Progress in CFD Discretizations Algorithms and Solvers for Aerodynamic Flows

Dimitri Mavriplis

University of Wyoming

What this talk covers:

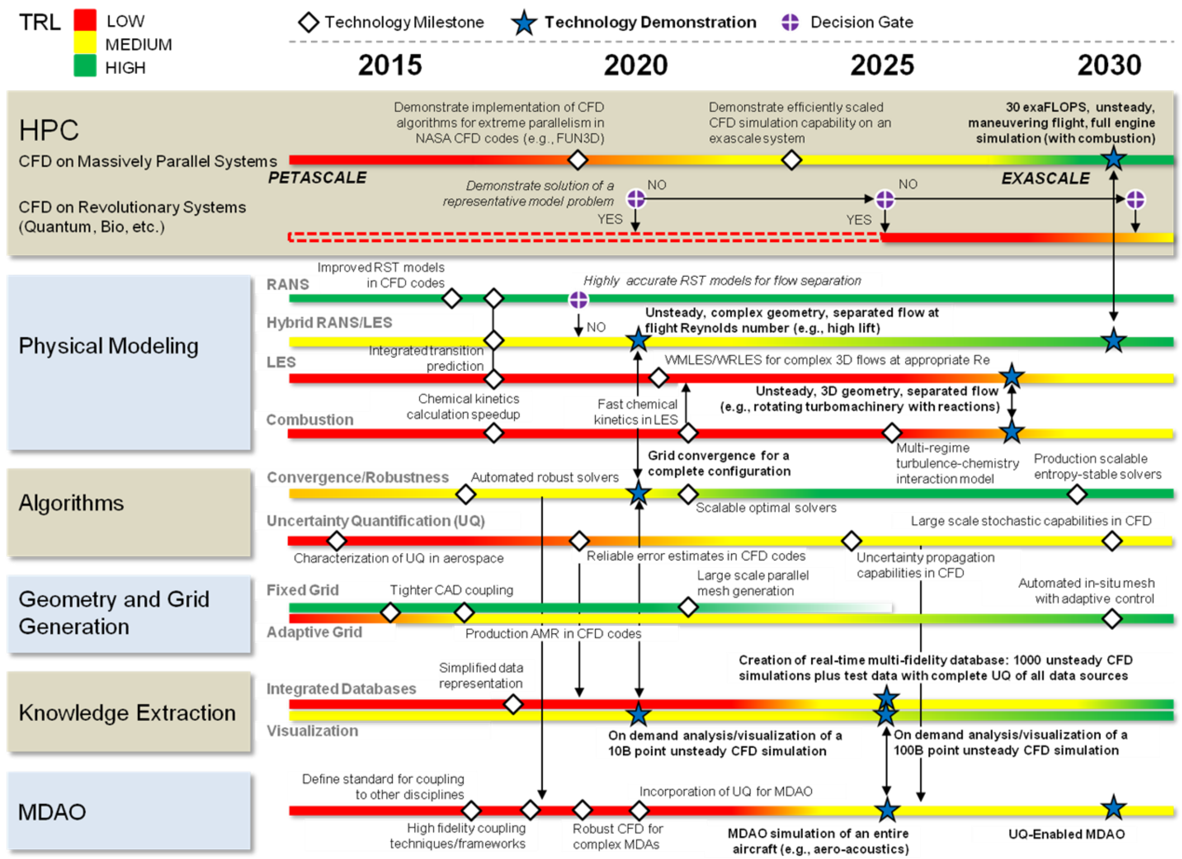

- **Algorithms:**
  - **Discretizations and Solvers**
  - **Steady/Unsteady Reynolds Averaged Navier-Stokes (RANS)**
  - **Scale resolving methods (LES, DES)**
  - Uncertainty quantification not covered

Other Areas:

HPC, Phys. Modeling, MDAO etc.

- **Only as they affect choices in Discretizations and Solvers**

Overview

- Community Efforts
- RANS Methods
  - Second-order accurate methods (FV and SUPG@p=1)
  - Higher-order accurate methods (DG and SUPG)
- Scale resolving methods
  - Second-order accurate methods
  - Higher-order accurate methods
  - Explicit vs Implicit
- Conclusions

Community Efforts

- Drag Prediction Workshop (DPW)
- High-Lift Prediction Workshop (HLPW)
- Aeroelastic Prediction Workshop (AePW)
- Hover Prediction Workshop (HPW)
- Certification by Analysis
- Geometry and Mesh Generation workshops (GMGW)
- International High-order Methods Workshop
- AIAA CFD solver discussion group
- Other…

Turbulence Modeling Resource

https://turbmodels.larc.nasa.gov

Langley Research Center

Turbulence Modeling Resource

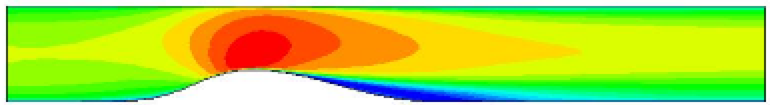

Purpose

The purpose of this site is to provide a central location where Reynolds-averaged Navier-Stokes (RANS) turbulence models are documented. This effort is guided by the Turbulence Model Benchmarking Working Group (TMBWG), a working group of the Fluid Dynamics Technical Committee of the American Institute of Aeronautics and Astronautics (AIAA).

The objective is to provide a resource for CFD developers to:

- obtain accurate and up-to-date information on widely-used RANS turbulence models, and
- verify that models are implemented correctly.

This latter capability is made possible through "verification" cases. This site provides simple test cases and grids, along with sample results (including grid convergence studies) from one or more previously-verified codes for some of the turbulence models. Furthermore, by listing various published variants of models, this site establishes naming conventions in order to help avoid confusion when comparing results from different codes.

The site should also help CFD code users to understand and compare the predictions of a variety of models on the fundamental flow problems in the validation database. Note that it is not the intention of this effort to provide validation of turbulence models for a wide range of complex flows for diverse applications. While this would undoubtedly be valuable, it is beyond the scope of what can be supported. Instead, the goal is to provide a set of test cases that illustrate the performance of models for flows that capture fundamental phenomena, in order to establish a consistent basis of comparison as a starting point from which a more thorough validation effort for flows of specific interest to users and developers can be conducted.

Finally, the site should serve as a forum for model developers to help disseminate new models to the CFD community.

It is anticipated that this site will be updated regularly as new models and/or verification/validation cases are incorporated and tested. If you have any questions or comments, please contact: Chris Rumsey of NASA Langley, Brian Smith of Lockheed-Martin, or George Huang of Wright State University.

https://turbmodels.larc.nasa.gov/onerawingnumerics_val.html

Langley Research Center

Turbulence Modeling Resource

Return to: Turbulence Modeling Resource Home Page

TURBULENCE MODEL NUMERICAL ANALYSIS

3D ONERA M6 Wing Validation Case

The purpose here is to provide a test case for a turbulent flow over a transonic wing. Over the years, the ONERA M6 experiment (Schmitt, V. and Charpin, F., "Pressure Distributions on the ONERA-M6-Wing at Transonic Mach Numbers," Experimental Data Base for Computer Program Assessment. Report of the Fluid Dynamics Panel Working Group 04, AGARD AR-138, May 1979) has been a widely used case for CFD "validation." This wing is used here primarily for numerical analysis of turbulence model simulations; e.g., convergence properties, effect of order of accuracy, etc.

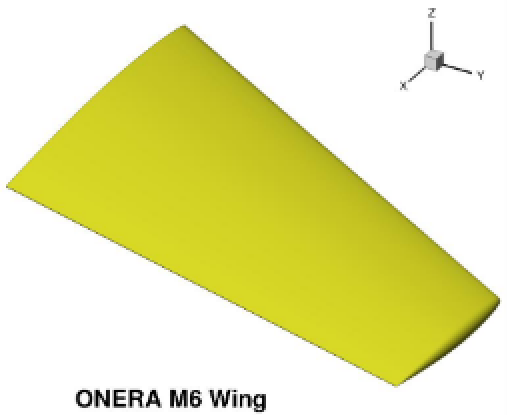

ONERA M6 Wing

Recently, a group at ONERA has looked into the M6 model and its past experiments in greater detail. See AIAA Papers 2015-1745 and 2016-1357. As part of this effort, the group has created a CAD geometry for the wing, provided below. In this geometry, the trailing edge of the wing has been made sharp for the purposes of this particular CFD exercise, as described in AIAA Paper 2016-1357. (The original ONERA M6 wing has a moderately thick trailing edge: see AileM6_with_thick_TE.igs.) Additional details are provided in: ONERA_M6_Test_Case_TMR.pdf

Overview

- Community Efforts
- RANS Methods
  - Second-order accurate methods (FV and SUPG@p=1)
  - Higher-order accurate methods (DG and SUPG)
- Scale resolving methods
  - Second-order accurate methods
  - Higher-order accurate methods
  - Explicit vs Implicit
- Conclusions

SUPG vs DG Characteristics

- SUPG/Continuous Galerkin
  - At p=1 same dofs as Finite-Volume
  - Mesh curving not required@p=1
    - Can run on same meshes as FV
- DG is element based
  - At p=1 More dofs on same mesh
- DG has nice block matrix structure
  - Nearest neighbor stencil for all p
  - Element based block
  - At high p well suited for
    - AMR (simple stencil)
    - Overset (simple stencil)
    - HPC (computationally intensive)

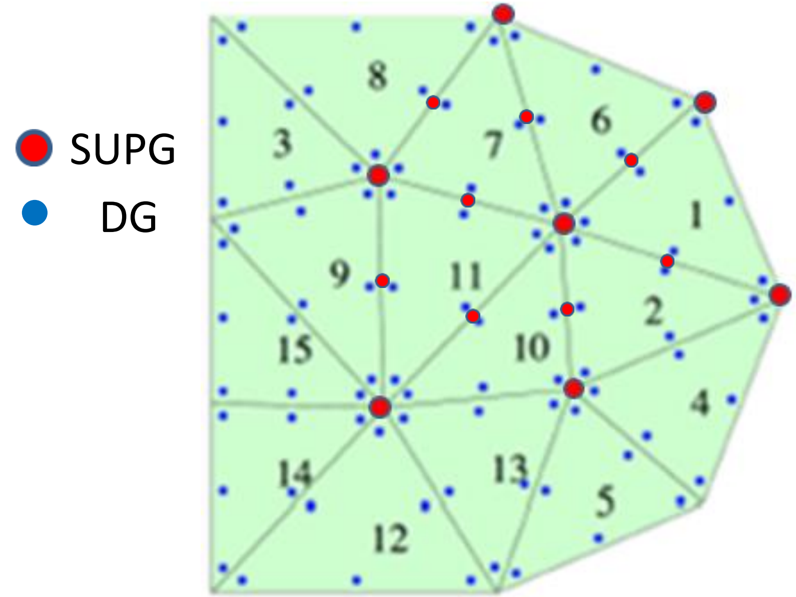

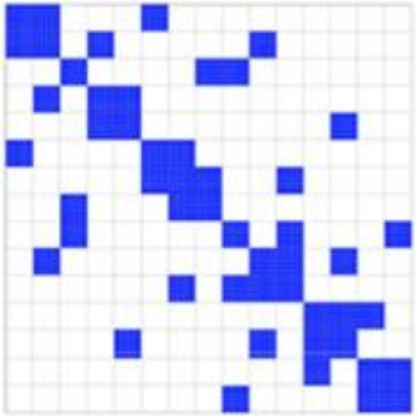

DG BLOCK MATRIX

SUPG vs DG at p=1,2,3

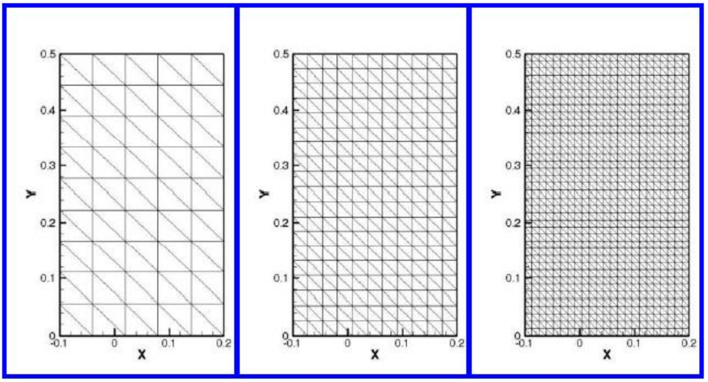

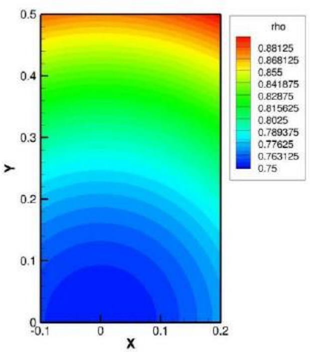

$\rho = 0.5(\sin(x^2 + y^2) + 1.5)$

MMS

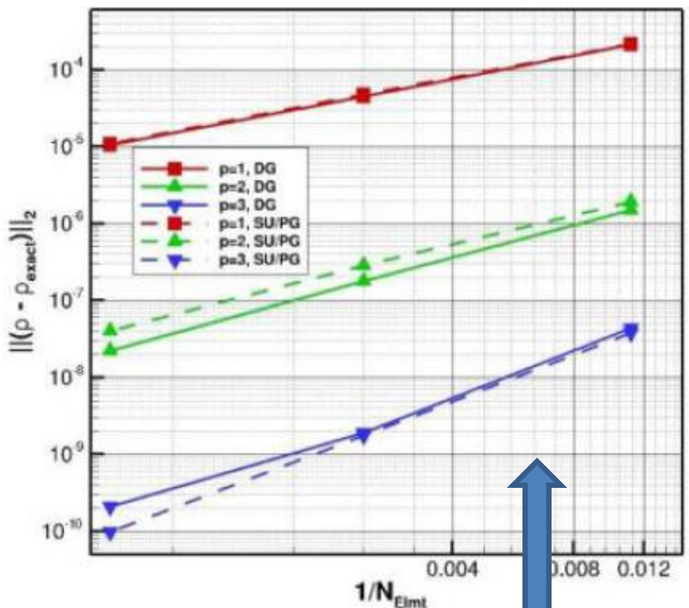

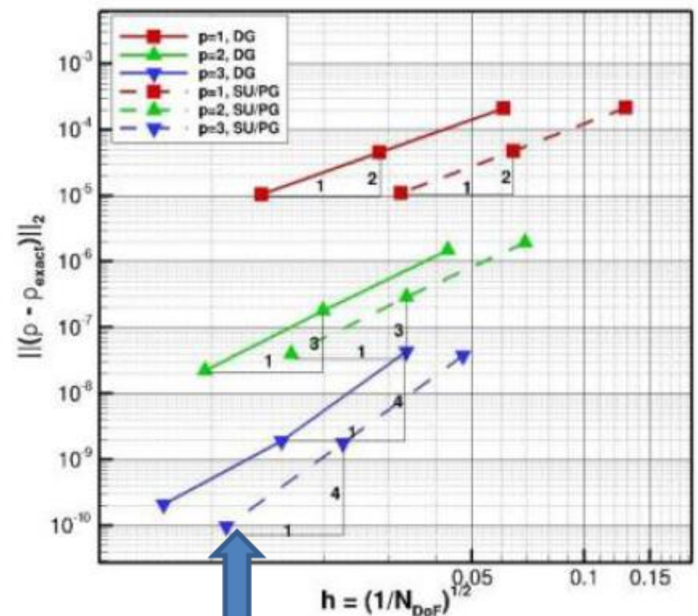

- Equivalent accuracy on same grid at same p-order
- DG has more degrees of freedoms (dofs) on same grid
  - More computational expense
- **Caveat: low Re test case**
  - **DG advantages for hyperbolic problems**
- At high p, DG has other advantages (more later)

Reproduced from: Glasby, Burgess, Anderson, Wang, Mavriplis, Allmaras, AIAA Paper 2013-0691

SUPG@p=1 Exhibits Higher Accuracy than Finite-Volume on Same Mesh

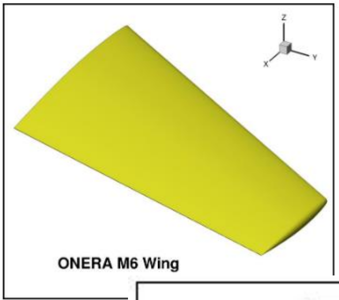

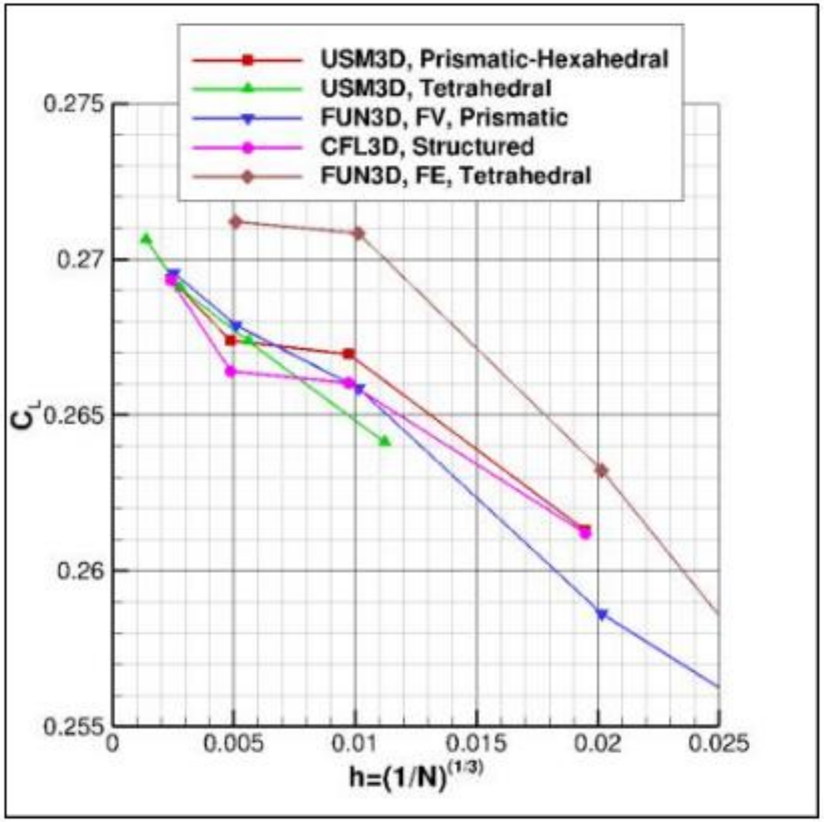

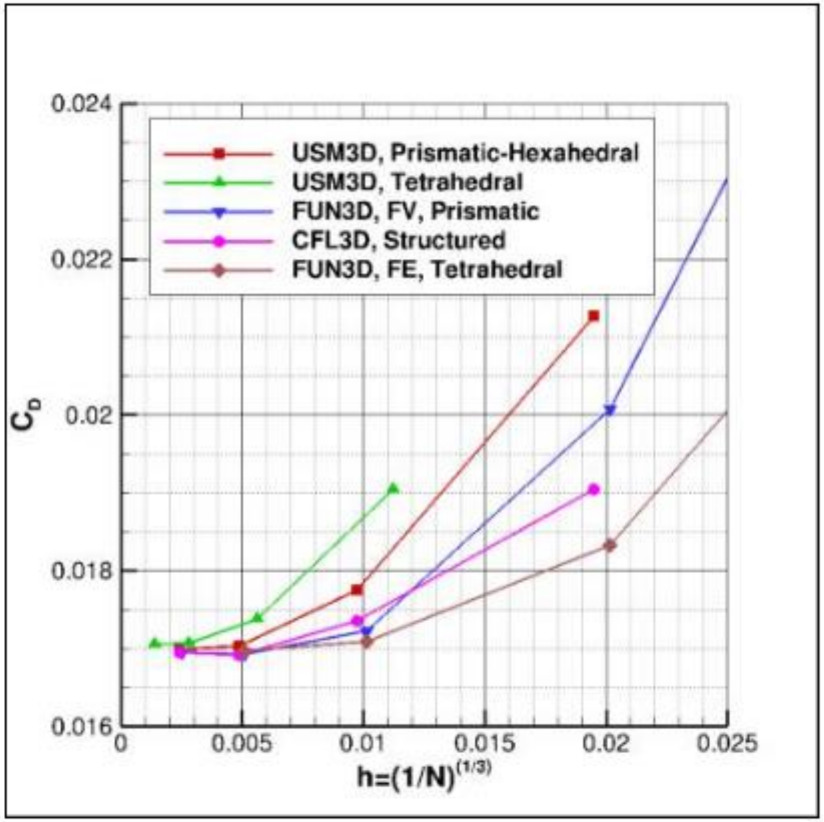

- ONERA M6 Results on NASA TMR web-site
  - Fun3D FE arguably better than FUN3D FV….

SUPG@p=1 Exhibits Higher Accuracy than Finite-Volume on Same Mesh

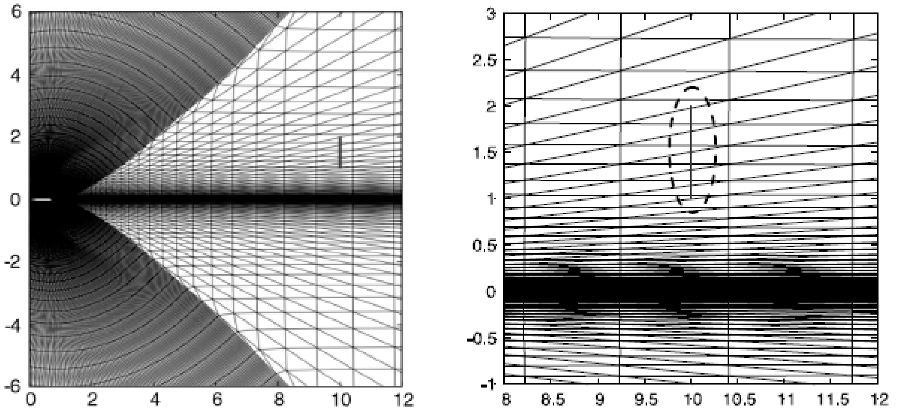

a) Medium-field view of profile location b) Near-field view of profile location

Fig. 14 Grid density in vicinity of profile for mesh used for quadratic elements with 230,000 DOF

AIAA Journal
Vol. 54, No. 9, September 2016

CrossMark

Finite Element Solutions for Turbulent Flow over the NACA 0012 Airfoil

W. Kyle Anderson,* Behzad R. Ahrabi,† and James C. Newman‡
University of Tennessee at Chattanooga, Chattanooga, Tennessee 37403

DOI: 10.2514/1.J054508

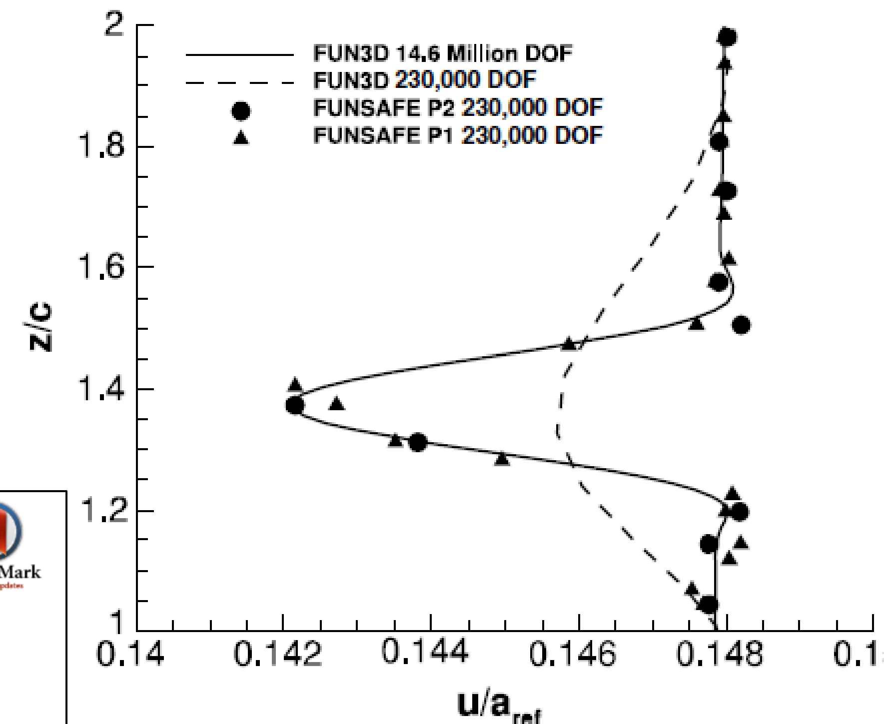

- Impressive wake resolution for SUPG@p=1 on coarse mesh
- However: SUPG often more expensive to converge
  - Importance of solver technology

Illustration of Solver Efficiency

Easy test case

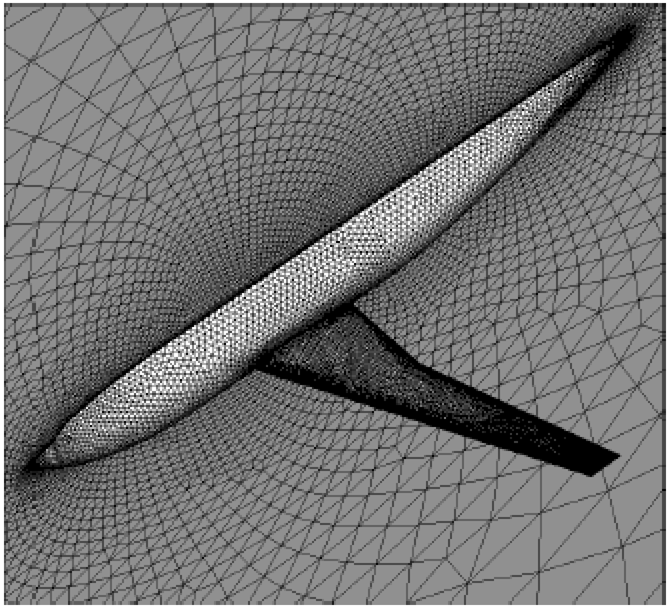

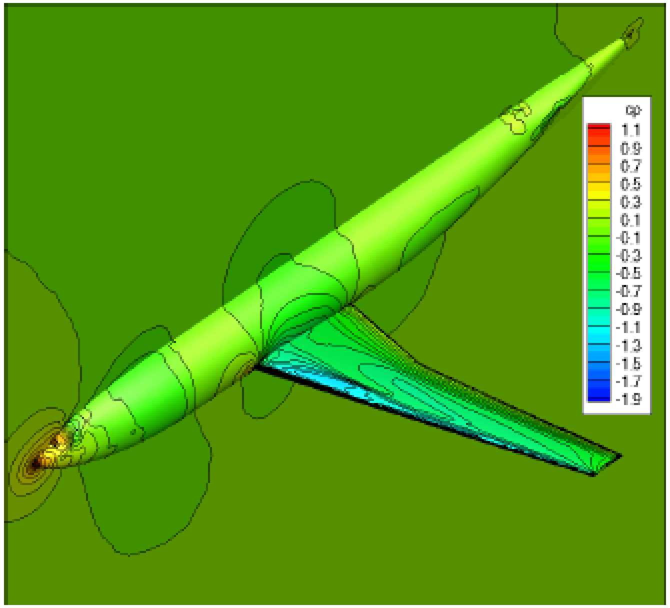

- F6 Wing-body (DPW3)
- Mach=0.75, Incidence=1deg, Re=3 million
- Prism-Tet Mesh: 1.2 million points (~3 million elements)

NSU3D Solutions for WB Test Case

1.2 million points on 128 cores

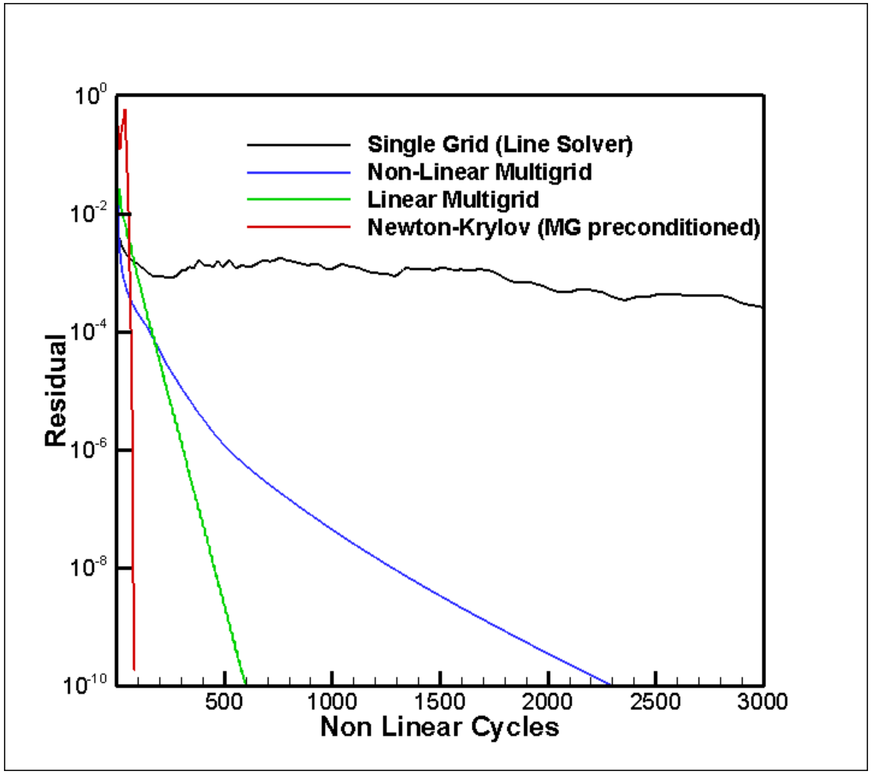

- Single grid solver is slow to converge
- FAS MG is much faster
- Linear MG is fastest
- Newton-Krylov takes only 88 nonlinear steps
  -

NSU3D Solutions for WB Test Case

1.2 million points on 128 cores

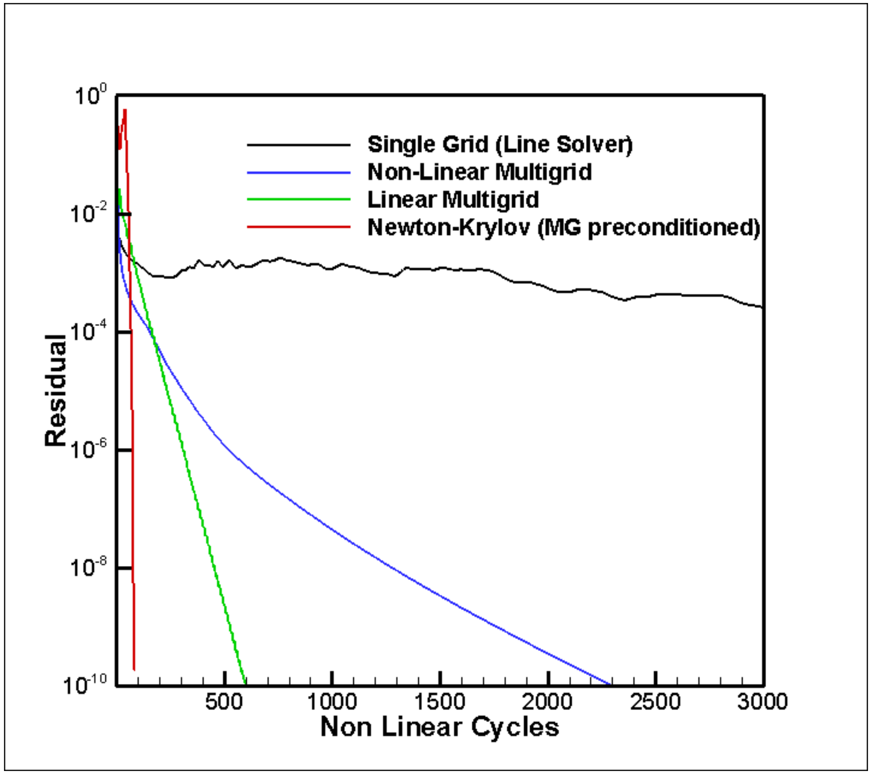

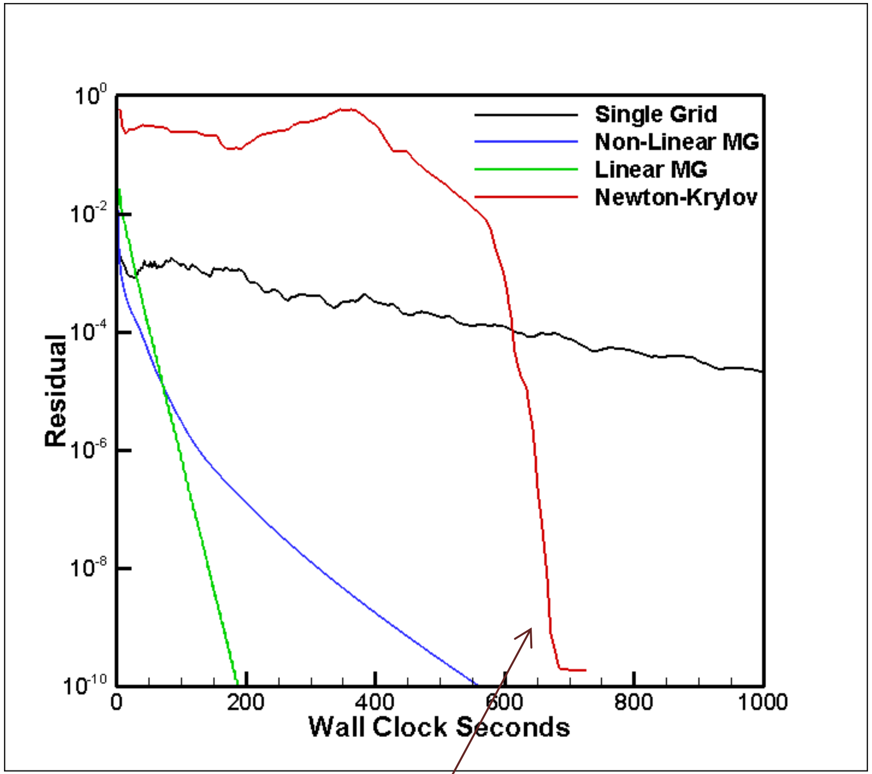

- Single grid solver is slow to converge
- FAS MG is much faster
- Linear MG is fastest
- Newton-Krylov takes only 88 nonlinear steps
  - But cost is higher due to slow initial convergence

NSU3D for HLPW2 Mesh Refinement Study
(More Difficult)

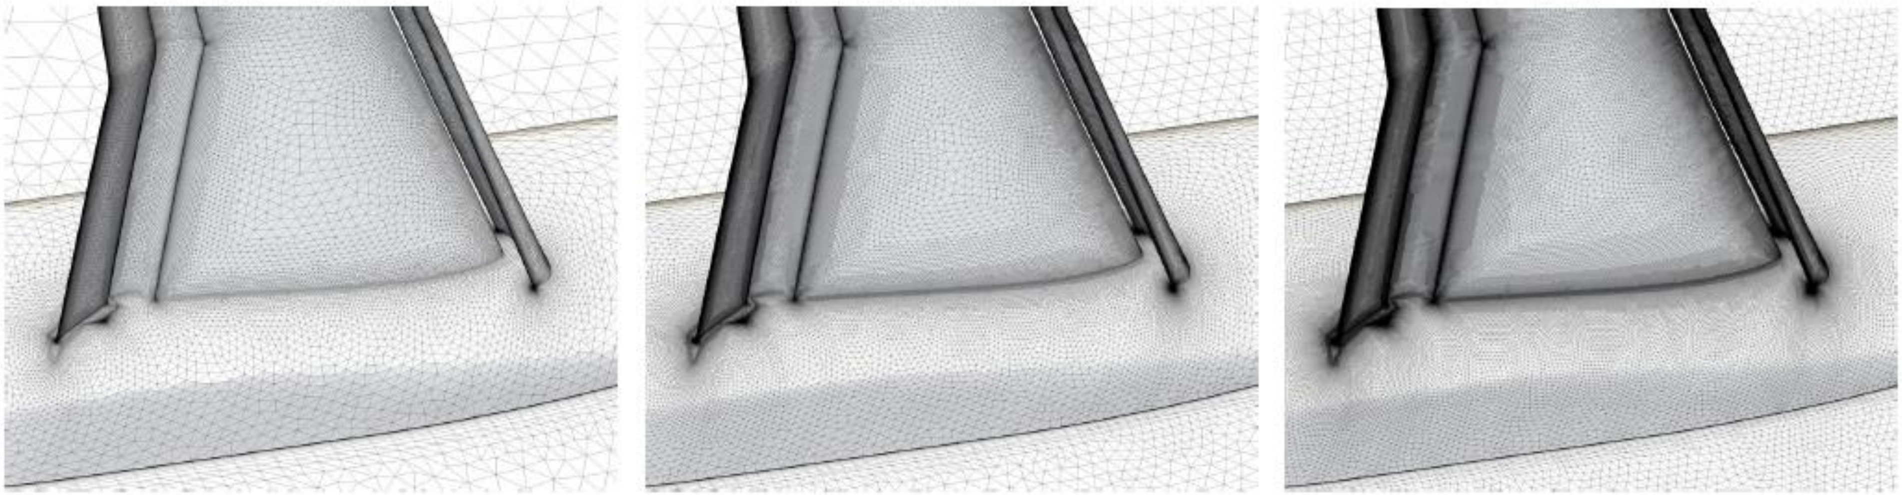

(a) Coarse (b) Medium (c) Fine

- Mach=0.175, Incidence=16deg, Re=15 million
  - Coarse Mesh: 10 million points
  - Medium Mesh: 30 million points
  - Fine Mesh: 75 million points

NSU3D for HLPW2 Mesh Refinement Study

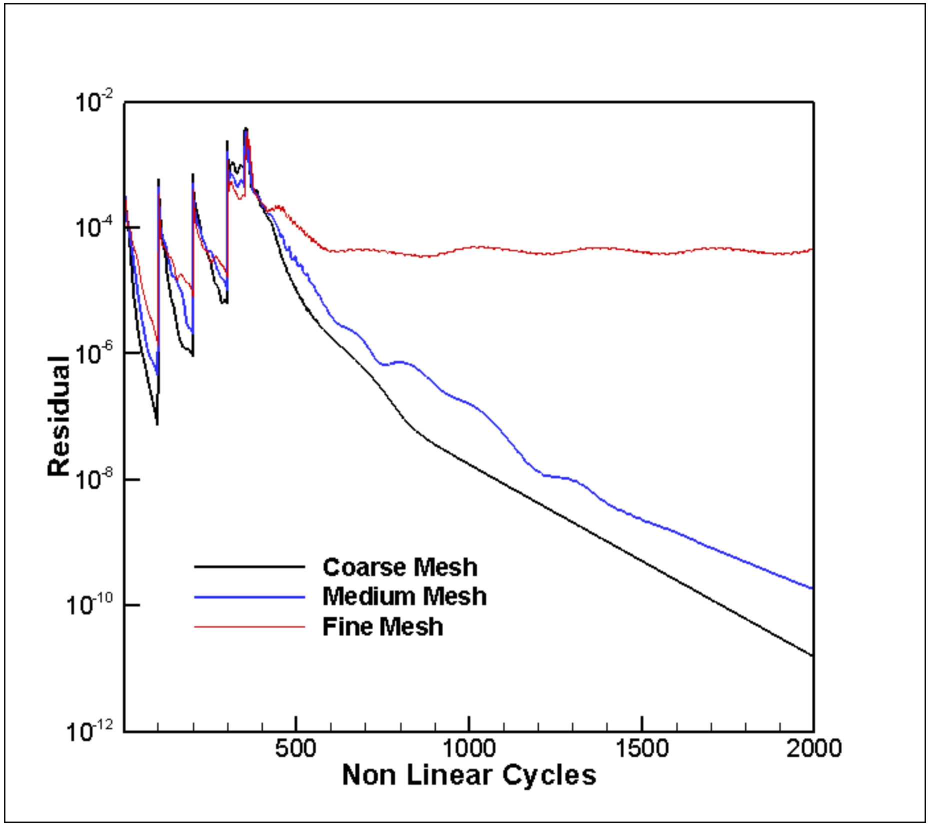

- FAS MG converges fully only on coarsest mesh
- Linear MG converges on coarse/medium, stalls on fine mesh
- Newton-Krylov converges fine mesh at considerable extra cost
  - Time-averaged forces from Linear MG on fine mesh very close to Newton final values

NSU3D for HLPW2 Mesh Refinement Study

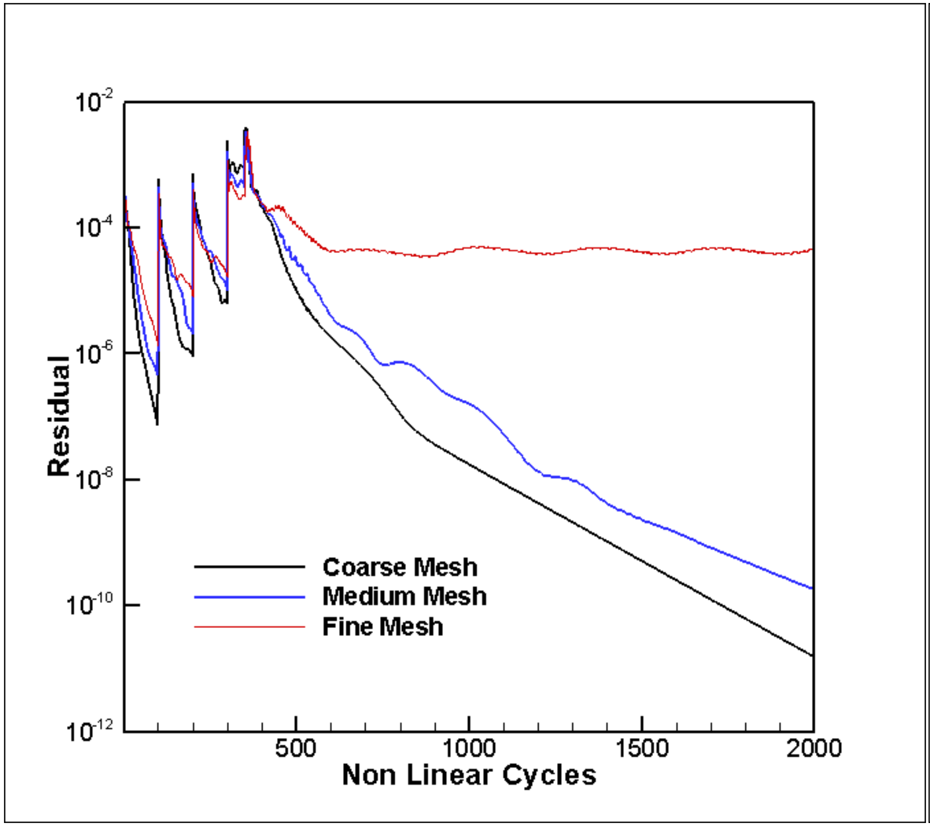

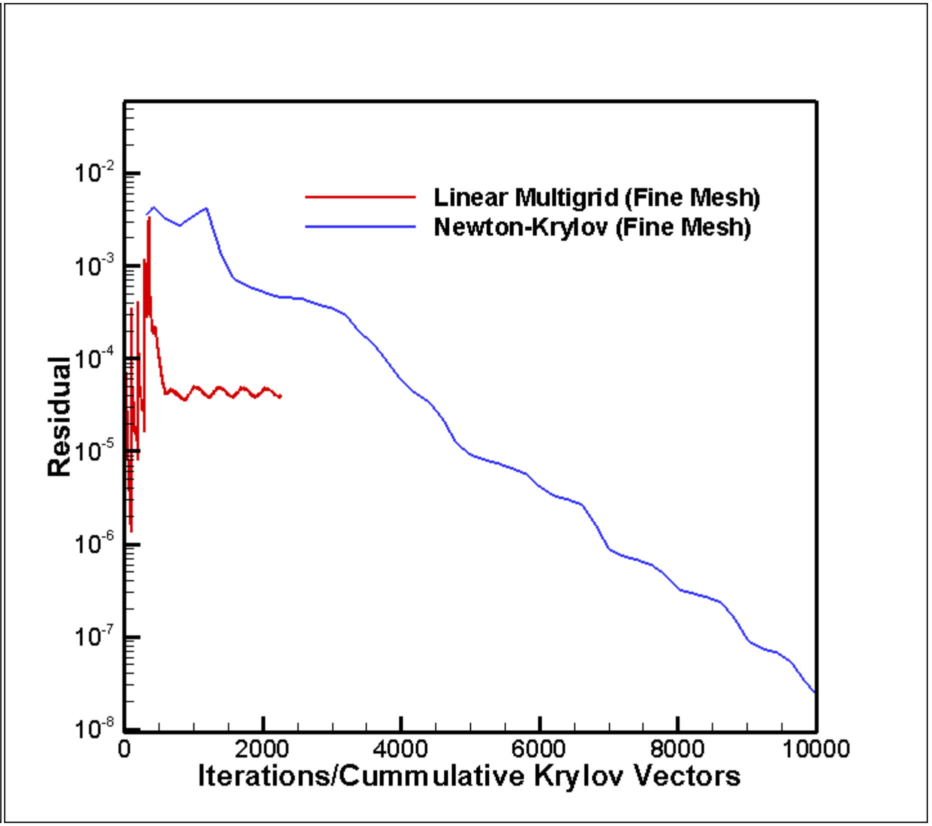

- FAS MG converges fully only on coarsest mesh
- Linear MG converges on coarse/medium, stalls on fine mesh
- Newton-Krylov converges fine mesh at considerable extra cost
  - Time-averaged forces from Linear MG on fine mesh very close to Newton final values

Hierarchy of Solvers

- FAS Multigrid
  - Fast when works
  - No tuning parameters
- Linear Iterative Solver (MG, GS, Lines, etc)
  - Somewhat more robust
  - Some tuning parameters
    - linear tol. , inner cycles, CFL ramping
- Newton-Krylov
  - Most robust
  - Considerably slower when other methods converge
  - Effective in final stages of convergence
  - Slow initial convergence
  - Forces/moments do not converge only at end !
- SUPG/DG methods only practical using Newton-Krylov methods
- Importance of improved solver technology
  - For ALL CFD DISCRETIZATIONS
  - For MDA/MDAO

Solver Technology

Linear and Nonlinear Solvers

- Linear Solvers (or preconditioners)
  - Components of nonlinear solver
  - Important to port efficiently to hardware
  - Easily abstracted into libraries and reused
  - Matrix factorization (ILU) used in SUPG/DG methods
    - Non-iterative
    - Robust (increases with more fill-in k>0)
    - Memory intensive (with larger fill in)
    - Not amenable to large scale parallelization (partition local)
  - Iterative Line solvers becoming more common
  - Algebraic Multigrid (AMG) libraries developed (DoE)
    - Work is on-going to assess AMG issues with CFD discretizations

Linear Solver Technology (Pandya et al. AIAA 2016-0860)

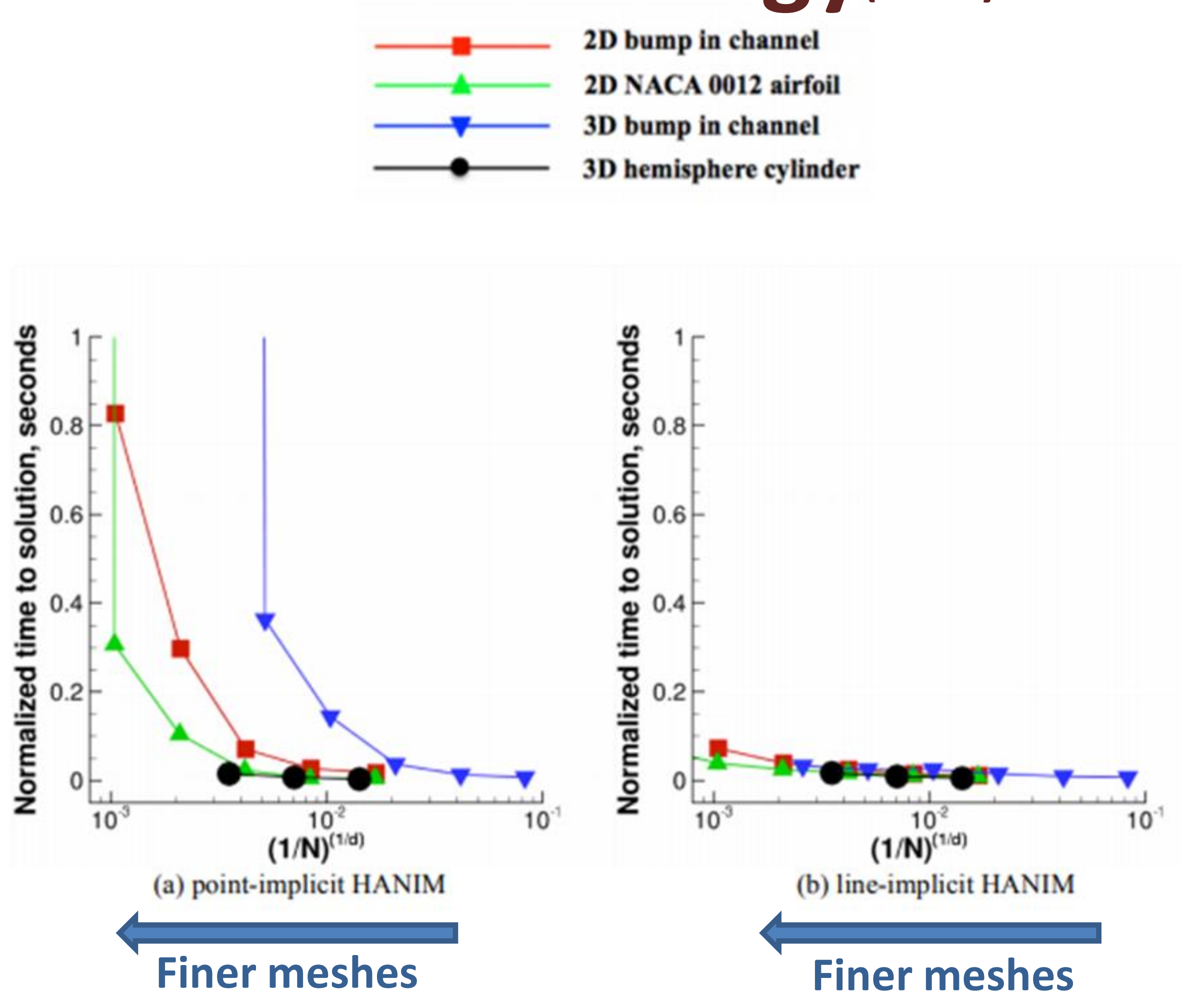

- Improved solution time in USM3D using line vs point linear solver in Newton-Krylov method

ILU(k) vs Iterative Line (Preconditioners)

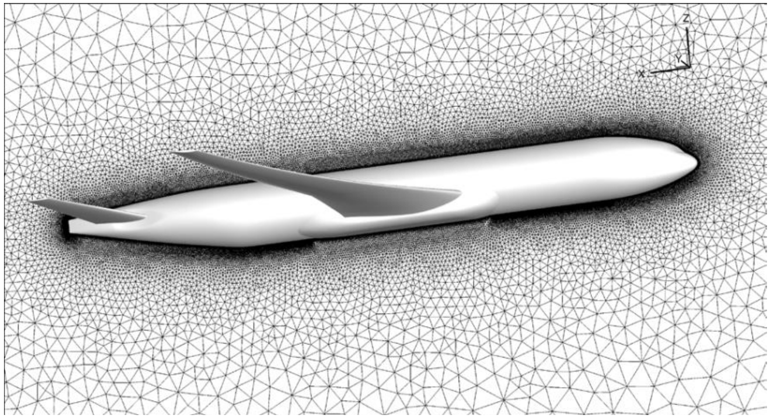

SUPG@p=1

From: Ahrabi and Mavriplis, AIAA 2017-0517

CRM WBT from DPW4
6.2 million grid points

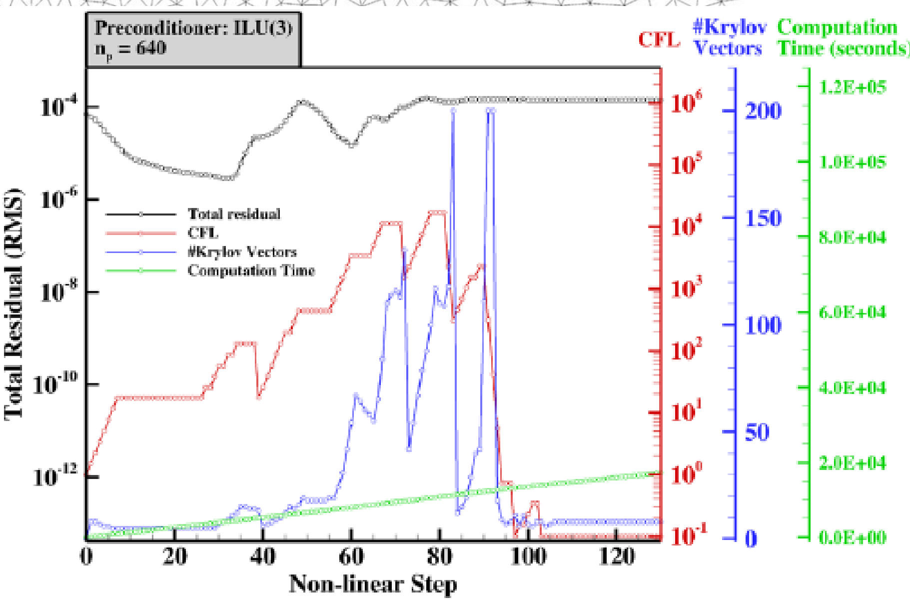

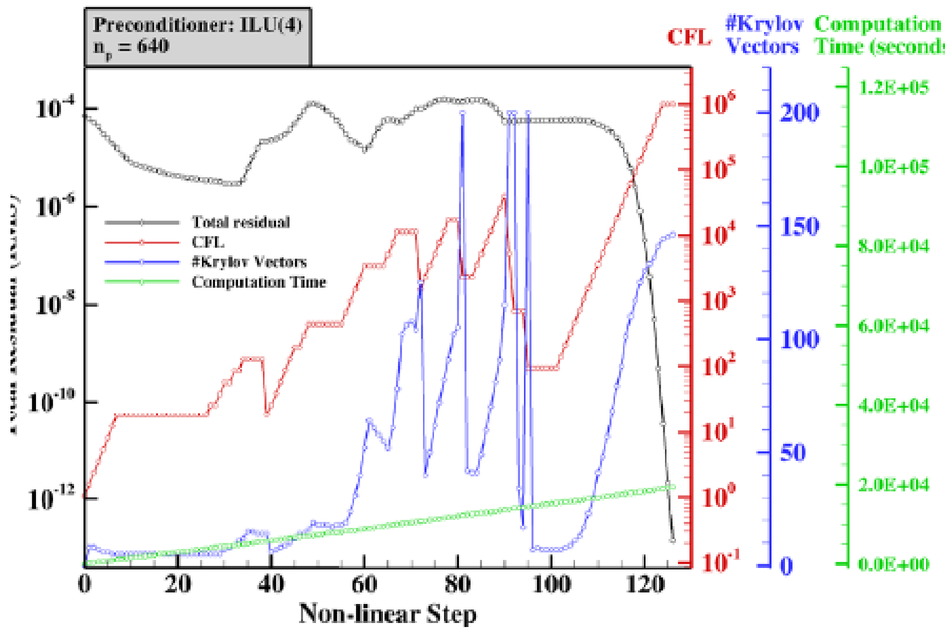

- ILU(3) fails to converge on 640 cpus

- ILU(4) converges on 640 cpus

ILU(k) vs Iterative Line (Preconditioners)

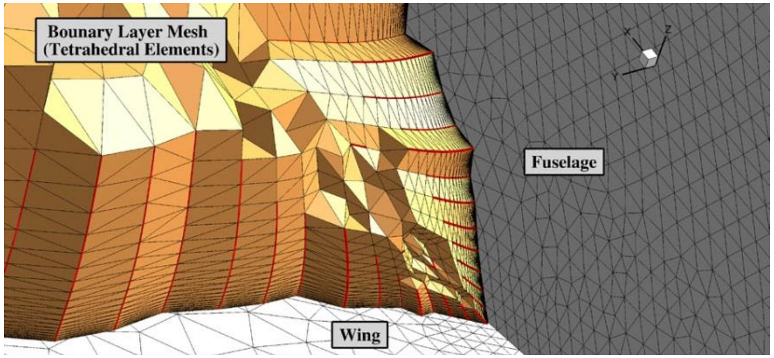

SUPG@p=1

From: Ahrabi and Mavriplis, AIAA 2017-0517

CRM WBT from DPW4
6.2 million grid points
Line structures in boundary layer

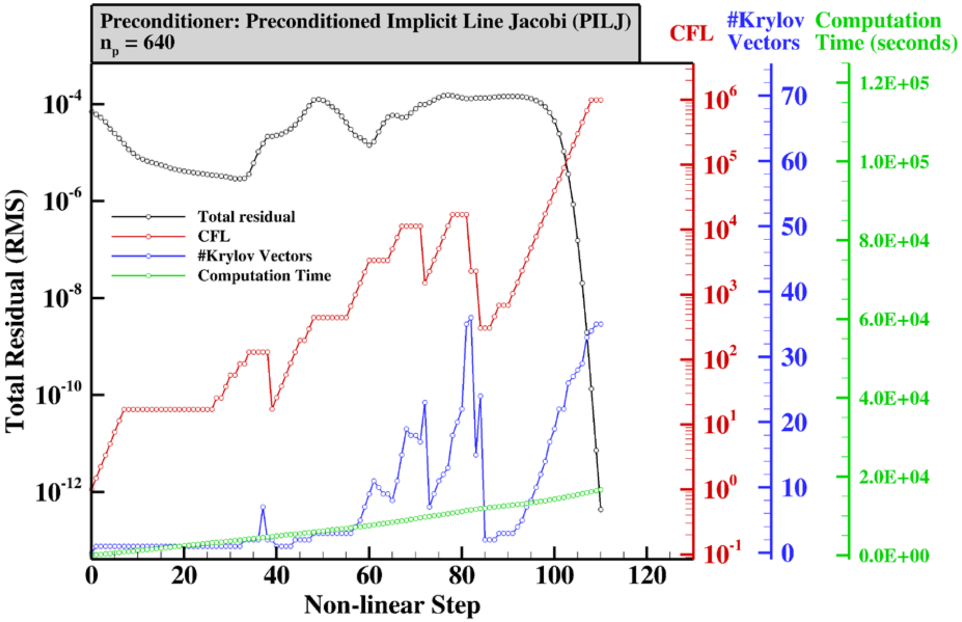

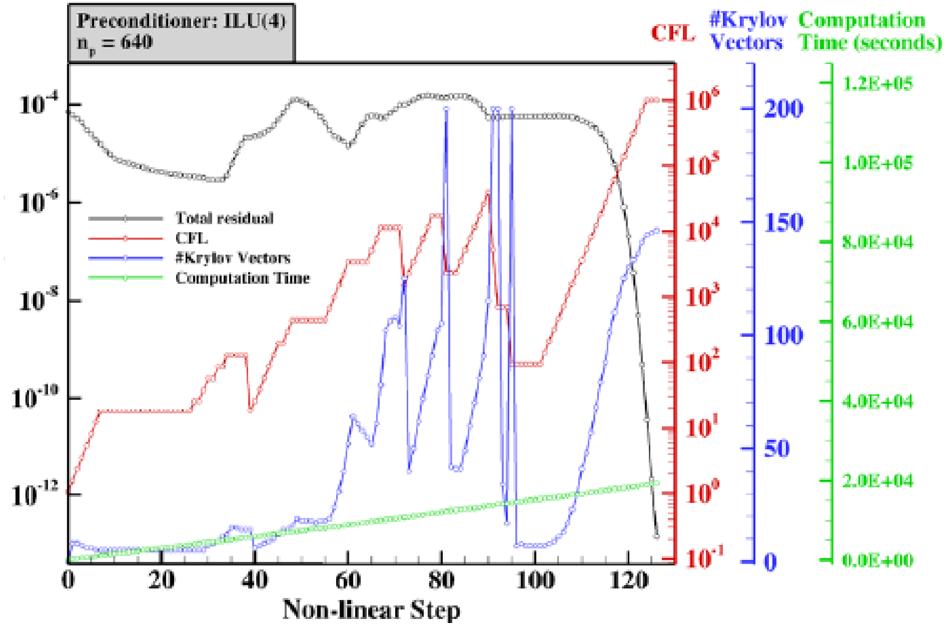

- Iterative line solver converges identically for any partitioning
  - Requires dual CFL for stability

- ILU(4) converges on 640 cpus

Limitations of ILU(k)

From: Ahrabi and Mavriplis, AIAA 2017-0517

	PILJ			ILU(2)			ILU(3)			ILU(4)		
n_p	n_{NL}	k_{max}	t_{tot}	n_{NL}	k_{max}	t_{tot}	n_{NL}	k_{max}	t_{tot}	n_{NL}	k_{max}	t_{tot}
80	110	36	1.202e+5	NM	NM	NM	NM	NM	NM	NM	NM	NM
160	110	36	6.203e+4	NC	NC	NC	NM	NM	NM	NM	NM	NM
320	110	36	3.188e+4	NC	NC	NC	124	200	3.204e+4	NM	NM	NM
640	110	36	1.671+4	NC	NC	NC	NC	NC	NC	126	200	1.936e+4

n_p = Number of processors (and partitions)

n_{NL} = Number of non-linear steps

k_{max} = Maximum of number of Krylov vector used in the solution

t_{tot} = Total elapsed time in seconds

NM = Not enough memory to run

NC = No convergence

- Numerical efficiency degrades on many cpus
- Increased k-fill incurs additional memory
- Can be mitigated with Overlap, Schur, shared mem parallelism…
- Investigating line and block iterative preconditioners for SUPG

Solver Technology

Nonlinear Solvers

- Non-Linear Solvers
  - FAS Multigrid
    - Very effective (optimal) but robustness issues remain
- Newton Methods more robust
  - Weakness is slow initial convergence
  - Accelerating initial convergence
    - Continuation methods (Mesh, homotopy, other…)
    - Residual smoothing
    - Multigrid (or coarse grid correction)

Accelerating Newton-Krylov Initial Convergence

From: Mavriplis et al. AIAA 2019-0100

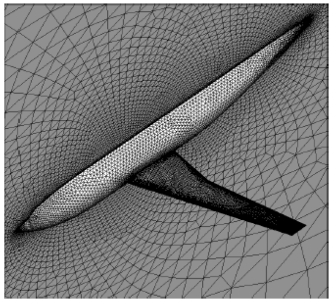

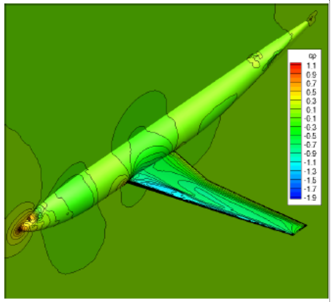

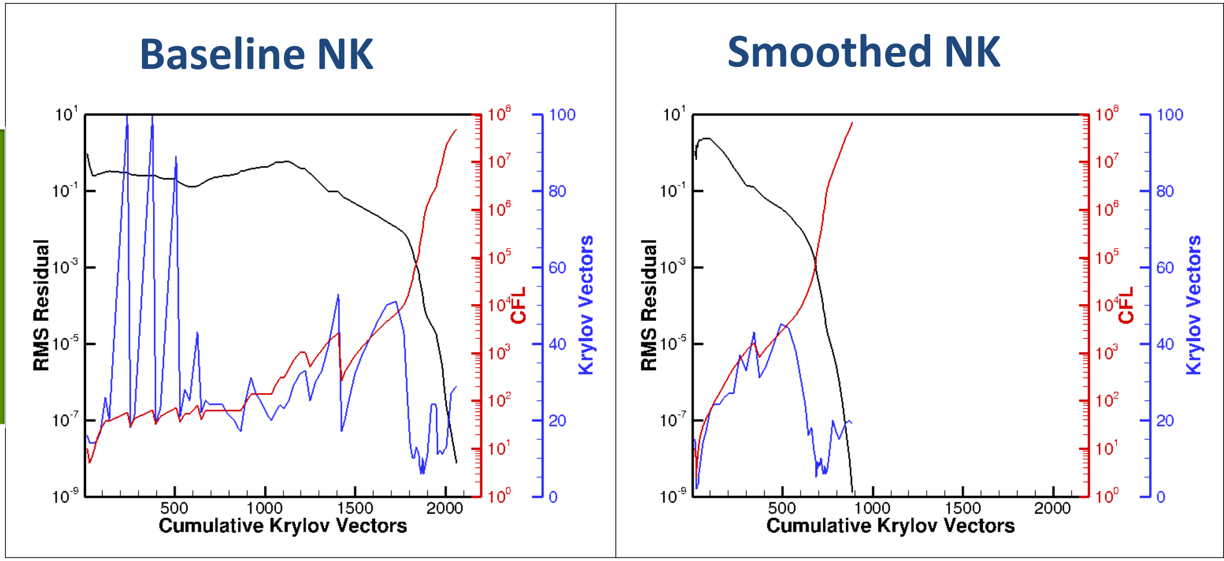

- Residual smoothing operator constructed from small number of local nonlinear passes
  - nonlinear point or line solver
- Produces significant gains in nonlinear convergence

Krylov Methods

- Linear solver that only requires matrix-vector product: $Ax=b$ $Av, A^2v, A^3v, \ldots$
- Provable monotone decrease in residual (may stall)
- Requires good preconditioner
  - Use linear solvers as preconditioners
  - Can wrap Krylov method around existing linear solver to produce better linear solver
- Becoming more ubiquitous in CFD applications
- Obvious choice for adjoint solvers (linear)
- Simple and effective for building strong multidisciplinary coupled solvers

Newton-Krylov for Monolithic (tightly coupled) Aero-structural Solver

Zhang and Zingg AIAA Journal, Vol. 56, No. 3, March 2018

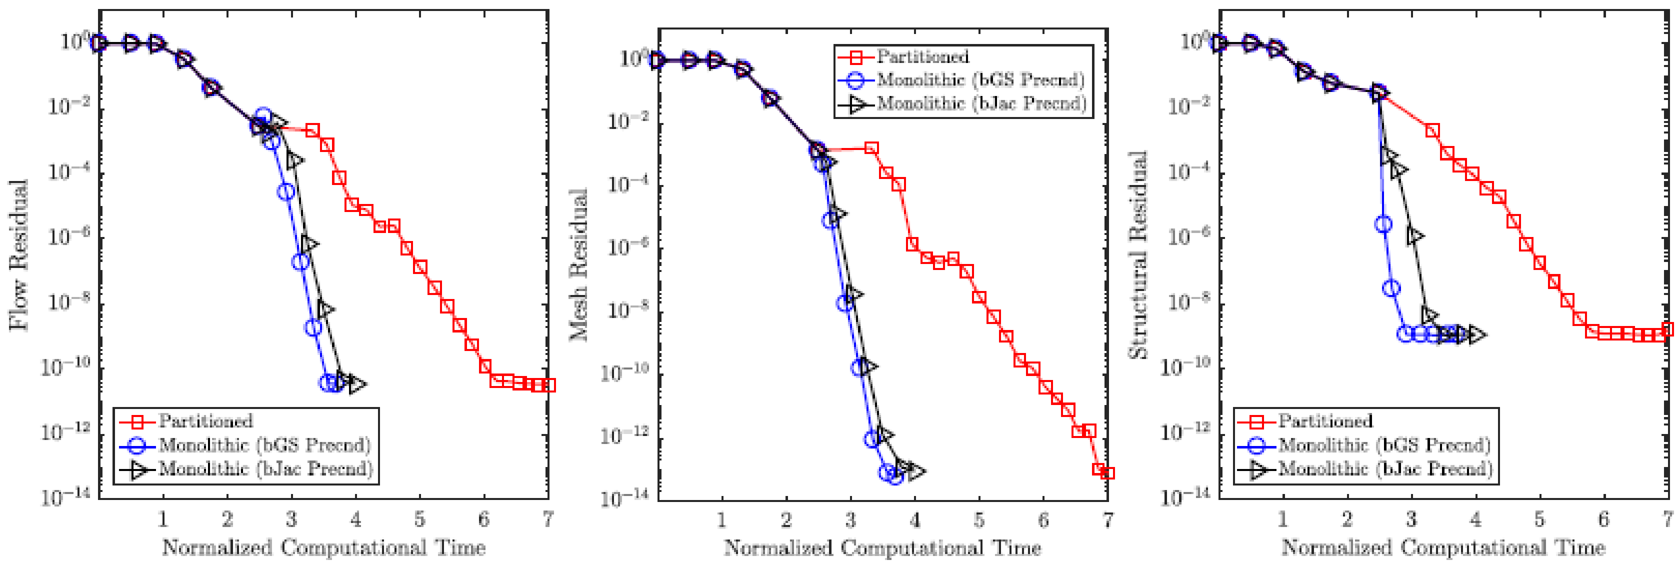

- Improved coupled convergence for Newton-Krylov solver applied to fully coupled aero-structural problem
- Improvement increases with dynamic pressure/more flexible structure

Overview

- Community Efforts
- RANS Methods
  - Second-order accurate methods (FV and SUPG@p=1)
  - Higher-order accurate methods (DG and SUPG)
- Scale resolving methods
  - Second-order accurate methods
  - Higher-order accurate methods
  - Explicit vs Implicit
- Conclusions

Discontinuous Galerkin for Simple Problem

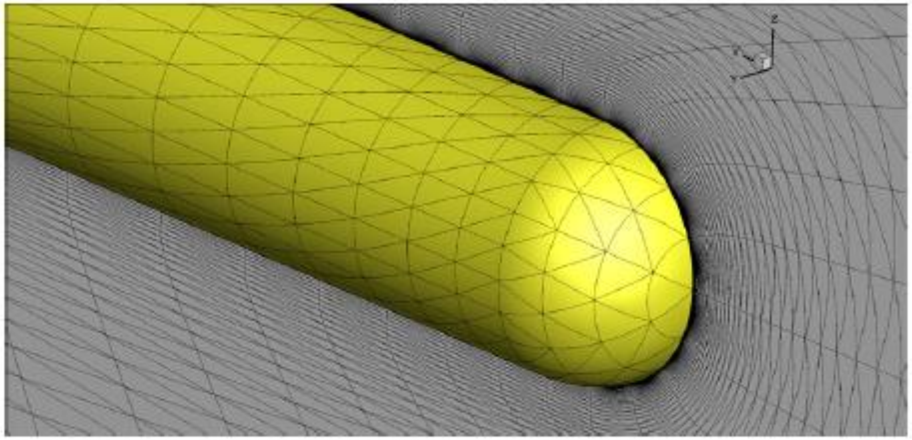

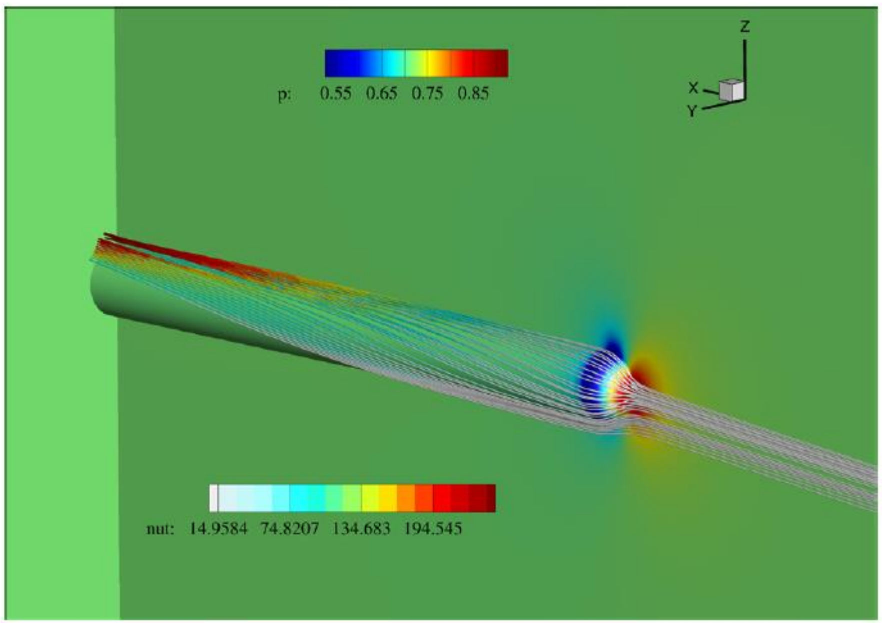

Hemispherical Cylinder
CFD Solver Discussion Group

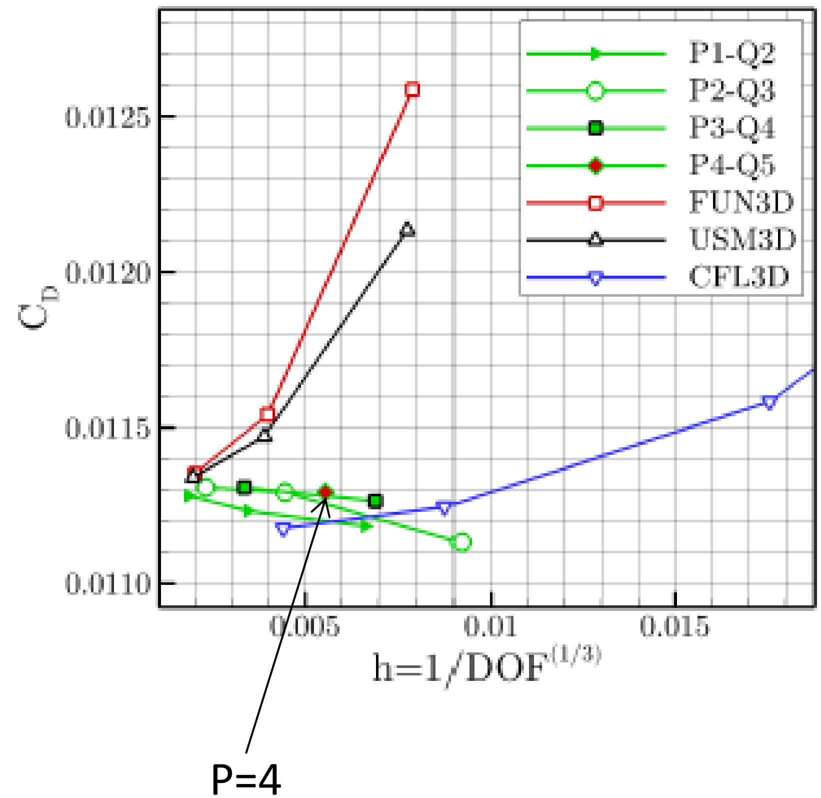

- Most accurate solutions using high p on coarse mesh
- Fast convergence for all cases

Discontinuous Galerkin for Simple Problem

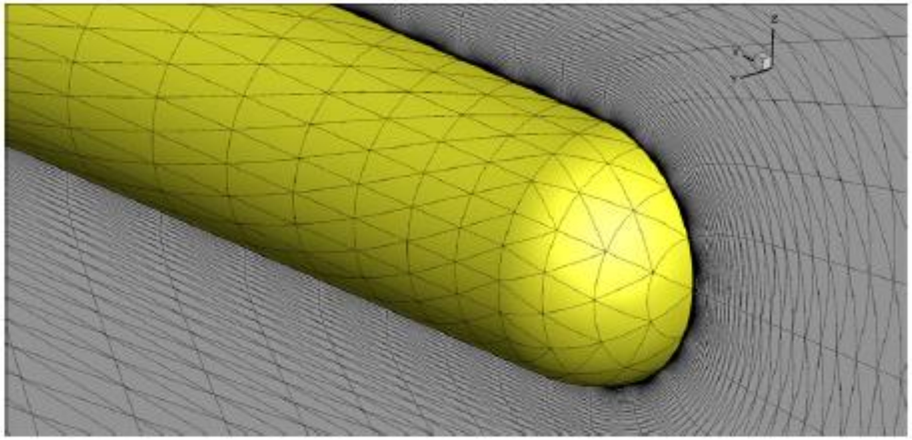

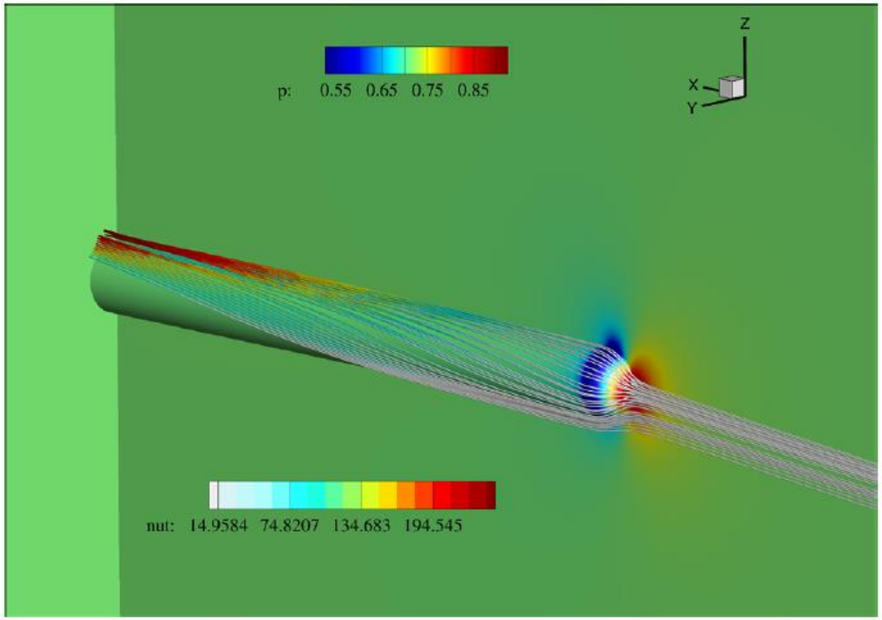

Hemispherical Cylinder
CFD Solver Discussion Group

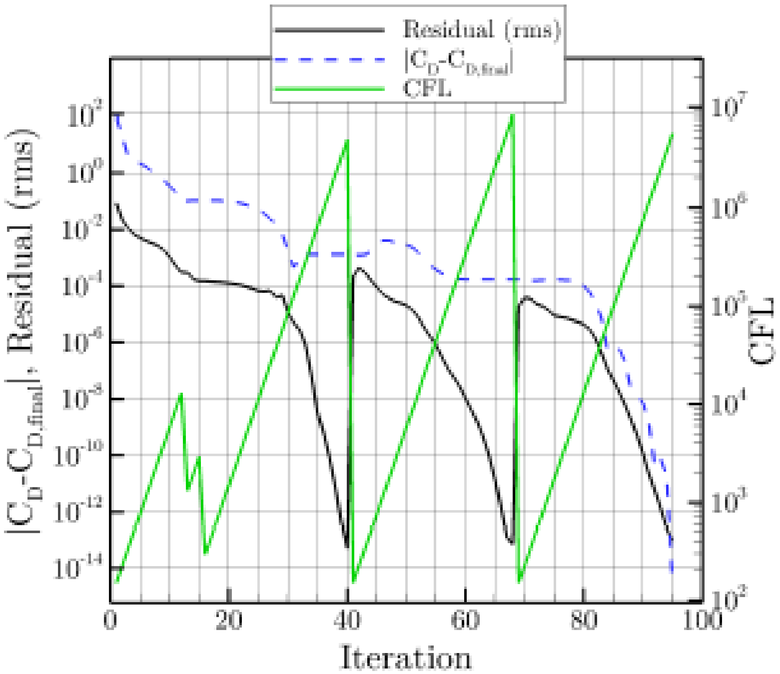

p-=1 p=2 p=3

- Most accurate solutions using high p on coarse mesh
- Fast convergence for all cases

Subsonic/Transonic CRM from HiOCFD5 (2018)

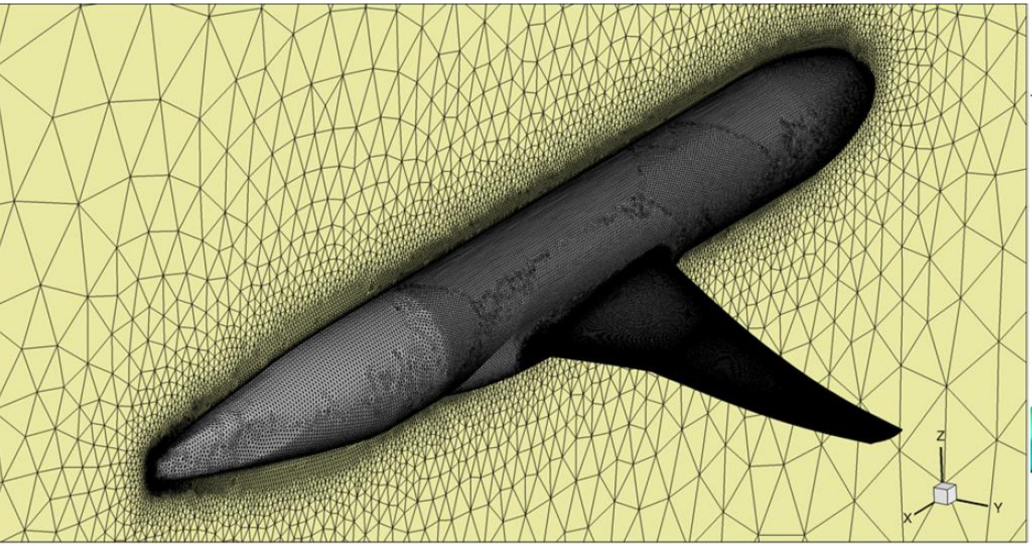

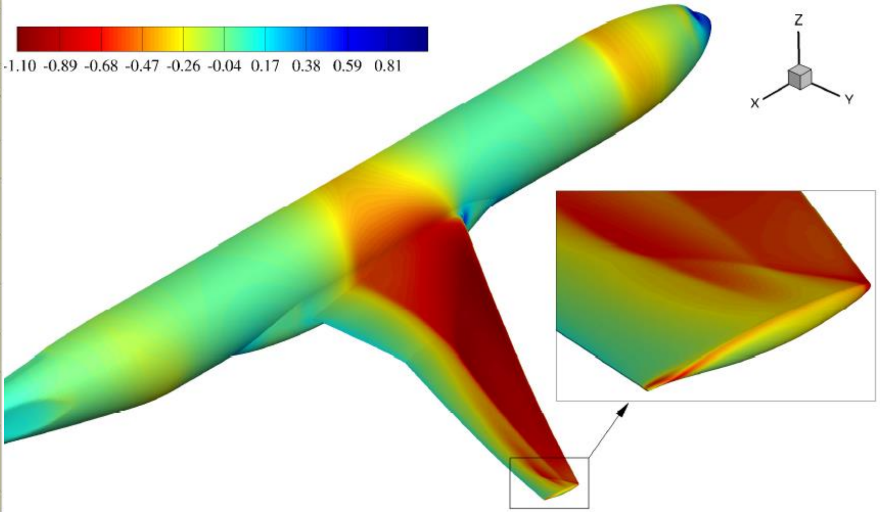

- Solutions for Mach=0.3, Mach=0.85
- SUPG p=1,2 and DG p=1,2,3 with mesh curving
- Grid sizes from 180,000 to 1.8M points
  - Coarse, but at p=2: 1.4M to 14.4M dofs
- Impressive accuracy gains going from p=1 to p=2

SG and SUPG Accuracy on CRM Test Case

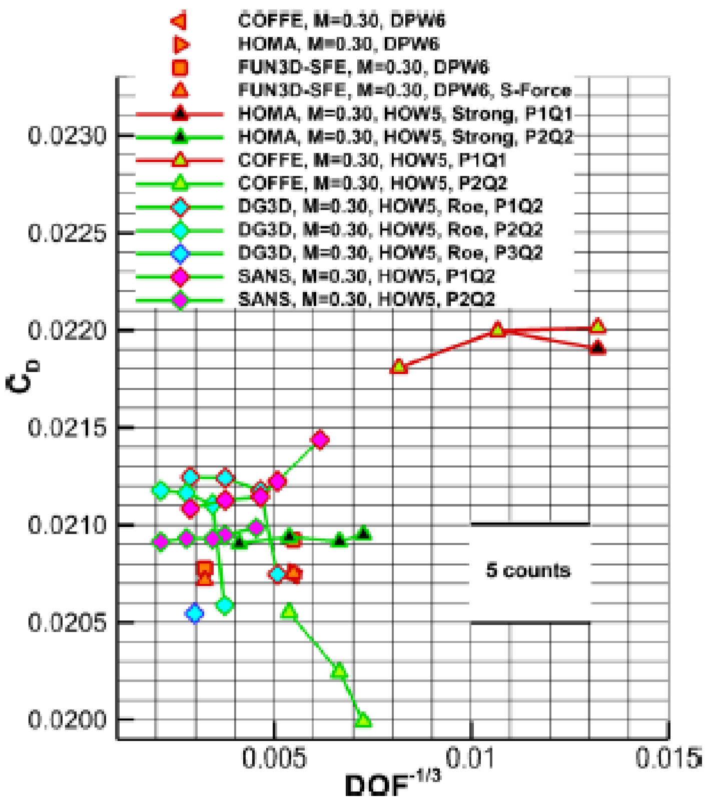

MACH=0.3

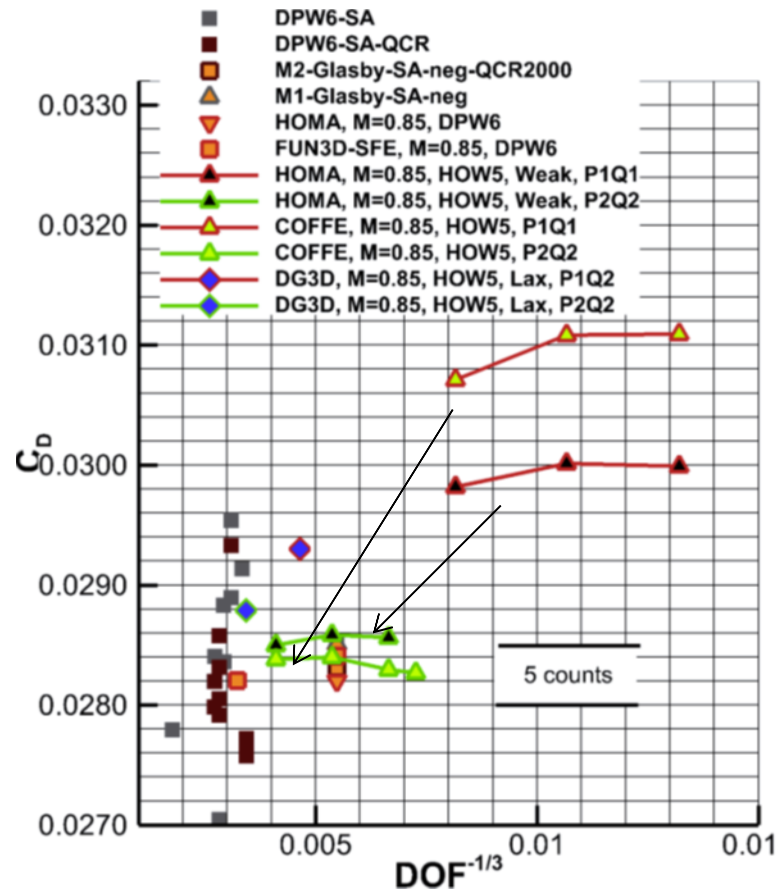

MACH=0.85

- Impressive Accuracy gains going from p=1 to p=2
- Requires curved mesh
- Expensive to solve

DLR-F11 (HLPW2) TEST CASE HiOCFD4 (2016)

Hartmann et al., Proceedings of the ECCOMAS Congress 2016

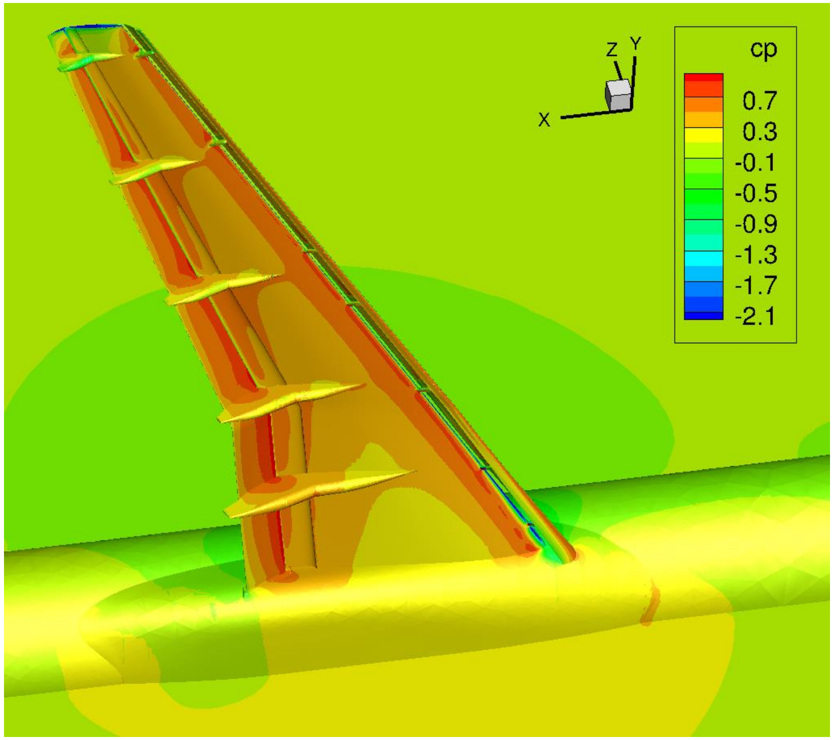

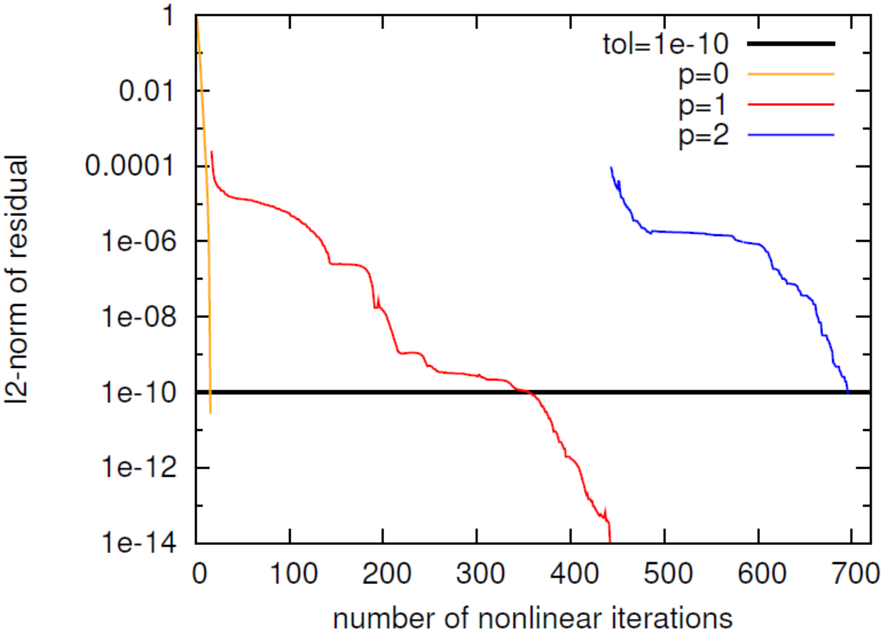

- DLR Page Code (DG discretization) (ILU(0)-GMRES Solver)
  - p=1: 14 million dofs
  - p=2: 35 million dofs
  - **One of the first high-order results on a “difficult” aerodynamic problem**

HL-CRM from HiOCFD5 (2018)

Ahrabi and Mavriplis AIAA-2019-0101

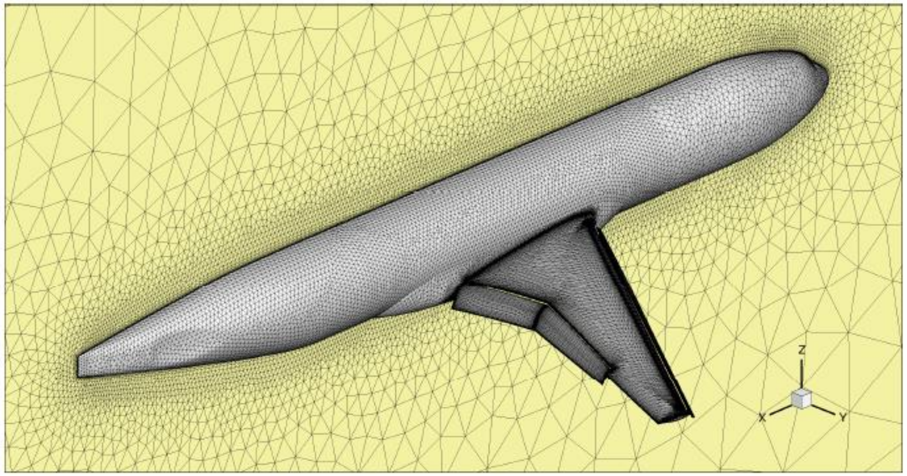

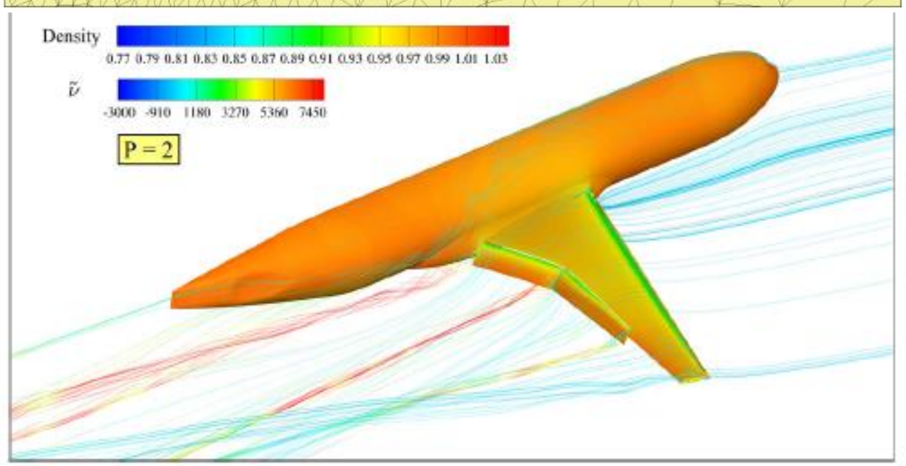

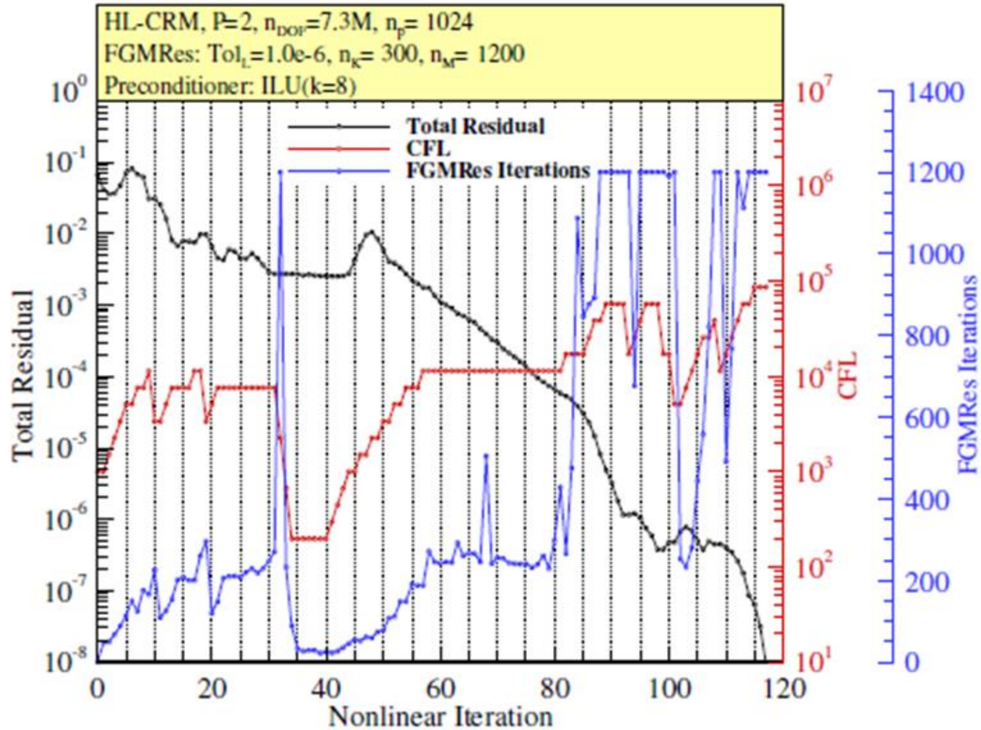

p=2 on curved mesh
7.2 million dofs

- p=2 convergence required ILU(8) preconditioner

HL-CRM from HiOCFD5 (2018)

Ahrabi and Mavriplis AIAA-2019-0101

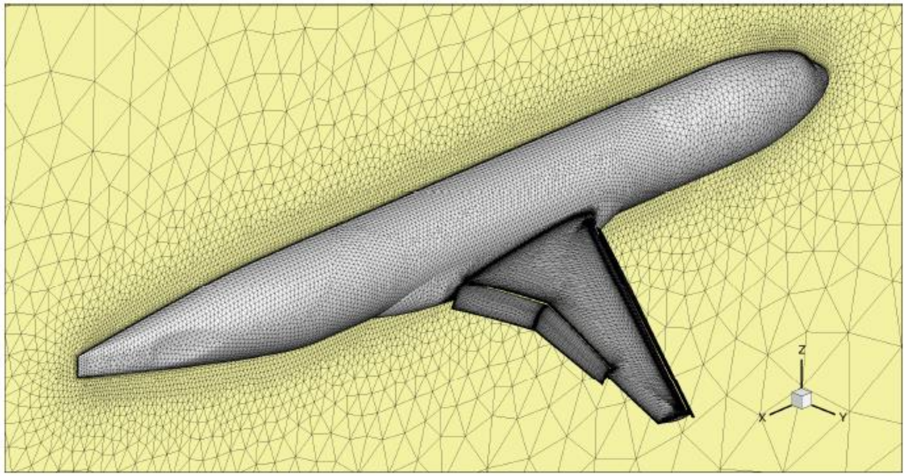

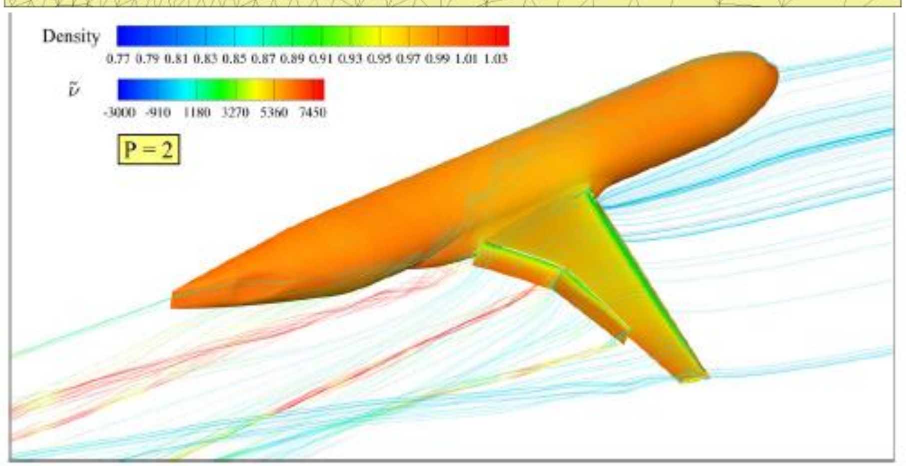

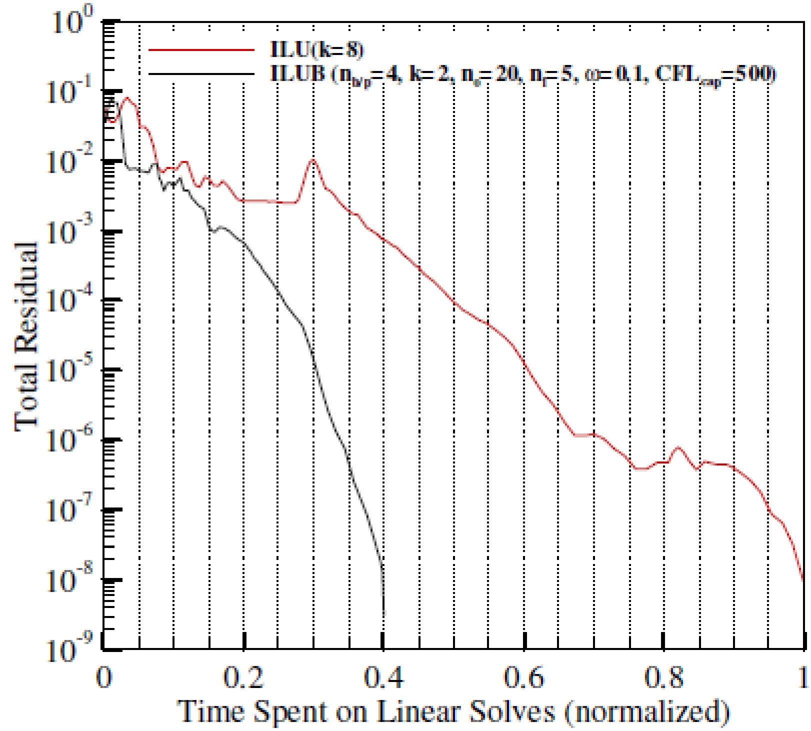

p=2 on curved mesh
7.2 million dofs

- Faster convergence using iterative block preconditioner
- Still very expensive

Overview

- Community Efforts
- RANS Methods
  - Second-order accurate methods (FV and SUPG@p=1)
  - Higher-order accurate methods (DG and SUPG)
- Scale resolving methods
  - Second-order accurate methods
  - Higher-order accurate methods
  - Explicit vs Implicit
- Conclusions

Scale Resolving Methods

- Issues are different
  - Explicit or implicit ?
    - Explicit: Solver is trivial
    - Implicit: Leverage RANS solvers (?)
  - Discretizations for
    - Accuracy
    - Low dissipation
    - Nonlinear stability (TKE preserving, entropy stable)
- As previously:
  - High-order discretizations must be demonstrably better than low-order methods on finer grids in order to be adopted
  - Some of the most consistent LES results today are using low-order methods (Spalart)

Lattice Boltzmann Method

- LBM is based on discretization of Boltzmann eqn
  - Converges to Navier-Stokes equations
    - Second-order accurate in space
    - First-order accurate in time (explicit)
    - Good at vorticity transport
  - Extremely fast per cell per time step
- LBM has shown good results for CL_{MAX} in HLPW
  - Low order, but many cells/time steps (cheap)
- Issues remain such as
  - Inability to fully capture thin boundary layers
  - Wall modeling, Turbulence modeling etc..

Second-Order Accurate NS Methods

- Mostly explicit solvers used
- Temporal and spatial scales are similar
- Other issues are shared with high-order discretizations
  - Low dissipation
  - Nonlinear stability
- Much of focus has been on physical modeling (not covered in this talk)
  - Wall modeling
  - Subgrid scale modeling
  - Transition for hybrid RANS-LES/DES

High-Order Scale-Resolving Methods

- Use of High-order methods for LES dates back to spectral/spectral element methods (1980's)
- Many inter-related discretizations
  - Discontinuous Galerkin
  - Spectral element
  - Spectral difference
  - Spectral volume
  - Flux reconstruction (FR)
  - Residual distribution
- Achieve high computational rates on modern HPC
  - Dense kernels, flops to mem ratio
  - Emergence of optimized libraries
    - pyFR, MFEM, BLAS3

High-Order Scale-Resolving Methods

- High-accuracy is not sufficient
  - DG methods are overly dissipative at smallest scales
  - Addition of sub-grid-scale (SGS) model only makes things worse
  - DG-ILES (implicit LES: no SGS model)
- Low dissipation, high accuracy schemes
  - Gassner et al.: TKE preserving DG
  - Carpenter et al., Murman et al.: Entropy Stable Schemes

Tensor Product DG

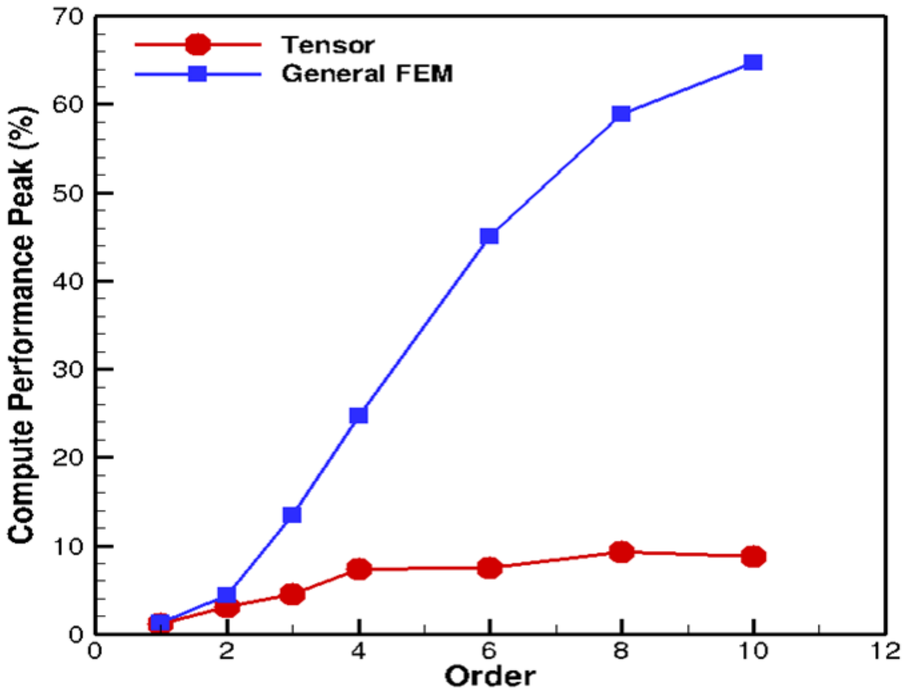

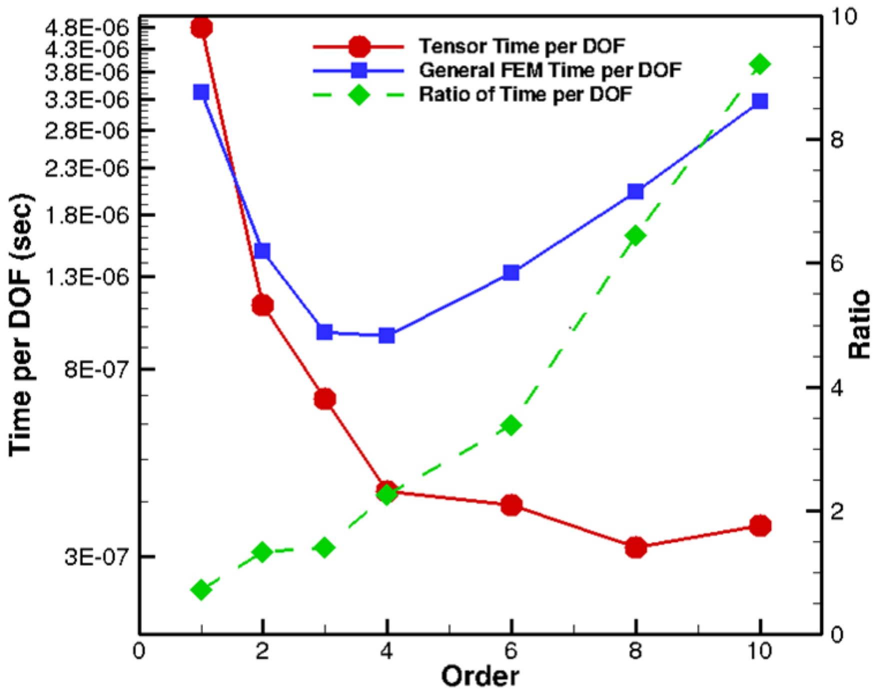

- Computational rate increases with p-order
- General formulation achieves 65% of peak at high p
- Tensor-product formulation has lower overall cost
  - $O(p+1)^6 \rightarrow O(p+1)^4$ in 3D
  - Lower cost per dof at higher p-order !

S-76 Rotor in Hover using h-p Refined DG in off-body region

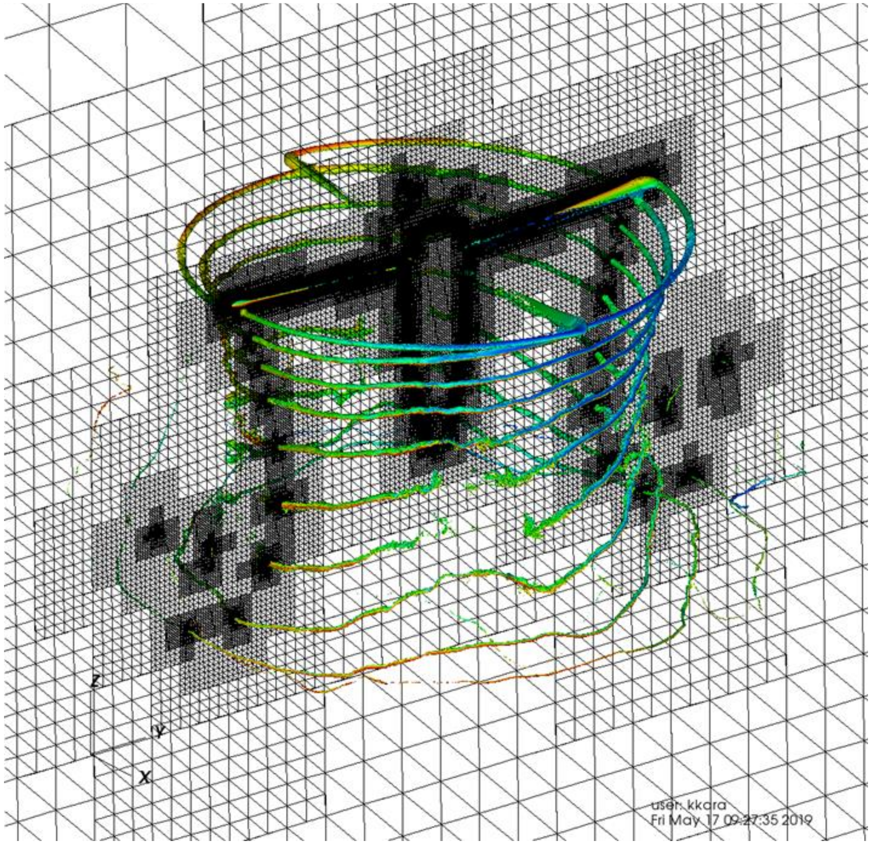

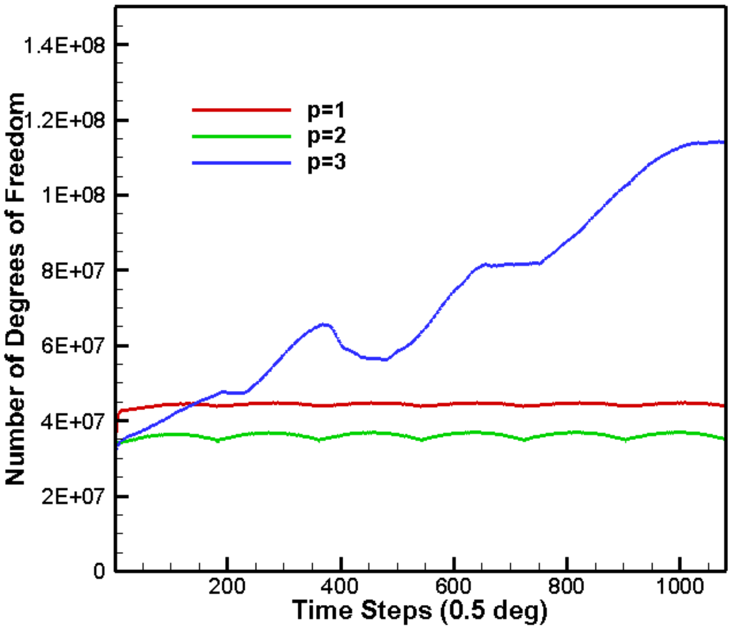

- After 7 rotor revolutions:
  - p=1,2 Simulation: 193M dofs: Wall-clock time: 53 secs/ΔT
  - p=1,2,3 Simulation: 210M dofs: Wall-clock time: 36 secs/ΔT

High-Order Implicit Methods

- Implicit time-stepping may still be desirable
  - Fast wave speeds (acoustic near incompressible limit)
  - Viscous/Diffusion time-step limit
  - Wide variation in cell sizes/p-order
- Dense Jacobian matrices destroy efficiency of Tensor-Product Formulation
  - Newton-Krylov methods with tensor-product preconditioners
  - P-multigrid using explicit steps on each level

Conclusions

- RANS Methods
  - For 2nd order discretizations
    - It is all about solver efficiency and robustness
    - Important as extend to
      - finer meshes
      - MDO and MDAO
    - SUPG@p=1 promising but limited by cost of solver technology
  - For high-order discretizations
    - SUPG seems to be more favorable over DG at lower p
    - Impressive accuracy gains even at just p=2
    - Work very well for simple problems (DG and SUPG)
    - Limiting issues are solver efficiency and robustness
      - Must be better than 2nd order schemes on finer meshes

Progress as Measured by HIOCFD

1st International Workshop on High-Order CFD Methods
Sponsored by Fluid Dynamics TC, AFOSR and DLR

January 7-8, 2012
at the 50th AIAA Aerospace Sciences Meeting, Nashville, Tennessee

Important Guideline (Read this first!)

Notes for all participants. If you have any questions, send e-mail to: hiocfd@gmail.com
Follow http://twitter.com/#!/hiocfd for updates!

Test Cases

C1. Easy, 2D
- C1.1 Internal inviscid flow over a smooth bump (Abgrall) (1/30/11), Mixed order and uniform order grids (5/24/11)
- C1.2 Transonic Ringleb flow (Huynh) (1/30/11), p4 quad grids (8/25/11)
- C1.3 Flow over the NACA0012 airfoil, inviscid and viscous, subsonic and transonic (May), p4 quad & triangular grids (far field boundary over 1000 chords away) (8/3/11).
- C1.4 Flat plate boundary layer (Bassi) (1/30/11), quad grids (9/23/11)
- C1.5 Radial expansion wave (van Leer) (Attention added on 9/30/11)
- C1.6 Vortex transport by uniform flow (Caraeni) (Updated on 10/05/11), Grids (4/11/11)

C2. Intermediate, 2D & 3D
- C2.1 Unsteady viscous flow over tandem NACA0012 airfoils with a smooth initial condition (Cary) (1/30/11)
- C2.2 Turbulent flow over a RAE airfoil (Deconinck), Linear and higher order grids (5/24/11)
- C2.3 Analytical 3D body of revolution (Kroll), Linear and high-order grids (5/24/11)
- C2.4 Delta wing at low Reynolds number (Hartmann) Grids (updated 8/15/11)

C3. Difficult, 2D & 3D
- C3.1 Turbulent flow over a multi-element airfoil (Wang) (1/30/11), Geometry (5/11/11)
- C3.2 Turbulent flow over DPW III wing alone (Fidkowski) (1/30/11), Grids (9/8/11)
- C3.3 Transitional flow over a SD7003 wing (Visbal) (1/30/11), Geometry posted on 5/11/11 (Points or IGES file)
- C3.4 2D laminar flapping wing case (Persson) (1/30/11)
- C3.5 Direct Numerical Simulation of the Taylor-Green Vortex at Re = 1600 (Hillewaert) (4/18/11), reference data (8/24/11)

HiOCFD5
5th International Workshop on High-Order CFD Methods

Home | Presentations | Test cases | News | Guidelines | Previous | Verification Resources | Agenda | Committee

Home » Test cases » Advanced » CR1 - Common Research Model

CR1 - Common Research Model

Home » Test cases » Computional challenges » MC1 - High-Lift Common Research Model

MC1 - High-Lift Common Research Model

- CRM is now “Advanced Test Case”
- HLCRM is “Challenge Test Case”

Conclusions

- Scale resolving methods
  - Battle between 2nd order methods and high-order methods continues
  - DG methods appear more favorable at very high order
    - Tensor product formulation
    - Block matrices
    - High computational rates
  - Main issues remain
    - Stability at low dissipation
      - TKE preserving, entropy stable, etc.
    - Physical modeling

Conclusions

- At industrial level, major impact from new discretizations/solvers still not felt
- However, gradual adoption of techniques is underway
  - SUPG@p=1 is in production
  - Krylov methods extending to adjoints, coupled MD problems
  - Spin-off technologies becoming ubiquitous
    - NK solvers
    - Negative SA model
    - Data bases for verification and validation
- Community efforts: Commendable
- Merging of these with others will provide progress towards CFD2030 milestones
  - 10B dof grids
  - Physical modeling
  - Better solvers
- Timeline ?

BACKUP SLIDES

FUN3D NonLinear Multigrid

(From: Diskin and Nishikawa AIAA 2014-0082)

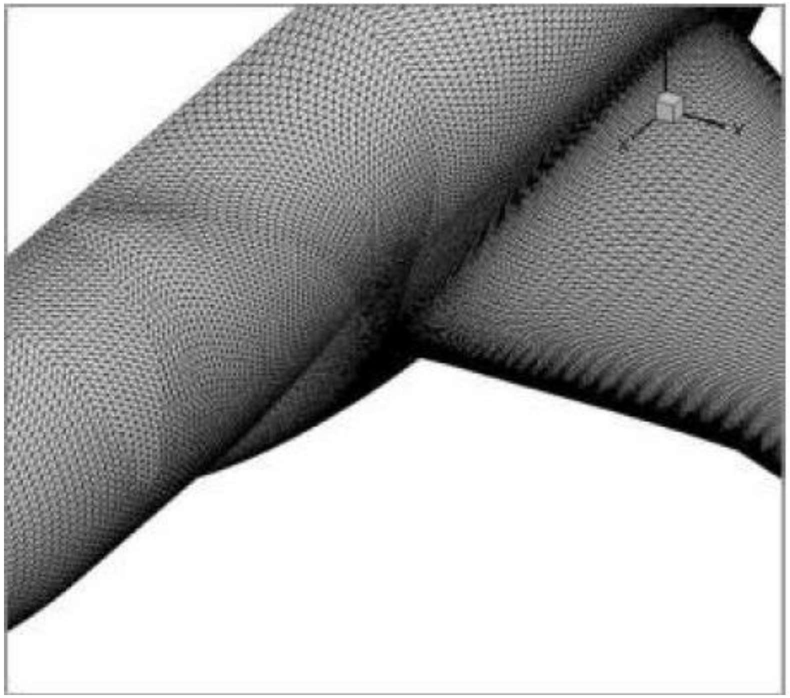

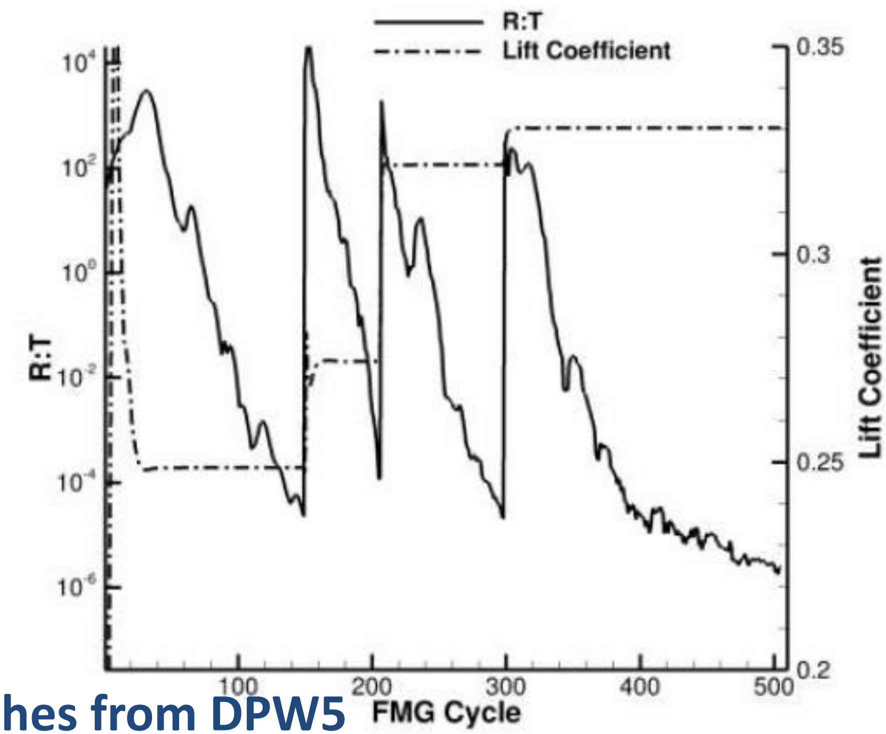

Sequence of Finer Structured Meshes from DPW5
Finest Mesh: 5.2 Million points

- Atypical approach:
  - Newton-Krylov method on each level
  - Bypasses slow initial convergence of single-grid NK method

FUN3D NonLinear Multigrid

(From: Diskin and Nishikawa AIAA 2014-0082)

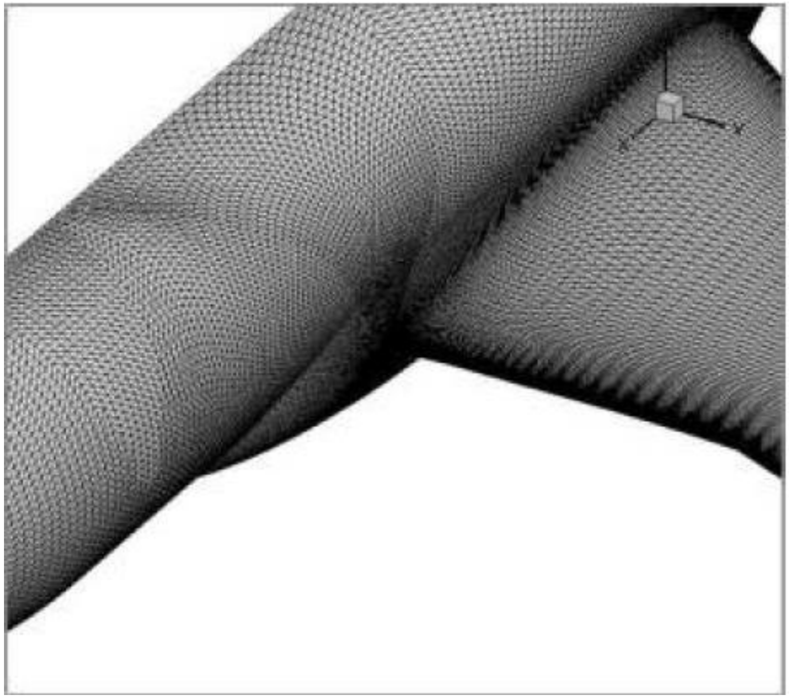

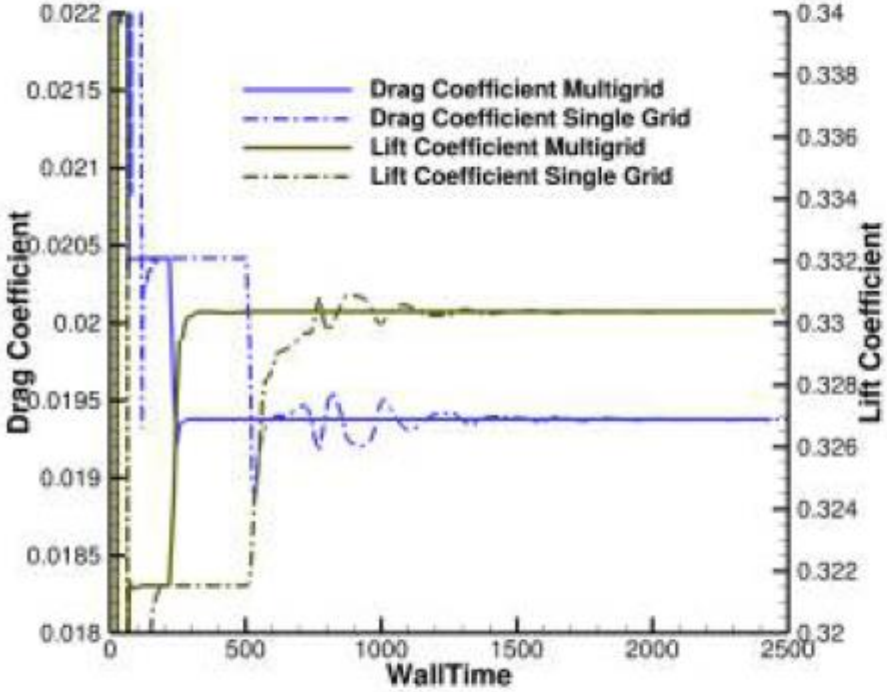

Sequence of Finer Structured Mes
Finest Mesh: 5.2 Million points

- Impressive speedup (order of magnitude)
- Relatively easy problem

DG and SUPG Convergence

(from Wyoming group results)

DG at Mach=0.3
ILU(0)-GMRES

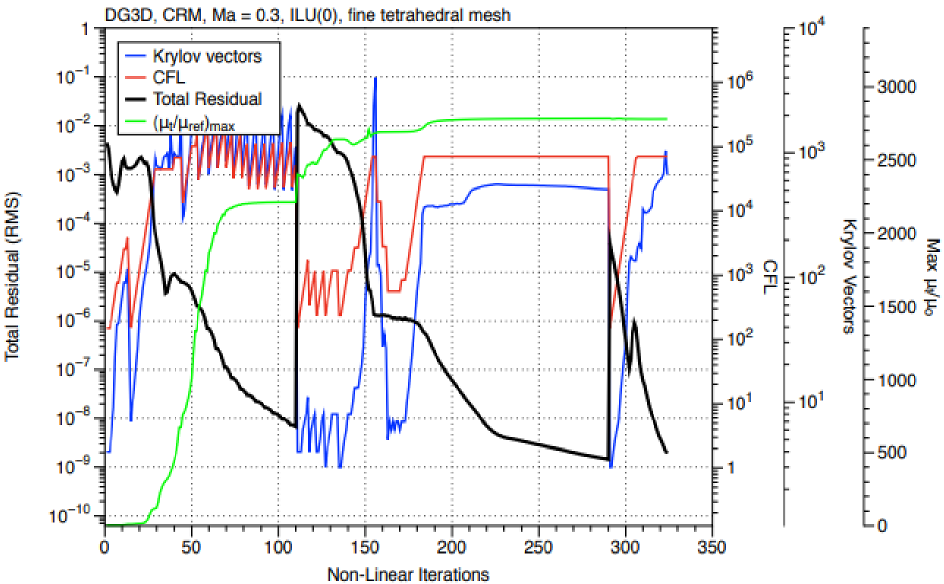

SUPG at Mach=0.85
Line Preconditioned GMRES

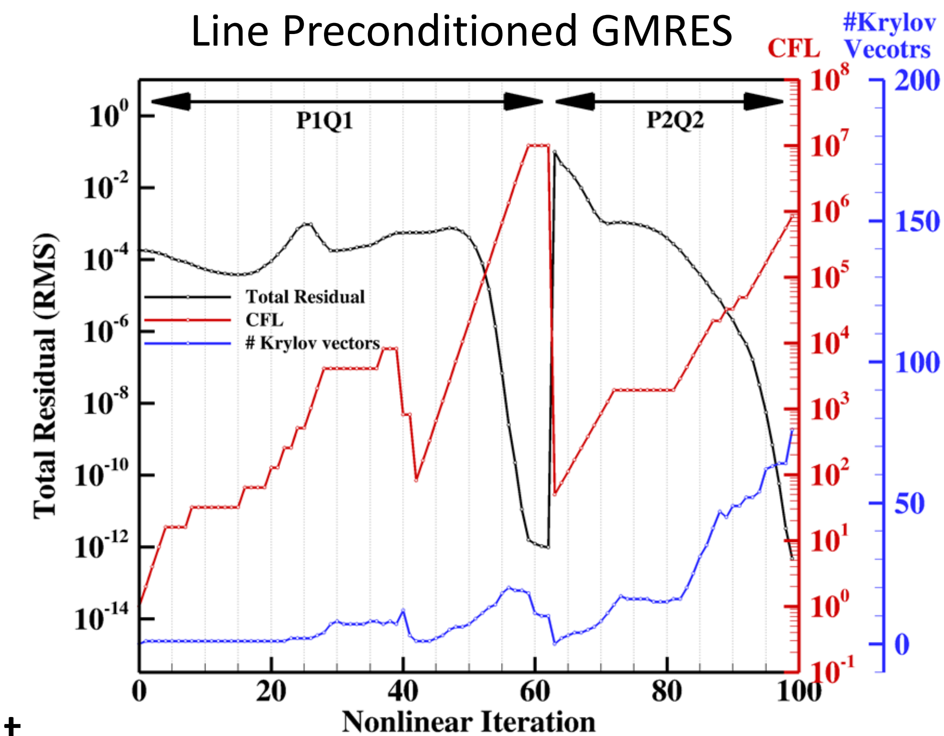

- Convergence at p>1 more difficult
  - Higher number of Krylov vectors
  - Higher number of dofs (on same grid)
  - Still likely more expensive than FV scheme on finer grid (same accuracy)
- Asymptotic properties (convergence and accuracy)

DG Solver Validation

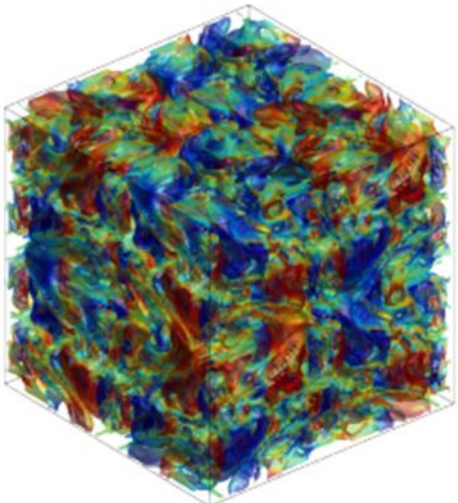

Taylor-Green Vortex Problem

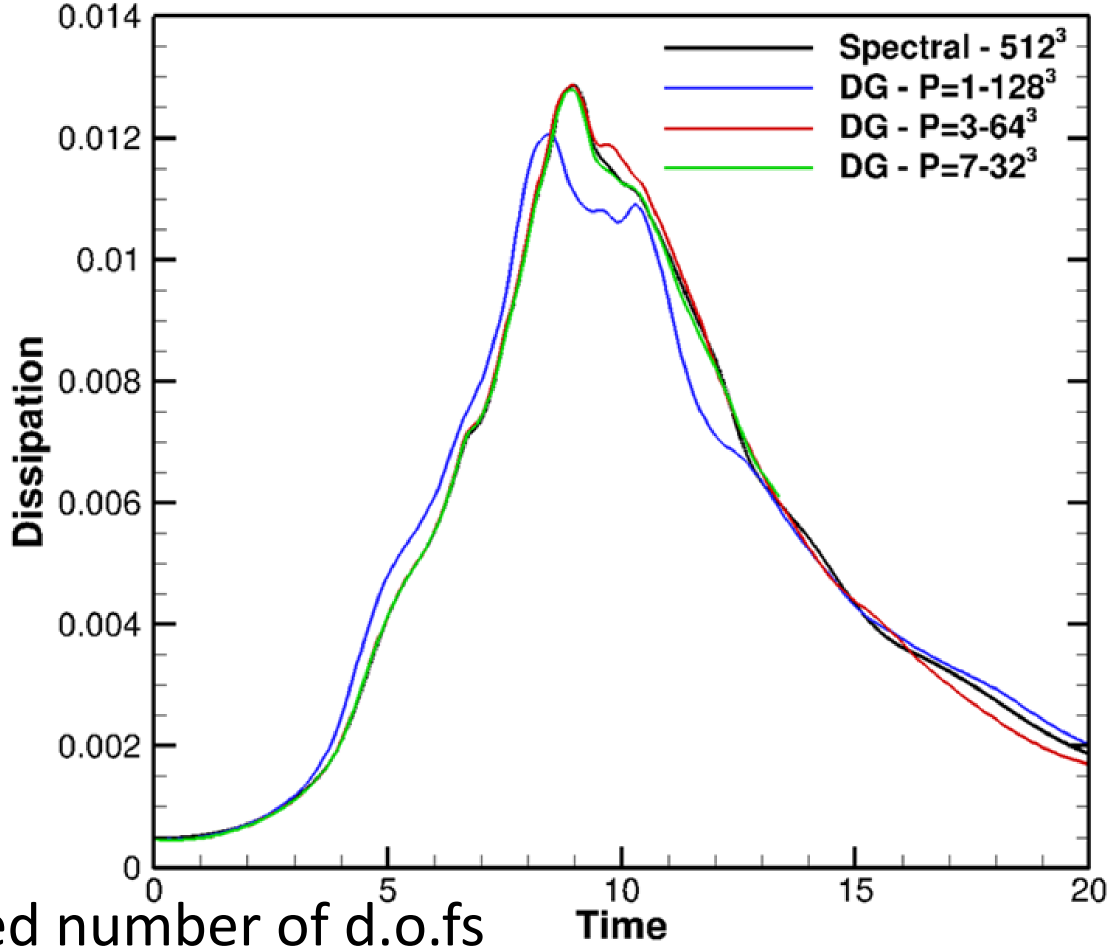

- Increasing p at fixed number of d.o.fs
  - Coarser meshes at higher p
  - Accuracy increases
  - Simulation cost decreases (per time step)