ExaWind at NREL: Upping the Ante

Nicholas Brunhart-Lupo
Ashesh Sharma
nicholas.brunhart-lupo@nrel.gov
ashesh.sharma@nrel.gov
National Renewable Energy Laboratory
Golden, Colorado, USA

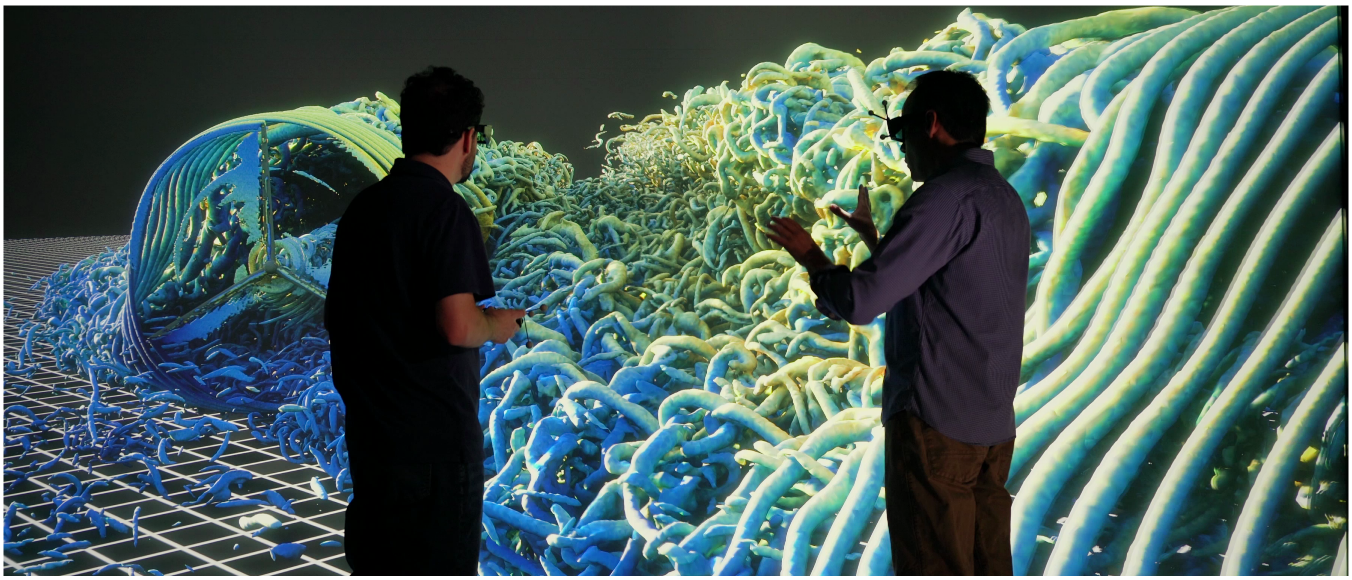

Figure 1: Results of a high-fidelity wind farm simulation artifact from the ExaWind solver suite. Iso-surface represents q-criterion, with shading derived from velocity.

ABSTRACT

The objective of the ExaWind component of the Exascale Computing Project is to deliver many-turbine blade-resolved simulations in complex terrain. These simulations bring new challenges to both compute and analysis of the resulting data. In this paper/video, we visually explore the impact of ExaWind on wind simulations through two studies of a small wind farm under two atmospheric conditions. We then turn to analysis and review tools that visualization researchers at NREL use to answer the challenges that ExaWind brings.

KEYWORDS

scientific visualization, high performance computing, exascale, wind simulation, renewable energy

ACM Reference Format:
Nicholas Brunhart-Lupo and Ashesh Sharma. 2023. ExaWind at NREL: Upping the Ante. In *Proceedings of Supercomputing Visualization Showcase (SC Showcase '23).* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

SC Showcase '23, November 14–16, 2023, Denver, CO, USA
2023. ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

Wind simulations are a critical element in wind energy research, providing insights ranging from materials and blade morphology to whole-farm design. These simulations are constantly challenged by a curse of scales; kilometers of context are required to capture large-scale phenomena such as atmospheric eddies, and at the same time, $O(\text{microns})$ of resolution is desired to properly express airflow dynamics near the turbine blades.

For accurate modeling, these simulations should be able to resolve the blade. Instead of actuator line models that approximate how a blade might affect the airflow, blade-resolved simulations use actual blade geometry, capturing the thin boundary layers and blade deflection under dynamic loads.

While compute hardware has advanced to allow single-turbine blade-resolved simulations in current petascale platforms, multi-turbine simulations require exascale-class compute, software, and strategy[12, 13]. The ExaWind component of the Exascale Computing Project has aims to build the required infrastructure and solve tooling to deliver these predictive simulations at extreme scale.

In this paper and accompanying video, we explore the advancements made in wind simulations and the new level of fidelity available to researchers. We also explore some new tools available to help disseminate this data.

2 BACKGROUND

The current ExaWind solver stack consists of AMR-Wind and Nalu-Wind. AMR-Wind[3] is an incompressible flow solver designed for large-scale wind simulations using adaptive mesh representations. AMR-Wind is suited to simulating large-eddy simulations (LES) of atmospheric boundary layer (ABL) flows and simulating wind farms using actuator disk or actuator line models for turbines. For blade-resolved simulations, AMR-Wind is typically used as a background solver, coupled with Nalu-Wind, which drives the near-body blade simulation. Nalu-Wind is also an incompressible flow solver[12], operating on unstructured meshes. Both Nalu-Wind and AMR-Wind share the objective of providing the highest-fidelity simulations of turbine flow fields on leadership class hardware.

The solving stack can provide resolutions down to 10 µm if required. Further, the stack makes use of adaptive refinement. This critical feature provides the kilometer scale needed to represent different atmospheric scenarios. While adaptive refinement increases the complexity of the solver and the analysis of the data, selective fidelity preserves full resolution only in regions where it is needed to fully capture flow dynamics (i.e., near turbine blades), lowering fidelity where only low-frequency phenomena should be preserved (such as in the periphery of the farm distant from the turbines), all while still capturing all the relevant flow physics.

In contrast, previous simulations were restricted. These simulations used actuator models or low-resolution adaptive refinement. In a linear actuator model, for example, a turbine blade is represented by a line, with forces projected from this line to the solver mesh[14]. While computationally efficient, this is a non-physical model and suffers from inconsistencies[11]. Figure 2 shows a visualization of one such simulation. Here, an isosurface of low-velocity is visible in blue, and an isosurface of q-criterion is visible in yellow. The low-velocity region shows the wake of the turbine, and the q-criterion provides information on the rotation of the air; the blade-tip vortices are visible, as well as the rolling of air off the blades. However, in this example, the simulation could only provide resolution of around 3 to 10 meters near the turbines. This low-resolution is evident in the 'blocky' appearance of the visualization, with large triangles and rough geometry obscuring flow features.

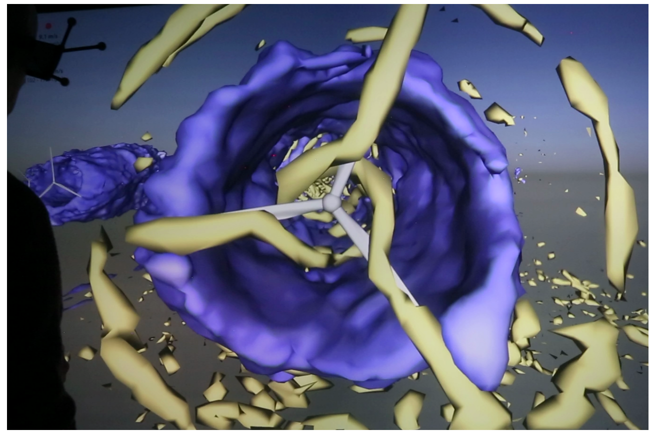

Figure 2: Example of previously-attainable simulation results. Triangles are easily visible with the lower resolution of the dataset.

3 SIMULATION

We demonstrate the applicability of the ExaWind solver to a four-turbine wind farm comprising four NREL 5 MW turbines under two different atmospheric boundary layer (ABL) conditions. The turbines are arranged in a square, with the wind directed in an approximate diagonal direction.

A domain of size 1.9 km by 1.5 km by 0.94 km is used. The background grid is block-structured with three levels of refinement for both stability states. All levels of refinement combine for a total of 615 million cells. An unstructured near-body grid with about 8 million cells is used. A mean flow of 10 m/s is enforced at the turbine hub height for both stability states simulated. A controller is not active; all turbines are rotated at a fixed number of rotations per minute (RPM) corresponding to the mean flow speed of 10 m/s.

3.1 Neutral ABL

For this study, we used an ABL corresponding to a neutral state, where the sun's influence on the atmosphere is minimal. The turbulent intensity in this regime is 5.5%. The inversion height of the neutral ABL is capped at 600 m. Upon accounting for the Coriolis force, the wind direction at hub height corresponds to 225° by meteorological conventions which dictate that north is positive y-axis and east is positive x-axis. The height of the capping is specified by manually defining the potential temperature along the height of the domain, including a zero surface potential temperature. The capping height was chosen to be representative of neutral ABLs in nature. For the presented results, the boundary layer developed over the course of 34 hours of physical time. We simulate the problem for 500 s of physical time, which corresponds to approximately two flow-through times, comprising 95 rotor rotations at an RPM consistent with that of $u = 10$ m/s. The isosurfaces presented account for the evolution of wake, turbine-wake interaction, and ABL-wake interaction within the wind farm.

Figure 3: Results from neutral ABL study. Shown is an isosurface of q-criterion, colored by the x-component of velocity. Note that while there is a considerable amount of 'noise' near the turbines, the wake maintains a straight-back structure.

3.2 Stable ABL

For our second demonstration study, we consider a strongly stable ABL. This kind of ABL corresponds to a cooling of the earth's surface and is particularly common in coastal areas. The turbulent intensity for this ABL is 1.3%. The wind speed and direction are consistent with the neutral ABL case. The height of the ABL is constrained by an applied inversion layer at 400 m. The capping of the ABL is based on the user-specified definition of potential temperature along the height of the domain. As in the neutral case, the capping inversion height was chosen to be representative of ABLs in nature. A surface heat flux is specified which cools the air near the ground, and causes the boundary layer to become stably stratified near the ground. The surface heat flux applied is commensurate with the desired potential temperature and stability profile. For the presented results, the boundary layer developed over the course of 28 hours of physical time. The long ABL development time is required due to the slower vertical transport of momentum under stable stratification conditions.

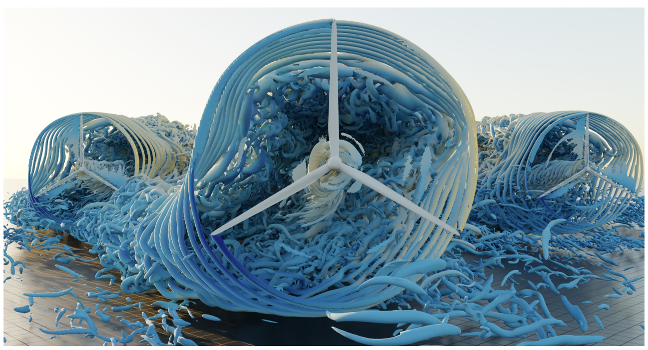

Figure 4: Results from the stable ABL study. Representation and coloring is the same in the neutral case. In contrast, however, note that the wake takes on a frame-left lobe as it progresses rearward. This is because of high veer resulting from the strong temperature stratification.

In both scenarios, the increase in fidelity is apparent. For these studies, micron-order resolutions are used near the blades, with resolutions on the order of 10 cm apparent in the isosurface artifact. Blade-tip vortices are more defined, and regions near the turbine hub show a more clearly articulated spiral. Outside this area, detail is vastly increased; in comparison to Figure 2, the surface-proximal isosurfaces are now visible with much more geometry.

Contrasting these two scenarios, we can see the wake of the neutral ABL is essentially straight. In the stable case, a lobe with a leftward trend is apparent. This is a result of the temperature boundary conditions. The more 'noisy' appearance of the neutral case results from the higher turbulent intensity.

In both cases, the wake completely obscures the fourth turbine. This is not an uncommon scenario. Turbine-wake interactions can negatively affect the power generation of that turbine (as the incoming flow already has significant energy extracted from it, or the flow is of higher turbulence) and also impacts turbine lifetime and maintenance burden due to the buffeting of turbulent flow [10]. Mitigating these effects remains an important topic in ongoing wind research [8].

4 VISUALIZATION

ExaWind data artifacts pose challenges for traditional visualization. The data is high-fidelity; while expected difficulties exist in processing large data, the higher resolution also realizes finer detail in more complex geometry presented to the viewer. Technical challenges also present themselves; this simulation is hybrid with AMR-Wind serving as the background solver and Nalu-Wind as the near-body solver. Hybrid visualization is supported by some packages, but, as these are rare code paths, this data tends to expose flaws in the software or suffer crashes. These tools also suffer from limitations in platform support.

These complexities frustrate more traditional techniques and platforms. Without effective visualization and analysis support, the fine details derived from these simulations will be wasted.

Therefore, to visualize these results, we made use of the immersive visualization suite at NREL's Insight Center [9] and consumer tablet devices, tied together with a new collaborative protocol developed at NREL, NOODLES [6].

4.1 Pre-processing

For this study, the data was first converted to an isosurface through PeleAnalysis[7] for AMR-Wind. PeleAnalysis is a collection of tools designed specifically for the analysis of block-structured, adaptively refined data at this scale. For Nalu-wind, a Paraview[2] script was used to extract the near-body information. Paraview was not usable for the extraction of the isosurfaces from the AMR-Wind data, due to ghost cell issues that are under resolution at the time of writing. These artifacts were then combined into a single scene. To reduce the load for interactive visualization, complex lighting components (such as ambient occlusion) were baked onto these surfaces using the open-source cinema renderer Blender.

The finalized scenes were several gigabytes in size (uncompressed), and were composed of $O(1 \times 10^7)$ triangles with multiple attributes per vertex. These scenes were then delivered to our visualization platforms using NOODLES.

NOODLES is an NREL-created collaborative visualization protocol designed to tie disparate software packages together to share in a visualization session[4, 5]. It provides a backbone for the composition of visualization tools and consists of a scene synchronization portion, a structured remote procedure call component, and alternate data delivery/representation portions. This protocol allows us to sidestep platform limitations of existing software.

For this work, a NOODLES server designed to serve arbitrary meshes (not specifically built for this study) published the study geometry, allowing arbitrary clients to subscribe to the visualization session.

4.2 Platforms

The first client was the immersive visualization suite at NREL. The suite is a custom-built tracked stereographic immersive virtual reality (VR) environment, comprising six projectors painting two surfaces. Recently upgraded to handle the demanding workloads of exascale and other next-generation data sources, each projector provides a 4k signal, and is backed by a high-performance render server—96 CPU cores, 2.0 TB of RAM, and six A6000 GPUS.

The immersive VR space provides an ideal framework to explore these data. Navigational difficulties common on a laptop or desktop are reduced to simple physical motion. likewise, the complex contours of these isosurfaces are easier to observe through intuitive head motion [9].

The immersive space uses a custom client application to connect to the NOODLES server, synchronize the scene, and represent that scene in the home-grown graphics engine that powers the uncommon rendering requirements of the immersive environment. Figure 1 shows the display in action.

The next platform was a consumer-grade tablet. Year after year, consumer devices continue to deliver higher performance with greater resources. These devices proliferate across the desks of researchers, and provide new opportunities for effective visualization. To demonstrate this, we used a 5th generation iPad Pro, with an M1 processor and 8 GB of RAM. Due to the restricted software environment, we used web technologies for the tablet client; Three.JS[1] powers the full client-side rendering; no server-side rendering offload is used.

Figure 5 shows the platform in use. While still a prototype and using non-native code, the display is interactive and provides the benefits of touch interaction in a smaller form factor.

As NOODLES facilitates collaboration, these two devices could be used simultaneously in the same session, with both platforms interacting, changing independent viewpoints simultaneously.

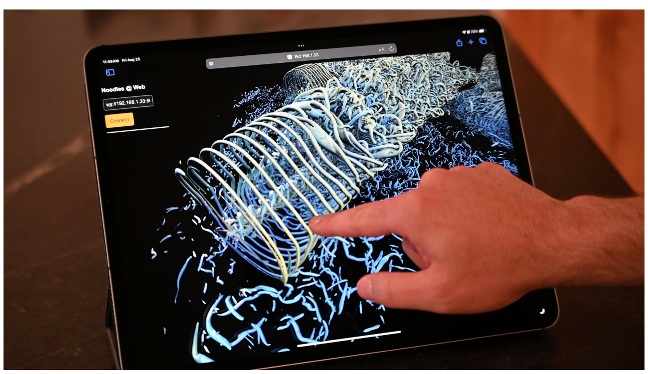

Figure 5: Example of the neutral ABL study on a consumer-grade tablet.

5 CONCLUSION

We demonstrated the ExaWind solver with a four-turbine wind farm, under two different ABL scenarios. We show the differences of fidelity between modern solvers and previous-generation solver artifacts. We explore the data in NREL's immersive space and on tablets using the NOODLES collaborative protocol.

For future work in the simulation domain, the next steps involve integrating AMR-Wind with Weather Research and Forecasting (WRF) inputs, which include surface heat flux, so that the boundary conditions are consistent with mesoscale modeling inputs.

For visualization, future work revolves around fleshing out the protocol specification as well as futhering client and server implementations. A tool using PeleAnalysis and VTK to deliver visualization results directly as a NOODLES server is in alpha stage, and will be used to automate isosurface extraction, lighting, and publication in the future (this tool was not available in time for this work). A native NOODLES client for mobile devices is in development, as well as other clients to support commodity virtual reality and augmented reality devices.

ACKNOWLEDGMENTS

This research was supported by the Exascale Computing Project (17-SC-20-SC), a joint project of the U.S. Department of Energy's Office of Science and National Nuclear Security Administration, responsible for delivering a capable exascale ecosystem, including software, applications, and hardware technology, to support the nation's exascale computing imperative. A portion of this research was supported by the U.S. Department of Energy Wind Energy Technologies Office.

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. This work was authored in collaboration with Sandia National Laboratories and Lawrence Berkeley National Laboratory.

A portion of the 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.

REFERENCES

[1] [n. d.]. ThreeJS. https://threejs.org.
[2] Utkarsh Ayachit. 2015. *The ParaView Guide: A Parallel Visualization Application*. Kitware.
[3] Michael Brazell, Shreyas Ananthan, Ganesh Vijayakumar, Lawrence Cheung, Michael Sprague, ExaWind Exascale Computing Project Team, and High Fidelity Modeling Project Team. 2021. AMR-Wind: Adaptive mesh-refinement for atmospheric-boundary-layer wind energy simulations. In *APS Division of Fluid Dynamics Meeting Abstracts (APS Meeting Abstracts)*. Article T29.007, T29.007 pages.
[4] Nicholas Brunhart-Lupo. 2022. NOODLES: Collaborative Visualization Protocol. https://github.com/InsightCenterNoodles.
[5] Nicholas Brunhart-Lupo. 2022. Noodles: Cooking Up Collaborative Visualization.
[6] Nicholas Brunhart-Lupo, Brian Bush, Kenny Gruchalla, Kristi Potter, and Steve Smith. 2020. Collaborative Exploration of Scientific Datasets using Immersive and Statistical Visualization. https://doi.org/10.5065/P2JJ-9878
[7] Marc Day, Lucas Esclapez, Bruce Perry, and Emmanuel Motheau. 2023. PeleAnalysis. https://peleanalysis.readthedocs.io/.
[8] P. Fleming, J. Annoni, M. Churchfield, L. A. Martinez-Tossas, K. Gruchalla, M. Lawson, and P. Moriarty. 2018. A simulation study demonstrating the importance of large-scale trailing vortices in wake steering. *Wind Energy Science* 3, 1 (2018), 243–255. https://doi.org/10.5194/wes-3-243-2018
[9] Kenny Gruchalla and Nicholas Brunhart-Lupo. 2019. The Utility of Virtual Reality for Science and Engineering. In *VR Developer Gems*, William R. Sherman (Ed.). Taylor Francis, Chapter 21, 383–402. https://doi.org/10.1201/b21598-21
[10] U. Hassan, G.J. Taylor, and A.D. Garrad. 1988. The dynamic response of wind turbines operating in a wake flow. *Journal of Wind Engineering and Industrial Aerodynamics* 27, 1 (1988), 113–126. https://doi.org/10.1016/0167-6105(88)90028-1
[11] Pankaj K. Jha and Sven Schmitz. 2018. Actuator curve embedding – an advanced actuator line model. *Journal of Fluid Mechanics* 834 (2018), R2. https://doi.org/10.1017/jfm.2017.793
[12] Michael Sprague, S Ananthan, G Vijayakumar, and M Robinson. 2020. ExaWind: A multifidelity modeling and simulation environment for wind energy. *Journal of Physics: Conference Series* 1452 (01 2020), 012071. https://doi.org/10.1088/1742-6596/1452/1/012071

[13] Michael A. Sprague, Stanislav Boldyrev, Paul Fischer, Ray Grout, William I. Gustafson, Jr., and Robert Moser. 2017. Turbulent Flow Simulation at the Exascale: Opportunities and Challenges Workshop: August 4-5, 2015, Washington, D.C. (1 2017). https://doi.org/10.2172/1338668
[14] Dan Zhao, Nuomin Han, Ernest Goh, John Cater, and Arne Reinecke. 2019. Chapter 7 - Aerodynamics of horizontal axis wind turbines and wind farms. In *Wind Turbines and Aerodynamics Energy Harvesters*, Dan Zhao, Nuomin Han, Ernest Goh, John Cater, and Arne Reinecke (Eds.). Academic Press, 431–461. https://doi.org/10.1016/B978-0-12-817135-6.00007-7

Received 30 August 2023; accepted 28 September 2023