The HiFi-TURB EU project: Vision and Progress of ML-based Turbulence Modeling

2022 NASA Symposium on Turbulence Modeling: Roadblocks and the Potential for Machine Learning

Charles Hirsch, Cadence-Belgium, *Fellow CFD*

July 2022

cādence®

Content

- The EU HiFi-TURB project
  - Structure and objectives
- Progress on ML driven EARSM new models (Stefan Wallin)
  - Approach
  - MEP
  - Auto-encoder
- Beyond RANS: the potential for WRLES in the near future
  - Some potential game changers

HiFi-TURB - High-Fidelity LES/DNS Data for Innovative Turbulence Models

The HiFi-TURB project has as objective to improve TU-models in presence of separation, based on the following:

- The exploitation of high-fidelity LES/DNS data for a range of reference flows that contain key flow features of interest
- The application of novel artificial intelligence and machine-learning algorithms to interrogate the data with the objective to identify significant correlations between representative turbulent quantities
- The guidance by world-renowned industrial and academic experts in turbulence
  - Philippe Spalart; Florian Menter; Stefan Wallin; Suad Jakirlic; Michael Leszchiner; Wolfgang Rodi; Stefan Hickel; Maria-Vittoria Salvetti
- HiFi-TURB Partners
  - Industry: Cadence-Belgium (formerly Numeca); Safran; Dassault Aviation
  - Research Centers: BSC; ONERA; DLR; CERFACS; CENAERO; ERCOFTAC
  - Universities: Bergamo; Imperial College; UCLouvain
  - Associate Partner: NASA-RCA (M; Malik, C; Rumsey)

cādence®

HiFi-TURB overview

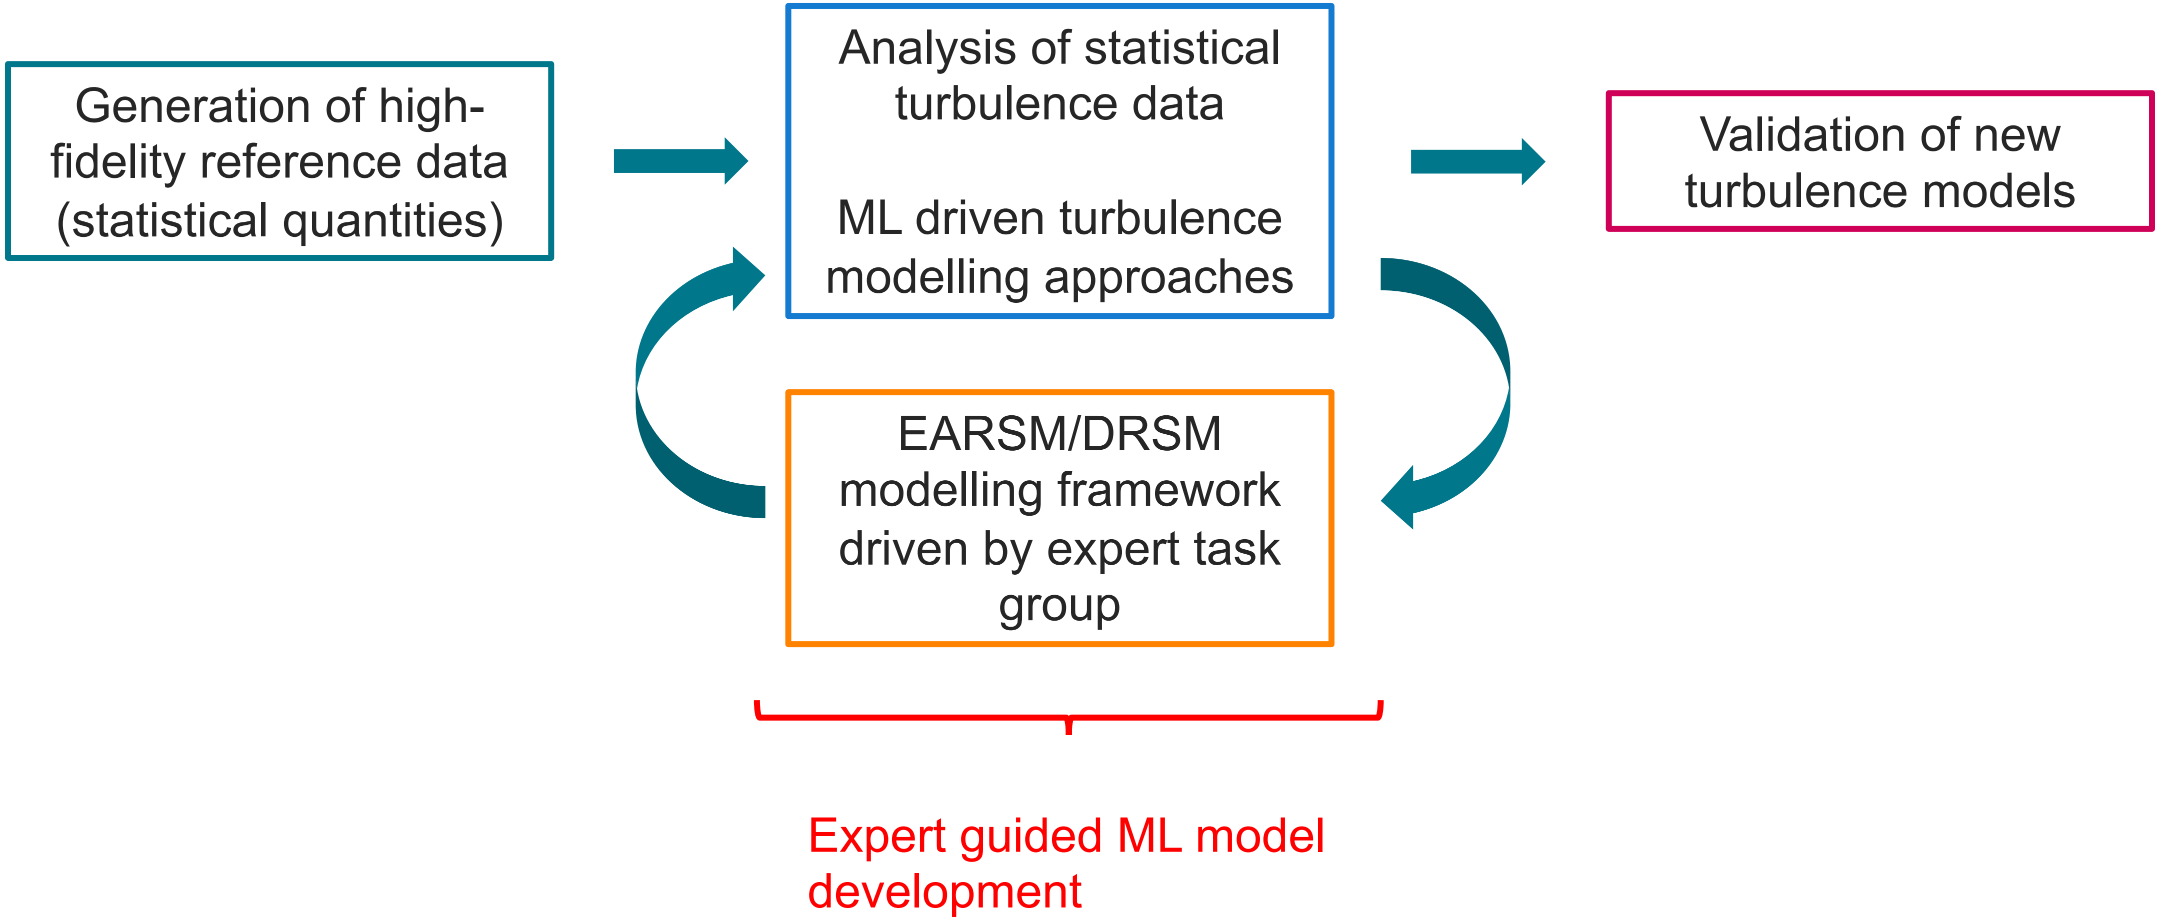

WP3 Test cases

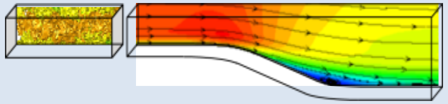	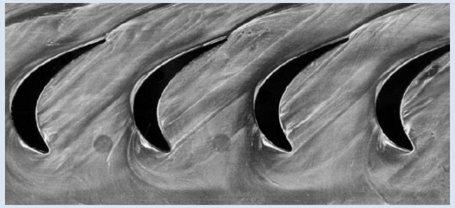	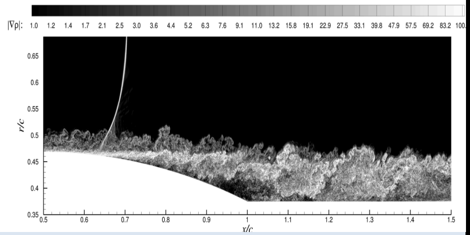 Uzun & Malik, 2018	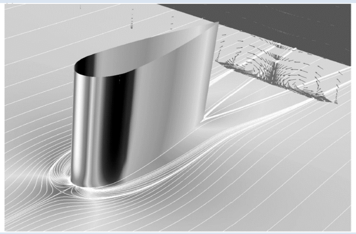 Apsley & Leschziner, 2001	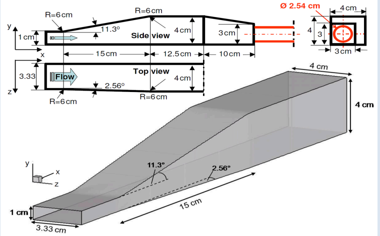 Cherry et al, 2008
TC01: **Adverse Pressure Gradient Experiment** $Re_\tau \approx 4179$ (related to channel)	**TC02:** **MTU T161 Cascade** Re = 200,000	**TC03:** **Bachalo-Johnson – axisymmetric bump** Re=1.6 million	**TC04:** **3D Wing-Body Junction** Re=115,000	**TC05:** **3D (Stanford) Diffusor** Re=10,000 to 30,000 ***www.kbwiki.ercoftac.org/w/index.php/abstr:ufr_4-16***
Physics to be modelled: Adverse pressure gradient; pressure-induced separation and reattachment	**Physics to be modelled:** End-wall boundary layers and corner flows	**Physics to be modelled:** Transonic shock-boundary layer inter-action; Shock/curvature induced separation.	**Physics to be modelled:** External Corner flow; 3D separation	**Physics to be modelled:** Internal corner flows; 3D separation
Test-case Team: **DLR**, UNIBG, CENAERO, DASSAV	**Test-case Team:** **ICL**, CENAERO, NUMECA, UNIBG	**Test-case Team:** **CENAERO**, ONERA, UNIBG	**Test-case Team**: **ONERA**, CERFACS, ICL, DASSAULT	**Test-case Team**: **NUMECA**, SAFRAN, BSC

WP3 Highlights

- For each of the testcases, all statistical averaged terms entering in the Re-Stress transport equations are stored, plus some basic additional turbulent quantities, for a total of ±175 quantities
- This represents a huge amount of data, to be analyzed by Big data, AI and ML technologies, to search for new meaningful correlations
- One important point is the definition and coordination of the computational campaign since several groups will investigate the same case, including quality criteria for the data sets in terms of spatial, temporal and statistical convergence.

WP4 Highlights

- WP4 is concerned with the analysis of the basic turbulence-averaged data coming from high-fidelity DNS datasets generated in WP3 by means of statistical analysis and machine learning technologies (ML)
- **Analysis of basic turbulence-averaged data via data driving methodologies focusing on RSM (DRSM, EARSM)**
- In this ML context, a main focus is respecting basic physics:
  - proper Re-number scaling,
  - Galilean invariance – including invariance to moving systems,
  - asymptotic behaviour at boundary-layer edge, log-law consistency,
  - locality vs non-locality

Hifi-TURB Highlights

- **Work Package 5: Turbulence modelling assessments and improvements – monitored by WP5 Task Group**
  - Task 1: Develop, improve and assess EARSM turbulence models
  - Task 2: Develop, improve and assess DRSM turbulence models
  - Task 3: Develop, improve and assess wall models for WMLES and Hybrid RANS-LES
- **Work Package 6: Validation of new turbulence models applied to representative and industrial relevant test cases**
  - Task 1: Validation on external flow configurations (High-Lift and Drag Pred. WS cases
  - Task 2: Validation on internal flow configurations for fixed (diffusor) and rotating. cases
  - Task 3: Assessment & recommendations
- **Work Package 7: Management of the LES/DNS databases for open accessibility (ERCOFTAC)**
  - Task 1: Definition of database criteria and implementation rules
  - Task 2: Creation and management of the LES/DNS database
  - Task 3: Integrating results of WP6 in the ERCOFTAC Wiki Knowledge Base

cādence®

WP6 Validation cases

	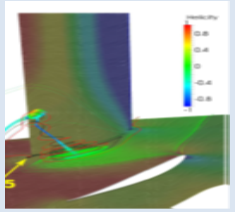	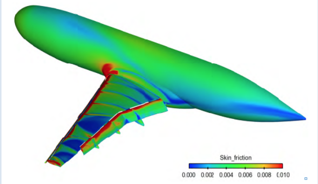 3rd High-lift prediction workshop, 2017	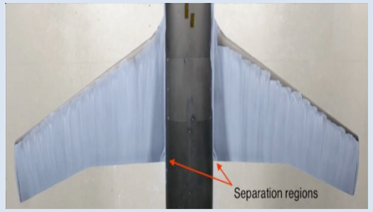 Rumsey, 2018	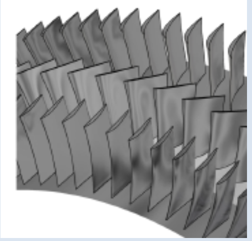	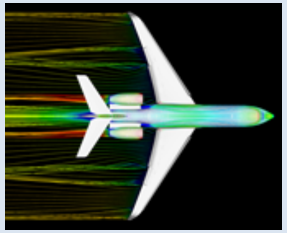
TC06: **Rotor 37**	**TC07:** **LMFA – NACA65**	**TC08a / TC08b: (a) NASA-CRM DPW, (b) High-Lift onfigurations**	**TC09:** **NASA Juncture Flow**	**TC10** - confidential: **1.5 axial stage compressor**	**TC11** - confidential: **Falcon business jet in cruise & landing configuration**
Physics to model Tip vortices, corner separation, shock-boundary layer interaction	**Physics to model** Compressor cascade; corner flow separation	**Physics to model** Turbulence model adapted to strong separation zone	**Physics to model** Wing-body corner separation	**Physics to model** Near-stall performance	**Physics to model** External flows incl. separation; whole aircraft
Test-case leader: NUMECA	**Test-case leader:** SAFRAN	**Test-case leader:** ONERA	**Test-case leader:** DLR	**Test-case leader:** SAFRAN	**Test-case leader:** DASSAV

cadence

Expert guided ML model development

Major novel approach for EARSM by Stefan Wallin

Computation of statistically averaged quantities from high-fidelity simulations

- Using the approach for EARSM the terms that need modelling are computed exactly from the high-fidelity data: Beta terms for EARSM and invariants

→ This forms **the data** on which the ML approaches act

- The data is either selected manually from the high-fidelity dataset or via the **Variational Auto-Encoder** tool

Machine Learning

1. Multiple Expression Programming (MEP)

- Let the computer find a mathematical model in function of invariants

2. Optimization tool for expressions

- Develop/tune handcrafted expressions vs the data
- Fine-tune MEP found expressions vs the data

3. Variational Auto-Encoder tool

- Quasi-exact representation of the high-fidelity data via a ML surrogate
- Physical insights through the nature of the 'embeddings'
- Extremely fast access to the data for extraction and a priori testing vs. heavy CPU intensive process without the VAE

Multi-Expression Programming (MEP)

- Multi Expression Programming (MEP) uses evolutionary algorithms to regress a mathematical expression from a set of data.
- Derived from the MEPX library [1], implemented in C/C++, with a high level python interface.
- In its most basic format:
  - Each chromosome stores an overall fitness: best fitness within a given complexity.
  - Fitness is evaluated for each element of the table.
  - The complexity associated with each operation and variable is set as an input parameter.
  - Constants do not add complexity to the expression.

Operations associated complexity:
+: 10
*: 5

MEP

id	op	addr1	addr2	Res	complexity
0	a	-	-	a	5
1	b	-	-	b	5
2	c (constant)	-	-	c	0
3	d	-	-	d	5
4	+	0	1	a+b	5+5+10 = 20
5	*	2	3	c*d	0+5+5 = 10
6	+	4	5	a+b+c*d	20+10+10 = 40
7	e	-	-	e	
8	log	7	-	log(e)	
9	sqrt	2	-	sqrt(c)	

Example of the computation of the complexity topdown from a chromosome (MEP table).

Variational Auto-Encoder (Cadence and partner Neural Concept)

- The Variational Auto-Encoder (VAE) is a type of Artificial Neural Networks (ANN)
  - The applied Geodesic Convolutional Neural Network (CNN) can be directly applied to an unstructured cloud of points, this is in contrast to standard CNNs which act on structured data such as they are presented by pixels in images
- Variational Auto-Encoders enforce smoothness constraints on the space of the embedding and allow for a physical interpretation

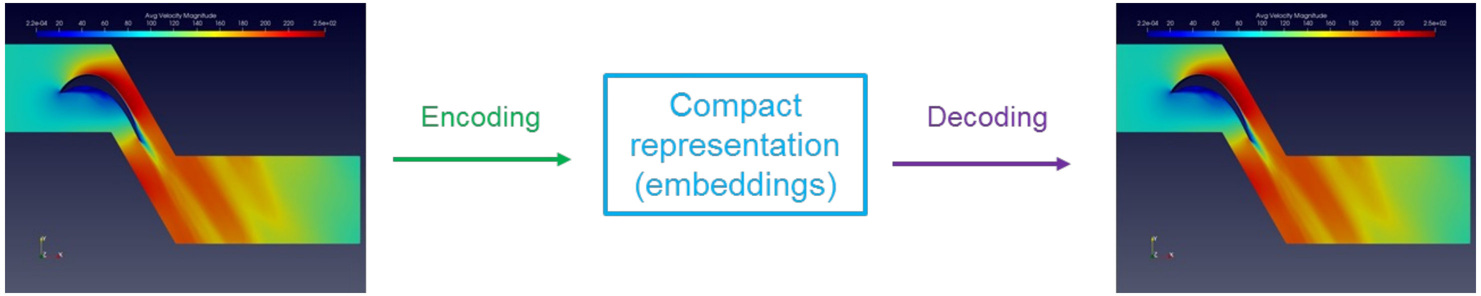

- An exact representation of the high-fidelity data can be constructed
  - Work easily with the data in a web-browser vs. HPC infrastructure
  - Analysis and data extraction based on the VAE
  - A priori model testing

Variational Auto-Encoder applied to Stanford diffuser

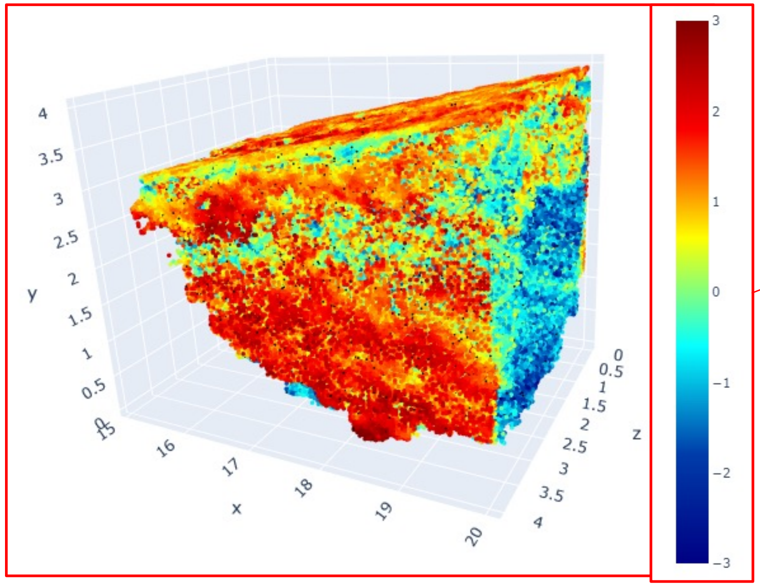

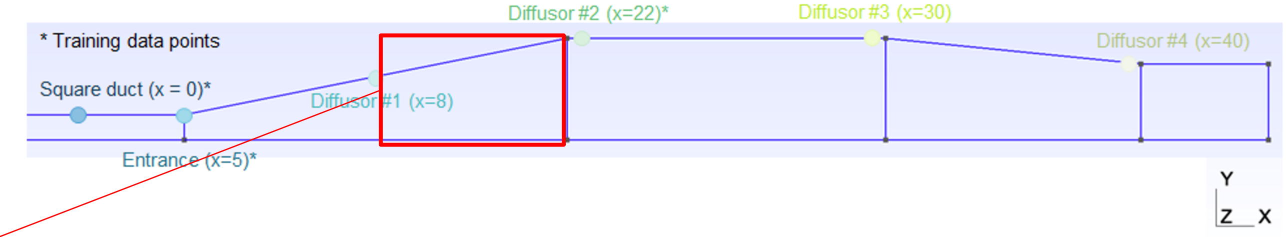

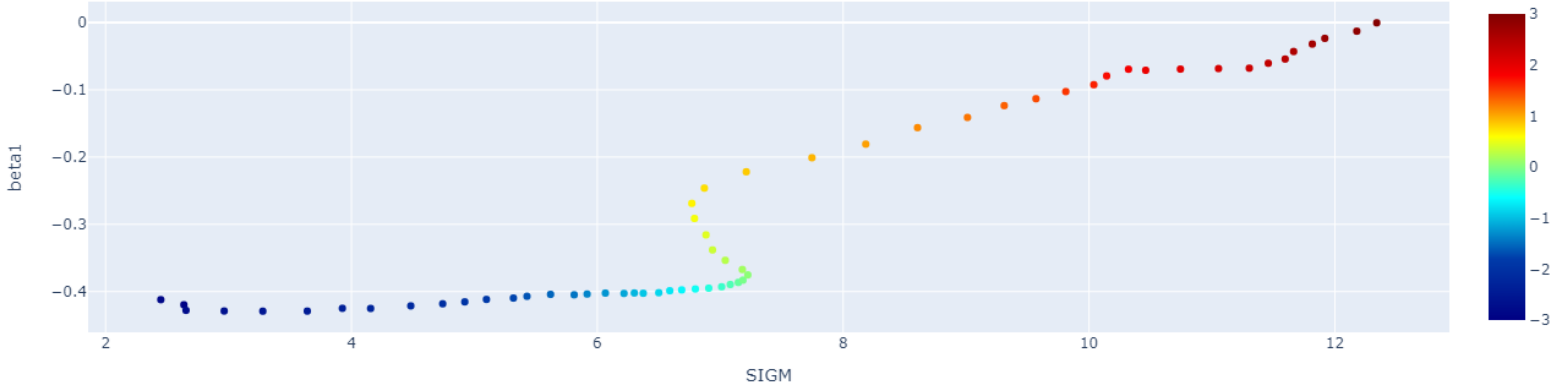

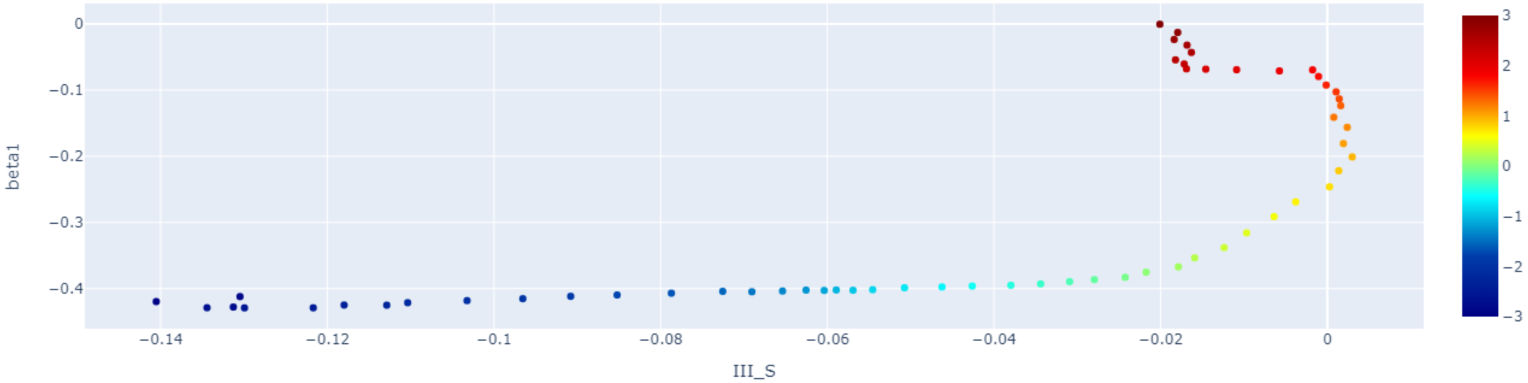

- Extraction of functional dependency of **beta1**
- The colors in the Cartesian plots correspond to the locations in the 3D view
- **→ Use this insight to select the data used for both manual and ML driven model development**

The novel EARSM framework---by Stefan Wallin

- The 5-term representation (Pope 1975; Wallin & Johansson 2000) has been found to be practical sufficient and not over-complex.
- Any algebraic relation of $a_{ij}(S_{ij}, \Omega_{ij})$ can be mapped on these 5 terms.

$$a_{ij} = \sum_{k=1}^{5} \beta_k T_{ij}^{(k)}$$

$$S_{ij}, \Omega_{ij} = \frac{1}{s}\left(\frac{\partial U_i}{\partial x_j} \pm \frac{\partial U_j}{\partial x_i}\right)$$
$$s = \|\mathrm{grad}(\boldsymbol{U})\|$$

where

$$T_{ij}^{(1)} = S'_{ij}, \qquad T_{ij}^{(2)} = S'_{ik}\Omega'_{kj} - \Omega'_{ik}S'_{kj},$$
$$T_{ij}^{(3)} = \Omega'_{ik}\Omega'_{kj} + \frac{1}{3}r\delta_{ij}, \qquad T_{ij}^{(4)} = S'_{ik}\Omega'_{kl}\Omega'_{lj} + \Omega'_{ik}\Omega'_{kl}S'_{lj} - \frac{2}{3}IV'\delta_{ij} + rS'_{ij},$$
$$T_{ij}^{(5)} = \Omega'_{ik}S'_{kl}\Omega'_{lm}\Omega'_{mj} - \Omega'_{ik}\Omega'_{kl}S'_{lm}\Omega'_{mj} - \frac{1}{2}r\left(S'_{ik}\Omega'_{kj} - \Omega'_{ik}S'_{kj}\right).$$

The novel EARSM framework---by Stefan Wallin

The strain- and rotation rate tensors are normalized with s

$$S'_{ij} = \frac{S_{ij}}{s}, \qquad \Omega'_{ij} = \frac{\Omega_{ij}}{s}.$$

Hence, the invariants are connected so that $II'_S - II'_\Omega = 1$ and we introduce r

$$II'_S = \frac{II_S}{s^2} = 1 - r, \qquad II'_\Omega = \frac{II_\Omega}{s^2} = -r,$$

and normalize also the higher invariants as

$$III'_S = \frac{III_S}{s^3} = S'_{ij}S'_{jk}S'_{ki}, \qquad IV' = \frac{IV}{s^3} = S'_{ij}\Omega'_{jk}\Omega'_{ki},$$

$$V' = \frac{V}{s^4} = S'_{ij}S'_{jk}\Omega'_{kl}\Omega'_{li} + \frac{1}{2}r(1-r).$$

cādence

The novel EARSM framework---by Stefan Wallin

- All invariants are, hence, normalized by the magnitude of the strain rate, s, and can be seen as "structure parameters" providing the principal directions. The consequence is also that the two 2nd invariants, II_S and II_Ω, can be replaced by one single variable r.

- One additional measure is then needed to complete the representation (as derived from the EARSM solution). That is the "equilibrium parameter" σ,

$$\sigma = \frac{sk}{\varepsilon},$$

which is then the only parameter dependent on ε. The variable σ will measure the state of equilibrium by relating the turbulence to mean-flow time scales.

Flow in equilibrium will have σ of order unity and the rapid distortion limit is characterized by $\sigma \to \infty$. Hence, any dependency on σ must approach a constant for $\sigma \to \infty$.

The novel EARSM framework---by Stefan Wallin

The different β_k are now functions of σ and the invariants, $\beta_k(\sigma, r, III_S, IV, V)$. The first term β_1 corresponds to the eddy-viscosity part with an effective coefficient

$$C_\mu^{\text{eff}} = -\frac{1}{2\sigma}\beta_1$$

Interestingly, the effective eddy viscosity becomes independent of ε

$$\nu_T^{\text{eff}} = -\frac{1}{2}\frac{k\beta_1}{s}$$

The production is related to the invariants as

$$-\frac{P}{sk} = \frac{a_{ij}S_{ji}}{s} = \beta_1(1-r) + \beta_3 IV + 2\beta_4 V.$$

In 2D mean flows $IV = V = 0$ and

$$-\frac{P}{sk} = \beta_1(1-r).$$

cādence

2D EARSM solution

- The 2D EARSM solution is here used for illustrating a possible functional behavior for P/sk and the β1,2 coefficients vs. the parameters σ and r.
- P/sk and β1 are related by definition and shown in Figure 1. β2 is not monotonic in r, but β2/(β1 σ) is. Hence, parameterization of β2/(β1 σ) might be more easily found than that of β2 directly.

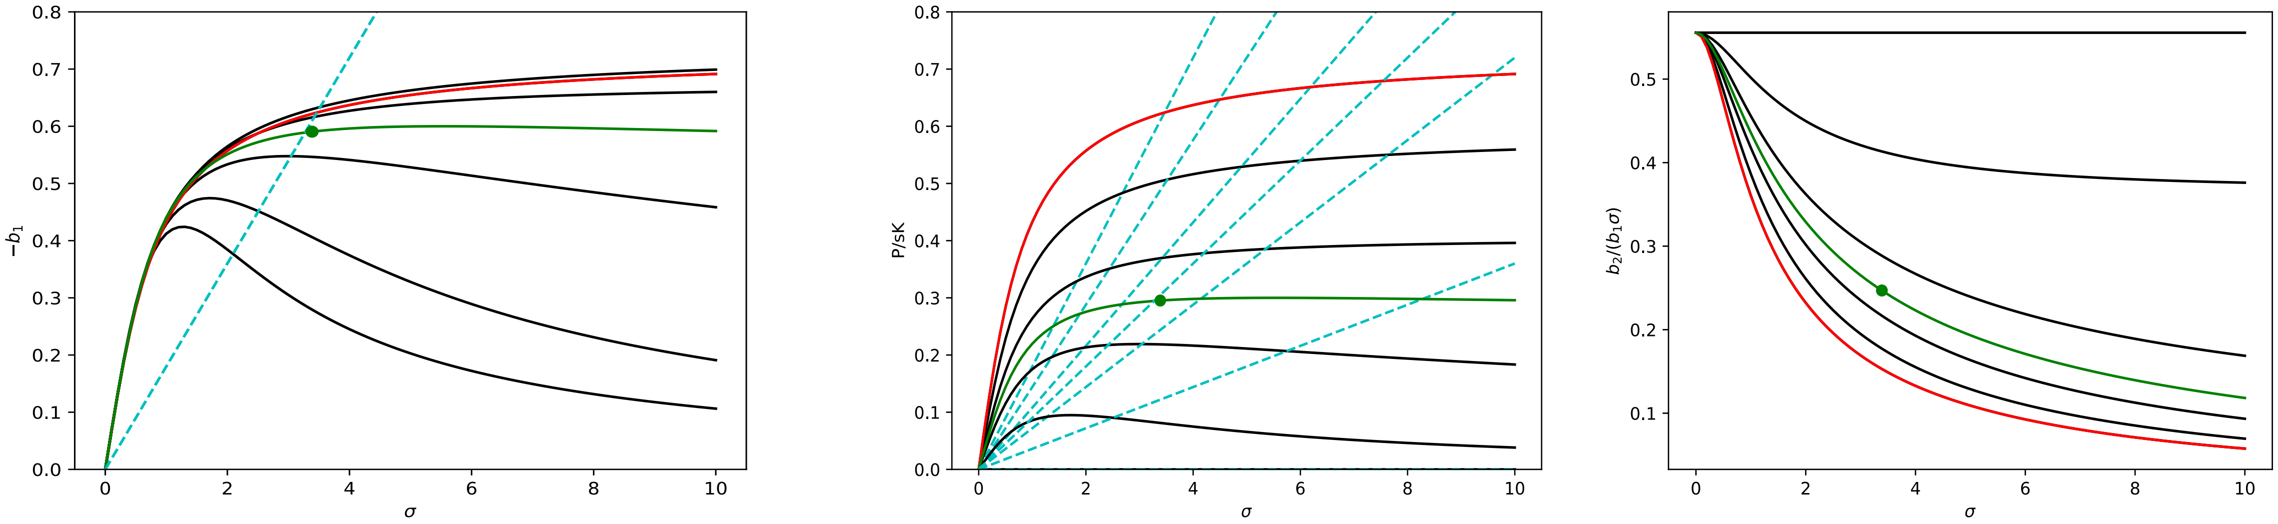

Fig. 1: The alternative EARSM solution for $\boldsymbol{\sigma}$ from 0-10 and $\boldsymbol{r}$ = 0.0, 0.2, 0.4, 0.5, 0.6, 0.8 and 1.0 (colours consistent with the curves). The green dot (•) is the log solution with $\boldsymbol{P} = \boldsymbol{\varepsilon}$. The green dashed lines are eddy-viscosity model with $\boldsymbol{C_\mu} = 0.09$. Left: β1; Middle: P/sk; Right: β2/(β1 σ)

Summary of New EARSM Model and ML Approach

- EARSM: novel framework within HiFi-TURB
  - $\boldsymbol{S}$ and $\boldsymbol{\Omega}$ are normalized by their magnitude s
  - Invariants become "structure parameters"
  - **limited** in magnitude, **independent of $\boldsymbol{\varepsilon}$**
  - $r = \boldsymbol{\Omega}^2/(\boldsymbol{S}^2 + \boldsymbol{\Omega}^2)$ – vorticity/swirl measure (0-1)
  - III, IV and V – measure of 3D effects (=0 in 2D)
  - Only one non-equilibrium parameter: $\sigma = sk/\varepsilon$
- Additional pros
  - Machine-learning problem much better posed
  - Invariants have a clear physical interpretation
  - Data completeness easily illustrated
  - Alternative parameters more easily added
- Cons
  - Singularities for $s \rightarrow 0$ (can be handled)

- Local data
  - Output: Five β-coefficients representing
    $$a_{ij} = \beta_k T_{ij}^{(k)}$$
  - Minimizing $\left\|\hat{\beta}_k - \beta_k\right\|$, β_k extracted from data
  - Using MEP/Auto-Encoder to find expressions
- Consequences
  - Physics more clearly decomposed on the different T_{ij}'s
  - β_k from data -> physical interpretation
  - Realizability outside of data range better kept
  - Big step forward (Sandberg and others)
  - Frame invariance enforced
  - However, resulting "model" not really general

cādence

The MEP-0 Model From BSC (Mirò, Lehmkuhl)

- Multi Expression Programming (MEP)
- Data: Stanford 3D diffuser (DNS by BSC)
- Resulting expression from MEP
- No additional tuning or fitting to “standard values” are made

$$\beta_1 = (-V + r)\sinh\left(\tanh(r - \sigma)\right)$$

$$\beta_2 = -\frac{0.269095\sigma^r}{\cosh\left(0.269095\sigma^r\right)}$$

$$\beta_3 = \log\left(-III_S + \tanh^{\tanh(V)+0.0863705}(V)\right)$$

$$\beta_4 = III_S - IV + V^{0.5}\sinh\left(e^r\right)$$

$$\beta_5 = VV^{IV}\left(-\sigma + \log(V)\right)$$

The model does not make much sense for a modeller …

cādence

Other Approach

- EARSM data generator
  - Use EARSM for generating synthetic data for meta training (training of training)
  - Invented and optimized approximate functions for synthetic data
  - Used to initiate MEP

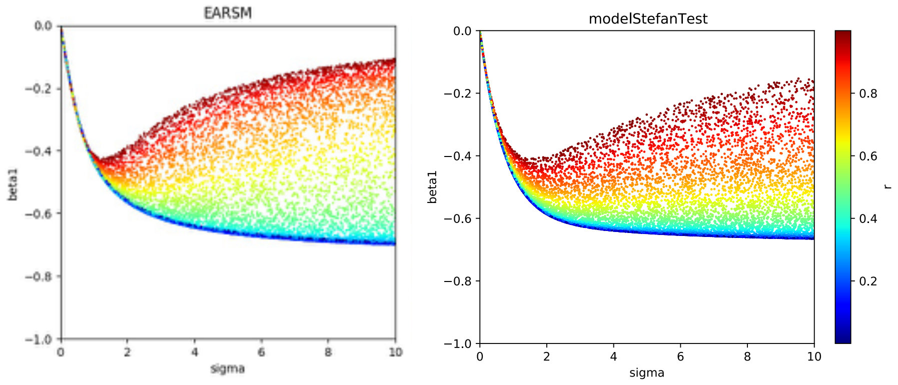

Example of person-made function (PMF)

$$\beta_1 = -0.591(f_1 + f_r f_2)$$

$$f_1 = 1 - \exp(-1.295\sigma)$$

$$f_2 = \sigma^2/(12.297 + \sigma^2)$$

$$f_r = -0.672(2r - 1) - 0.502(2r - 1)\text{^}2$$

Preliminary Results

- Analytical model checking (upper fig)
  - Limit of $\sigma \to 0$: $\beta_i \to 0$
  - RDT limit of $\sigma \to \infty$: $\partial/\partial\sigma \to 0$
  - Log layer fix point $\mathrm{P}/\varepsilon = 1$, $r = 0.5$
  - Homogeneous shear-layer fix point $\mathrm{P}/\varepsilon = 1.8$
  - Bradshaw assumption captured by MEP-0 $-a_{12} = 0.3$
  - Non-realizable red zones avoided by MEP-0
- A-posteori testing (lower fig)
  - Automatic model implementation
  - MEP model formulation as input
  - Channel flow solved within seconds – can be used ”in the loop” of MEP
  - MEP-0 model gives $\kappa = 0.38$

Surprisingly close to reference!

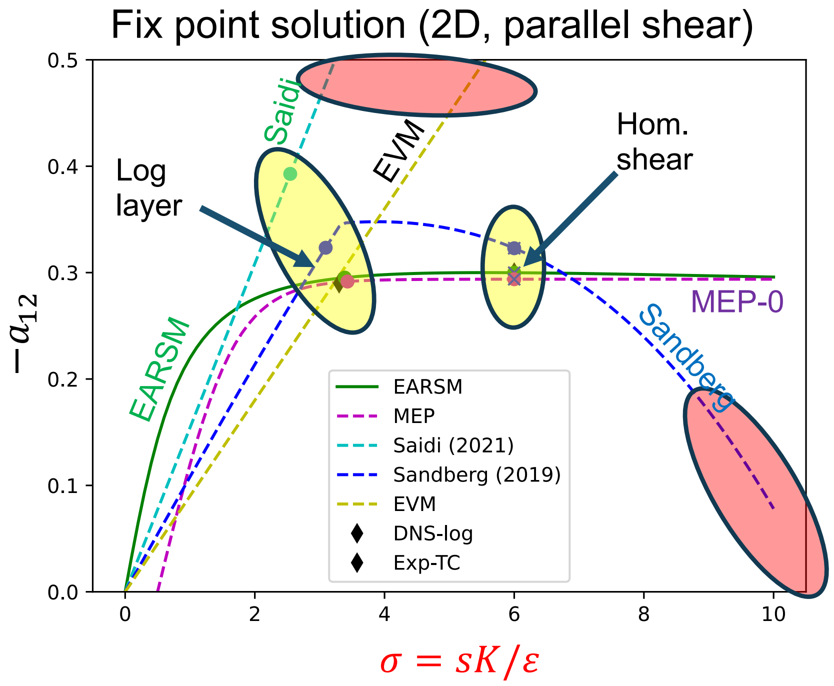

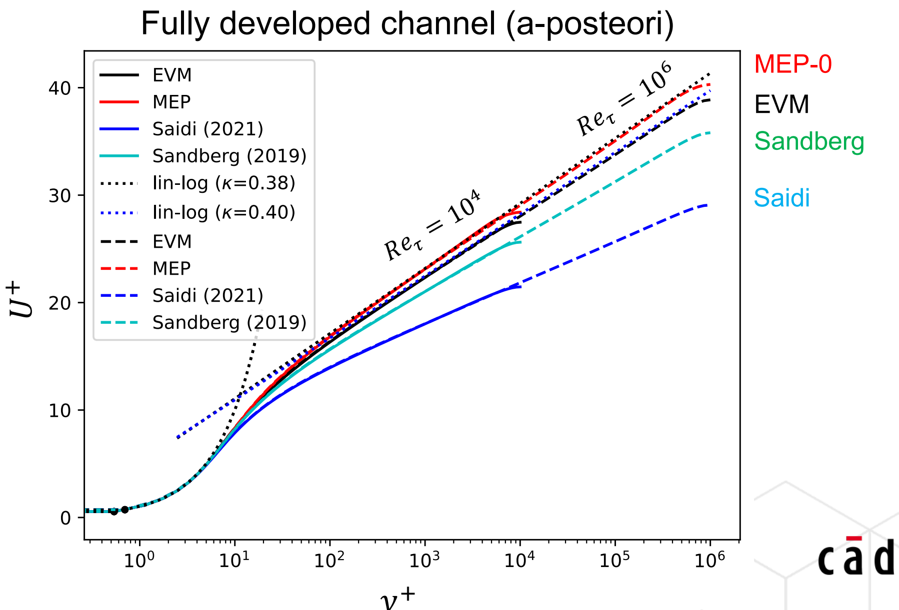

Requirements on Data

- Data should contain **all five $\boldsymbol{\beta}$'s** being non-zero
- Data should contain **all invariants** in their complete ranges
  - r – relative rate of rotation (0-1)
  - III, IV and V – measure of 3D effects (=0 in 2D)
  - $\sigma = sK/\varepsilon$ – non-equilibrium parameter
- We don't see sufficient variation in our data
  - Data lives in $r \approx 0.5$, III, IV, $V \approx 0$ and $T_{ij}^{(4)}, T_{ij}^{(5)} \approx 0$.

- **Data from separated curved step and from 3D Stanford Diffuser are mostly around $\boldsymbol{r \approx 0.5}$**

- Lack of data variation – problem is underdetermined
  - Std. approach taken by others – problem is "hidden"
  - With the present approach - problem is visible and can possible be handled
- ML strategy
  - Unintelligent approach will fail
  - ML approach must be guided by physical knowledge
  - Can we use synthetic data from EARSM to guide through the wilderness?

cādence®

First Validations, MEP-0 and manually derived model from A-E (Dirk Wunsch-Cadence)

Case description

Solver description:

- Finite-volume density-based solver (Open^{TM}) with low-speed preconditioning, local-time stepping, residual smoothing and full-multigrid acceleration with central scalar scheme.
- Three models tested:
  - S-BSL EARSM of Menter et al (2012) as reference
  - BSL MEP0 NLEVM from Wallin et al (2022)
  - BSL DIRK NLEVM from Wunsch et al (2022)
- During coarse-grid initialization, CFL = 1 and, on the fine grid, CFL = 2.
- 3 grid levels for the multi-grid.

Computation settings

- Turbulent flat plate with a length of 4 m
- Freestream velocity: 68.79 m/s
- Kinematic viscosity: $3\ 10^{-5}\ m^2/s$
- Reference data: Weighardt and Tillman (1951), Klebanoff (1955)

Results: Turbulent Fluctuations

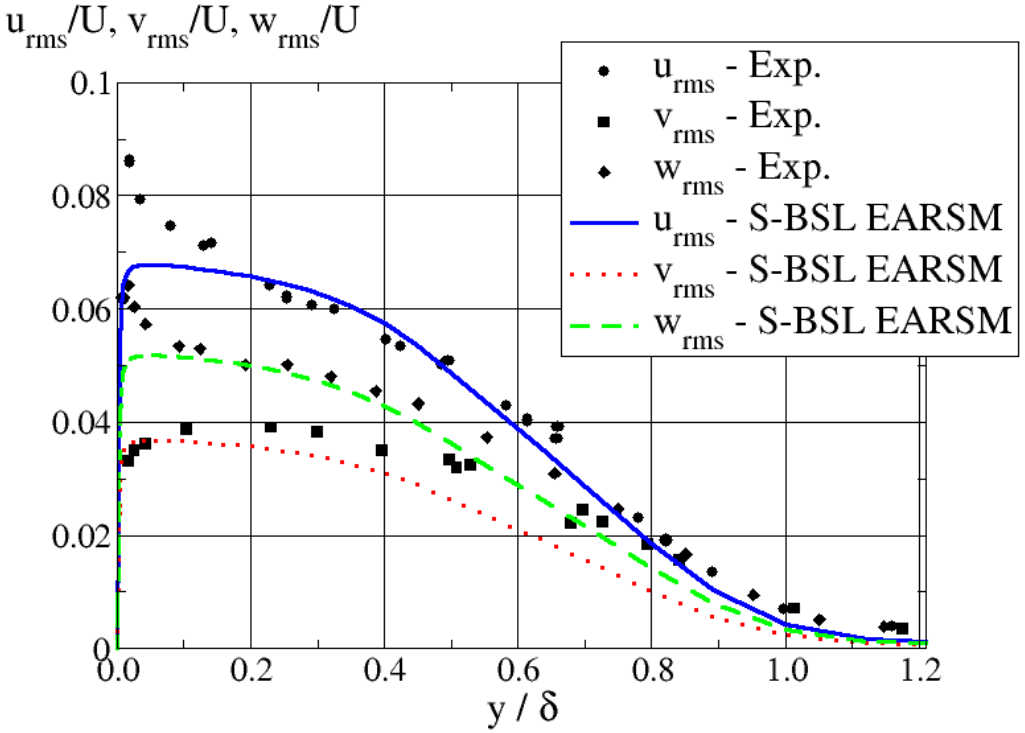

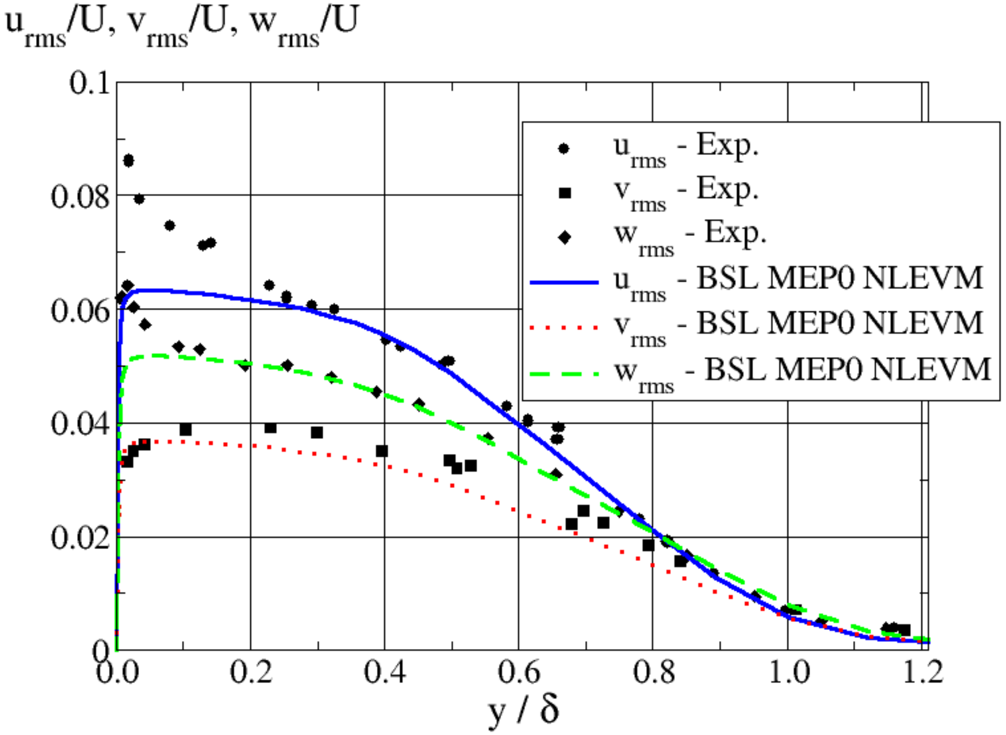

Results: shear stress

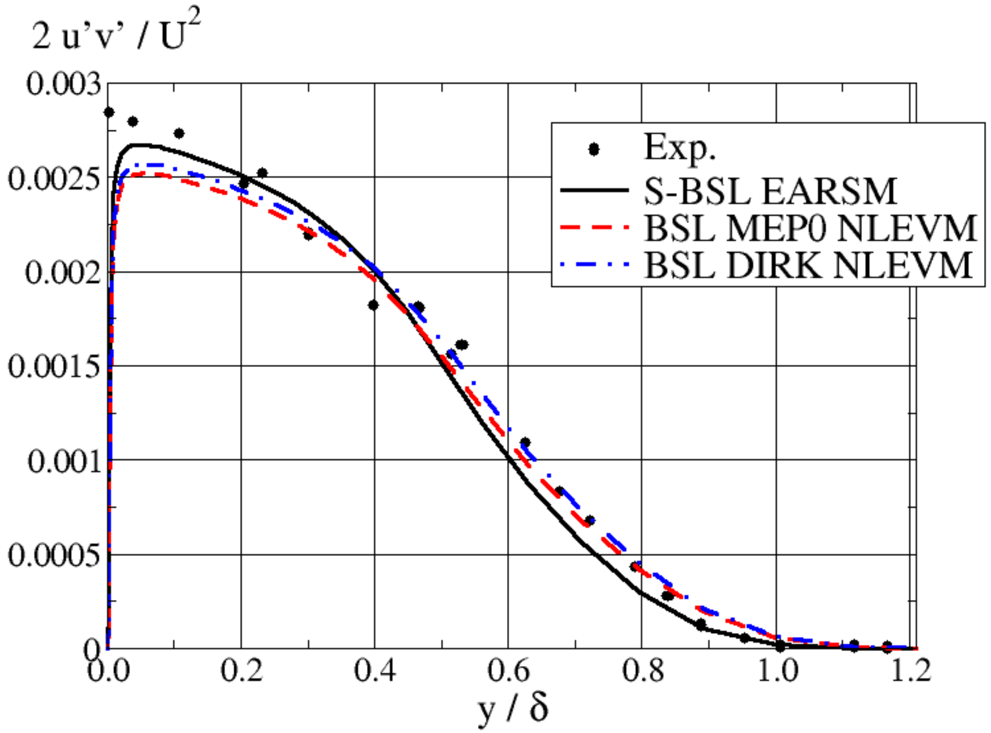

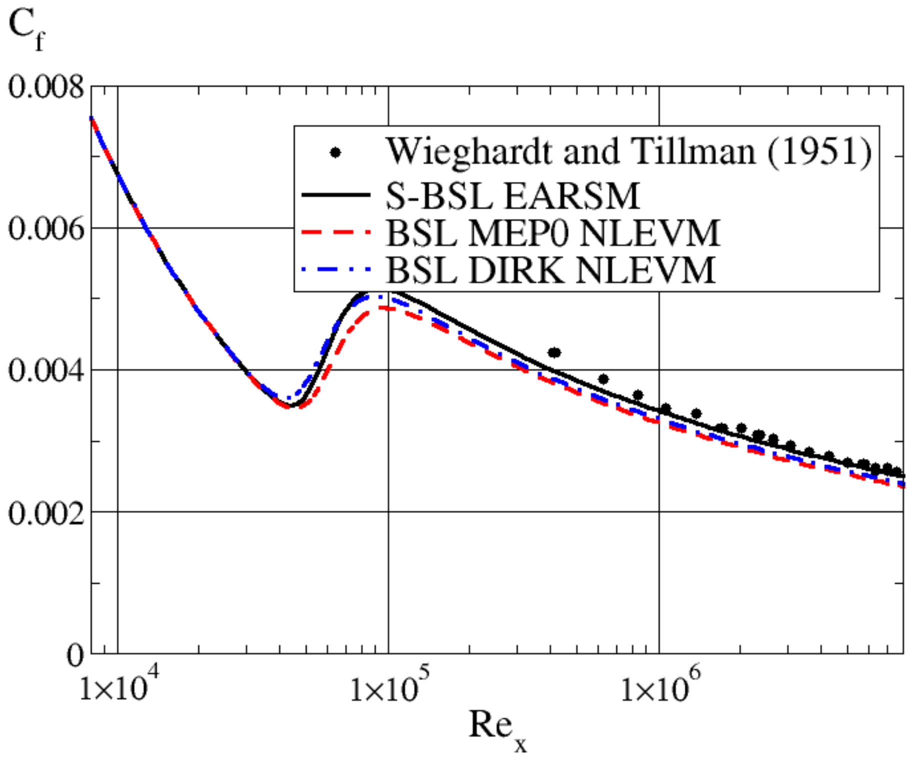

Results: streamwise velocity profiles

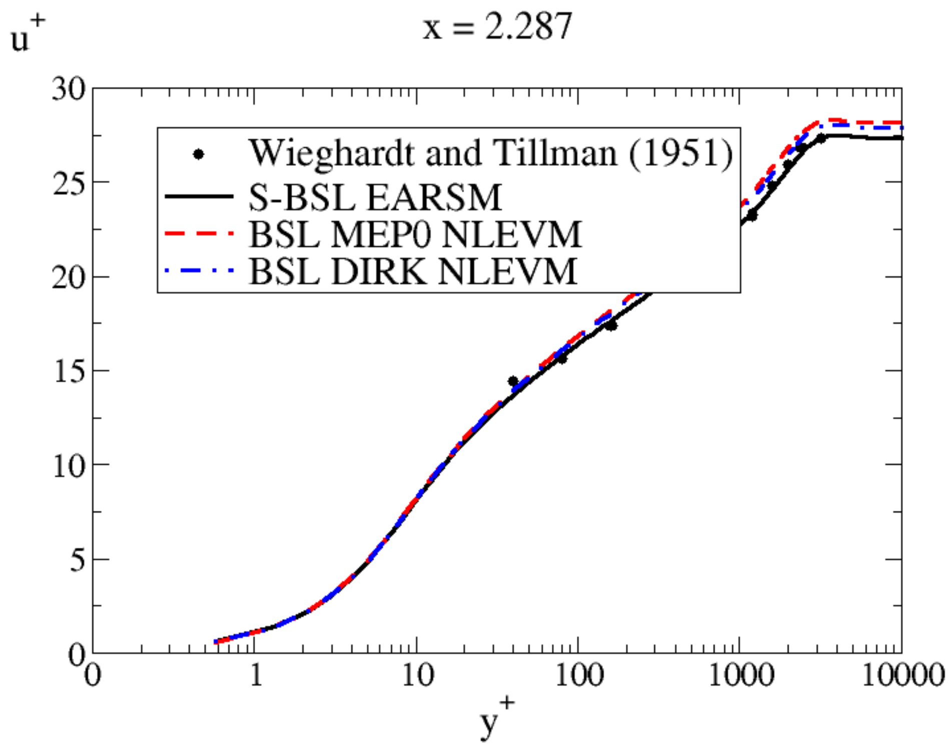

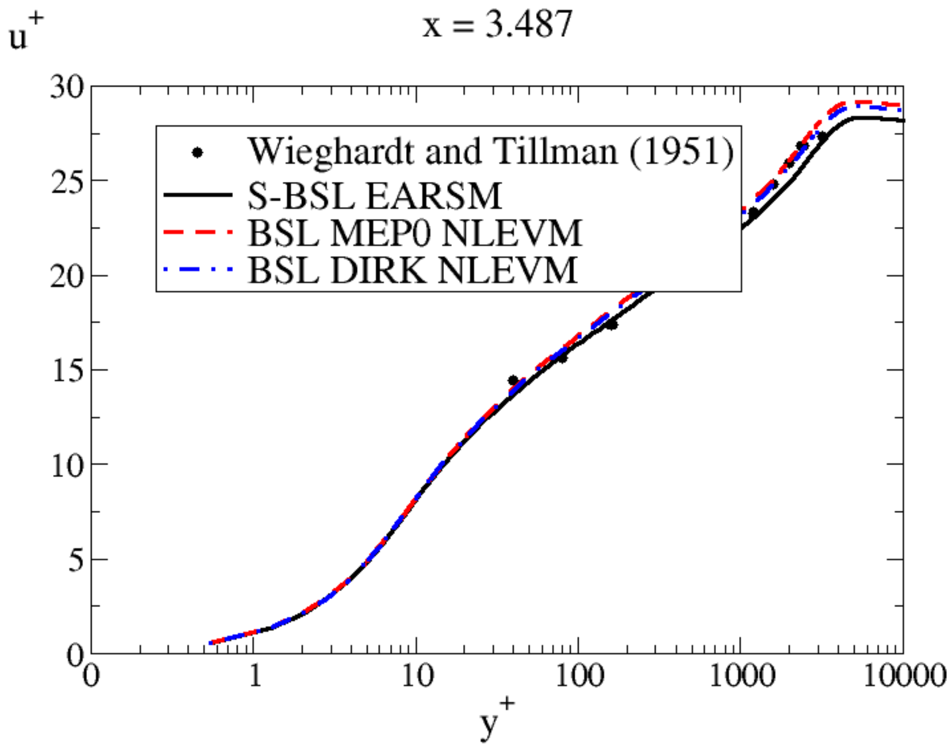

Comments from Stefan Wallin: *I find it extraordinary that the basic log layer is so accurately captured from the rather low Re 3D diffuser data with all its complexity. There are limited log regions in the data though, but one could suspect that it will be drowned in the vast amount of data with high complexity.*

More tests are underway.

The road to Scale Resolving Industrial Simulations

The road to Scale Resolving Industrial Simulations

From Spalart & Strelets (2019)

- *Traditional turbulence modeling is challenged from two sides:*
  - *Turbulence-resolving simulations*
    - *These are promising, but far from industry practical. We need **many** exaflops*
    - *The flow fields have some very "interesting" features…*
    - *We contend that DES is cleaner, and will deliver well before WMLES and VLES*
  - *Artificial intelligence*
    - *We contend that this work is still in its infancy, and much of it is simply unsuccessful*
    - *A lot of "adult supervision" is needed*
    - *Did we the "adults" fail to explain modeling (too bad Wilcox's book is now rare)?*

Cost estimates for WRLES/DNS

- Numerous papers have been published on the grid point requirements for WRLES and DNS
- More recent ones from Yang & Griffin (2021); Choi & Moin (2012) and various presentations by Spalart and Strelets
  - These estimates rely on assumptions on the number of grid points per cube ($n_x n_y n_z$) and time steps estimates assuming explicit time integration
- *According to M. Strelets N_cubes ~ 6 10^6 would be required for a swept wing, and on a high-lift wing, with 3 LE, this would lead to 2 10^7 cubes. And assuming 20^3 DOF per cube, leads to 160 billion points for the boundary layer. Not counting some 10^4 time steps*.
- Note: such a high resolution is perhaps not required for an engineering accuracy of 1 count (E-4)

Moreover two additional elements might be a game changer

- High Order Methods
- GPU technology, beyond Moore's law

The disruptive numerical features of HO

- Low order schemes (less than 3rd order) fail in predicting wave propagation dominated flow cases, on acceptable mesh sizes
- This can be seen from the convection of a vortex, comparing second order to 4th order accuracy with same number of data points (16000).
    - After having been transported 47 times through the width of the passage, the vortex is completely dissipated by the second order scheme, while it is still unchanged by the 4th order method, on the same number of data points
- Another view: A second order code requires 40 to 50 points per wavelength to represent accurately a propagating wave or signal over longer times, while a 4th order code requires only 5 points per wavelength
- This difference means that a second order code will require a mesh of 500 to 1000 million points to reach the accuracy of a fourth order code on a mesh of 1 million points.

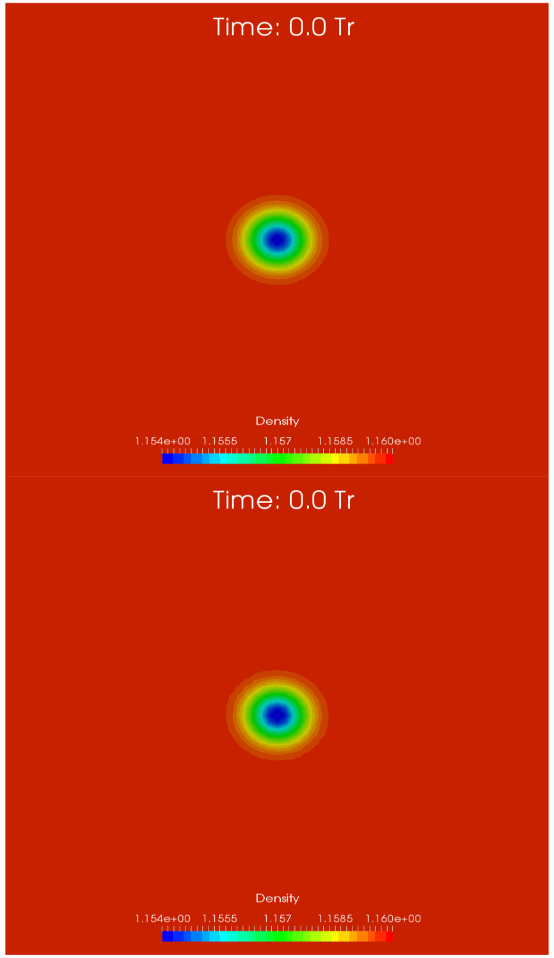

HPC Capacities in the Near Future

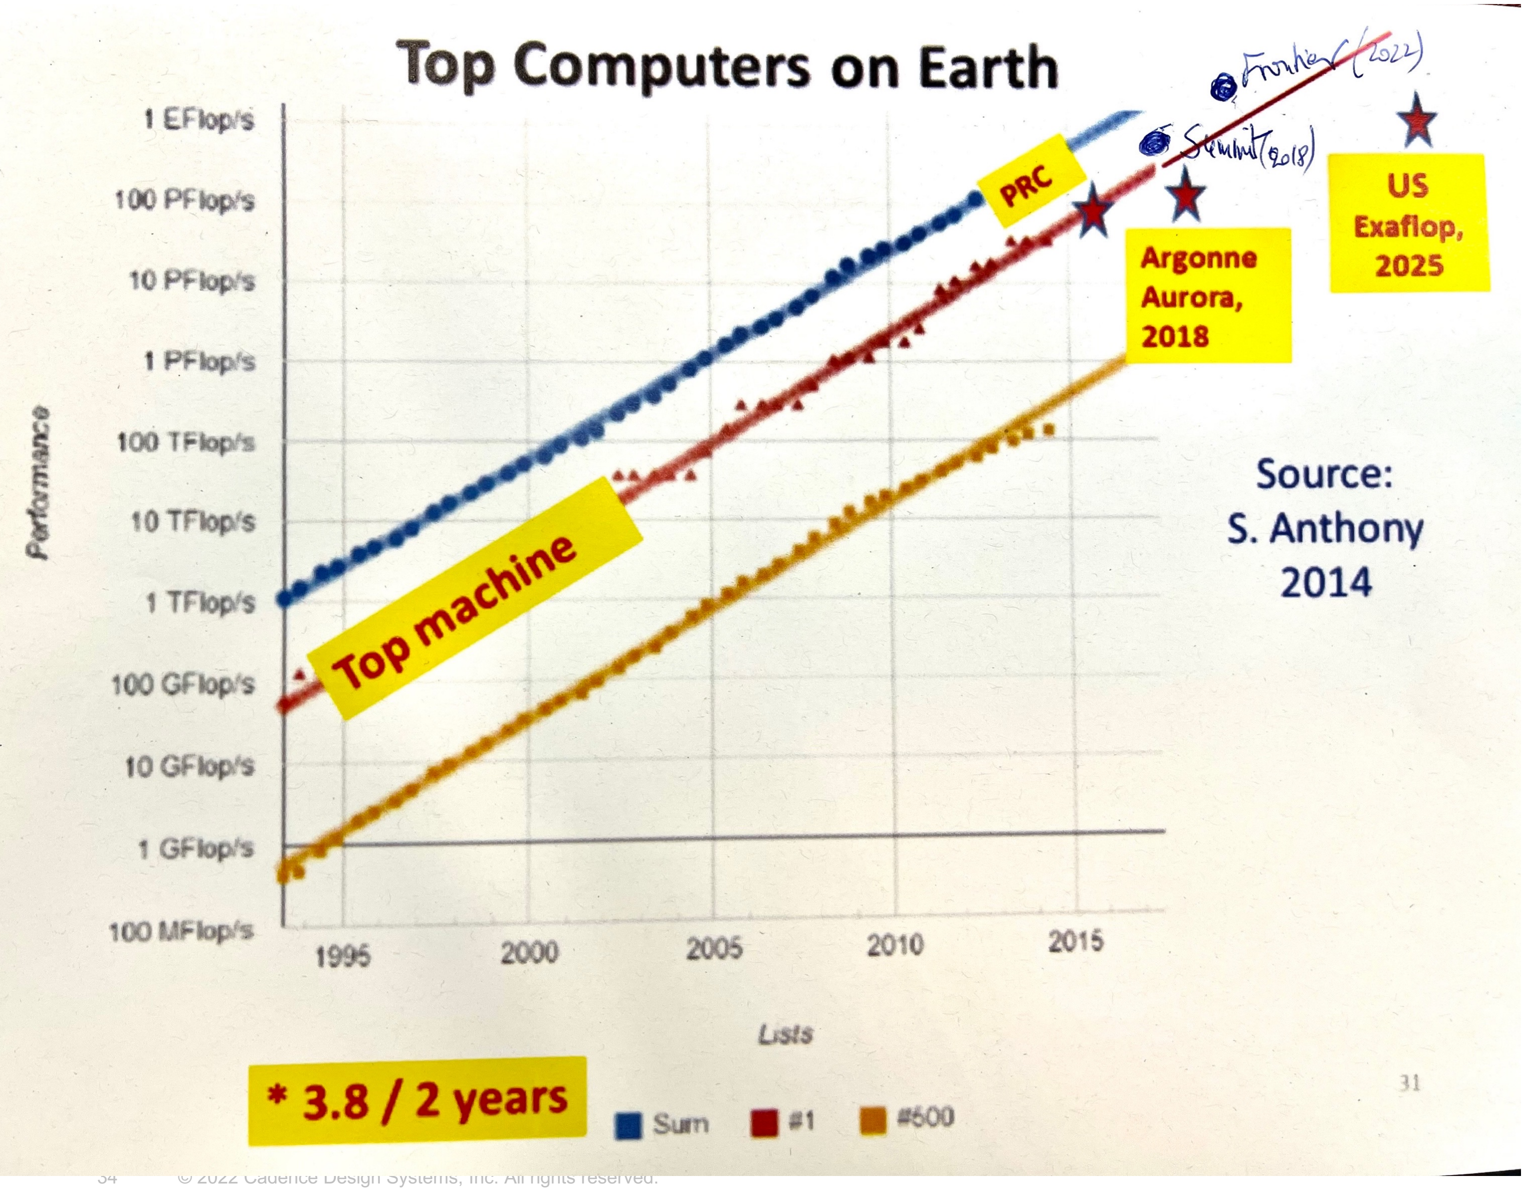

H1OO NVIDIA GPU

Exascale High-Performance Computing

The NVIDIA data center platform consistently delivers performance gains beyond Moore's Law. And H100's new breakthrough AI capabilities further amplify the power of HPC+AI to accelerate time to discovery for scientists and researchers working on solving the world's most important challenges.

H100 triples the floating-point operations per second (FLOPS) of double-precision Tensor Cores, delivering 60 teraFLOPS of FP64 computing for HPC. AI-fused HPC applications can leverage H100's TF32 precision to achieve one petaFLOP of throughput for single-precision, matrix-multiply operations, with zero code changes.

H100 also features DPX instructions that deliver 7X higher performance over NVIDIA A100 Tensor Core GPUs and 40X speedups over traditional dual-socket CPU-only servers on dynamic programming algorithms, such as Smith-Waterman for DNA sequence alignment.

cādence

Fidelity/HO New Scale Resolving Industrial Product

High-Order iLES

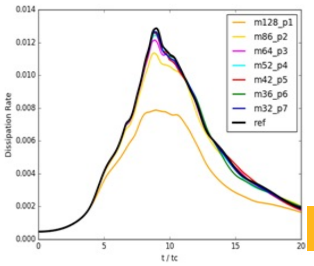

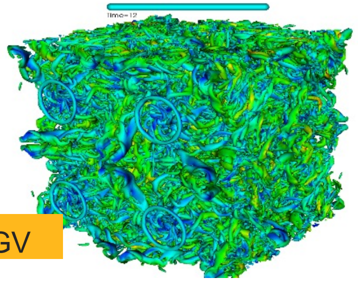

TGV

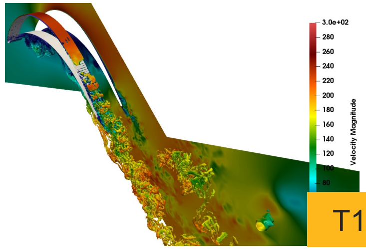

T106-C

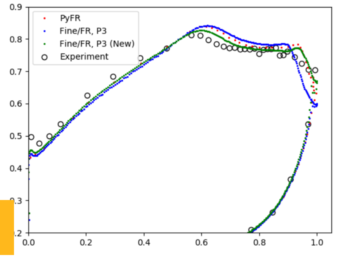

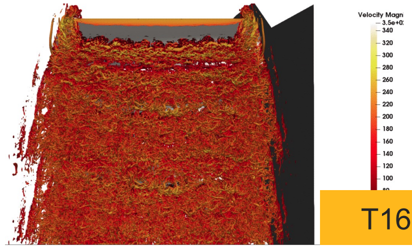

T161

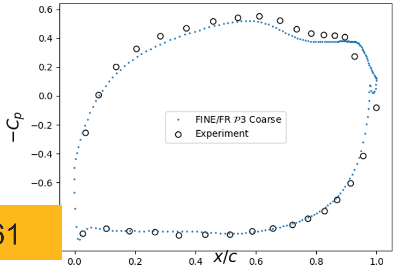

State-of-the-art algorithms (Flux Reconstruction)

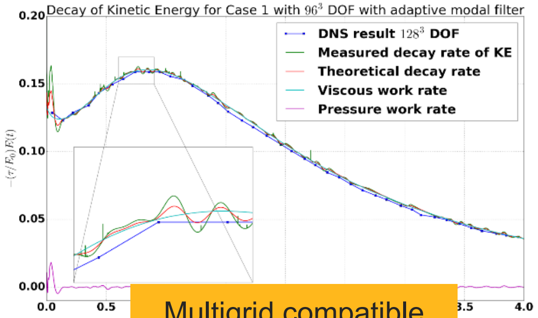

Multigrid compatible regularization

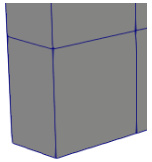

HN support

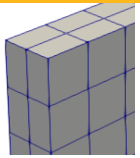

Efficient convergence acceleration for steady & unsteady via dual-time-stepping

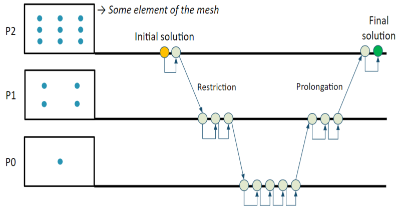

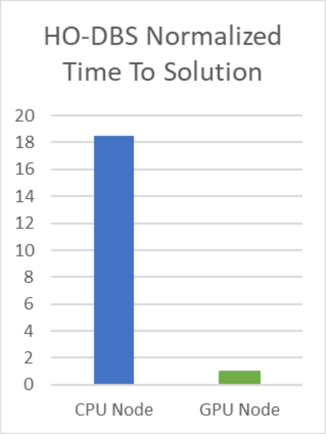

Scalable distributed parallel, GPU compatible implementation.
1 GPU >= 250 CPU cores

Integrated Workflow

CAD & Geom preparation

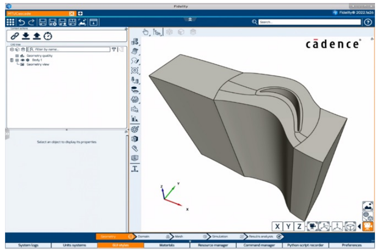

High-Order curved mesh generation/import

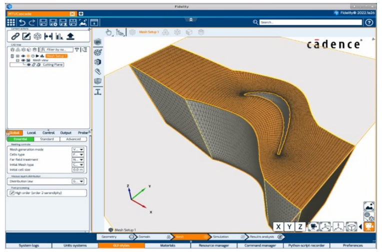

Precursor + HO simulation

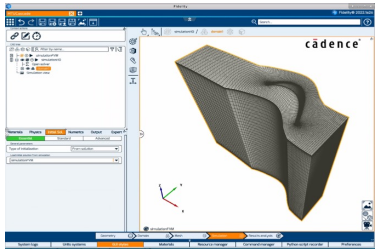

Postprocessing

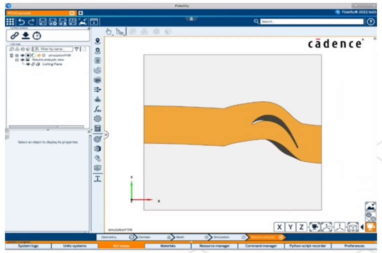

cādence

Automatic curved meshing

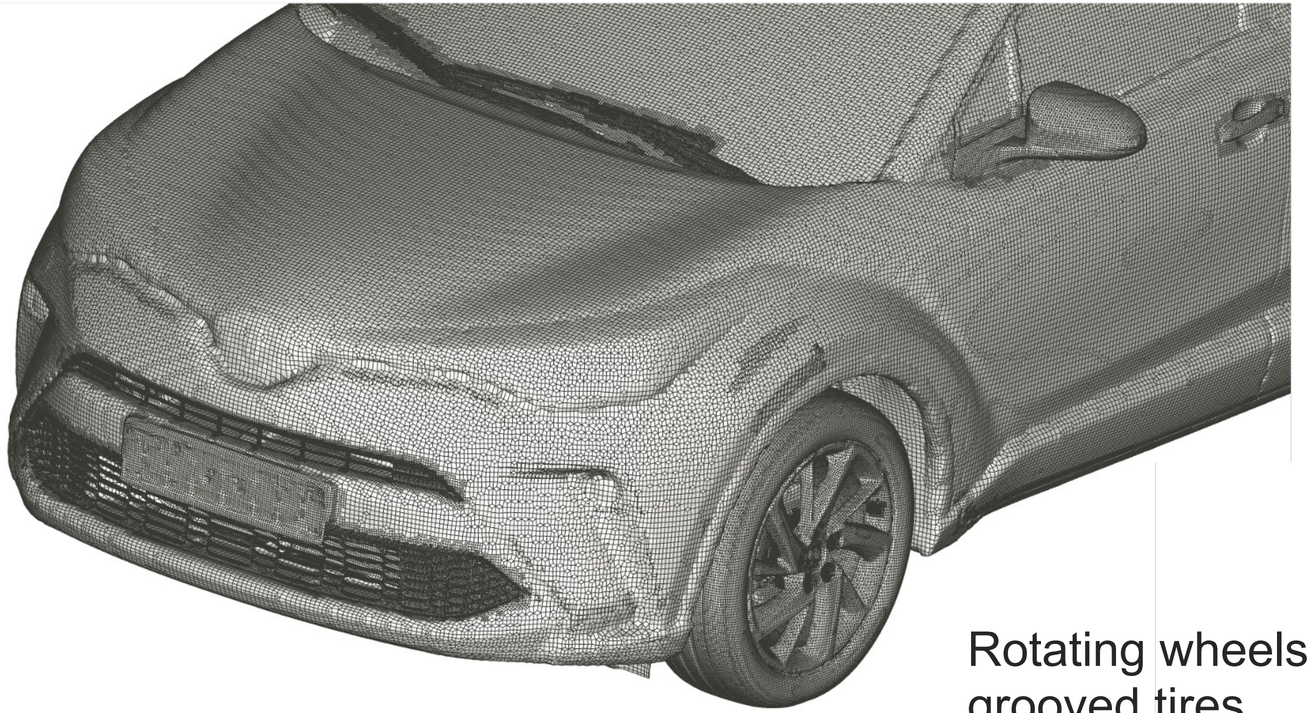

Mesh size around 5M points

Rotating wheels with grooved tires

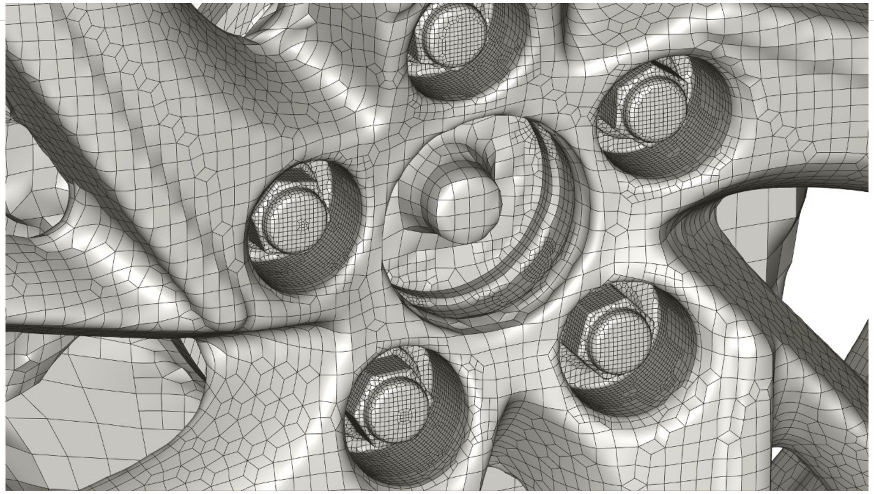

Curved mesh

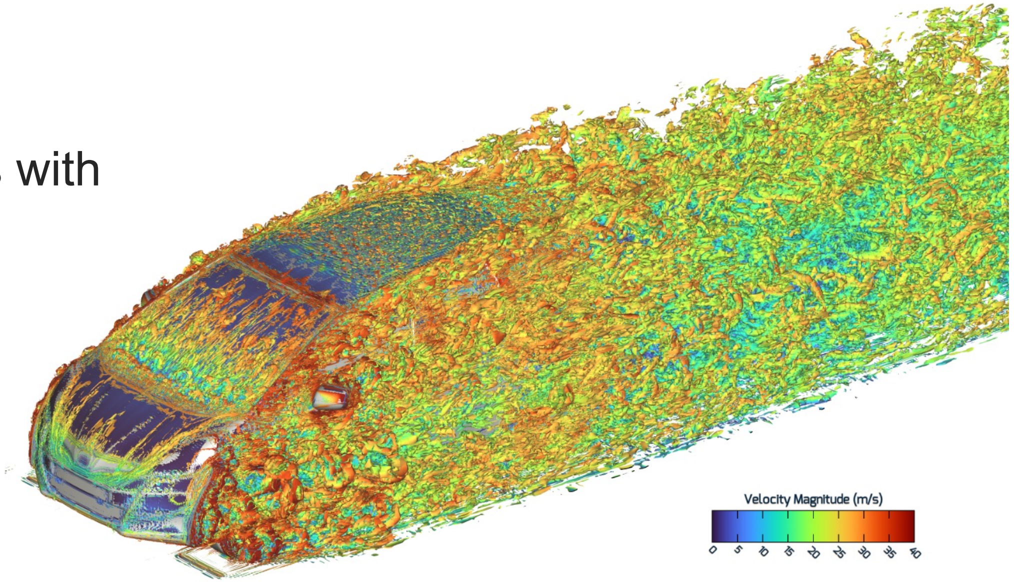

Massive GPU acceleration

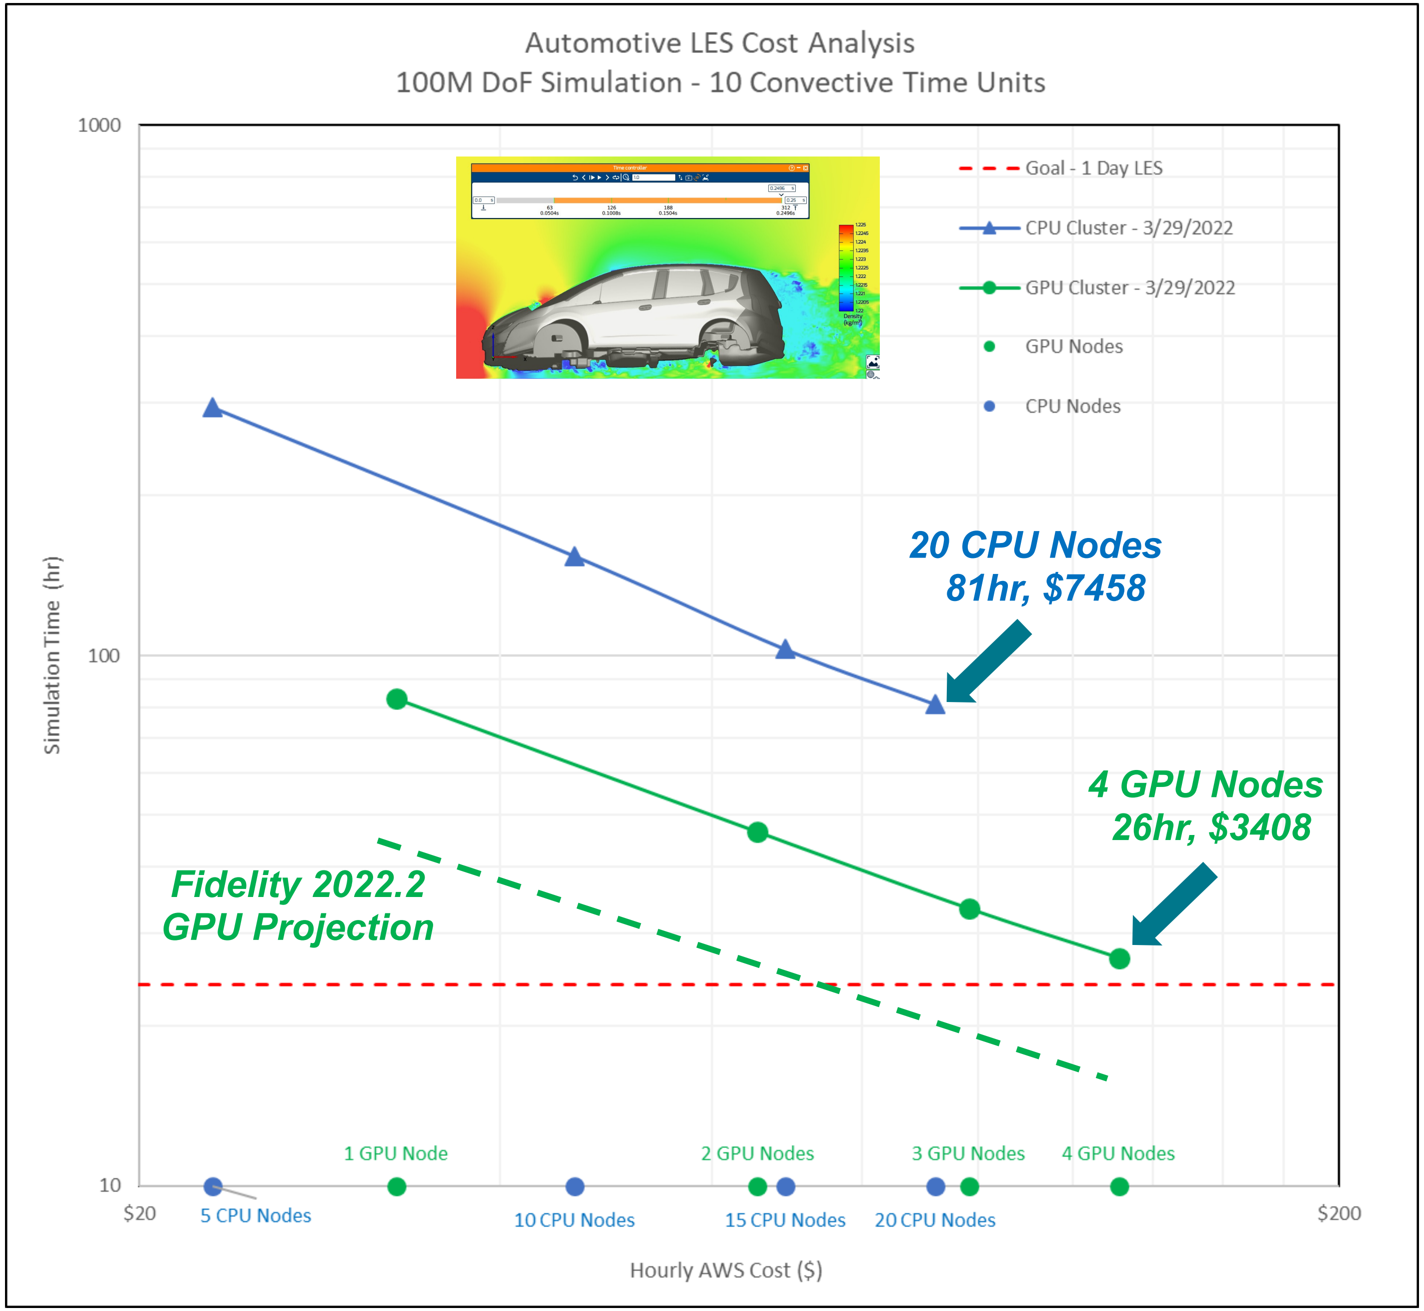

- The numerical nature of HO allows for **massive GPU acceleration**
- The current version of the solver achieves speed-up of **1 GPU equivalent to 250 CPU's**
- **1 GPU node contains 8 NVIDIA A100 GPU's and 64 CPU's**
- This enables **1-day return time** on a few GPU nodes for LES on industrial automotive configurations; accuracy of 2-3 counts (E-3) for ±150 MDoF
- At a fixed simulation time, GPU computing is up to 3X cheaper

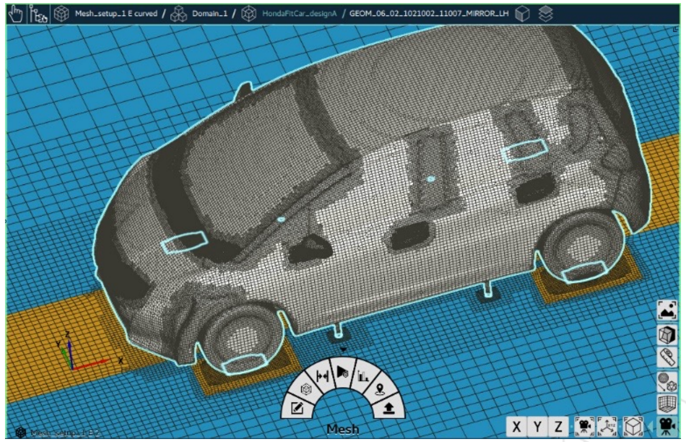

Summary on industrial WRLES in the near future

- High Order solvers allow a gain of 3 orders of magnitude in required number of points
- Full GPU optimization should allow another gain of 2 orders of magnitude in the ExaFlops projections
- Restricting to engineering accuracy, we estimate another potential gain of 2 orders of magnitude (to be confirmed)
- Our HO industrial solver is parameter free, based on first principles

The Challenge:

- Define appropriate meshing best practices for coarse curved grids, providing ILES (WRLES) accuracy on various industrial configurations
- Current Applications: Automotive and Turbomachinery
- Aeronautical applications for H1-2022, with DES and ILES options to be demonstrated