SAND2018-7468C

Sandia National Laboratories

The suitability of hybrid meshes for low-Mach large-eddy simulation (LES)

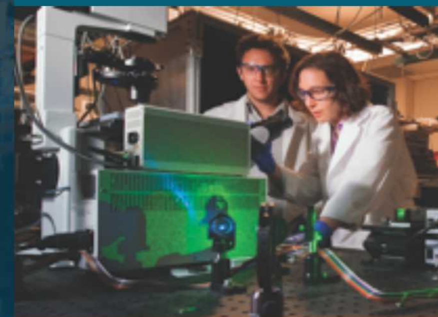

PRESENTED BY

Stefan P. Domino

Computational Thermal and Fluid Mechanics

Sandia National Laboratories SAND2018-TBD

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.

Presentation Objectives

- Overview of the low-Mach application of interest
- Numerical Methods (discretizations and coupling)
- Validation and Verification on general mesh topologies
- Comments on Typical Validation Efforts
- NGP Activities (ExaWind-centric)
- Conclusion

Consider the Abnormal/Thermal Environment

- Characterized by a highly sooting, turbulent, reacting flow with Participating Media Radiation (PMR) and Conjugate Heat Transfer (CHT) mutiphysics coupling

Time-averaged (inset transient)

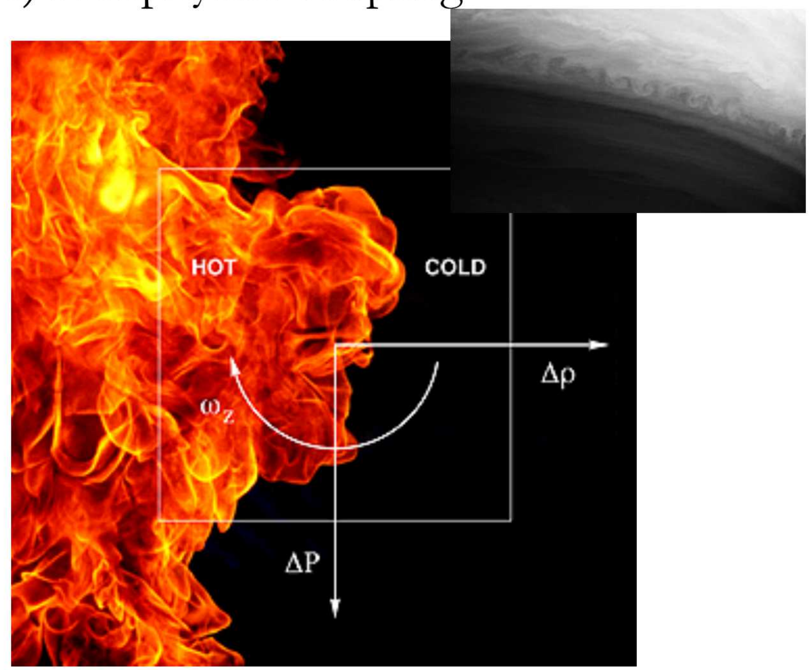

Vorticity generation

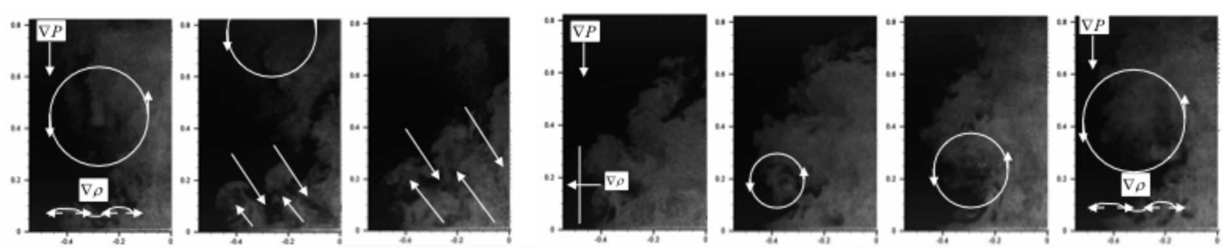

Helium plume puffing cycle

Evolution of a Mindset….. Cross Flow

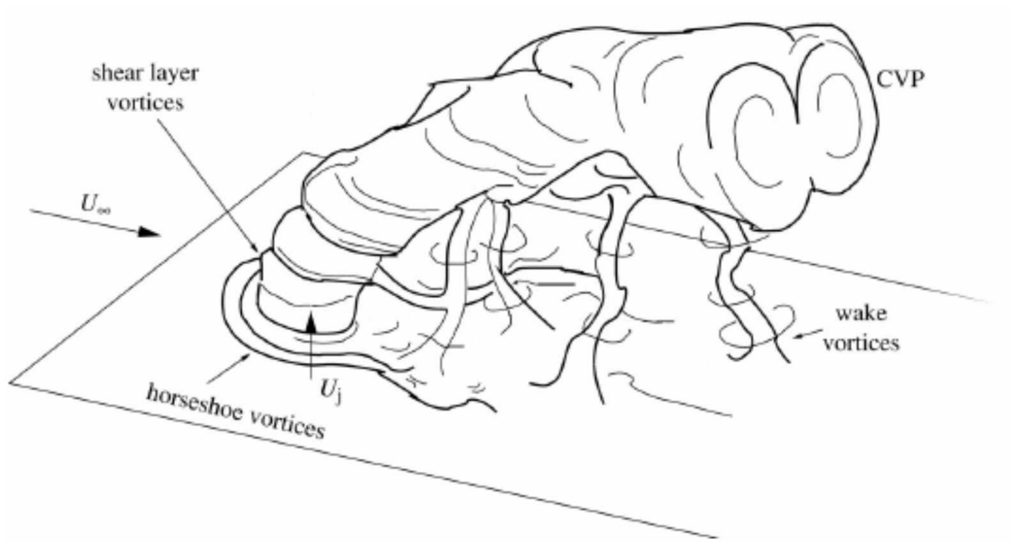

LES of pulsed jet in cross flow; Coussement et al, JFM, 2012

- Conclusion: The inclusion of a cross-flow wind profile couples vorticity of the pool and streamwise momentum which drives the formation of column vortices, increases the importance of mixing and, therefore, convective loads on the object become more important
- Change in mindset: Invest in validation use cases to highlight the importance of fire accident scenarios in the presence of an external momentum field

Ten meter (top) experiment and three meter (bottom) simulation

Evolution of a Mindset…..Whirling-like Flow

Brush fire (Curtin Springs, Australia)

Fire whirls from a 3-meter diameter pool in the Fire Laboratory for Accreditation of Modeling by Experiment, or FLAME, facility at Sandia National Laboratories. (Photo by Richard Simpson; A. Hanlin, lead experimentalist)

Evolution of a Mindset….. Modeling Whirling-like Flow

- Idealized chamber in which swirl is provided by selective wall placement in the experimental design
- Gap varied between 10, 20, and 40 cm
- Objective: Can the onset of swirl be predicted? What is the strength?

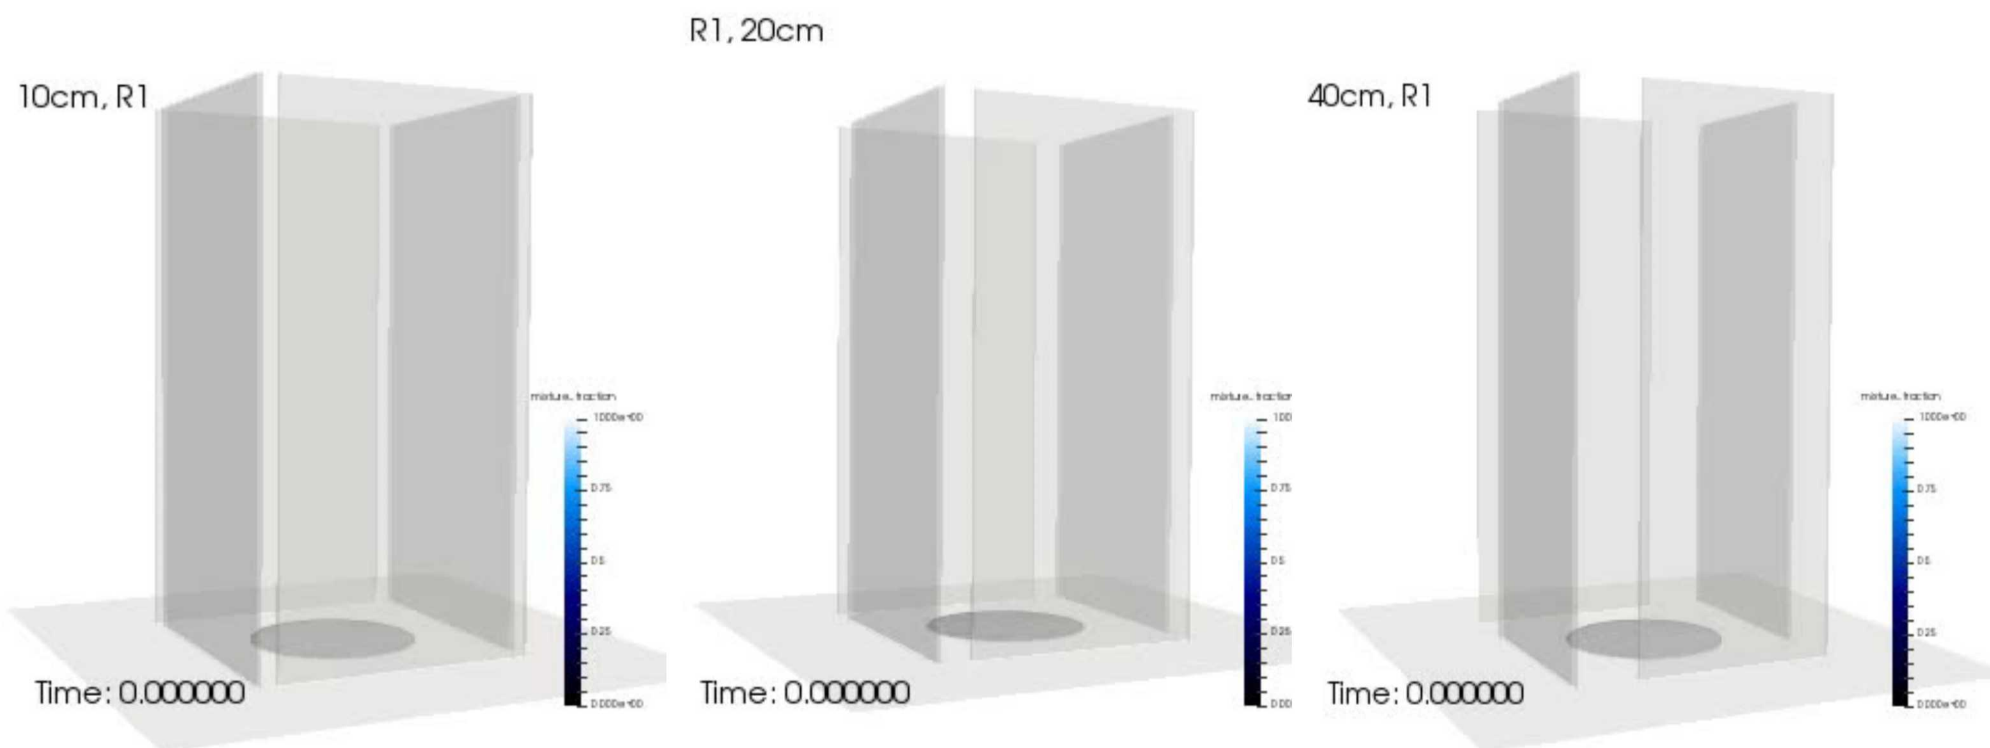

Volume rendered mixture fraction

Disparity in Time and Length Scales, Fire

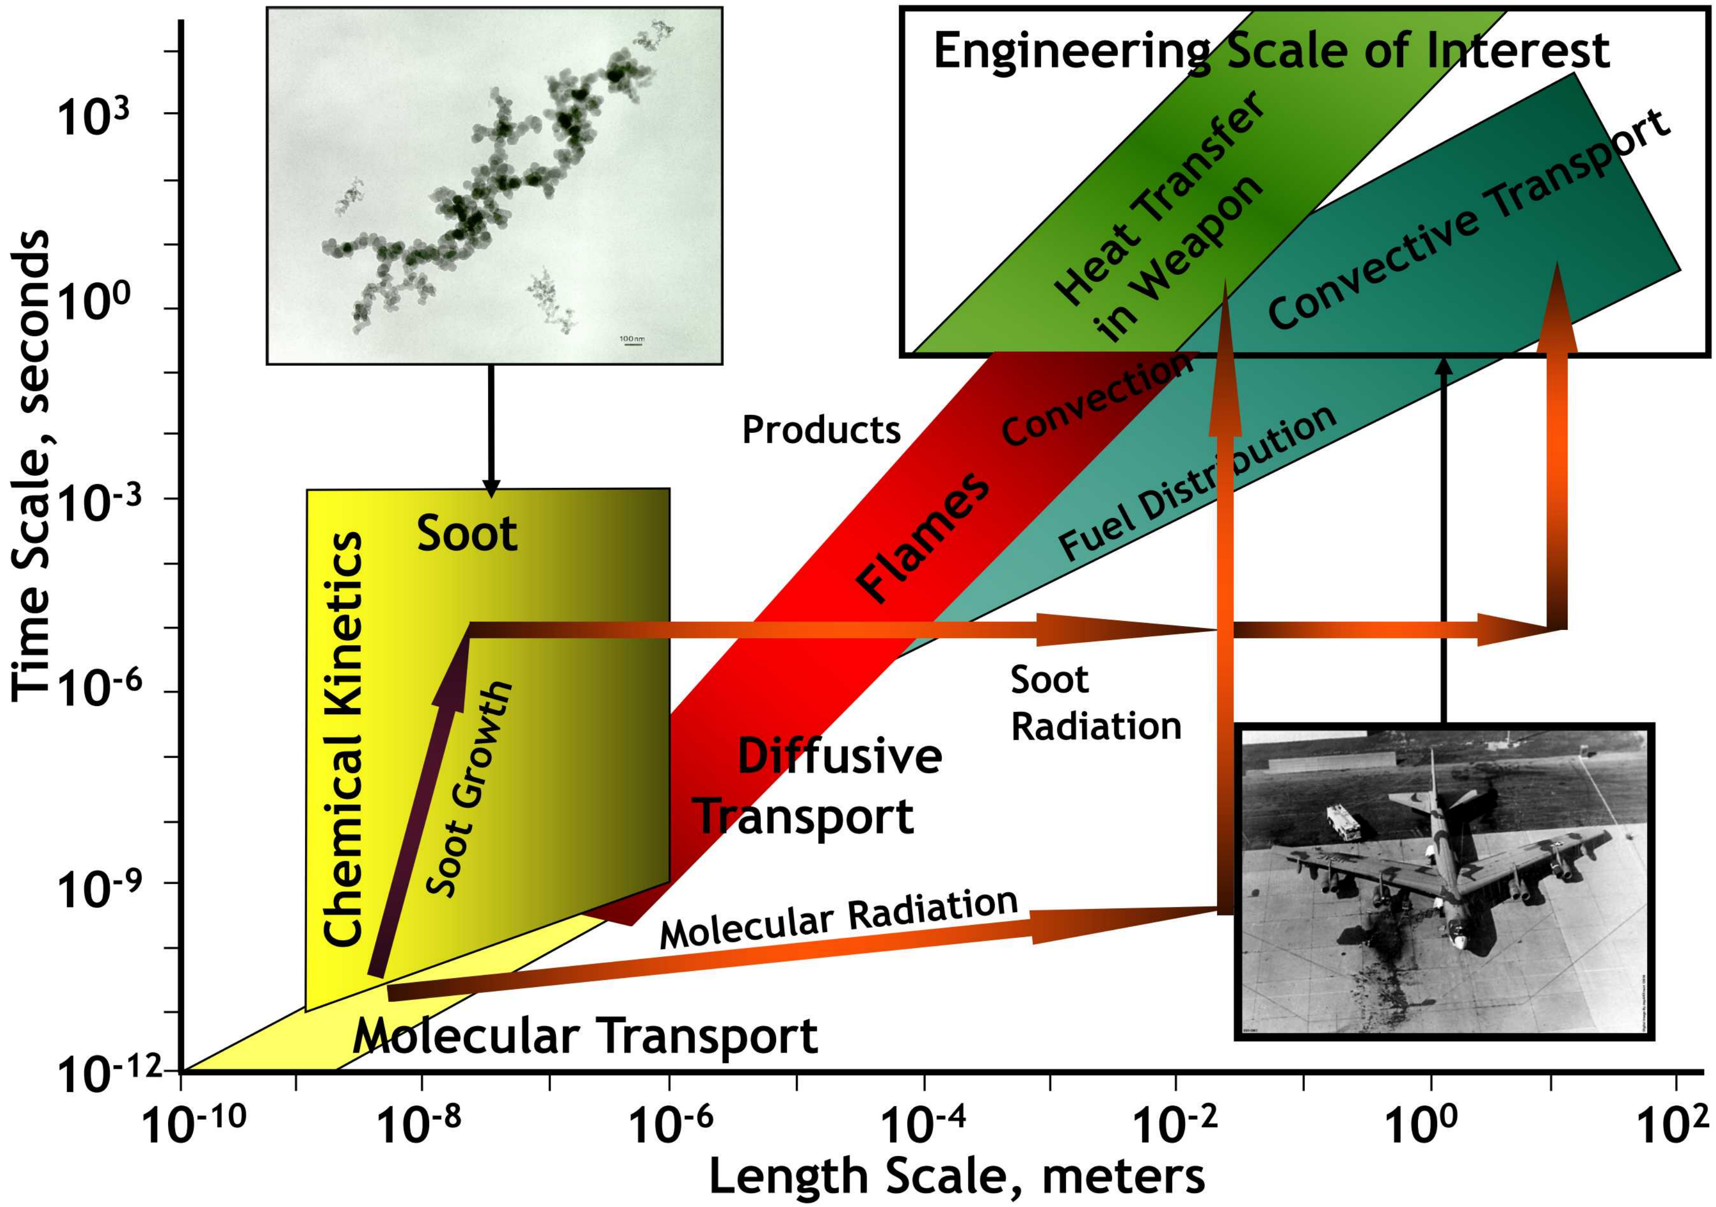

Presentation Objectives

- Overview of the low-Mach application of interest
- Numerical Methods (discretizations and coupling)
- Validation and Verification on general mesh topologies
- Comments on Typical Validation Efforts
- NGP Activities (ExaWind-centric)
- Conclusion

Equal Order Interpolation Edge-Based Vertex-Centered (EBVC) Finite Volume

- All primitives are collocated at the vertices of the elements with equal-order interpolation
- A dual mesh is constructed to obtain flux and volume quadrature locations
- Classic two-state, "L" and "R" approach provides spatially second-order accuracy

▲ Surface quadrature point (area summed to edge)

◆ Volume quadrature point (sub-vol summed to node)

- Same grad-op as CC:

$$\frac{\partial T}{\partial x_j} \approx \frac{(T_R - T_L)n_j}{\|\Delta x_j\|} + O(\Delta x)$$

- Iterate nodes for volume contributions
- Iterate edges for surface flux contribution

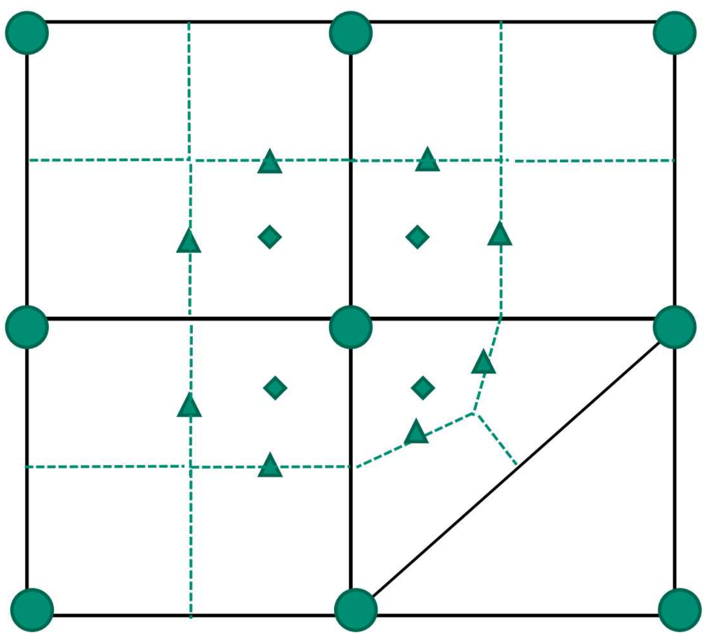

Dual-volume definition

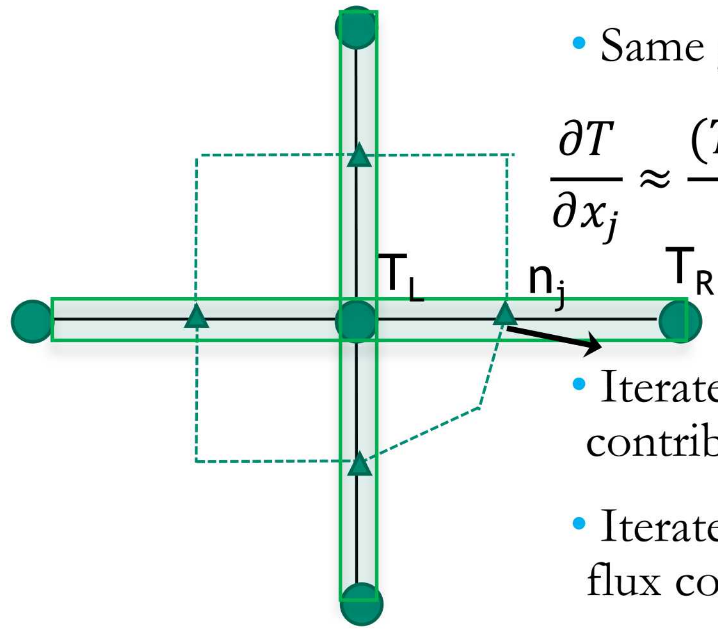

Edge-based stencil

Typical Failings for Two-State Discretization Methods

- With two points, only a linear basis can be used.
- Therefore, unstructured CC and EBVC are limited to second-order spatial accuracy
- Non-orthogonality is problematic for gradient-operator

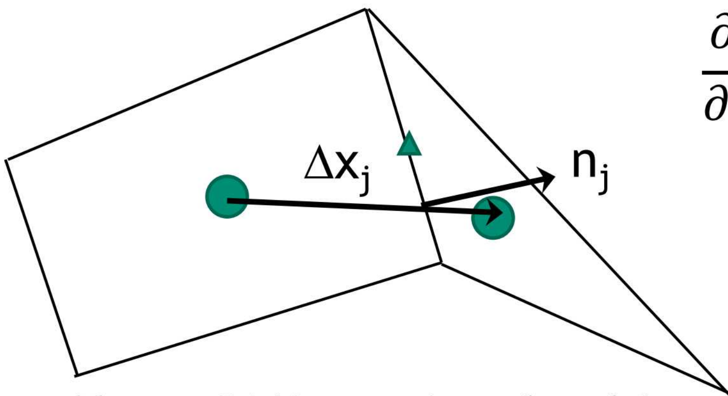

$$\frac{\partial T}{\partial x_j} = G_j T + \left[(T_R - T_L) - G_k T \Delta x_k\right] \frac{A_j}{A_l \Delta x_l}$$

With area vector defined by: $A_j = n_j dS$

- Above, $G_j T$ is a projected nodal gradient at the cell-center, or vertex center:
- Non-orthogonality is simply defined as the mis-alignment of the distance vector between the two "L" and "R" states and the surface normal

$$G_j T = \frac{\int T A_j}{\int dV}$$

- Both edge- and cell centered-based schemes show degraded accuracy on typical production meshes
- Several non-orthogonality approaches are available, for the best source, see Jasak
  - Jasak, "Error analysis and error estimation for the finite volume method with applications to fluid flow", Imperial College Dissertation, 1996

The Hybrid Control-Volume Finite Element Method (CVFEM)

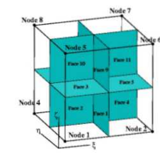

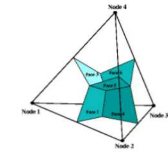

- A combination between the edge-based vertex-centered and FEM is the method known as Control Volume Finite Element ()
- A dual mesh is constructed to obtain flux and volume quadrature locations
- As with FEM, a basis is defined:

$$T(x_k) \approx \sum_{i=1}^{npe} N_i(x_k) T_i \qquad \frac{\partial T(x_k)}{\partial x_j} \approx \sum_{i=1}^{npe} \frac{\partial N_i(x_k)}{\partial x_j} T_i$$

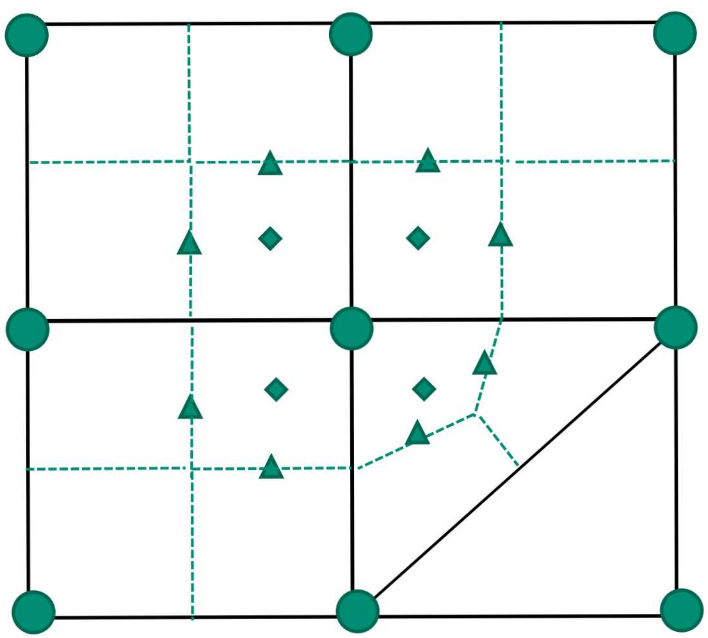

Dual-volume definition

- Integration-by-parts over test function w:

$$\int w \rho C_p \frac{\partial T}{\partial t} dV + \int \frac{\partial w}{\partial x_j} \lambda \frac{\partial T}{\partial x_j} dV - \int w \, \lambda \frac{\partial T}{\partial x_j} n_j dS = 0$$

- However, define a test function, w, as a piece-wise constant function (Heavyside) to be 1 inside the dual volume and 0 outside. Gradient is a Dirac-delta function:

$$\frac{\partial w}{\partial x_j} = -n_j \delta(x_j - xIP_j)$$

- Leading to: $\int \rho C_p \frac{\partial T}{\partial t} dV - \int \lambda \frac{\partial T}{\partial x_j} n_j \, dS = 0$

Control-Volume Finite Element Method Attributes

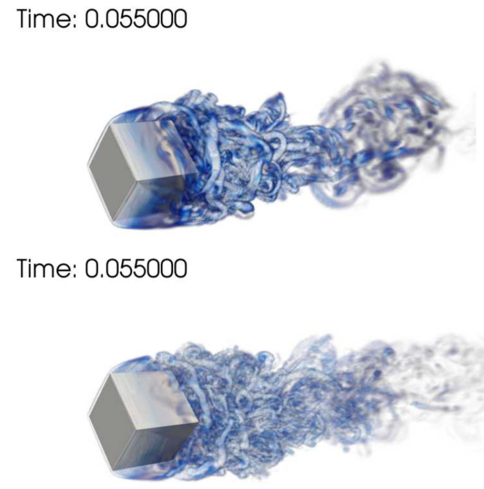

Rotating cube (Re 4000, RPM 3600)
P=1 (top) and P=2 (bottom)

- CVFEM is a locally conservative finite volume scheme
- Gradient operator, like its FEM counterpart, is absent of any error due to non-orthogonality, however, suffers from high AR mesh and monotonicity
- CVFEM can be viewed as Petrov-Galerkin method
- The method can also be promoted in polynomial space (higher efficiency on NGP due to increased local work)
- However, suitability of higher-order for LES is an open argument (I mean research opportunity)

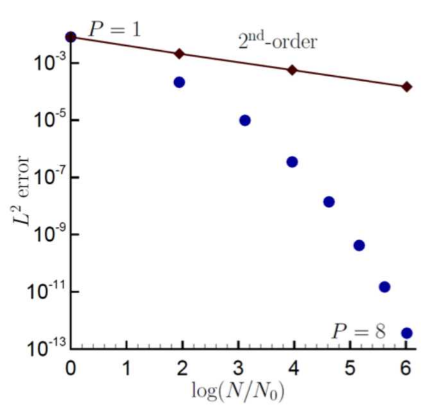

Spectral convergence

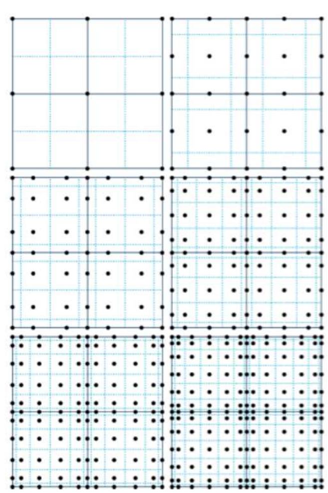

Dual-volume for promoted quad4

P=1 (left) and P=4 (right)
Helium plume (VR-density)

CVFEM Allows for Mix/Match of FEM and FV

- Sliding mesh using a hybrid CVFEM/DG interface (Domino, 2018)

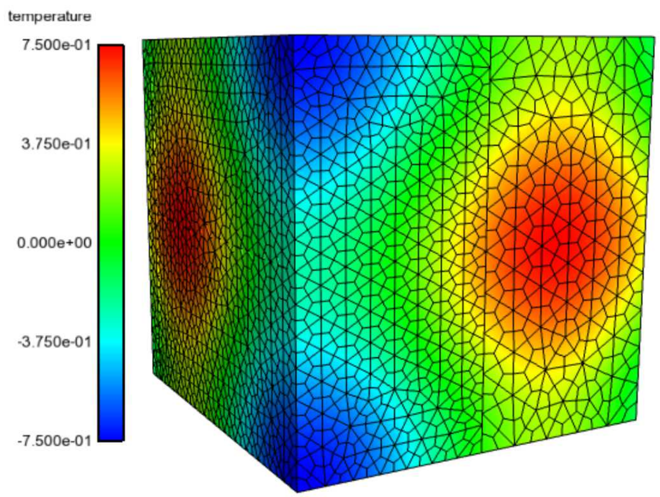

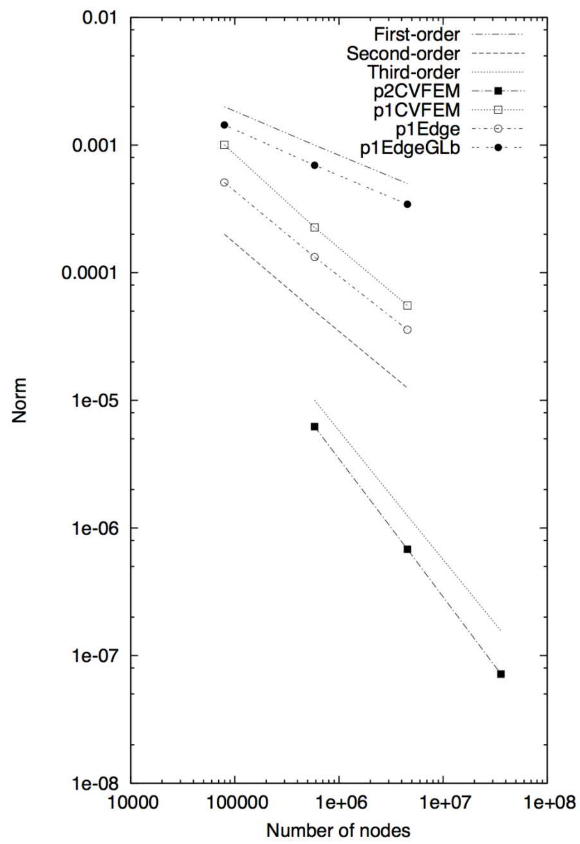

- Nonlinear Stabilization Operator (NSO) approaches (combining DCO of Shakib and entropy/visc of Guermond)

$$\sum_e \int_\Omega \nu(\mathbf{R}) \frac{\partial w}{\partial x_i} g^{ij} \frac{\partial \phi}{\partial x_j} d\Omega \quad \Longrightarrow \quad -\sum_e \int_\Gamma \nu(\mathbf{R}) g^{ij} \left(\frac{\partial \phi}{\partial x_j} - G_j \phi\right) n_i dS$$

With… $$\nu(\mathbf{R}) = \sqrt{\frac{\mathbf{R_k R_k}}{\frac{\partial \phi}{\partial x_i} g^{ij} \frac{\partial \phi}{\partial x_j}}}$$

Examples of Various Supported Topologies

Hex8

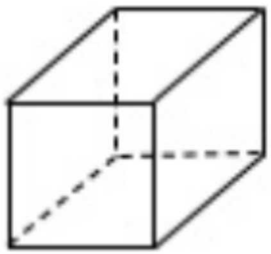

Tet4

Pyramid5

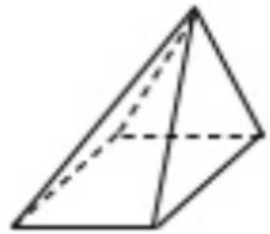

Wedge6

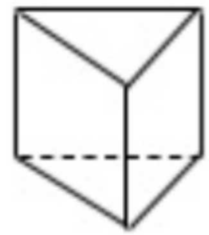

Arbirtrary

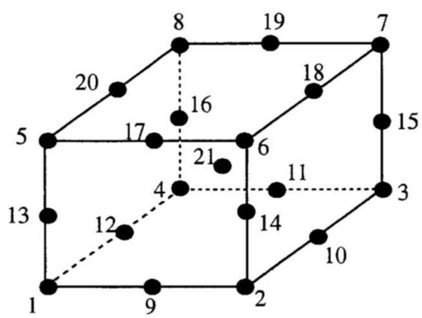

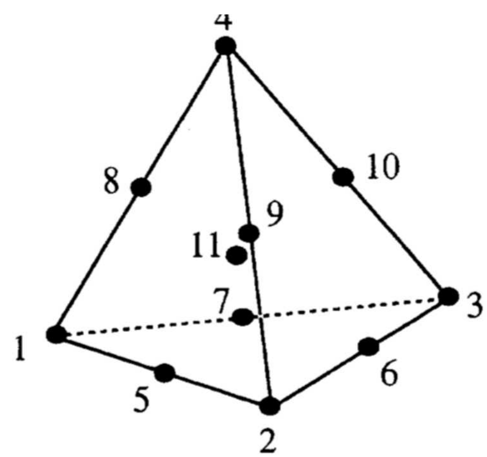

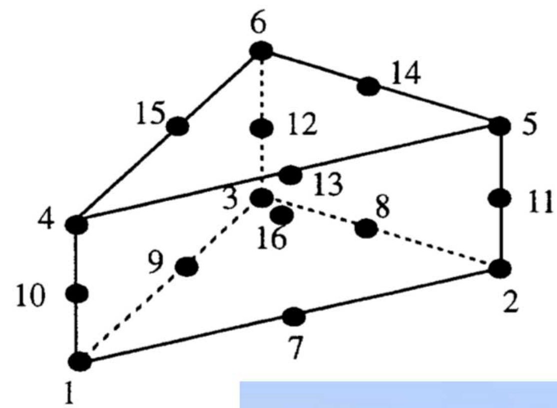

Higher-order promoted elements (Hex27, Hex64, etc.)

Presentation Objectives

- Overview of the low-Mach application of interest
- Numerical Methods (discretizations and coupling)
- Validation and Verification on general mesh topologies
- Comments on Typical Validation Efforts
- NGP Activities (ExaWind-centric)
- Conclusion

Reality: Meshing time for complex applications remains a significant bottleneck!

- Many applications of interest to SNL contain complex geometries
- low-Mach fluids users interested in high-quality simulation results tend towards hexahedral-based topologies (if possible)
- However, if a scheme is "design-order" accurate, any topology may suffice as it is simply a matter of mesh size and efficiency – not unlike the active discussion on low- vs higher-order
- Sometimes, the penetration of a low-Mach fluids physics addition for a DSW analysis is high as the meshing can be prohibitively complex, i.e., "inside the skin"

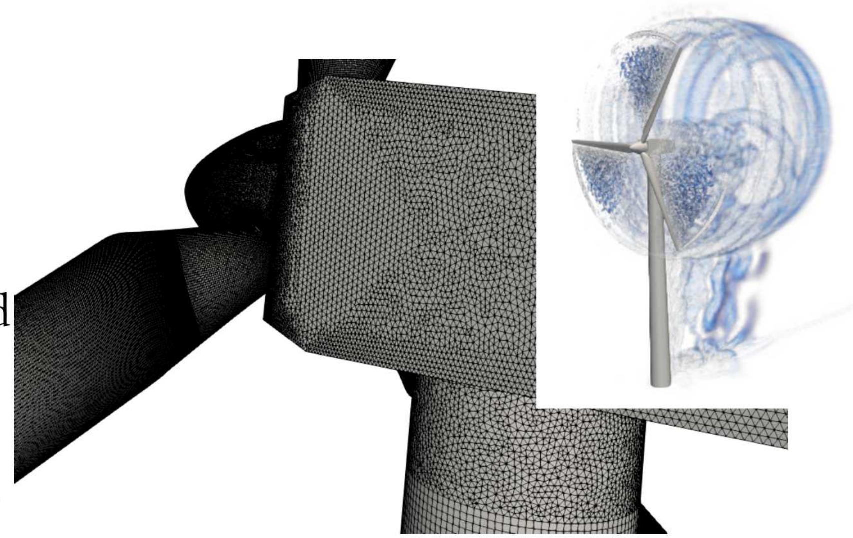

UUR Example: Vestas V27 225 kw
Hybrid low-order hex/tet/pyr/wedge

Reality: Meshing time for complex applications remains a significant bottleneck!

- Many applications of interest to SNL contain complex geometries
- low-Mach fluids users interested in high-quality simulation results tend towards hexahedral-based topologies (if possible)
- However, if a scheme is "design-order" accurate, any topology may suffice as it is simply a matter of mesh size and efficiency – not unlike the active discussion on low- vs higher-order
- Sometimes, the penetration of a low-Mach fluids physics addition for a DSW analysis is high as the meshing can be prohibitively complex, i.e., "inside the skin"

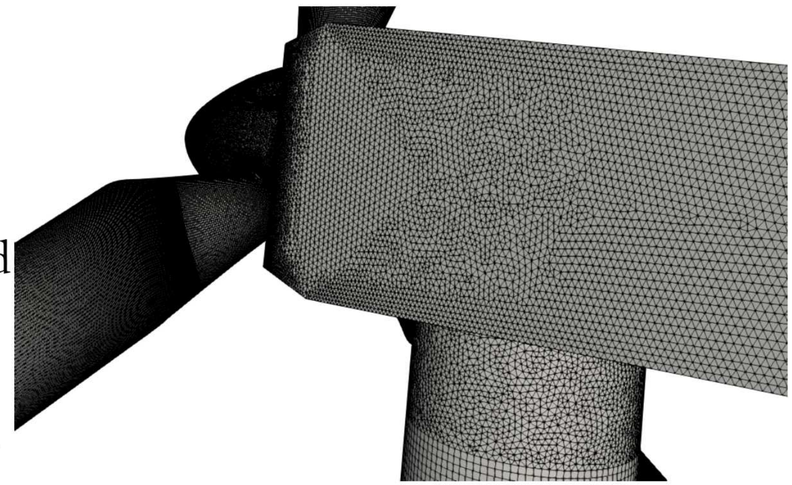

UUR Example: Vestas V27 225 kw
Hybrid low-order hex/tet/pyr/wedge

UUR sample of a somewhat complex geometry

ASC/SNL is focused on reducing meshing time, increasing code usability, and providing a ***Next Generation Simulation*** Capability

- NNSA Advanced Simulation and Computing (ASC) project, has initiated a *Next Generation Simulation* foundational research project that seeks to improve the throughout, effectiveness, and credibility of multi-physics simulation analysis
  - Development of advanced discretization schemes, error indicators, embedded VVUQ, and efficient parallel mesh generation
  - Paradigm shift from insisting on "friendly" meshes (using discretization schemes that excel) to a fast-meshing generation archetype whose time scale can support penetration of the design cycle
  - Algorithms Next Generation Platform (NGP) ready
- Calling all Numerical Architects, sorry, we must **reverse the current paradigm**!!
- Rather than prescribing the mesh type to the user community, we need to work within the constraints of what types of meshes can be generated quickly and deploy discretizations & algorithms that effectively run on those meshes

- Current meshing investments include:
  - Tetrahedral-based (possibly with transition elements), Arbitrary-polyhedra, etc.
- **Discretization Architects Lament**: "Wait a minute - please…."
  - How about a set of foundational studies on the suitability of different element topologies - **first**!
- FY18 Goal: start easy (laminar) and build up to turbulent flow (T-G, HIT, PCF, etc.)

A Review of the Typical V&V Workflow

- Verify on ***nice*** Hexahedral-based elements (at SNL, Thex)

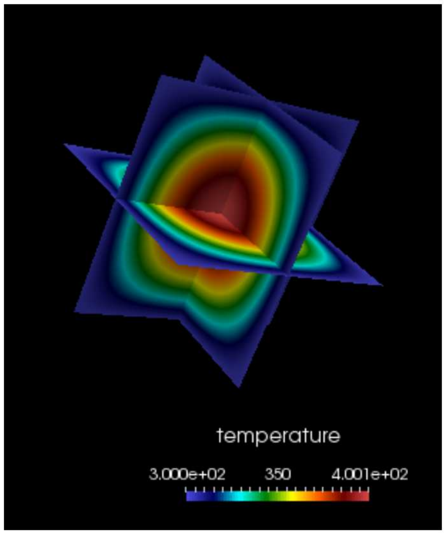

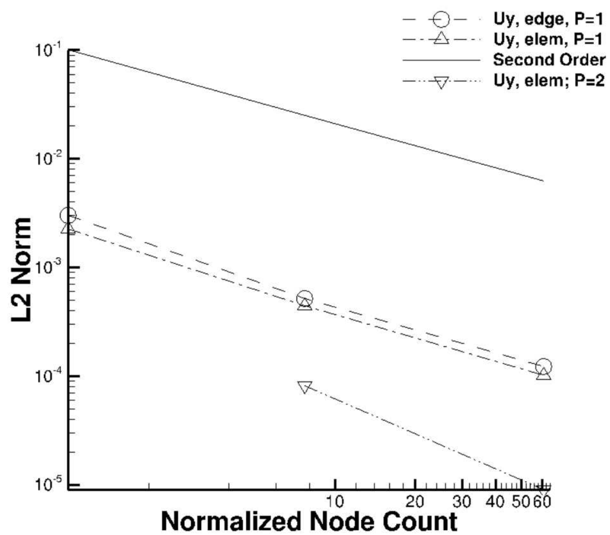

Variable-density low-Mach MMS (Domino, 2016)

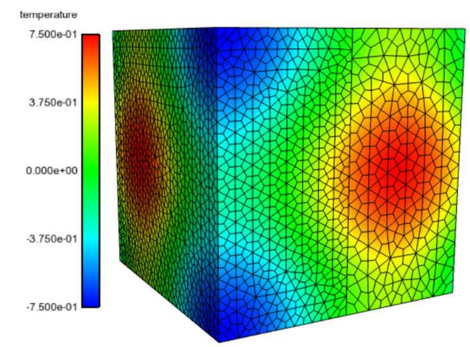

Thex-based non-conformal interface

- Validate on ***perfect*** meshes on a variety of canonical flows

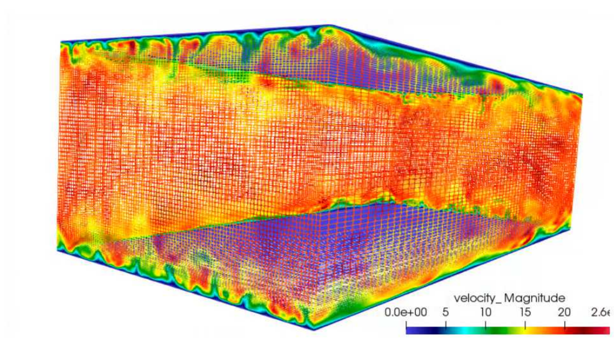

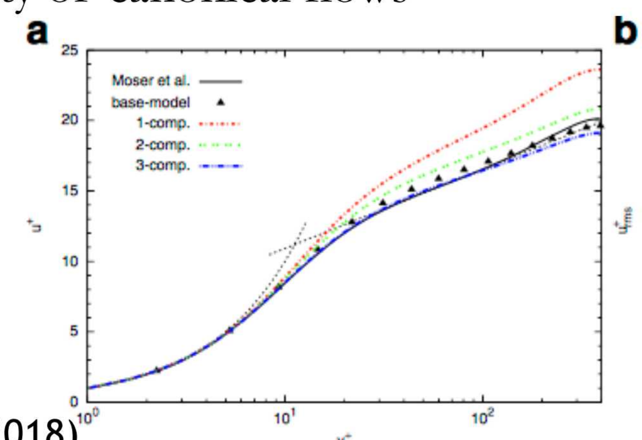

Re^{τ} 395 plain-channel (Jofre, Domino, Iaccarino, 2018)

Case 1: Laminar 1x2x10 Channel

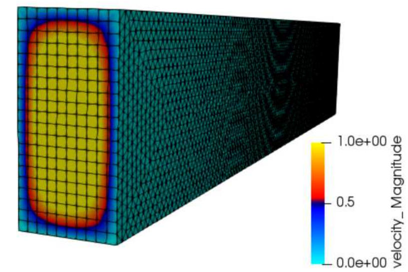

- 1x2x10 specified pressure drop duct flow (Domino et al, 2007):

The material properties are a fluid with a density of 1E-3 kg/m^3 and a viscosity of 1E-4 kg/m. A specified pressure drop is provided to be 1.60e-3 Pa/m.

The axial velocity distribution is

$$v = -\frac{1}{2\mu}\frac{dp}{dz}\left[b^2 - y^2 - \frac{4}{b}\sum_{n=0}^{\infty}(-1)^n \frac{1}{m^3}\frac{\cos(my)\cosh(mx)}{\cosh(ma)}\right], \tag{42}$$

where

$$m = \frac{(2n+1)\pi}{2b}. \tag{43}$$

The half-width in the x-direction is a and the half-width in the y-direction is b. The axial direction down the duct is z.

- Perform simulations on Hex8, Hex27, Tet4, Pyramid5, Werdge6 and Hybrid

Case 1: Laminar 1x2x10 Channel - Results

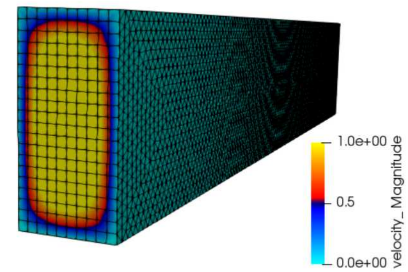

Findings:

- All topos result in design-order accuracy
- Hex8 and Wedge6 surprisingly provide the same accuracy
- Tet4 results in ~3.5x higher error and matches pyramids
- Higher-order (P=2) notes ~an order of magnitude lower error for the first mesh refinement level (promoted Hex8 P0 mesh)
- Tet4/Pyramid5 O(5)x slower than Hex8/Wedge6 to meet the same accuracy (resource limited model)

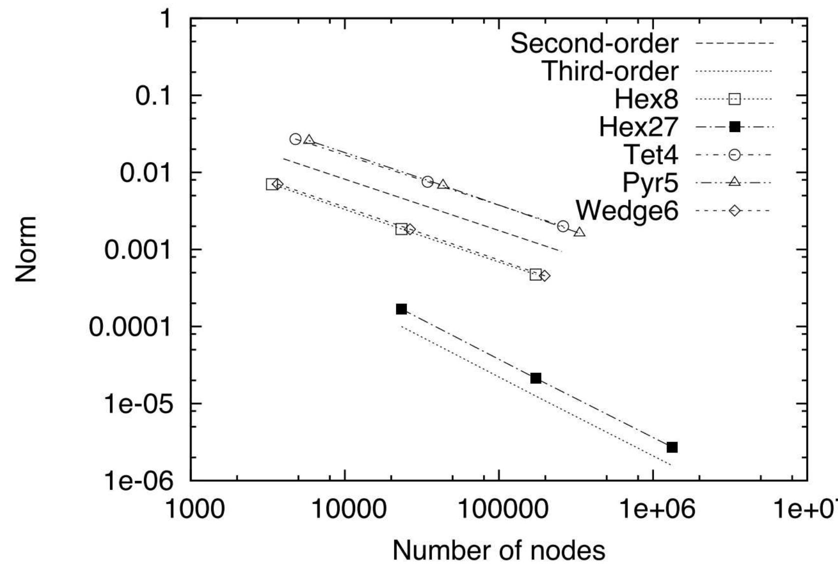

Mesh	R0 Node/Elem	R1 Node/Elem	R2 Node/Elem
Hex8	3,366/2,500	23,331/20,000	173,061/160,000
Hex27	23,331/2,500	173,061/20,000	1,331,721/160,000
Tet4	4,782/23,051	34,418/184,408	260,459/1,475,264
Pyramid5	5,866/15,000	43,331/120,000	333,061/960,000
Wedge6	3,672/5,510	264,99/45,760	197,862/367,240

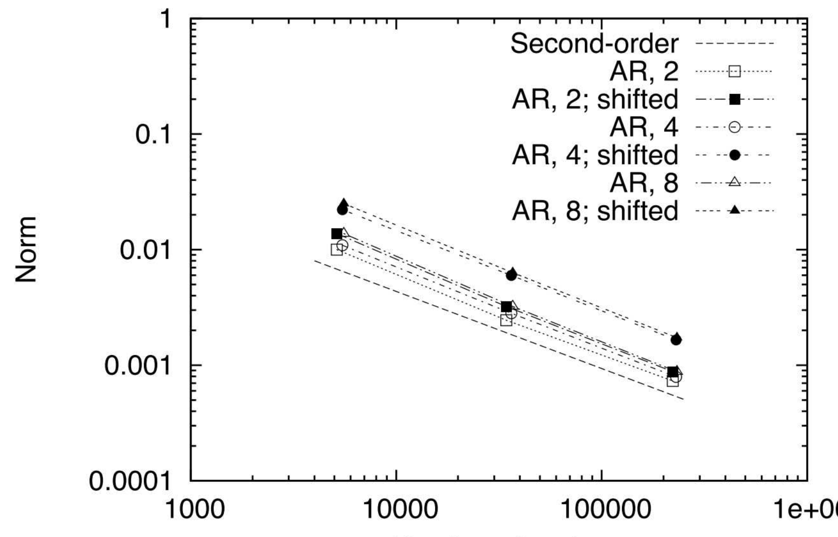

Case 2: Taylor-Green Vortex (Re 1600)

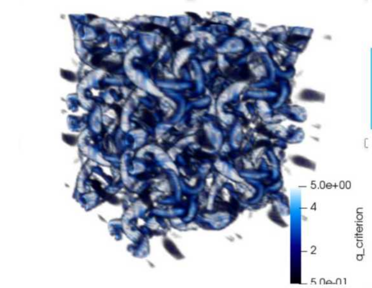

- Well studied problem that is part of the following numerical benchmark:
  - "C3.5 DNS of the T/G Vortex at Re = 1600"
- QoI: turbulent kinetic energy vs time and dissipation rate vs time
- First phase: small viscous effects, small-scale structures are laminar and organized [easy]
- Second phase: viscous (diffusion) dominates, structures are distorted [harder]
- Break-up phase: flow is dissipating fully [easy]

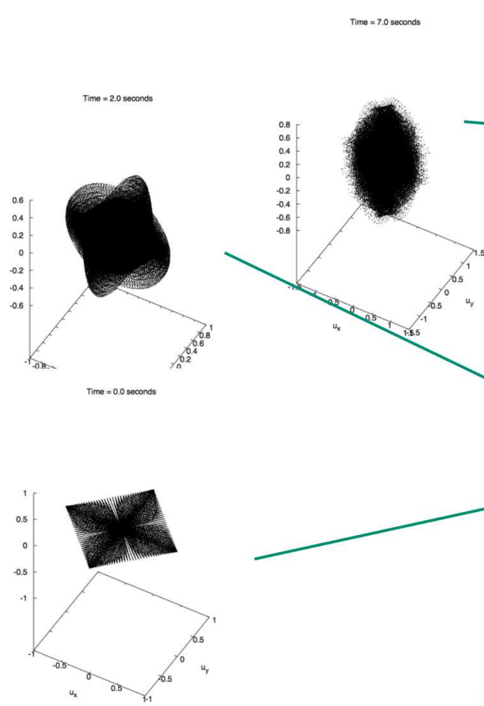

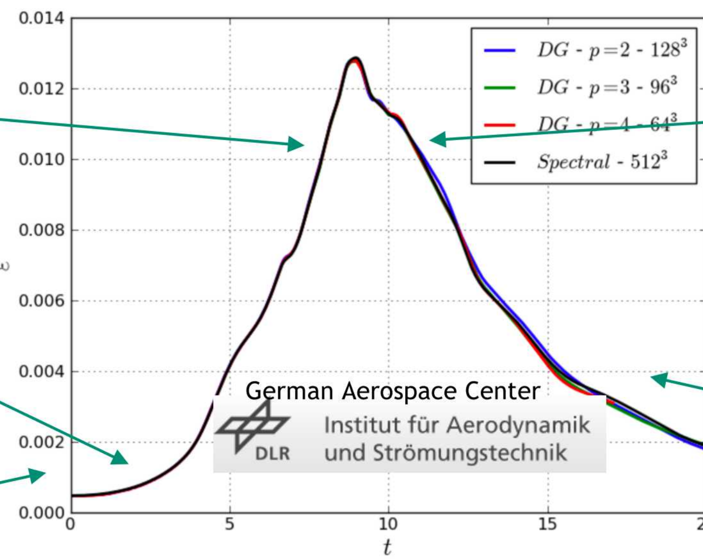

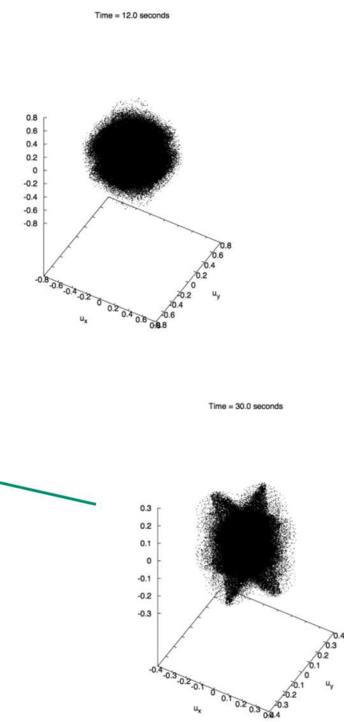

Figure 2: Evolution of the dimensionless energy dissipation rate as a function of the dimensionless time: results of pseudo-spectral code and of variants of a DG code.

Case 2: Taylor-Green Vortex (Re 1600) - Results

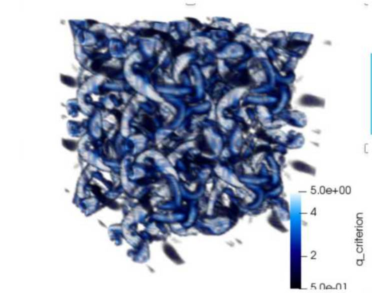

Findings:

- All topos converge to baseline DNS
- Pyramid5/Wedge6 (at R1) looks very good

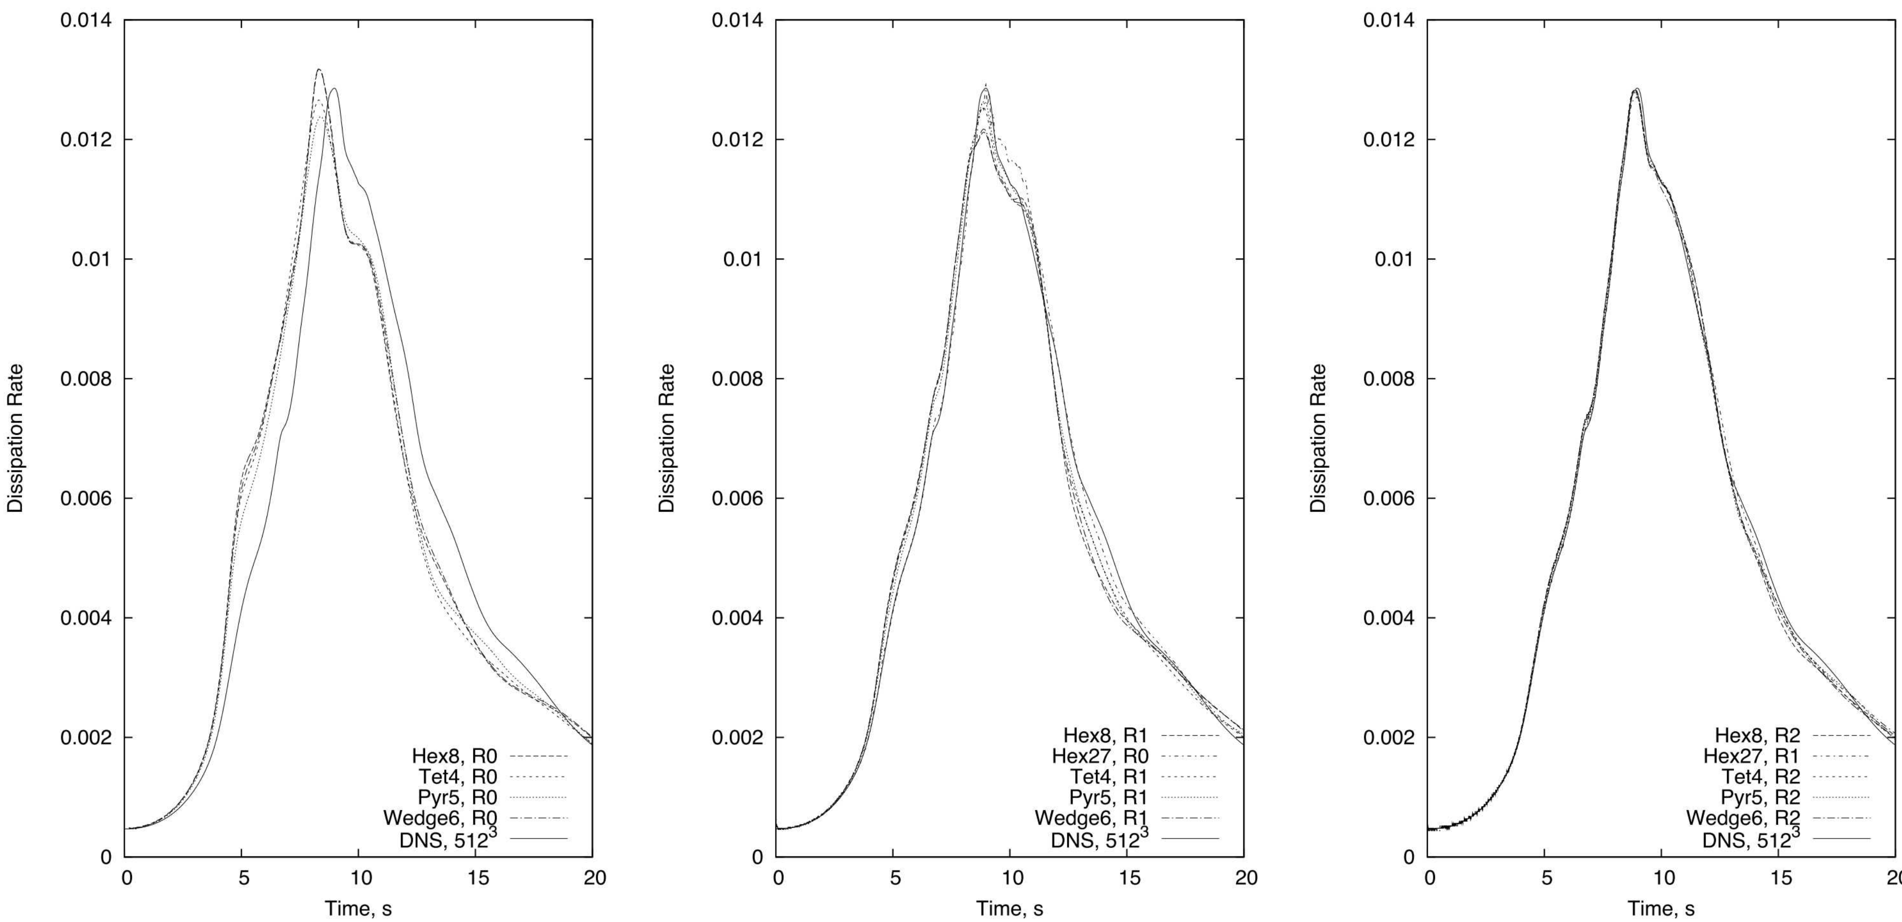

Case 2: Taylor-Green Vortex (Re 1600) - Results

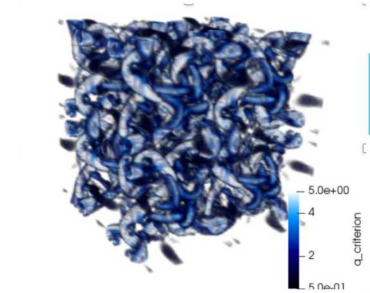

Timing Findings:

- At R2 mesh resolution, Wedge6, Pyramid5 < ~1.5x timing of Hex8; Tet4 ~2x
- Again, based on a resource-limited model (penalizes high element counts)

Mesh	**R0 Node/Elem**	**R1 Node/Elem**	**R2 Node/Elem**
Hex8	1,030,301/1,000,000	8,120,601/8,000,000	64,481,201/64,000,000
Hex27	8,120,601/1,000,000	64,481,201/8,000,000	n/a
Tet4	1,558,290/9,185,501	12,370,858/73,484,008	98,500,835/587,872,064
Pyramid5	2,030,301/6,000,000	16,120,601/48,000,000	128,481,201/384,000,000
Wedge6	1,180,185/2,296,800	9,313,737/18,374,400	74,002,545/146,995,200

	Hex8	Tet4	Pyr5	Wedge6
R0	1	2.24/1.0	1.92/1.0	1.15/1.0
R1	2.58	1.75/2.0	1.56/2.10	1.3/2.9
R2	5.48	2.2/5.4	1.59/2.16	1.3/6.25

Case 3: Plain Channel Flow (Re^{τ} 395)

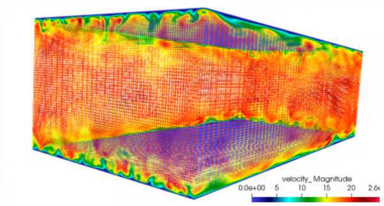

- What about a real wall-bounded turbulent flow?
- Plain channel flow is a well studied flow consisting of a double-periodic flow with upper and lower walls
  - Body force provided
- Simulations run in "Wall-Resolved LES", or WRLES using WALE model
  - No wall functions
  - Yplus < 1 requirement

Mesh Procedure

- Hex8 the mesh (with bias); base mesh from Jofre et al, 2018 study
- Hex27 obtained by promotion of Hex8
- HTet4 the Hex8 mesh
- Tri3 surface extruded for Wedge6

Mesh	R1 Node/Elem	R2 Node/Elem
Hex8	813,345/786,432	6,398,529/6,291,456
Hex27	6,398,529/786,432	n/a
HTet4	813,345/4,718,592	6,398,529/37,748,736
Wedge6	813,345/1,572,864	6,398,529/1,258,2912

Case 3: Plain Channel Flow (Re^{τ} 395) - Results

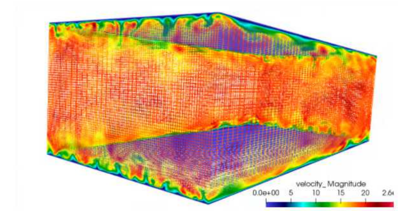

Findings:

- R0 low-order mesh simulation demonstrates reasonable prediction wrt Moser's DNS; promoted Hex27 is well resolved
- R1 low-order mesh shows good results using the Tet4 and Wedge6 mesh
- Timing shows that the increased assembly time is offset by the decreased linear solver time; Hex8=Tet4=Wedge6
- Comparison of higher-order statistics is in progress

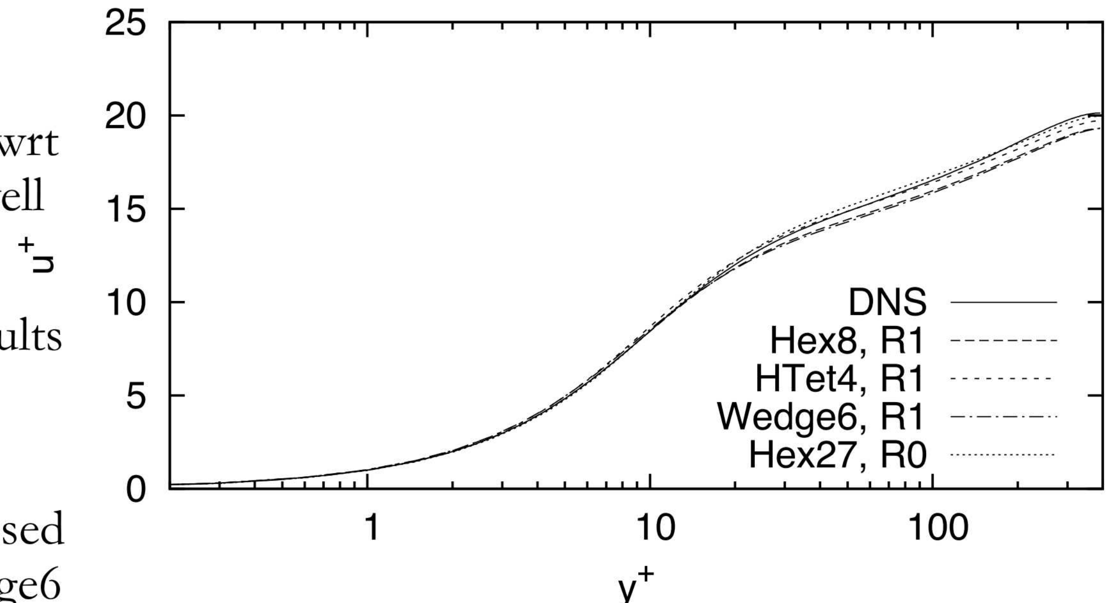

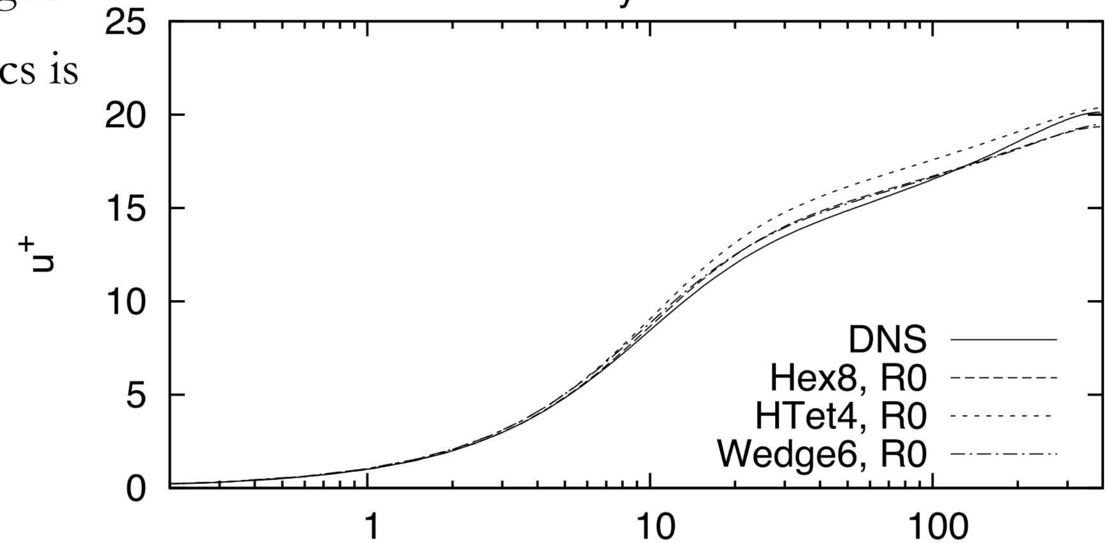

Case 3b: Plain Channel Flow (Re^{τ} 180) - WIP Results

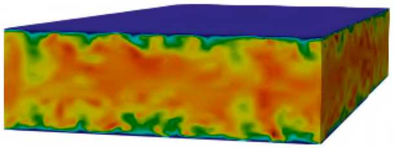

- Hybrid mesh study based on Ham and Iaccarino (2006) found that simulations were extremely sensitive to mesh topology
- Non-symmetric time mean flow found for cell-centered; better for the CTR node-centered formulation
- Simulations shown are absent of LES model

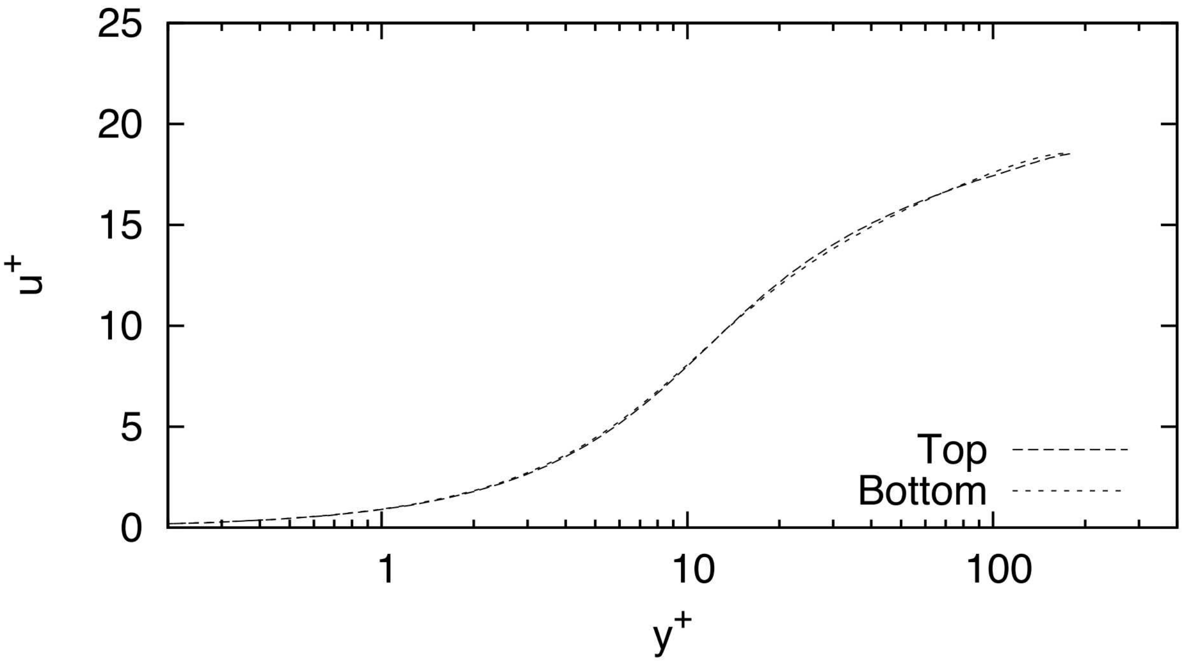

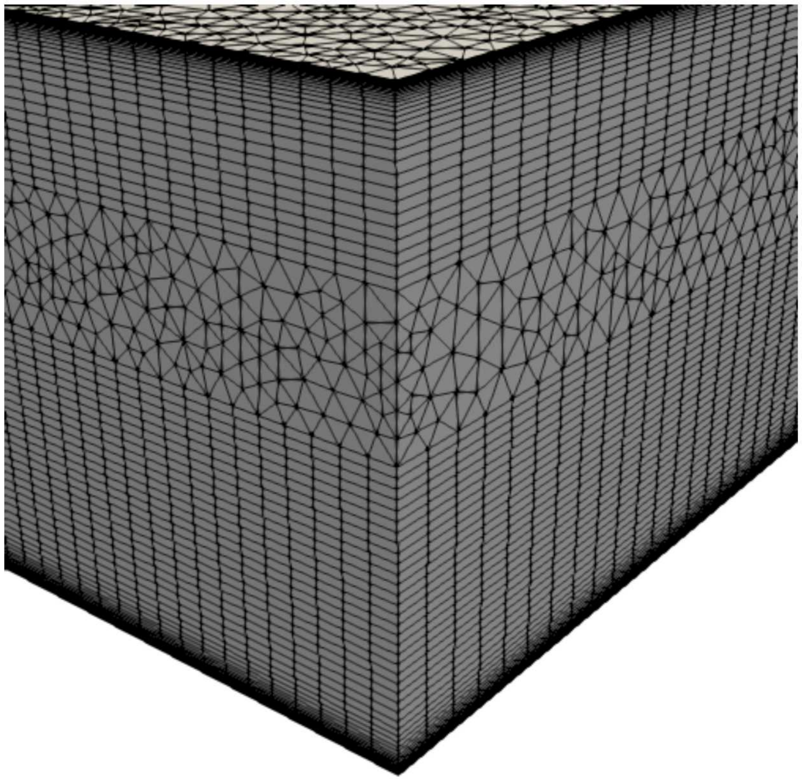

- Again, higher-order statistics and LES ongoing

Flow past water channel

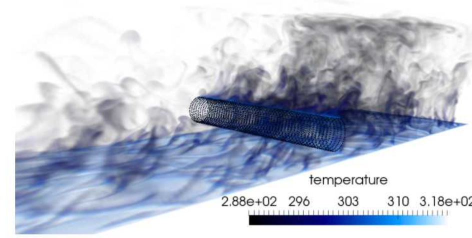

- Model Configuration: SNL-based Sean Kearney Experiment, "Experimental investigation of a cylinder in turbulent convection with an imposed shear flow", AIAA, 2005
- RANS Conclusion: The presence of the heated bottom wall significantly challenged ability to predict the QoI; q"

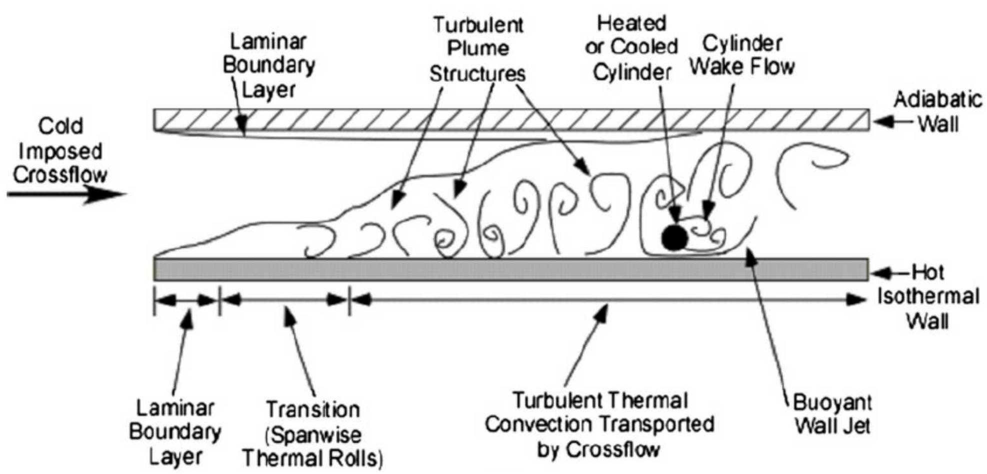

Kearney experimental configuration, 2005

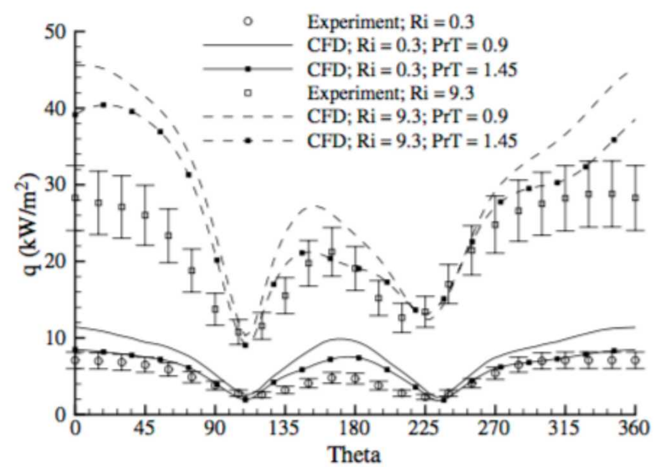

Fig. 13. Effect of turbulent Prandtl number on cylinder heat flux predictions for cases 3 (cooled cylinder) and 4 (heated cylinder).

Laskowski et al., 2007

RANS-based simulation (v2-f, k-e) study conducted by Laskowski et al., AIAA 2007

Flow past water channel - Results

- WRLES Ksgs model demonstrates decent mesh convergence
- low- and high-order results very similar (Domino, 2016)
- Tet4 simulations are also predicting the flow well at reduced time.

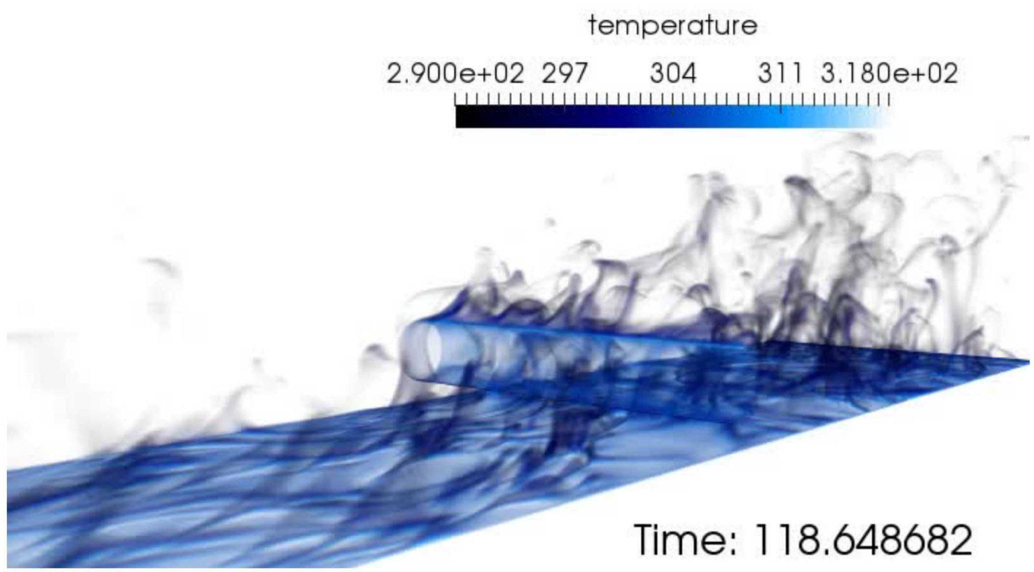

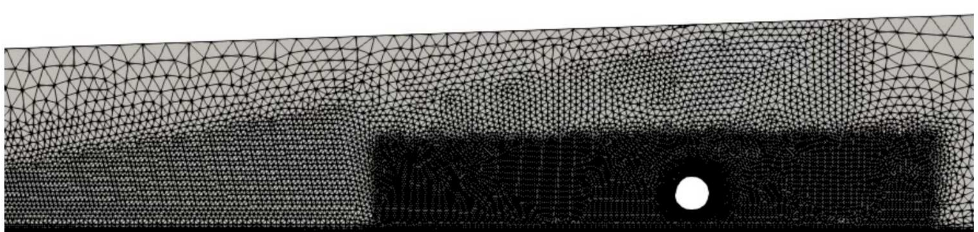

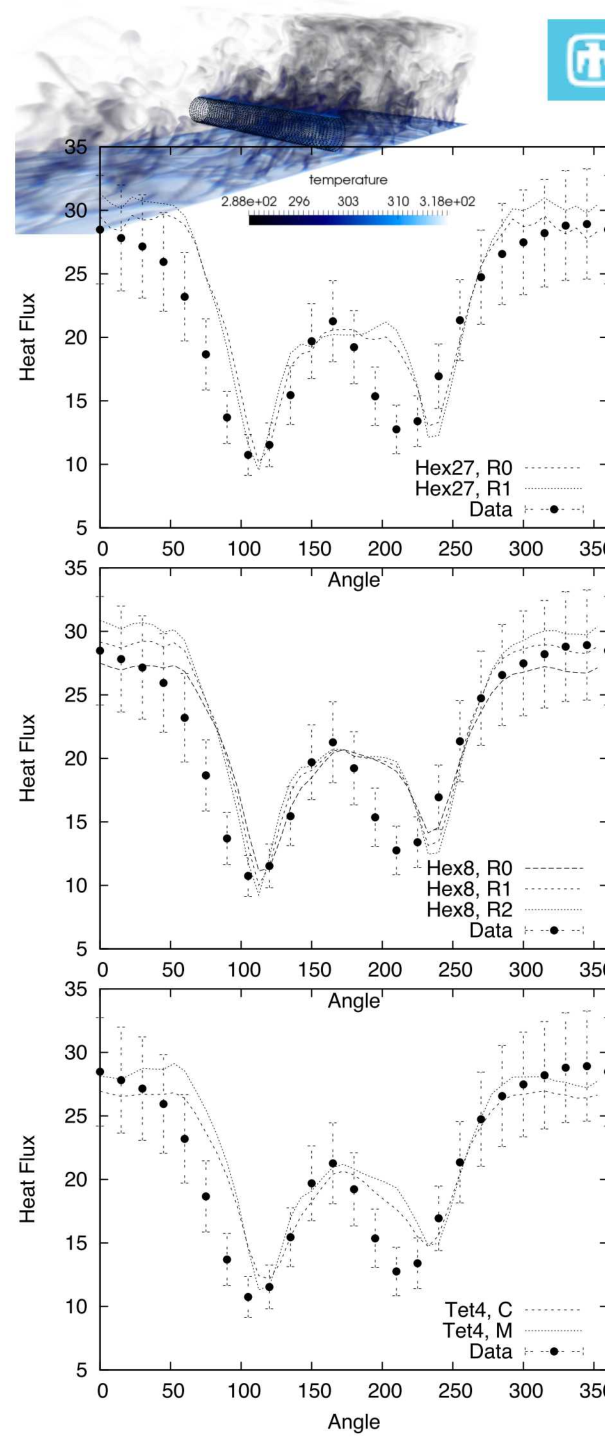

Wind Energy Applications Including Blade-Resolved Simulations

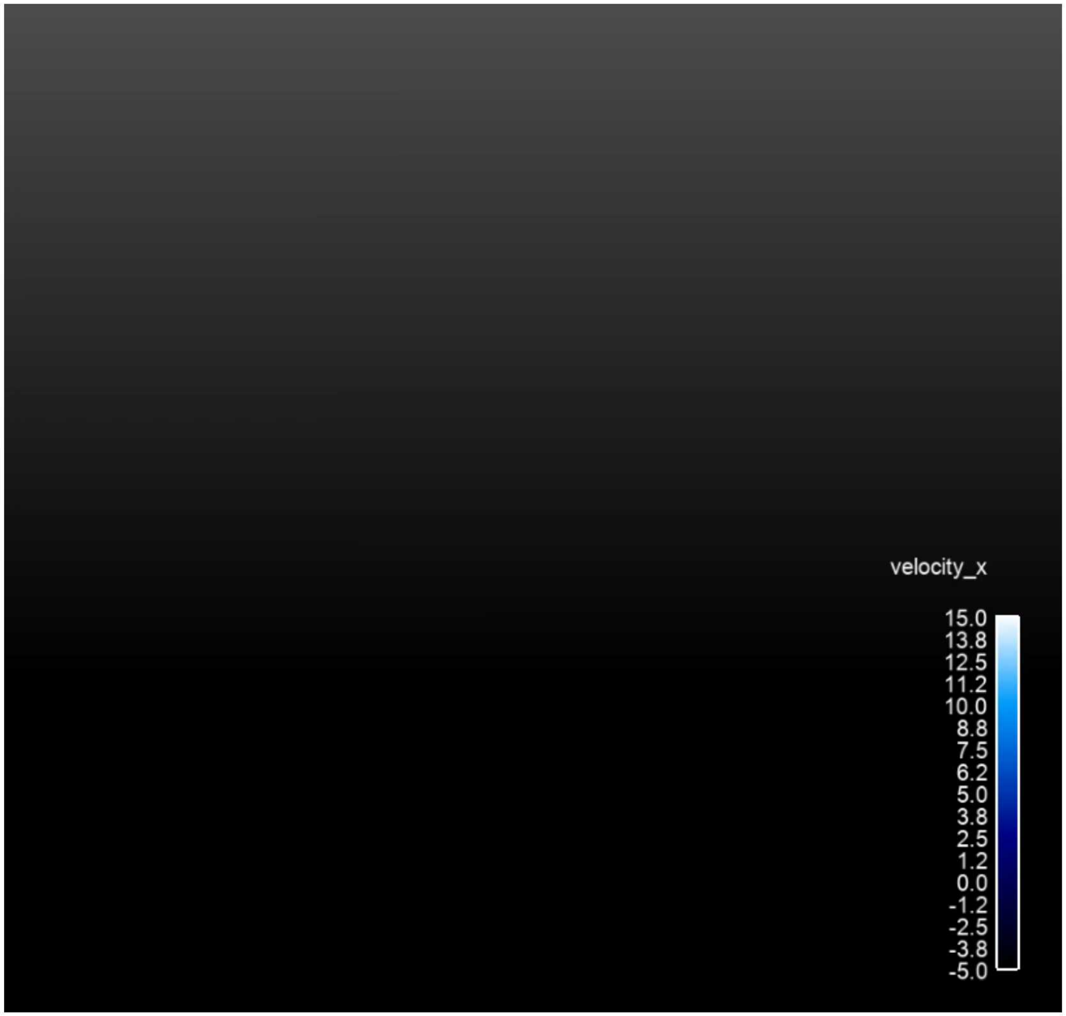

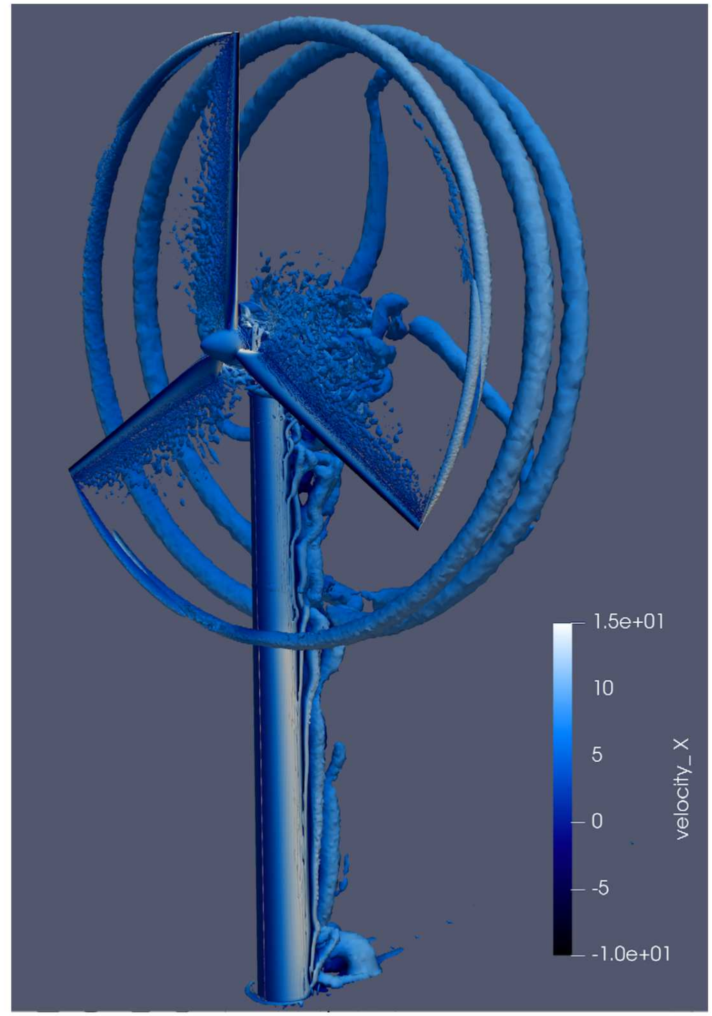

EXASCALE COMPUTING PROJECT

Vestas V27 225 kw wind turbine simulation

NALU LOW-MACH

Disparity in Time and Length Scales, Wind

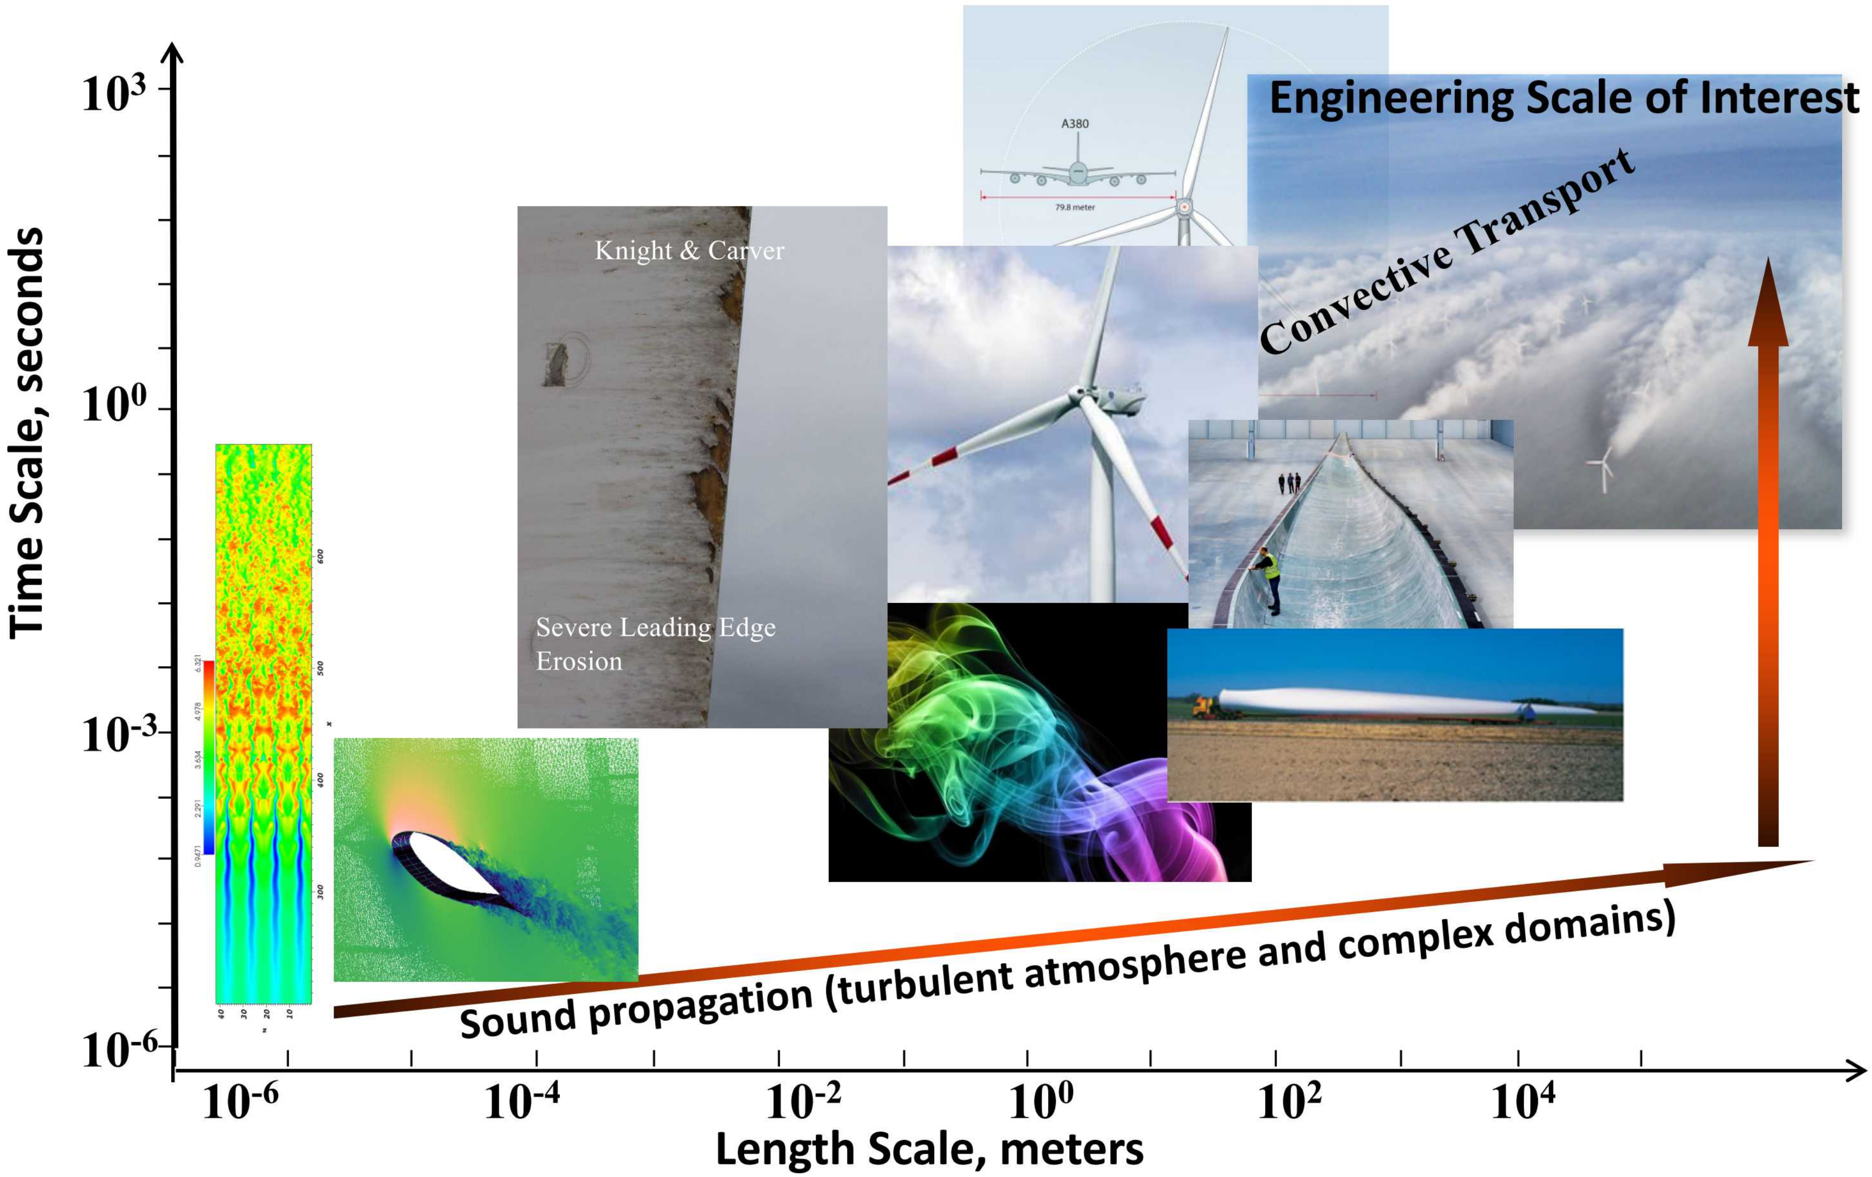

Vestas V27 225 kw full turbine – Preliminary Results

- O(200) million element low-order hybrid mesh
  - hex8, tet4, pyr5, and wedge6
  - Novel hybrid, design-order CVFEM/DG in use (IP)
- Usage of Nalu/SIMD/Kokkos-based Kernels
- Preliminary deployment of Nonlinear Stabilization Operator
- WALE-based LES model
- Thrust well predicted relative to previous characteriza

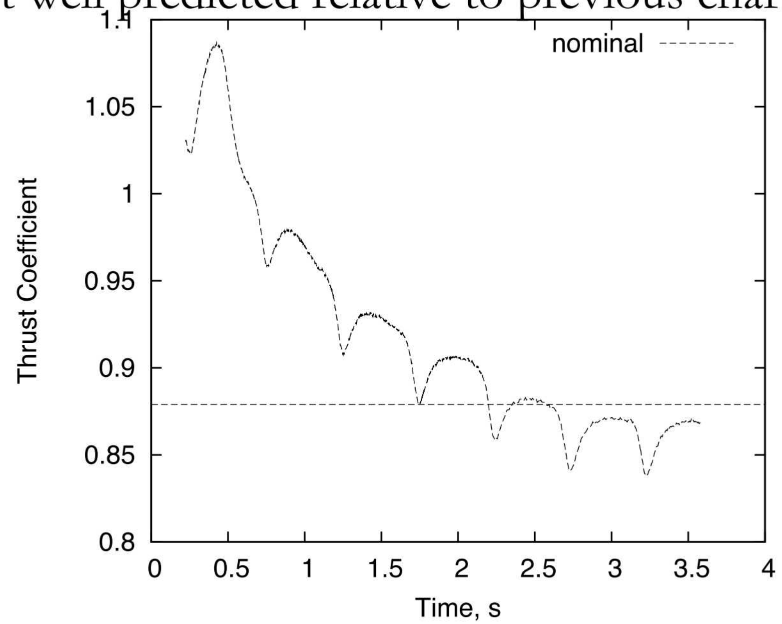

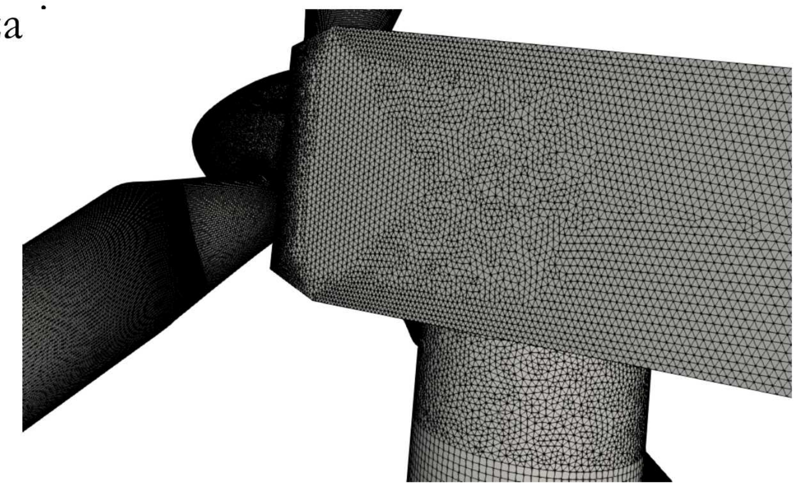

Volume-rendered Q (top)
Hybrid mesh (bottom)

Presentation Objectives

- Overview of the low-Mach application of interest
- Numerical Methods (discretizations and coupling)
- Validation and Verification on general mesh topologies
- Comments on Typical Validation Efforts
- NGP Activities (ExaWind-centric)
- Conclusion

Conceptual Challenge: Distinguishing Between Model-form Error, Discretization Error, and Code Error

- One mesh, one model, no code pedigree…

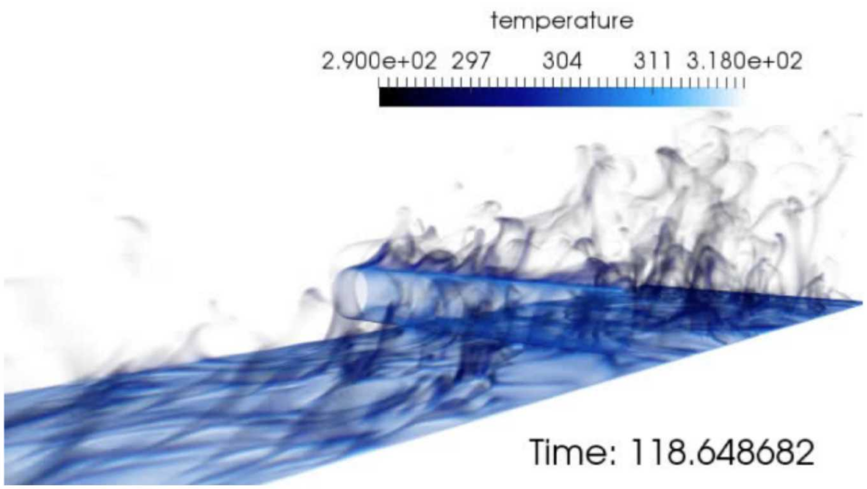

Heat flux to the cylinder
Volume-rendered temperature

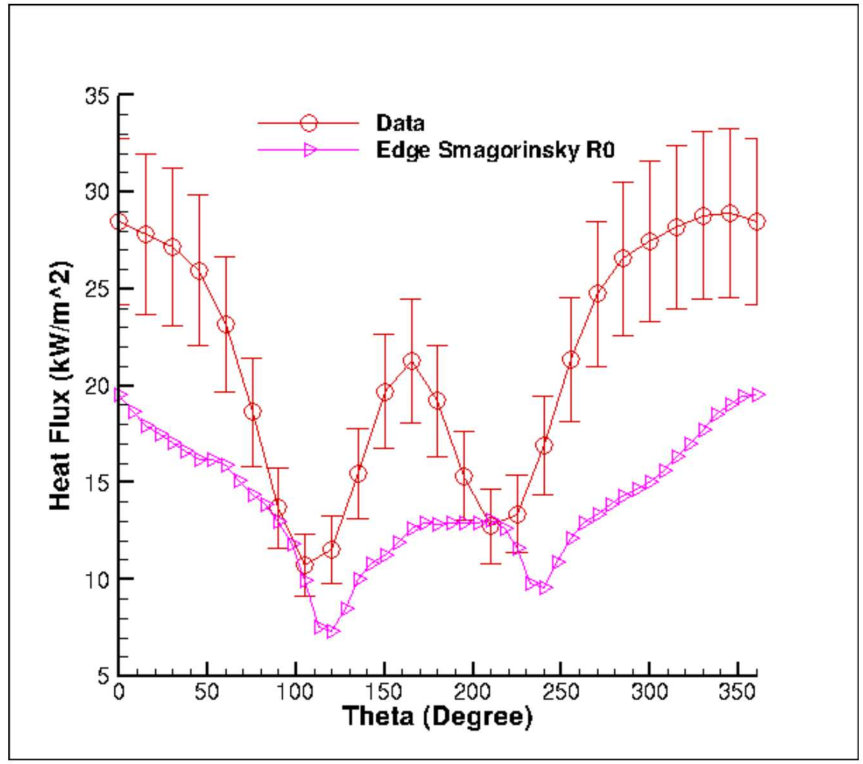

Time-averaged heat flux to cylinder

- What credible scientific hypothesis can be tested in this context?
- Was a Phenomena Identification and Ranking Table (PIRT) was conducted:
  - Process that defines 1) what you know, 2) what you think you know, and 3) what "you know not of"

Solution Verification and Due Diligence for Model-form (Structural) Uncertainty

	Import	Adequacy			
	Phen	Mod	Code	Val	Mats
Convective Processes					
Convective heat transfer	M	M	M	L	

- Three meshes, three models; code verification
- The heat flux results also show error bars due to time and spatial averaging over a line-of-site

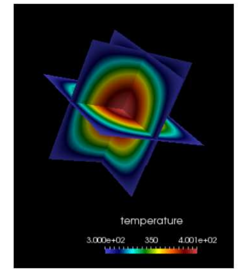

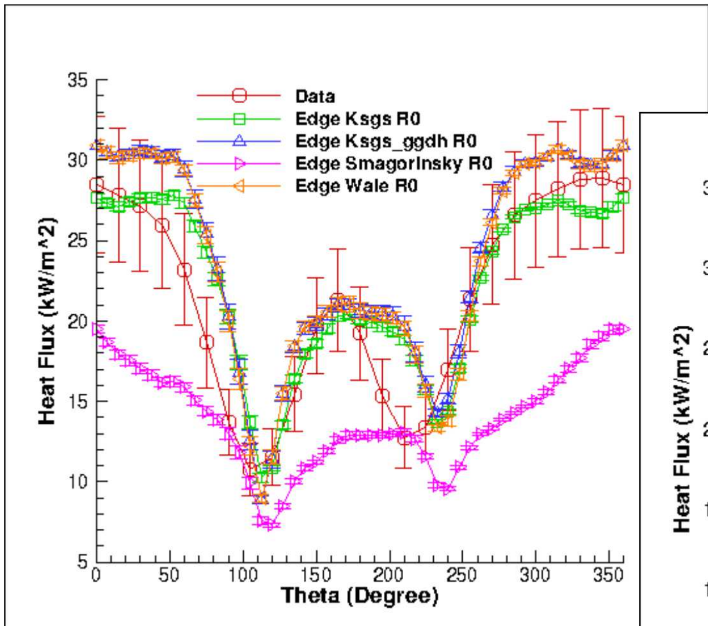

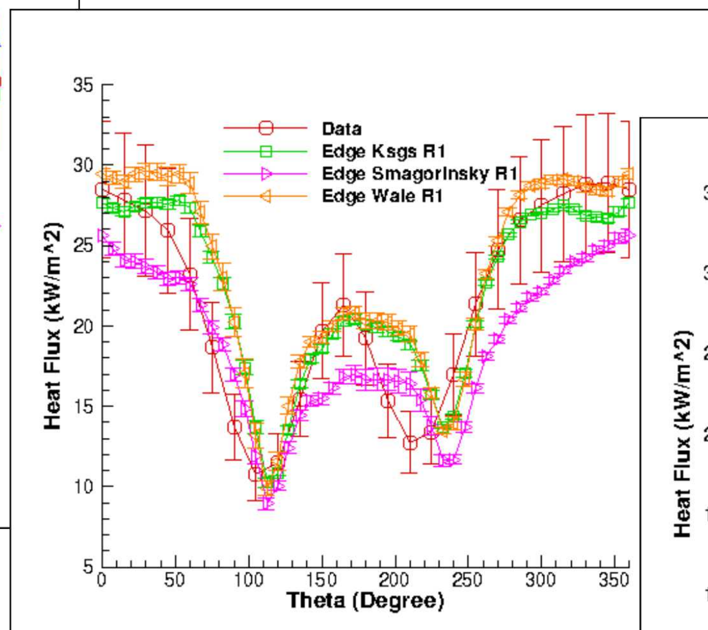

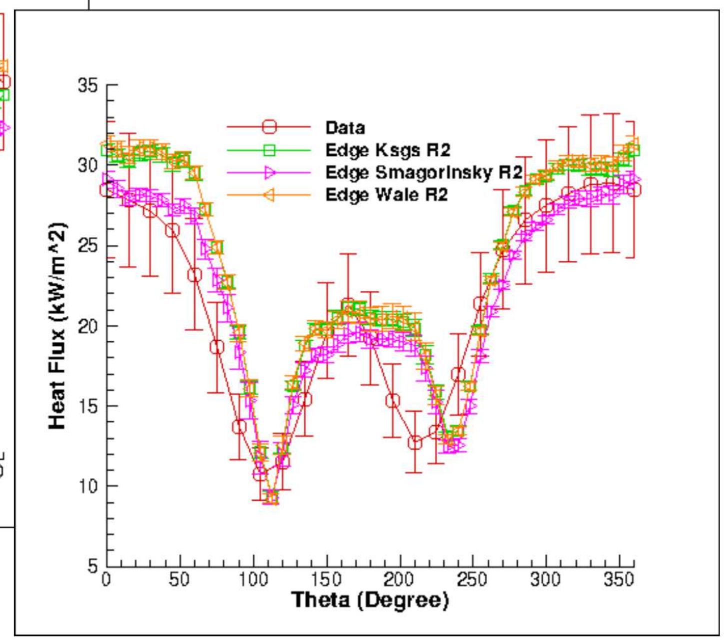

- Better…
- However, multiple models implemented and verified, maintained, transitioned, etc.

More Effective/Efficient Structural Uncertainty

- In the previous high-quality LES validation (cylinder in x-flow), three models were implemented and tested
- Is there a more efficient approach? Yes! Eigenvalue perturbation of the SGS stress

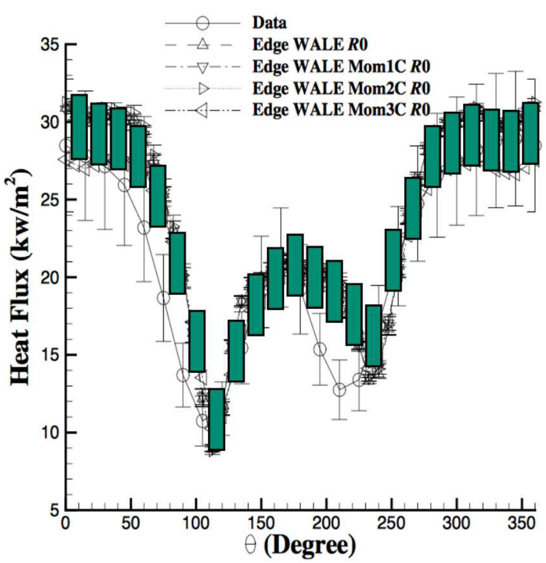

$$\tau_{ij}^{sgs} - \frac{\tau_{kk}^{sgs}}{3}\delta_{ij} = -2\nu_{sgs}\overline{S}_{ij},$$

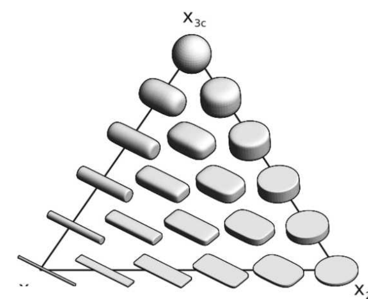

$$a_{ij}^{res} = \frac{1}{\overline{u_k u_k}}\left(\overline{u}_i\overline{u}_j - \frac{\overline{u}_k\overline{u}_k}{3}\delta_{ij}\right) = v_{in}^{res}\Lambda_{nl}^{res}v_{jl}^{res}$$

$$a_{ij}^{sgs} = \frac{1}{\overline{u_k u_k}}\left(\tau_{ij}^{sgs} - \frac{\tau_{kk}^{sgs}}{3}\delta_{ij}\right) = v_{in}^{sgs}\Lambda_{nl}^{sgs}v_{jl}^{sgs},$$

$$\overline{u_i u_j}^* = \overline{u}_i\overline{u}_j + \tau_{ij}^{sgs*} = \overline{u}_i\overline{u}_j + \overline{u_k u_k}^* a_{ij}^{sgs*} + \frac{\tau_{kk}^{sgs*}}{3}\delta_{ij},$$

$$\text{with} \quad \overline{u_k u_k}^* = \overline{u}_k\overline{u}_k + \tau_{kk}^{sgs*} \quad \text{and} \quad a_{ij}^{sgs*} = v_{in}^{sgs*}\Lambda_{nl}^{sgs*}v_{jl}^{sgs*}.$$

- See “Framework for characterizing structural uncertainty in LES”, Jofre et al, 2017

Turbulent open jet structural uncertainty

- Re 25k turbulent open jet
- Partnership with G. Iaccarino and L. Jofre (Stanford PSAAP-2)

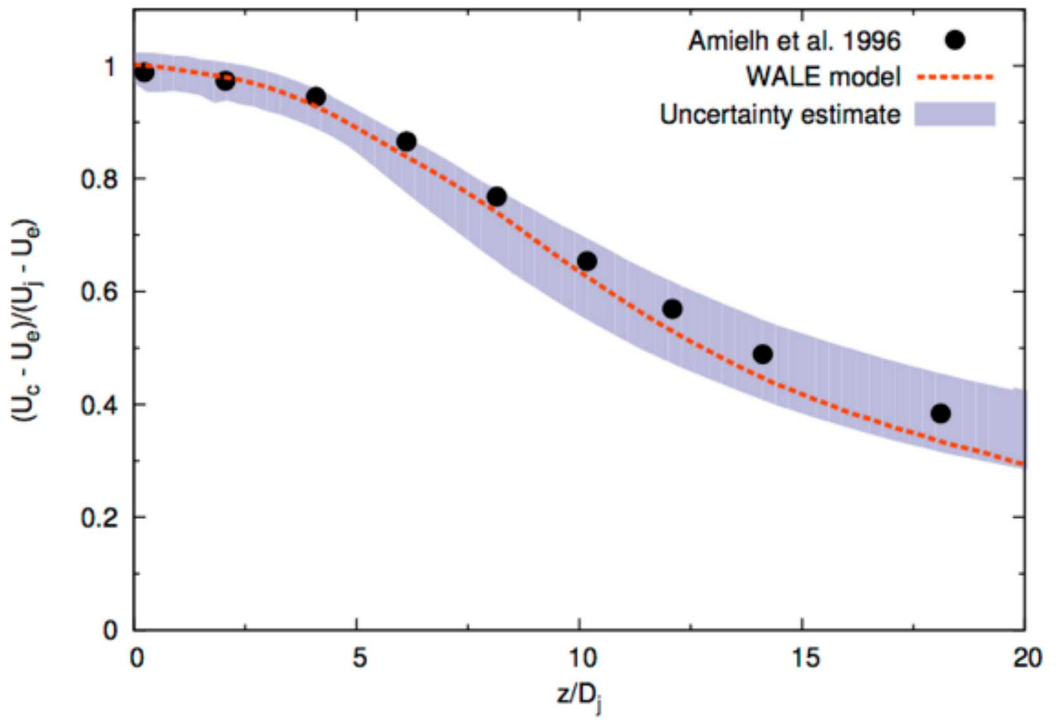

Centerline velocity decay

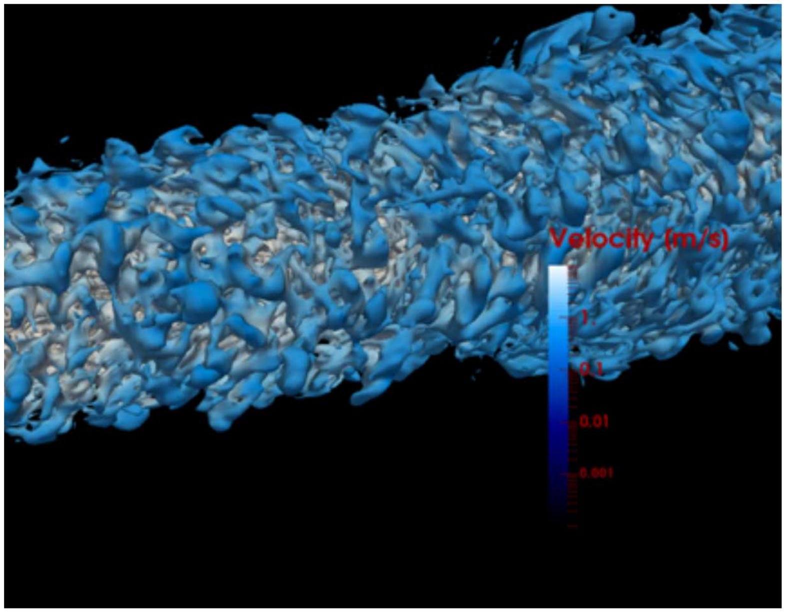

Velocity iso-surface from internal DNS

Presentation Objectives

- Overview of the low-Mach application of interest
- Numerical Methods (discretizations and coupling)
- Validation and Verification on general mesh topologies
- Comments on Typical Validation Efforts
- NGP Activities (ExaWind-centric)
- Conclusion

Goal: Beyond 32-bit Computing

- Circa 2013, many scientific production codes were limited to 32-bit
- Therefore, maximum simulation size for entities, e.g., node, edge, face, element, etc., was ~2.2 billion
- Next Generation Platforms were advocated to overcome poor MPI scaling and power needs to support Exascale computing (10^{18} floating point operations/second)
  - Platform architectures are not yet known

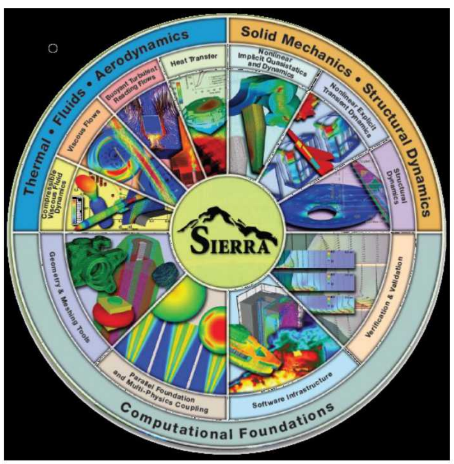

+ ASC IC Investments

()

Sierra Toolkit/Trilinos (open-source)
MPI+X parallelism
Support for new architectures

ExaWind Sample of Research Topics

- Low-/higher-order tradespace for LES
- Sliding mesh and/or overset
- Advanced stabilization techniques for nonlinear PDEs
- Increased solver performance at scale; O(100) billion elements
- Matrix storage reduction techniques for higher-order (static condensation)
- AMG coarsening strategies
- In situ matrix modification
- Efficient parallel searches
- Kokkos integration
- NGP focused for Exascale on open-source (BSD)

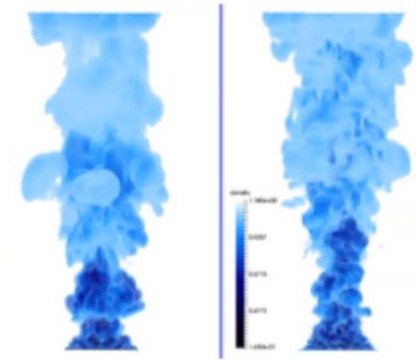

Low/high-order (with NC interfaces)

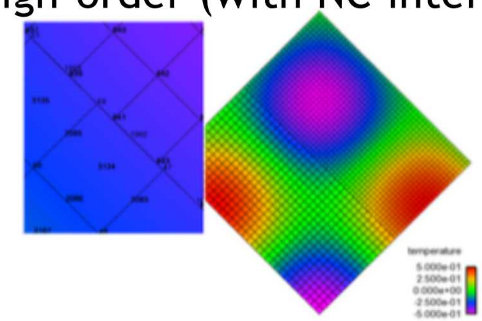

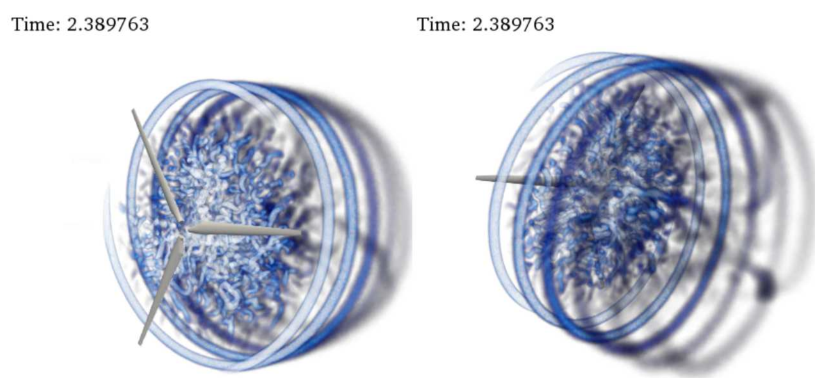

Sliding mesh/overset

Value Proposition: Deployment of HPC towards HF Wind Plant Modeling

Current Wind Plant Flow Physical Models **+** **Scalable flow solver technology** **=** ***Paradigm shift* in wind plant design & operation**

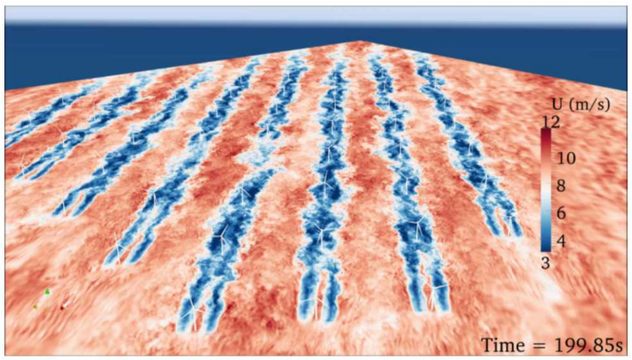

Current State of the Art: Wind plant wake simulation using NREL SOWFA tool

Millions of low-order elements

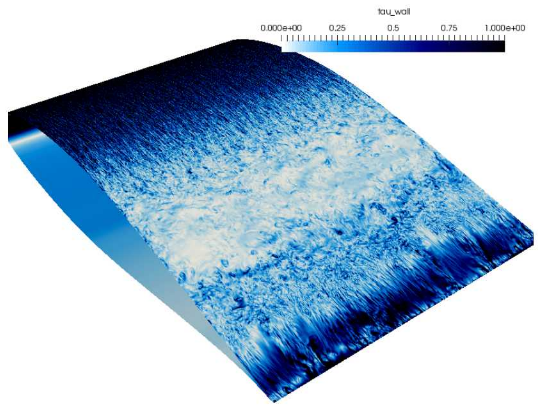

Near-blade simulation using SNL Nalu code on Trinity
Billions of elements using P=1, 2, …

Enabling SNL technology

- Wind plant simulation that *simultaneously resolves wakes and near-blade flow.*
- Allows the engineer to fully characterize linkages between *turbine design, site characteristics,* and *plant performance.*

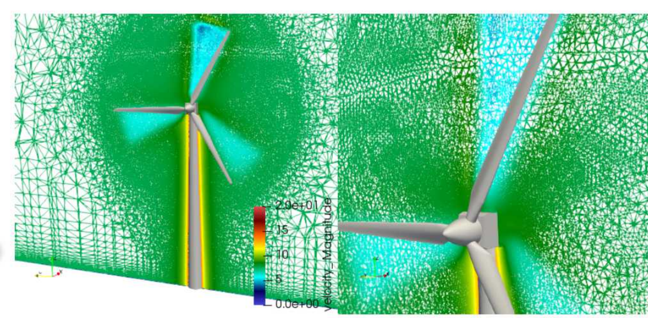

Circa FY17 View of Nalu

- MPI-based
- Lots of std::vector resize!!
- Physics within specialized heterogeneous algorithms, e.g., AssembleMomentumSolver
- Built upon NGP infrastructure Sierra Toolkit (STK) and Trilinos Tpetra
- Very much , a non-DSL design, i.e., the developer manages rhs(), lhs(), with rho*u_j*n_j()*dS

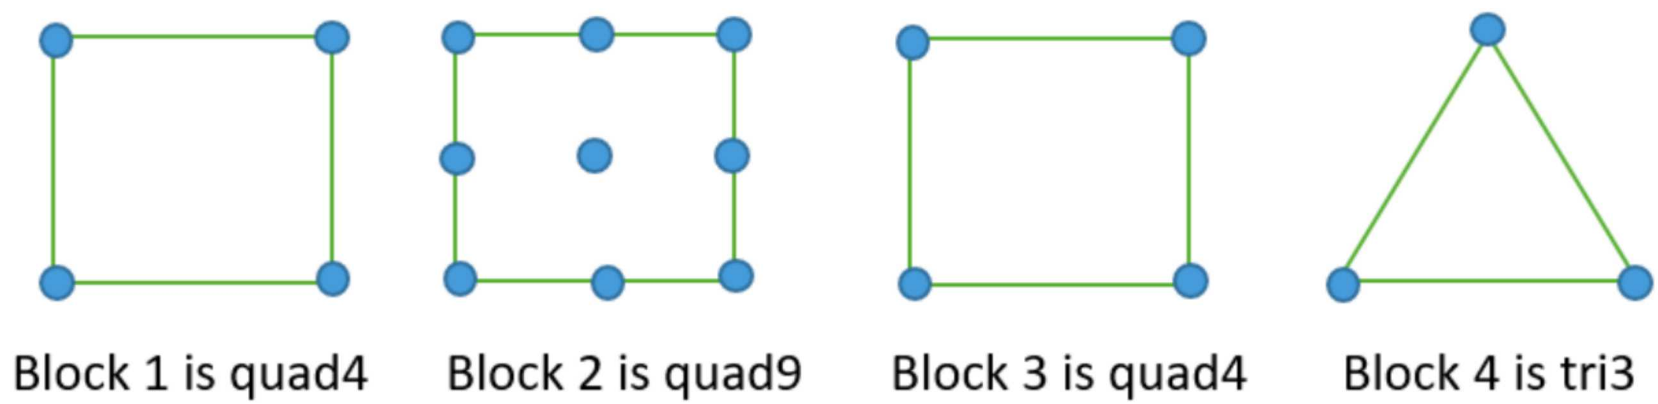

Heterogeneous mesh use case

Lowest Level; Nalu::Kernel

- Nalu::Kernel is templated on AlgTraits, e.g., integration rule such as nodesPerElement_, numIntgPoints_, etc.

```
template<typename AlgTraits>
MomentumNSOElemKernel<AlgTraits>::MomentumNSOElemKernel(
  ElemDataRequests& dataPreReqs)
{
  // define master element rule for this kernel
  MasterElement *meSCS
    = sierra::nalu::MasterElementRepo::get_surface_master_element(AlgTraits::topo_);

  // add ME rule
  dataPreReqs.add_cvfem_surface_me(meSCS);

  // add fields to gather
  dataPreReqs.add_coordinates_field(*coordinates_, AlgTraits::nDim_, CURRENT_COORDINATES);
  dataPreReqs.add_gathered_nodal_field(*velocityNp1_, AlgTraits::nDim_);

  // add ME calls
  dataPreReqs.add_master_element_call(SCS_GIJ, CURRENT_COORDINATES);
}
```

Attributes of a Nalu::Kernel

```
template<typename AlgTraits>
void
MomentumNSOElemKernel<AlgTraits>::execute(
  SharedMemView<DoubleType**>& lhs,
  SharedMemView<DoubleType *>& rhs,
  ScratchViews<DoubleType>& scratchViews)
{
  SharedMemView<DoubleType**>& v_uNp1
    = scratchViews.get_scratch_view_2D(*velocityNp1_);
  SharedMemView<DoubleType***>& v_gijUpper
    = scratchViews.get_me_views(CURRENT_COORDINATES).gijUpper;

  for ( int ip = 0; ip < AlgTraits::numScsIp_; ++ip ) {

    // determine scs values of interest
    for ( int ic = 0; ic < AlgTraits::nodesPerElement_; ++ic ) {

    // assemble each component
    for ( int k = 0; k < AlgTraits::nDim_; ++k ) {

      // determine scs values of interest
      for ( int ic = 0; ic < AlgTraits::nodesPerElement_; ++ic ) {

        // save off velocityUnp1 for component k
        const DoubleType& ukNp1 = v_uNp1(ic,k);

        // denominator for nu as well as terms for "upwind" nu
        for ( int i = 0; i < AlgTraits::nDim_; ++i ) {
          for ( int j = 0; j < AlgTraits::nDim_; ++j ) {
            gUpperMagGradQ += constant*v_gijUpper(ip,i,j);
          }
        }
      }
    }
  }
}
```

Thread-local scratch arrays using A Kokkos SharedMemView

Templated

MD-array rather than error-prone pointer arithmatic

Nalu Algorithm Abstraction: Kokkos integration, parallel_for()

- Classic SNL, Team-based nested Kokkos::parallel_for() model

```
Algorithm::execute()
{
  int bytes_per_thread = //scratch bytes per element
  auto team_exec = get_team_policy(... bytes_per_thread, ...);
  Kokkos::parallel_for(team_exec, buckets, [&](team)
  {
    ScratchViews scratchViews(topo, meSCS, dataNeeded ...);
    Kokkos::parallel_for(team, bucket.size()) {
      fill_pre_req_data(dataNeeded, elem, ..., scratchViews);
      for(kernel : computeKernels) {
         kernel->execute(elem, lhs, rhs, scratchViews);
      }
    }
    apply_coeff(lhs, rhs, ...);
  }
}
```

Nested parallel_for() thread-team parallelism.

Each bucket is processed by a team, inner loop over elements is split among SIMD/threads in a team.

scratchViews is created in outer loop, contains views into scratch-memory allocation (no heap-alloc is done here).

fill_pre-req_data() fills thread- local storage for a given element (shared over all kernels)

Nalu Supported Physics

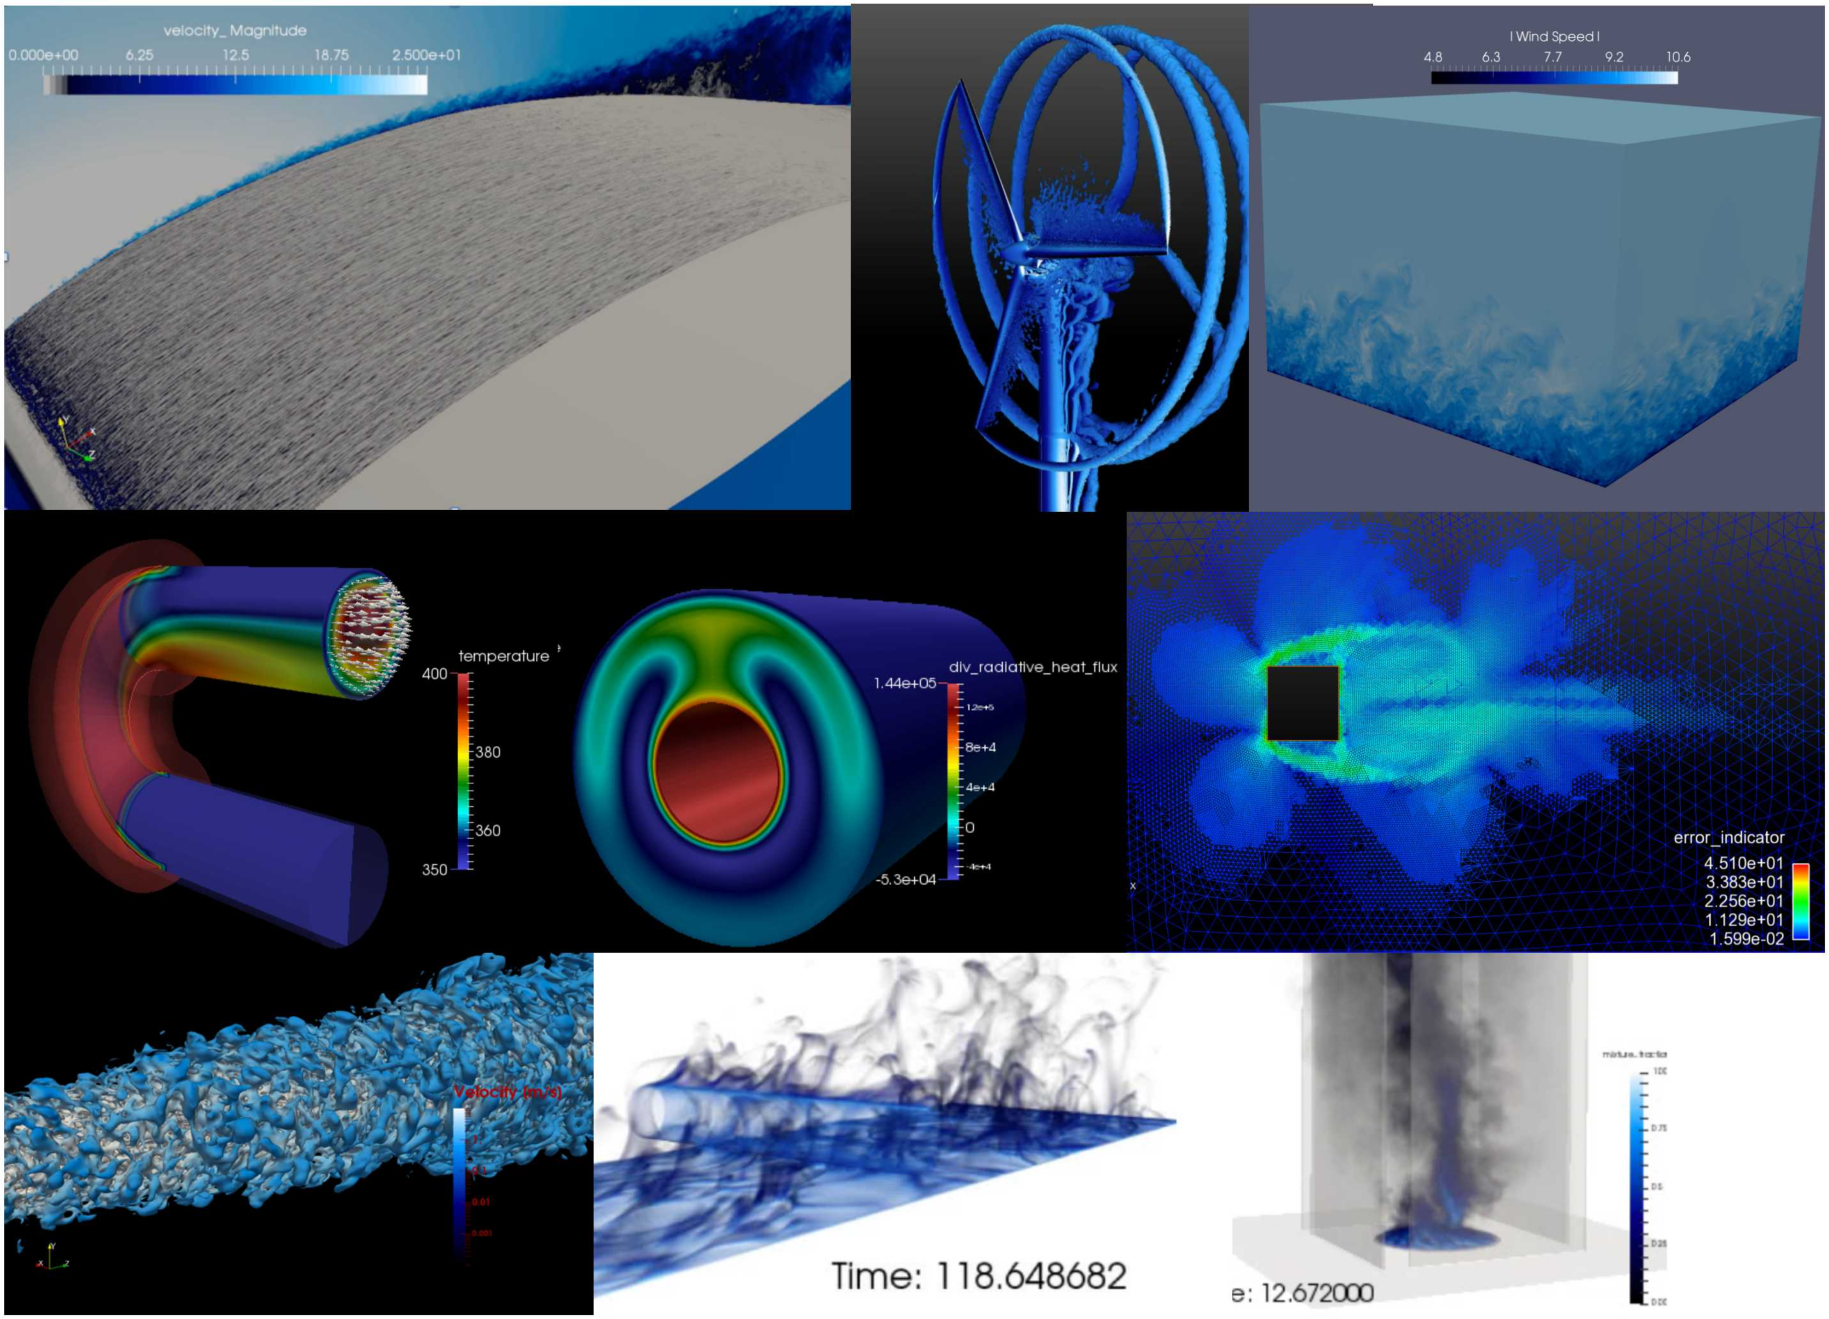

Conclusions

- Typical multi-physics applications include a wide range of time and length scales
- Complexity in desired use-case creates mesh burden to the user
- The ASC/NGS project has been created to provide foundational advances in providing a credible simulation study result
- Validation of several canonical turbulent flows indicate that current element-based discretization approach found in CVFEM provides acceptable accuracy for low-order hybrid element types
- In most cases, atypical element topologies are competitive on complex flows
- SGS decomposition approaches can provide an efficient manner in which structural uncertainty is obtained
- New NGP efforts incorporate Kokkos
- GPU activities are underway within the open source ExaWind Nalu code base

Acknowledgements

- Atypical mesh work draws on select works from Matt Barone (SNL Aerosciences) and Philip Sakievich (SNL Thermal/Fluids)
- ASC Next Generation Simulation (NGS) funding stream
- ECP ExaWind

Abstract

The suitability of hybrid meshes for low-Mach large-eddy simulation

Recent advances in computational hardware/core count have allowed for the large-eddy simulation (LES) technique to be deployed to a variety of multi-physics applications. In many cases, the complexity of the underlying solid geometry drives the extensive usage of *hybrid* meshes. Ideally, such meshes seek to utilize predominantly hexahedral topologies, however, in practice, production meshes frequently include Hex, Tet, Pyramid, and Wedge elements. Although hybrid meshing strategies are employed, the human-cost of mesh generation frequently remains the bottleneck within a general product construct cycle that typically consists of idea/conception, prototyping/improving, and, finally, the design/deployment iteration phase.

Very recently, Sandia National Laboratories, under the Advanced Simulation and Computing (ASC) project, has initiated a *Next Generation Simulation* foundational research project that seeks to improve the throughout, effectiveness, and credibility of multi-physics simulation analysis in support of Stockpile Stewardship. Key research aspects of this project include the development of advanced discretization schemes, error indicators, VVUQ, and efficient high-quality mesh generation. Whereas formally the computational scientist has tacitly prescribed the desired mesh type, i.e., hexahedral-based, a paradigm shift that drives accurate and scalable methods development on next generation platforms (NGP) that conforms to the fast-meshing generation archetype must be realized.

In this seminar, a wide range of simulations of interest to the low-Mach LES application space are presented that employ Hex8, Hex27, Hex64, Tet4, Wedge6, and Pyramid5 topologies. In addition to verification and validation use-cases, the talk will overview recent advances in quantifying structural uncertainty in LES that include eigenvalue perturbation of the decomposed LES subgrid scale model. Finally, next generation platform (NGP) transition efforts will be overviewed within the open-source low-Mach code, Nalu, with an emphasis on describing the transition of critical low-Mach physics algorithms to a nested team-based Kokkos parallel_for() paradigm.