High order in Exawind

SAND2019-9280PE

Robert Knaus, Sandia National Laboratories

6 August 2019

EXASCALE COMPUTING PROJECT

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.

Exawind project overview

Goals/motivation for predictive simulations

- Advance our fundamental understanding of the flow physics governing whole wind plants
- Predict the response of wind farms to a wide range of atmospheric conditions

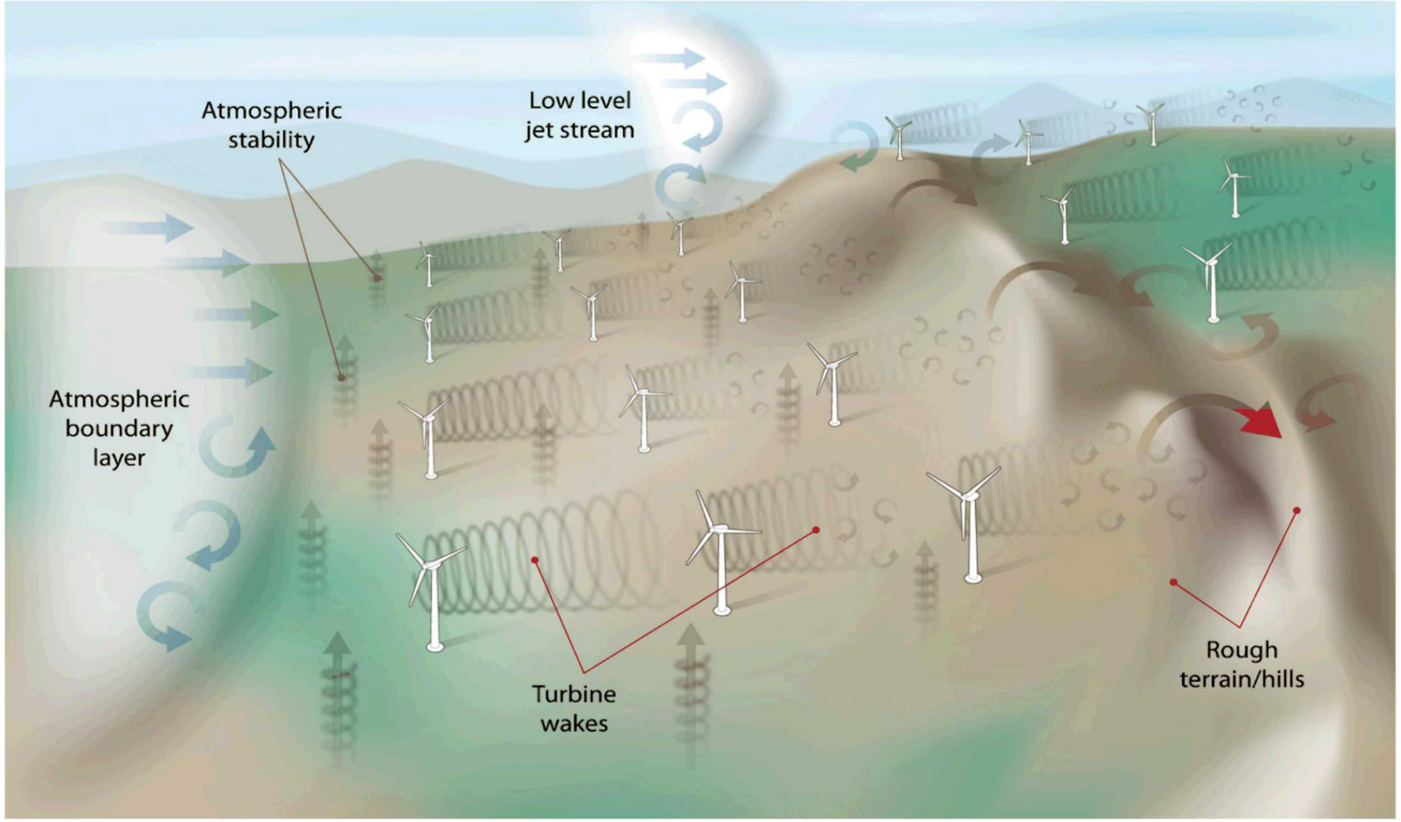

Nalu-Wind

- Open source: https://github.com/Exawind/nalu-wind
  - Builds with Spack
- Unstructured finite volume, low-Mach flow solver
  - 2nd-order node-centered, low-dissipation FV scheme
  - **Arbitrary order accurate element-based continuous finite volume scheme**
  - Fully implicit
  - Mixed-order interfaces
- C++. Built on "Trilinos"
  - Kokkos for performance portability
  - Tpetra, MueLu, ShyLU, Ifpack2, Belos packages
    - Also support hypre
  - STK package
    - IOSS for IO

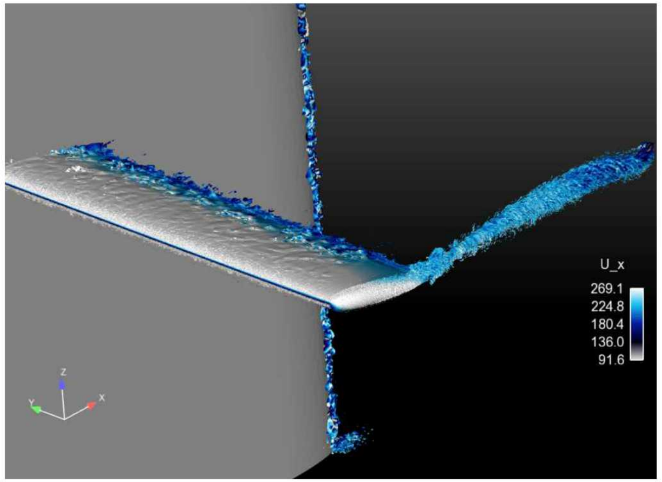

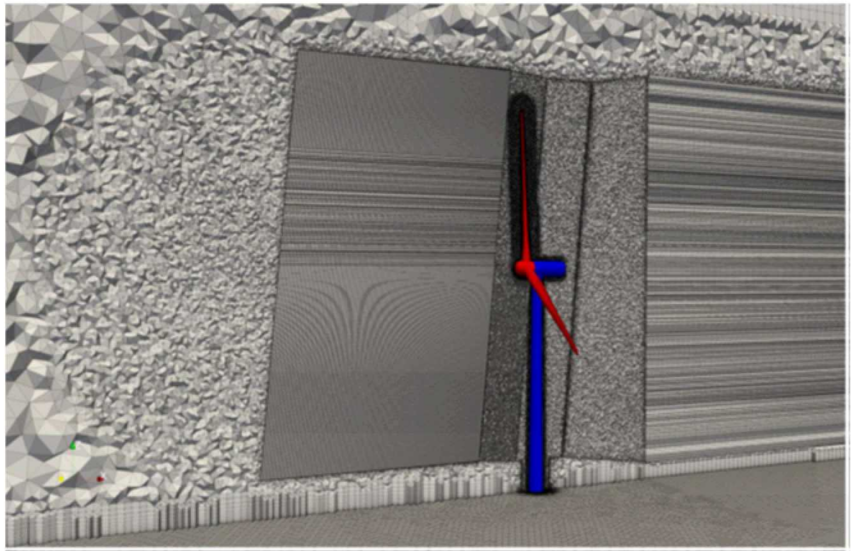

Nalu-Wind physics capabilities

This image cannot currently be displayed.

- Atmospheric boundary layer modeling
  - Monin-Obukhov wall models
  - Boussinesq, Coriolis forcing

- Full turbine modeling
  - Sliding mesh, overset technologies
  - Hybrid RANS-LES models
    - ”TAMS” model being developed at UT Austin

- Actuator line modeling
  - Coupling with the OpenFAST code

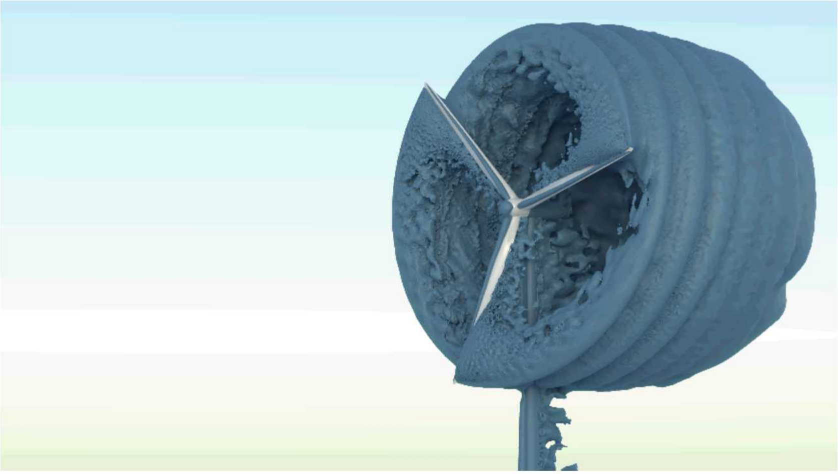

EXASCALE COMPUTING PROJECT

Control volume finite element

- Node-centered FV (EBVC) production discretization in nalu-wind
  - Design is semi-flexible in terms of discretization
  - Linear CVFEM traditionally the "main" discretization
- Basic idea for CVFEM: define a test space of indicator functions
  - Other names "finite volume element", "covolume method", etc.

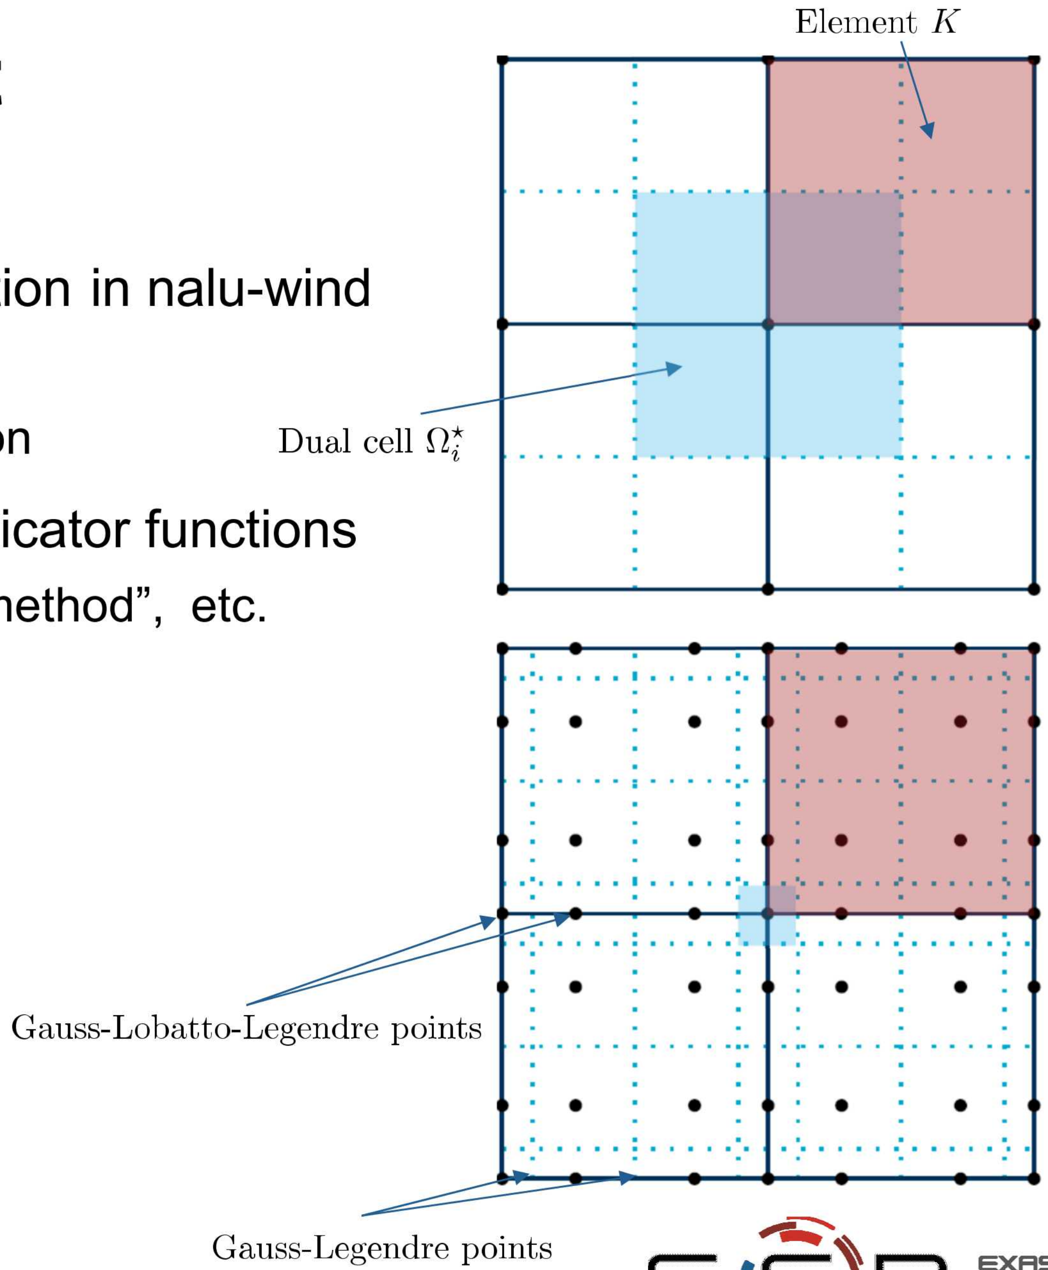

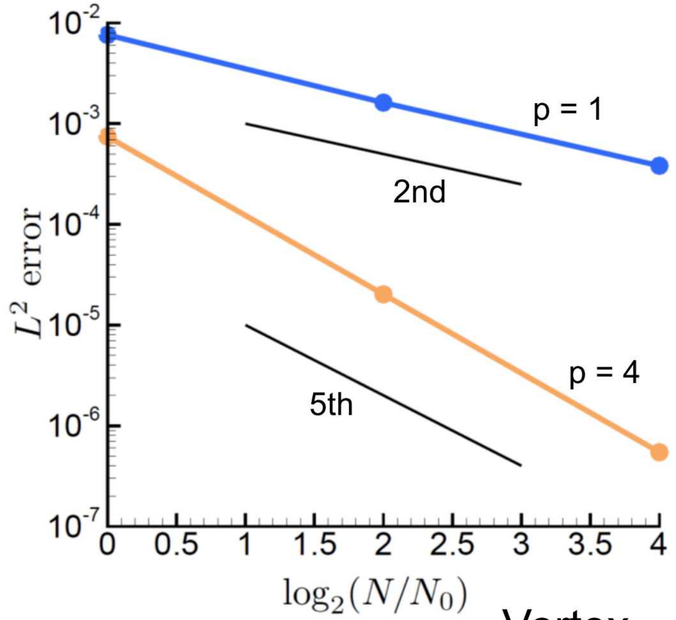

Vortex transport

CVFEM Operators

- In one-dimension:

- Operators:

$$\tilde{I}_{ij} = \ell_j\left(\hat{x}_i^{\mathrm{GL}}\right), \quad \tilde{D}_{ij} = \ell_j'\left(\hat{x}_i^{\mathrm{GL}}\right)$$
$$W_{ij}^{-1} = h_j\left(\hat{x}_i^{\mathrm{GLL}}\right), \quad D_{ij} = \ell_j'\left(\hat{x}_i^{\mathrm{GLL}}\right)$$
$$\widetilde{\Delta} = \begin{cases} -1 & i = j \\ +1 & i = j+1 \\ 0 & \text{otherwise} \end{cases}$$

where $h_j(\hat{x}) = -\sum_{k=1}^{j} d_k'(\hat{x})$ and d_k is the k-th Lagrange interpolant between $\{-1, \hat{x}_1^{\mathrm{GL}}, \ldots, \hat{x}_p^{\mathrm{GL}}, +1\}$
This construction is done so $\int_{-1}^{1} \mathrm{d}x \, 1_{\Omega_i^\star}(x) h_j(x) = \delta_{ij}$. ℓ_k is the k-th interpolant between $\{\hat{x}_j^{\mathrm{GLL}}\}_{j=1}^{p+1}$

- For a tensor-product element in 3D, we have

$$S^{\hat{x}} = \tilde{\Delta} \otimes W \otimes W \qquad \tilde{I}^{\hat{x}} = \tilde{I} \otimes I \otimes I \qquad \mathbf{D}^{\hat{x}} = \begin{pmatrix} \tilde{D} \otimes I \otimes I & \tilde{I} \otimes D \otimes I & \tilde{I} \otimes I \otimes D \end{pmatrix}$$
$$S^{\hat{y}} = W \otimes \tilde{\Delta} \otimes W \qquad \tilde{I}^{\hat{y}} = I \otimes \tilde{I} \otimes I \qquad \mathbf{D}^{\hat{y}} = \begin{pmatrix} D \otimes \tilde{I} \otimes I & I \otimes \tilde{D} \otimes I & I \otimes \tilde{I} \otimes D \end{pmatrix}$$
$$S^{\hat{z}} = W \otimes W \otimes \tilde{\Delta} \qquad \tilde{I}^{\hat{z}} = I \otimes I \otimes \tilde{I} \qquad \mathbf{D}^{\hat{z}} = \begin{pmatrix} D \otimes I \otimes \tilde{I} & I \otimes D \otimes \tilde{I} & I \otimes I \otimes \tilde{D} \end{pmatrix}$$

Poisson Equation

- Over a dual cell:

$$\Delta u = f \Rightarrow \int_{\partial \Omega_i^\star} \nabla u \cdot n \,\mathrm{d}S = \int_{\Omega_i^\star} f \mathrm{d}x$$

- This becomes on an element,

$$\sum_i S^{\hat{x}_i} \left[\mathbf{G}^{\hat{x}_i} \circ \mathbf{D}^{\hat{x}_i} u_e \right] = (W \otimes W \otimes W) \left[\det \mathbf{J} \circ f_e \right]$$

$$\mathbf{G}^{\hat{x}}_{mjk} = \det \mathbf{J}_{mjk} \left(\mathbf{J}^T \mathbf{J} \right)^{-1}_{mjk} \mathbf{e}_{\hat{x}}$$
$$1 \le m \le p \quad 1 \le j, k \le p + 1$$

- Use explicit SIMD types and sum factorization for apply
- Uses Trilinos-Belos package for the Krylov solver implementation
  - No matrix required

Matrix-free Preconditioning

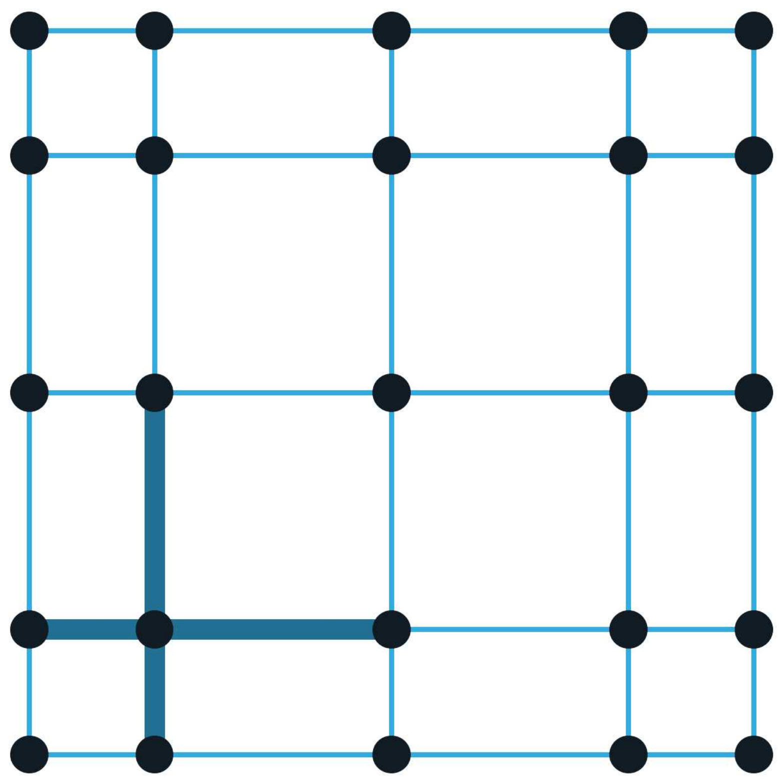

- Just Jacobi for momentum
  - Does OK at low (~1) Courant numbers

- Continuity uses AMG with a sparsified “edge” Laplacian
  - Sparse system is passed to “MueLu", which creates the preconditioner

Case (144^3 dofs)	Avg Continuity GMRES iterations
Order 1 full system	5
Order 8 no preconditioner	129
Order 8 full sparsified system	15
Order 8 edge sparsified system	9

Re=1600 Taylor Green Vortex Breakdown

$$u(t=0,x,y,z) = +u_0 \sin(x)\cos(y)\cos(z)$$
$$v(t=0,x,y,z) = -u_0 \cos(x)\sin(y)\cos(z)$$
$$w(t=0,x,y,z) = 0$$

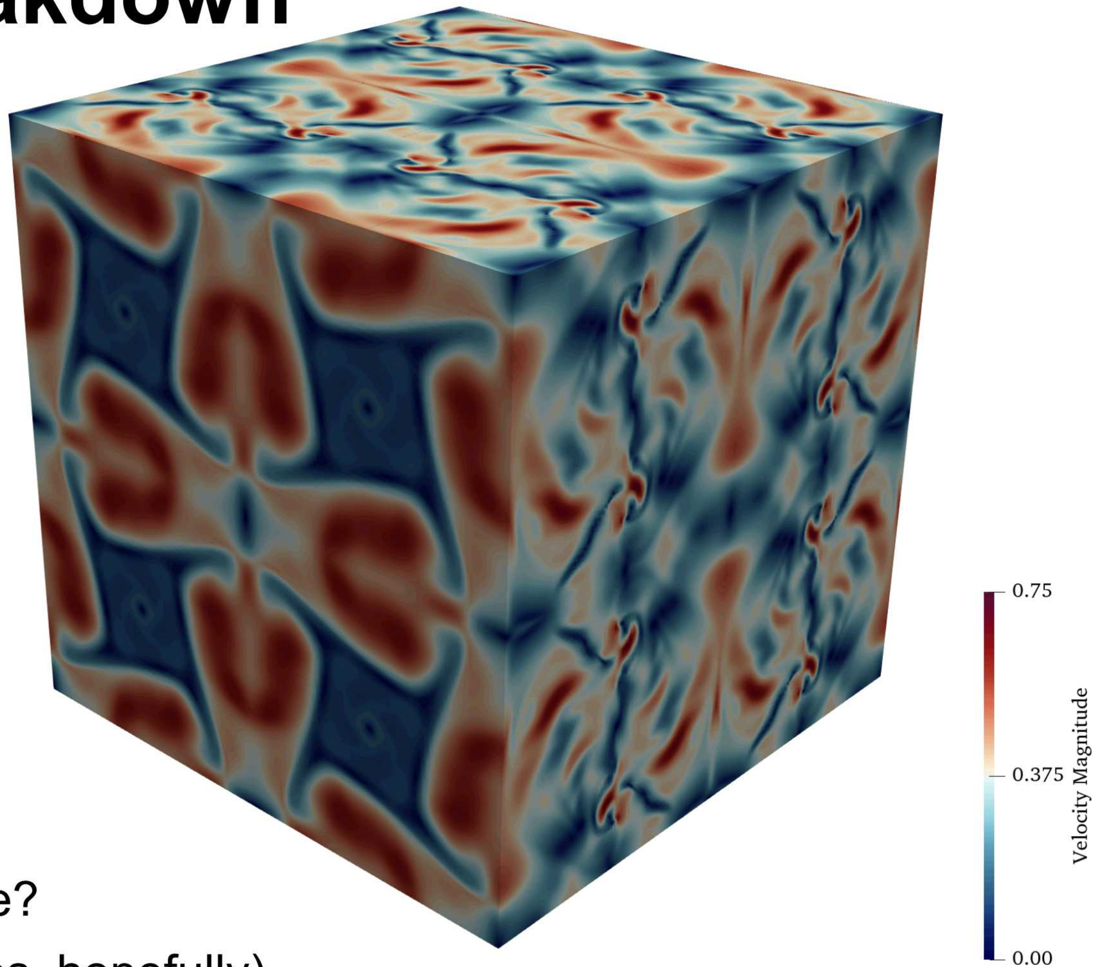

- Standard, turbulent-like test case
  - Just trigonometric functions to set up
  - Three resolutions $(96/p)^3$, $(144/p)^3$, and $(192/p)^3$
- Cases are run underresolved (256^3 for DNS)
  - How does the scheme behave outside the asymptotic range?
  - LES will be underresolved (in localized spatiotemporal zones, hopefully)

Results

- High order scheme captures initial laminar breakdown and scale generation and behaves “OK” when resolution insufficient
  - Need some stabilization beyond the pressure stabilization

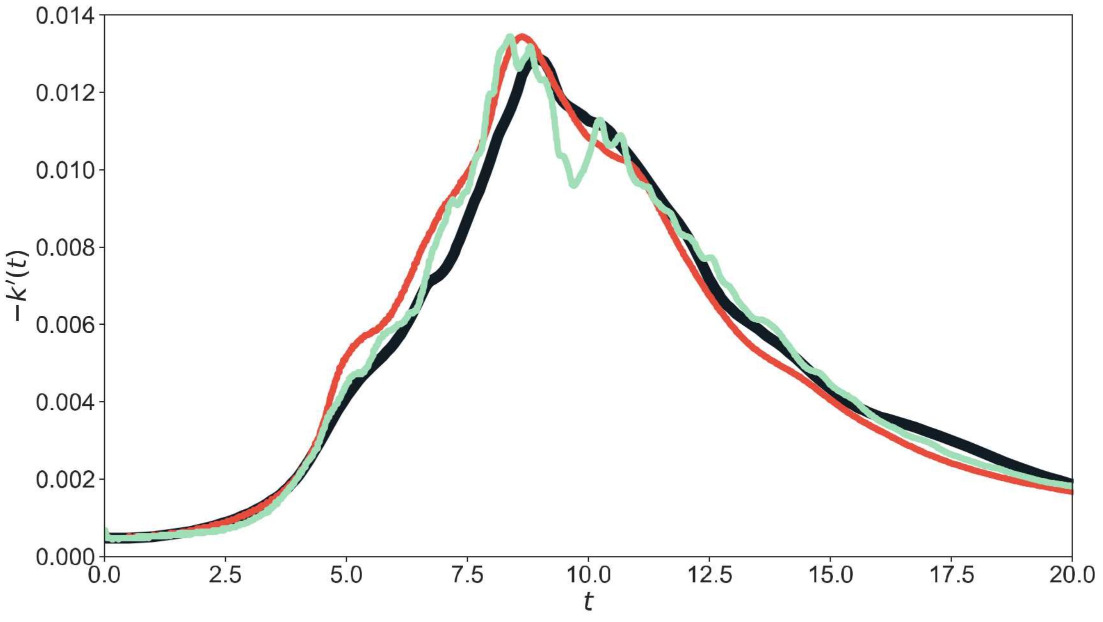

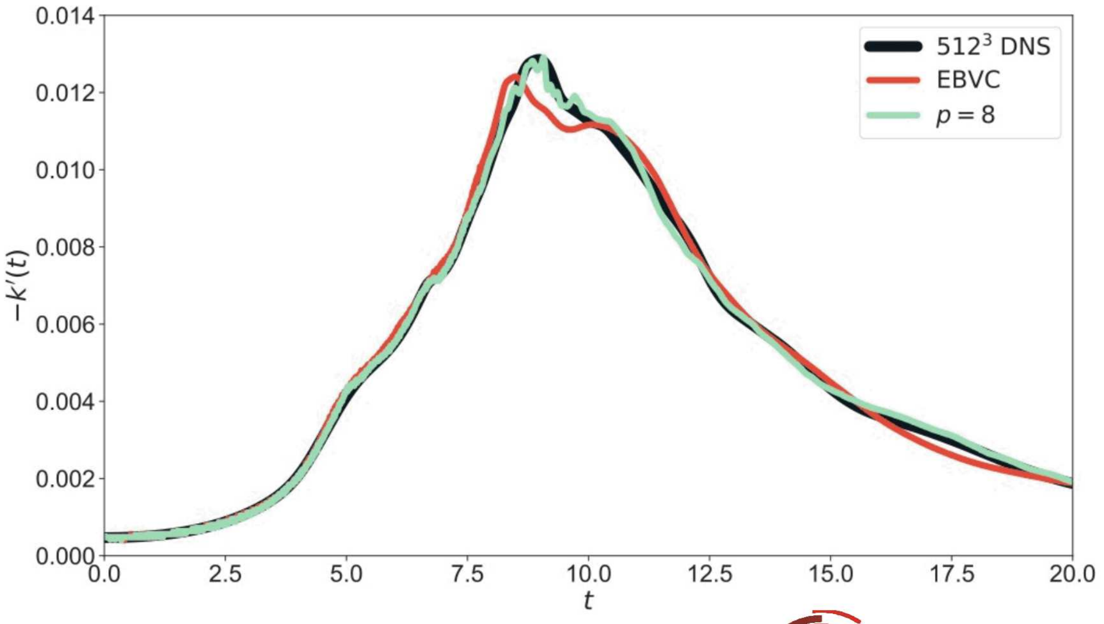

EXASCALE COMPUTING PROJECT

Overall

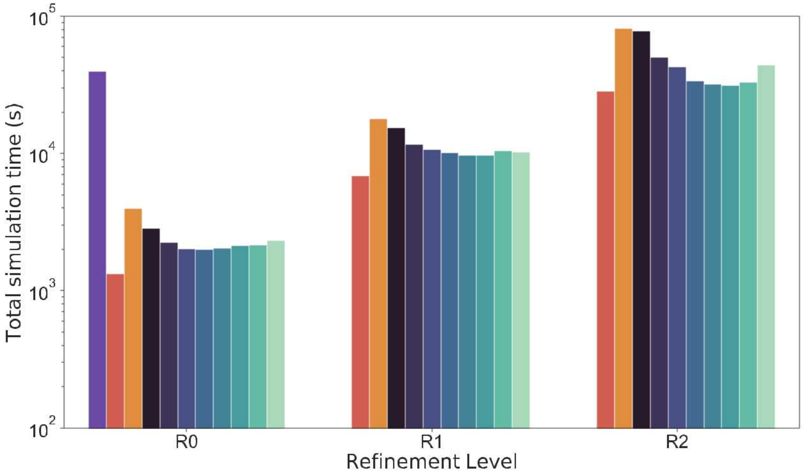

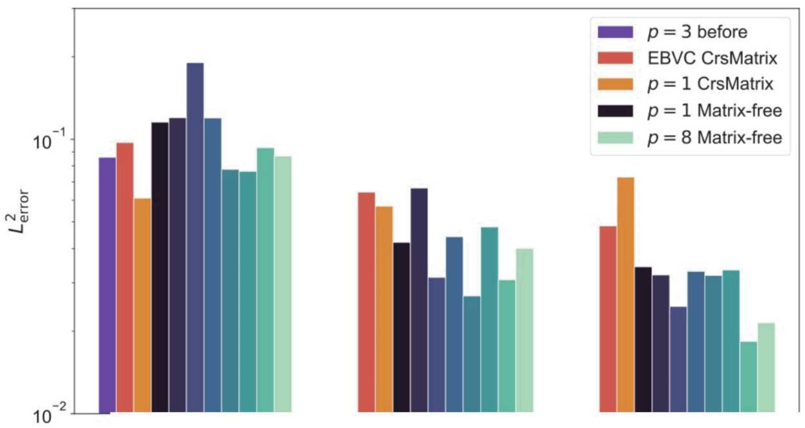

- Computational cost is roughly independent of polynomial order for polynomial orders 4-7
  - Slower than the edge-based finite volume
  - Slow down is due to extra equation essentially
  - Still faster than the default element scheme
- Maybe some benefit in marginally underresolved case
  - Resolved region performs very well, affecting rest of result

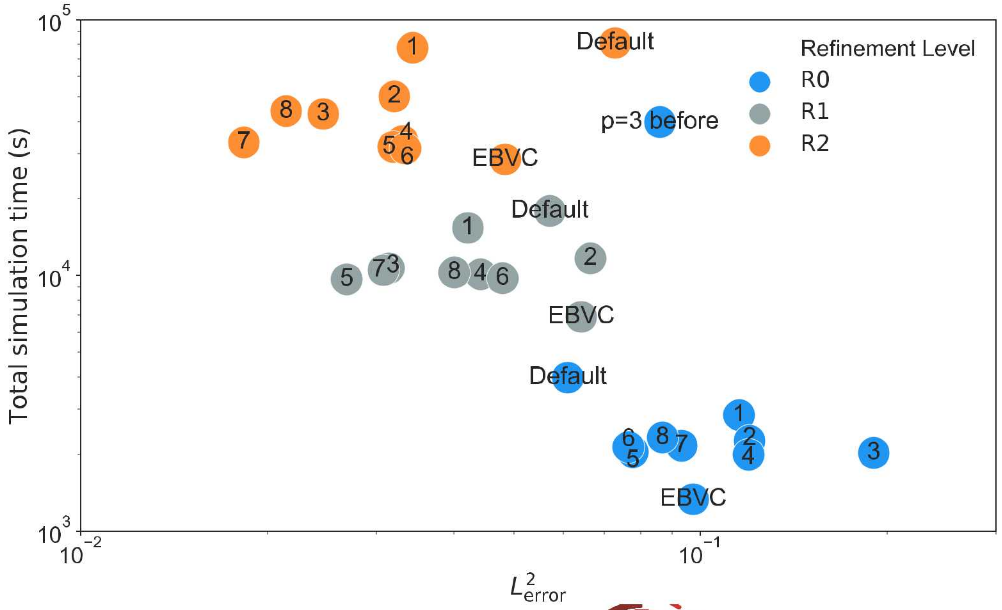

Wall-resolved channel flow LES

- Kim, Moser, Mansour (1999) $\mathrm{Re}_\tau = 590$ case
- "Wall adapting local eddy-viscosity" model of Nicoud and Ducros (1999)
- Run 64x64x48 nodes up to 192x192x128 in 3/2 increments
  - DNS: 384x257x384
  - Constant pressure gradient

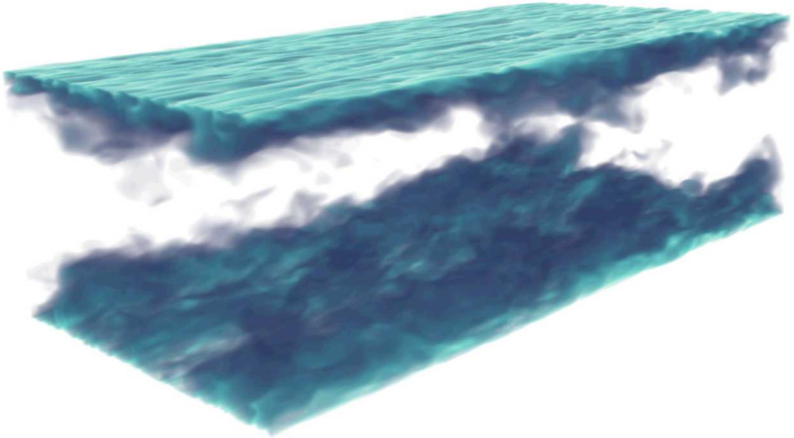

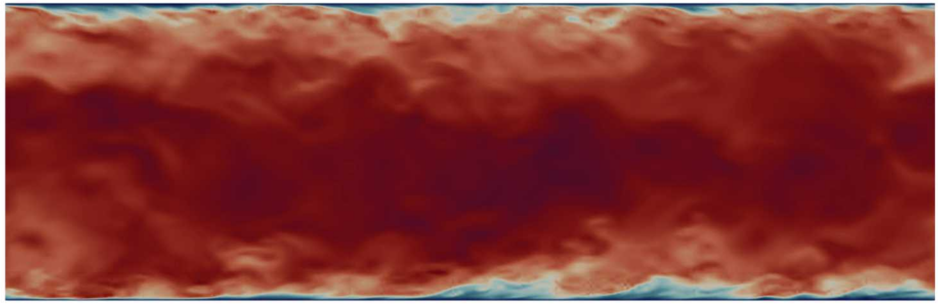

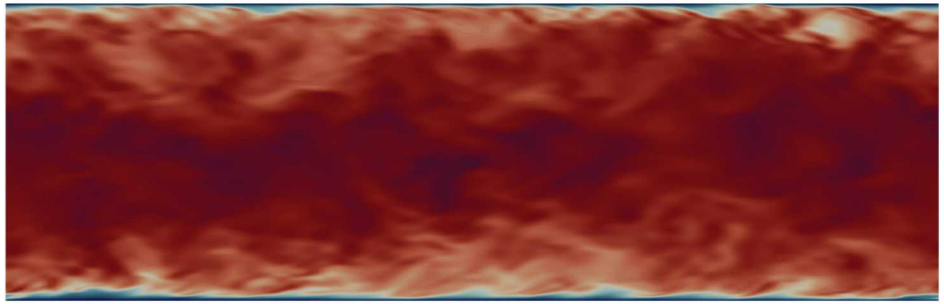

$p = 7$

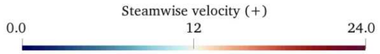

$EBVC$

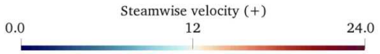

Effect of resolution on predictions

Mean

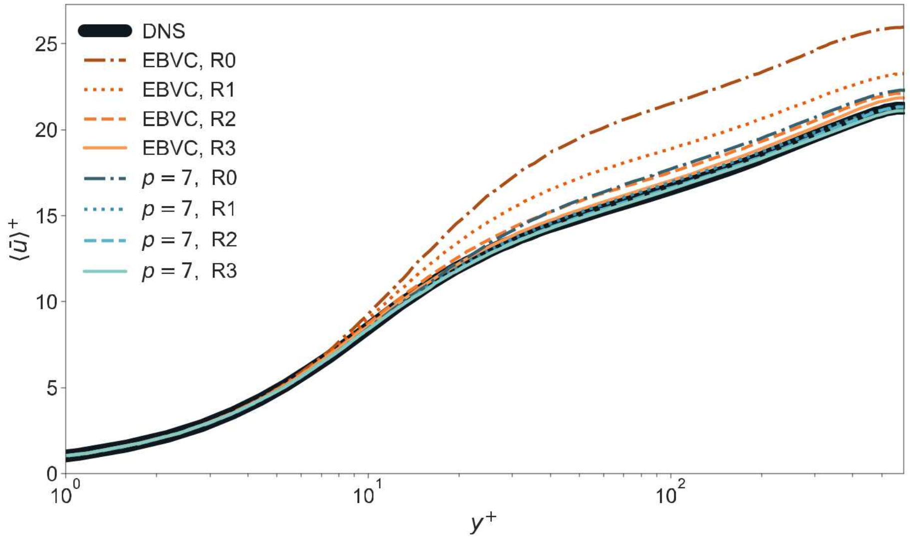

Resolved Reynolds stress

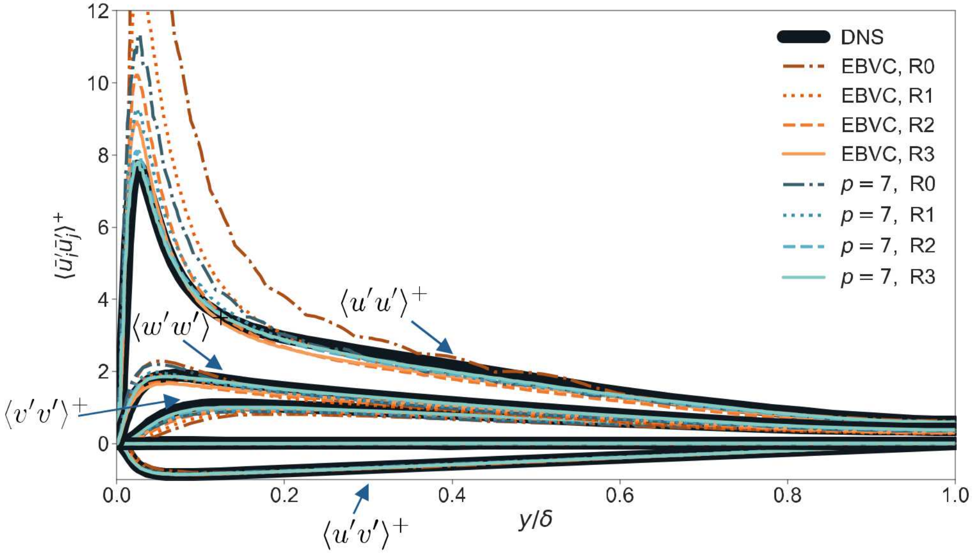

Estimated turbulent shear stress

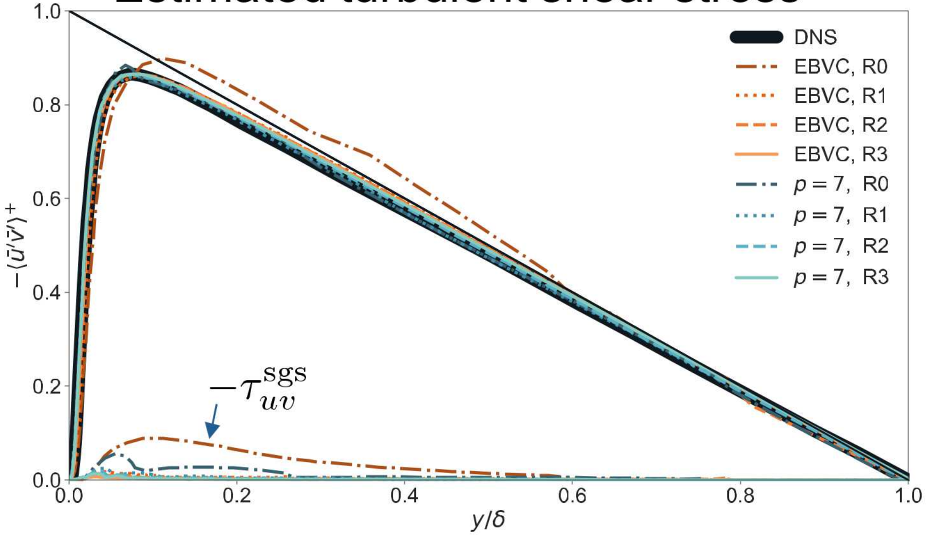

EXASCALE COMPUTING PROJECT

Effect of polynomial order on mean velocity prediction

R0

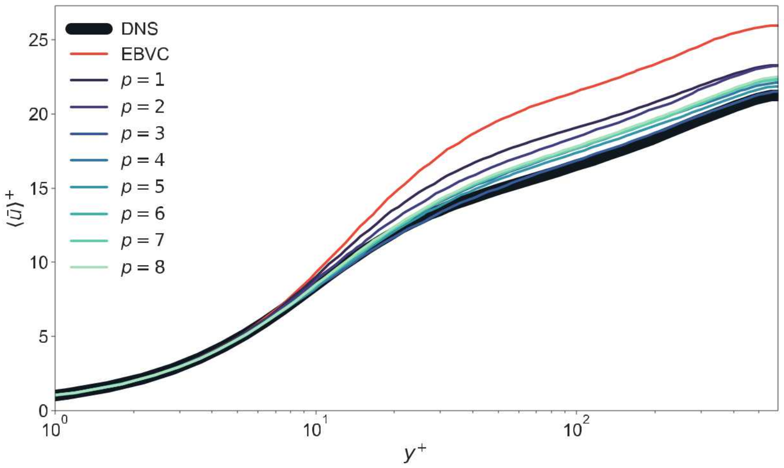

R1

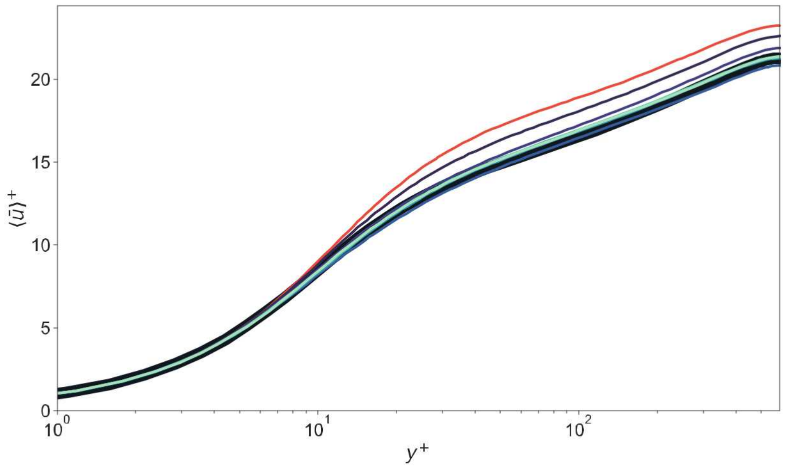

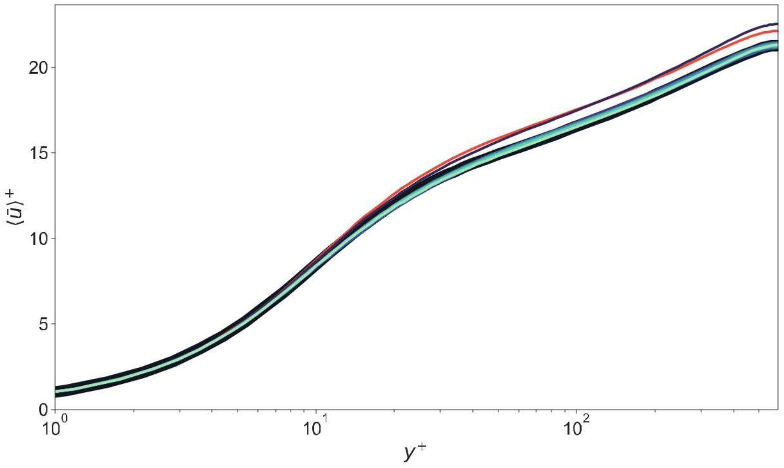

R2

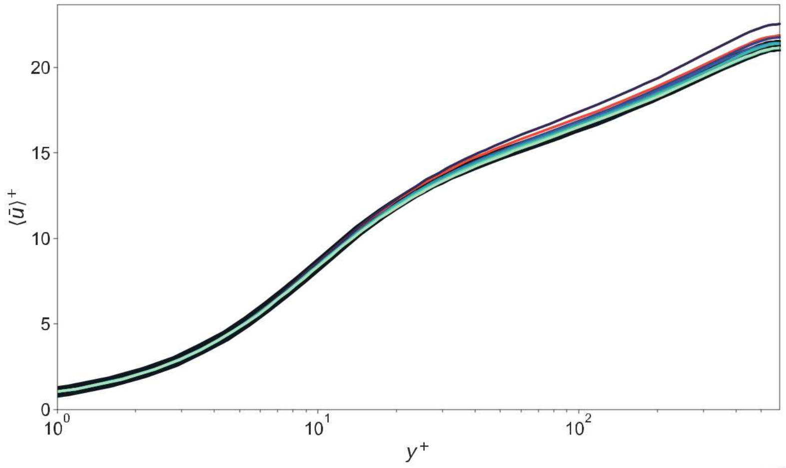

R3

Effect of order on the resolved Reynolds stress

R0

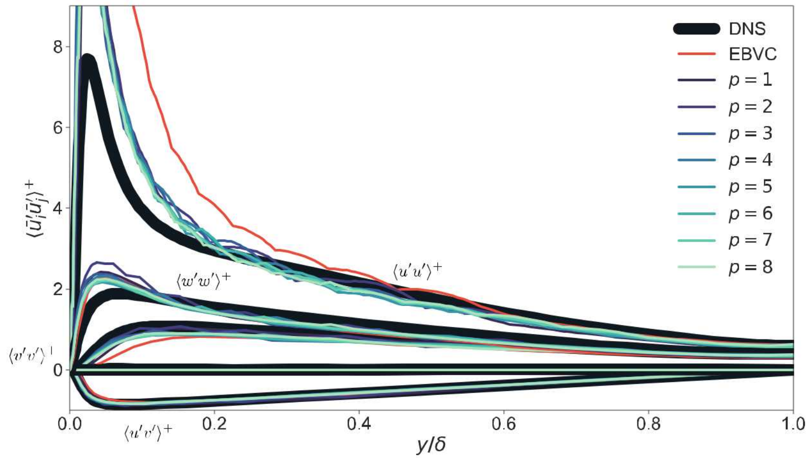

R1

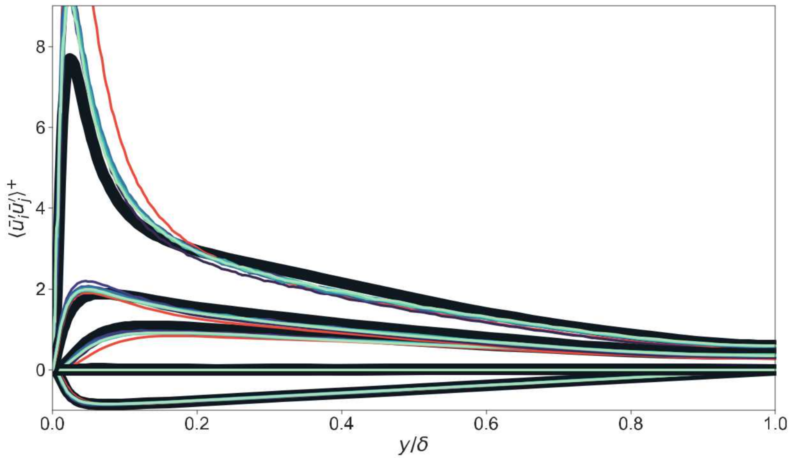

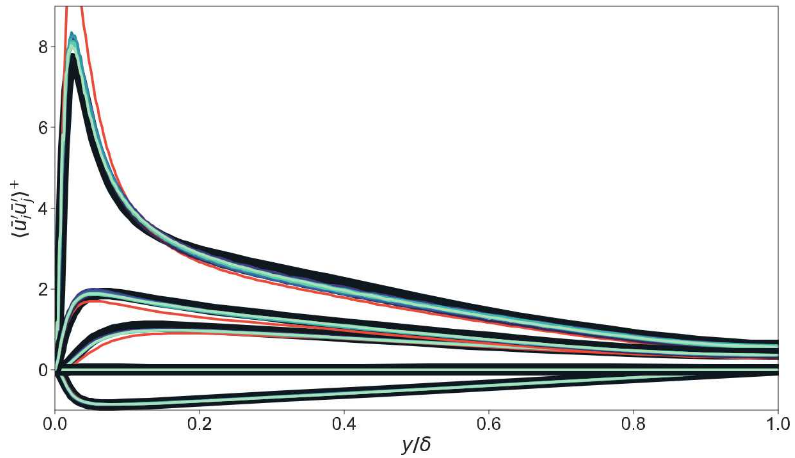

R2

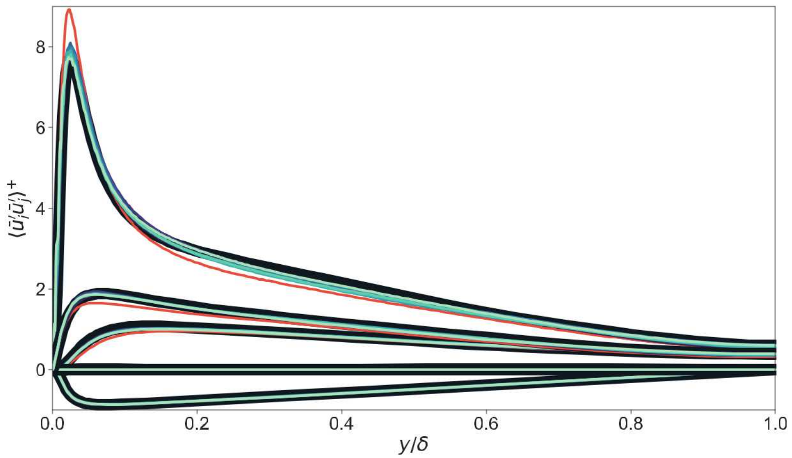

R3

Overall

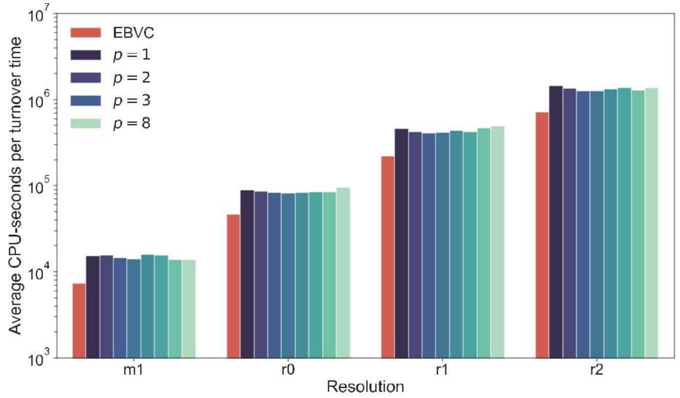

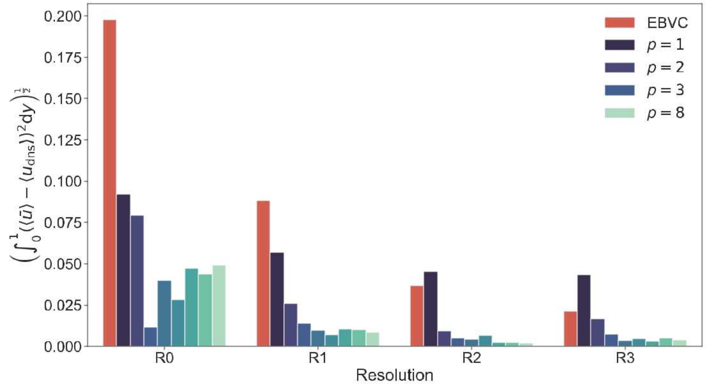

- High order performs better
  - Some variation for lower-high order
  - Turbulence model does well but could be adapted

- EBVC is relatively faster than the Taylor-Green
  - Preconditioner somewhat worse, momentum particularly
  - Partially due to switch to Haswell

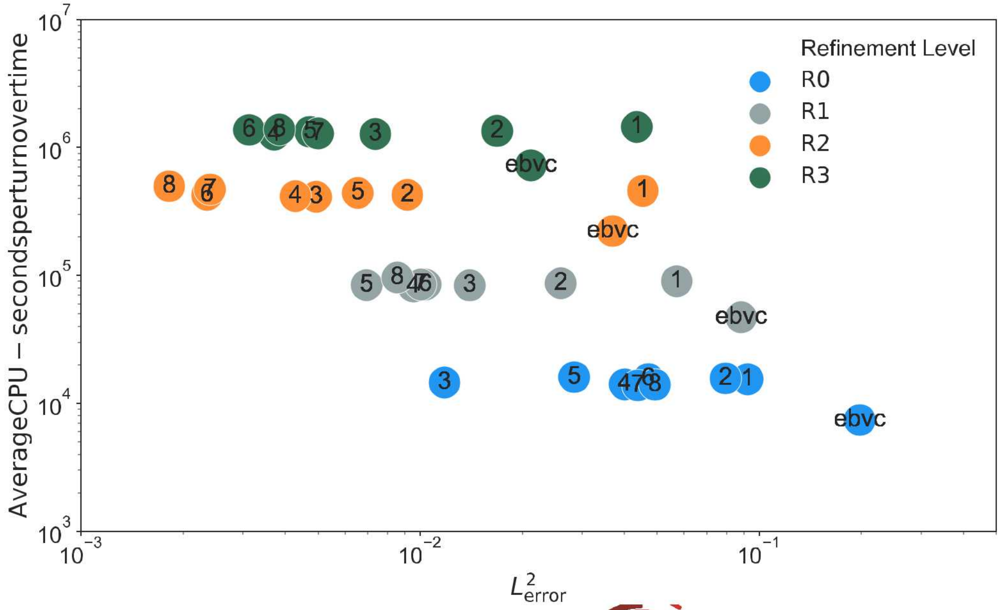

Mixing layer

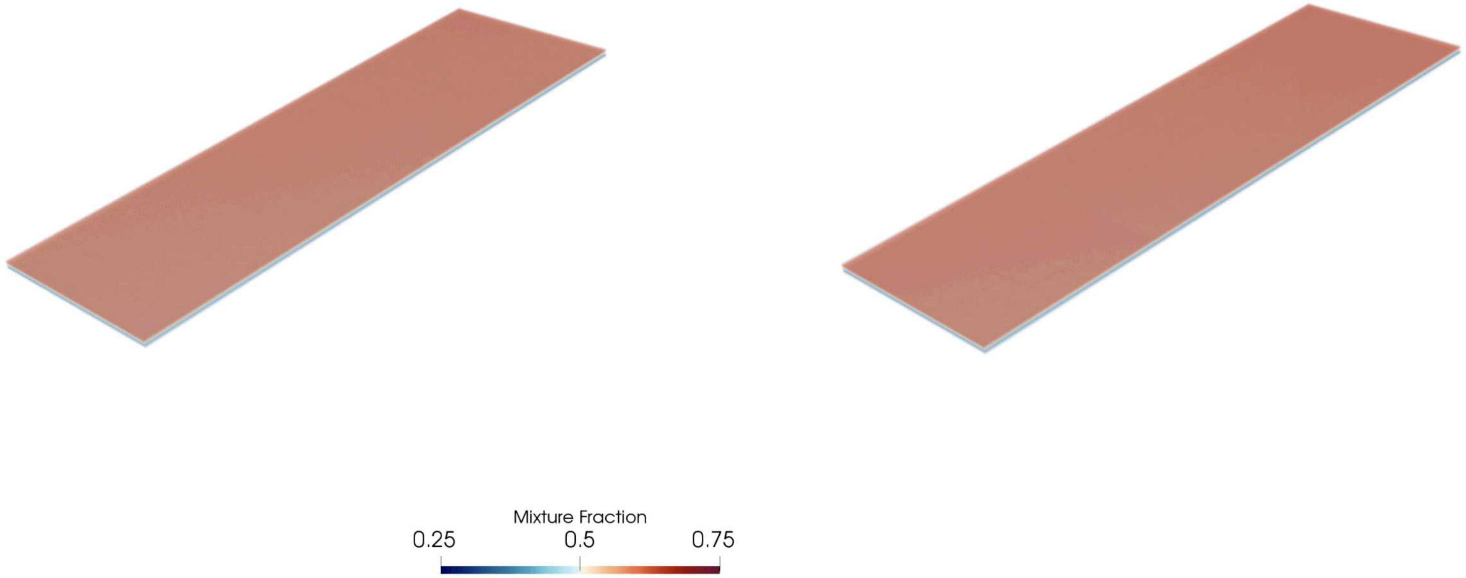

Closing

- Accuracy vs time looks good for LES flows
  - Still places for improvement
- Not working on the GPU
  - Written using Kokkos heavily but still work to do
- Looking at a flow more directly relevant to wind
  - Looking at mixed discretization for near-turbine and off-turbine regions

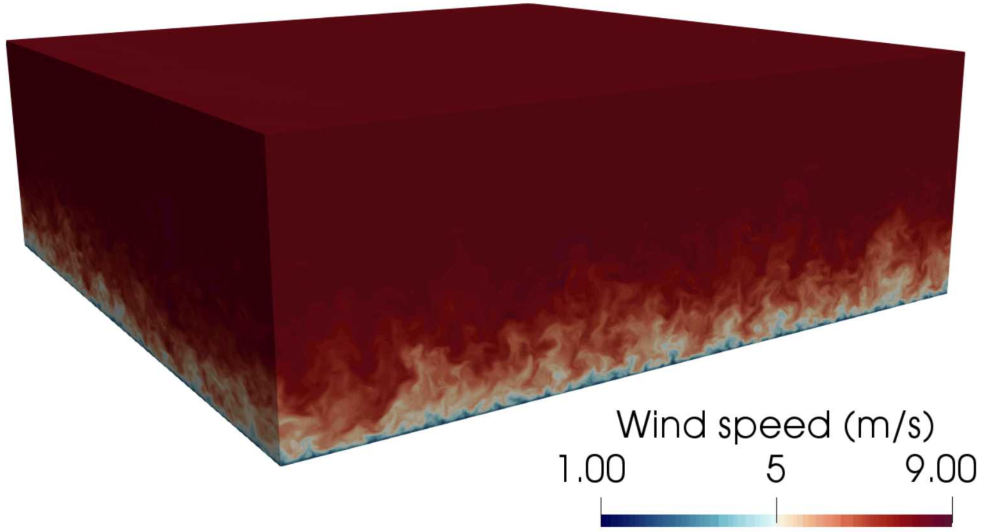

Acknowledgement

- 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.

Citations

- Nicoud, Franck, and Frédéric Ducros. "Subgrid-scale stress modelling based on the square of the velocity gradient tensor." *Flow, turbulence and Combustion* 62.3 (1999): 183-200.
- Moser, Robert D., John Kim, and Nagi N. Mansour. "Direct numerical simulation of turbulent channel flow up to Re τ= 590." *Physics of fluids* 11.4 (1999): 943-945.

EXASCALE COMPUTING PROJECT