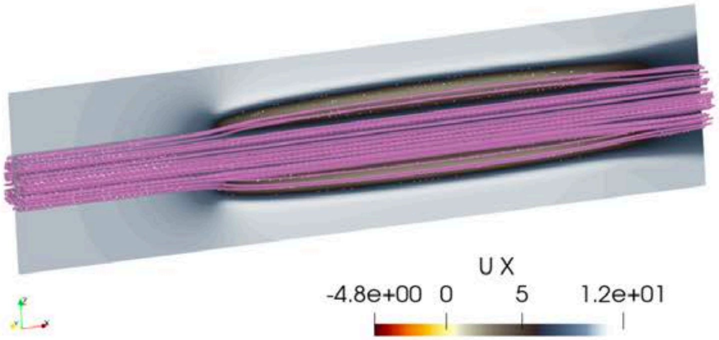

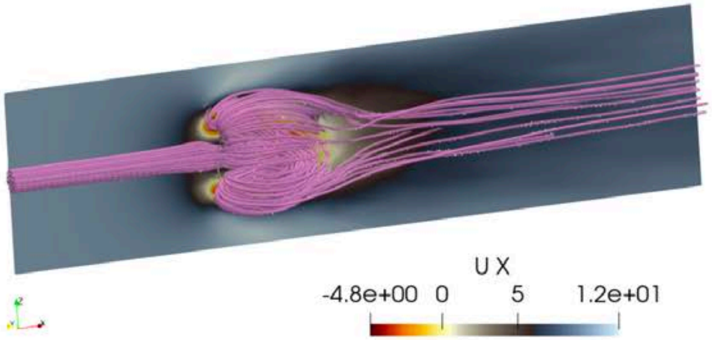

Team: Luis A Martínez-Tossas, Emmanuel Branlard, Kelsey Shaler, Ganesh Vijayakumar, Shreyas Ananthan, Philip Sakievich, Jason Jonkman

What would the pioneers of wind energy say?

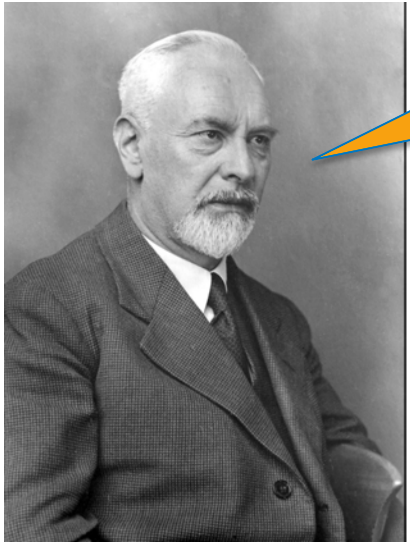

Ludwig Prandtl
(1875 –1953)

Albert Betz
(1885 –1968)

Progress will depend on a generation of scientists educated deeply in their own specialty as well as in the breadth of wind energy science.
Grand challenges in the science of wind energy, Science 2019

Next frontier for ExaWind: Offshore wind

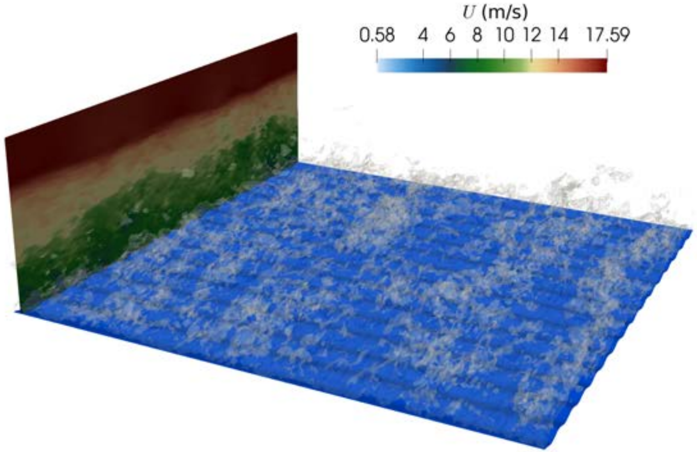

Simulation of wind over waves using the ExaWind/Nalu-Wind fluid solver. (G. Deskos)

Q&A

www.nrel.gov

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. Funding provided in part by the U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Wind Energy Technologies Office. Funding was also provided 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. 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. A portion of this research was performed using computational resources sponsored by the Department of Energy's Office of Energy Efficiency and Renewable Energy and located at the National Renewable Energy Laboratory. This research also used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725.

NREL
Transforming ENERGY