University of Stuttgart

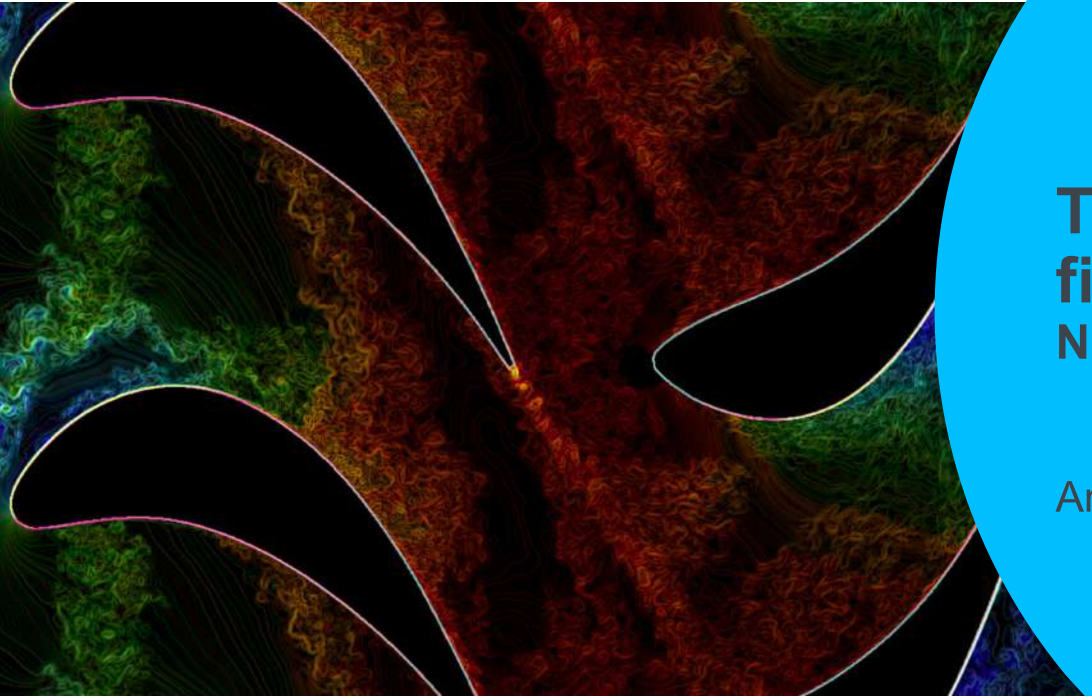

Towards data-driven high fidelity CFD
NHR4CES Workshop 23

Andrea Beck

Multiscale - Multiphysics

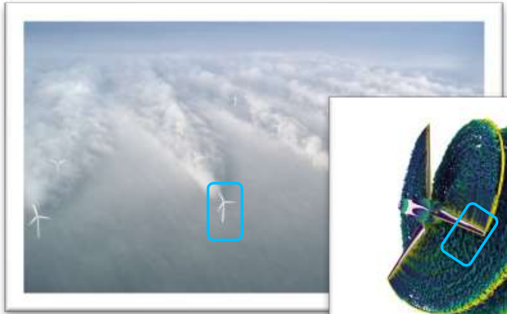

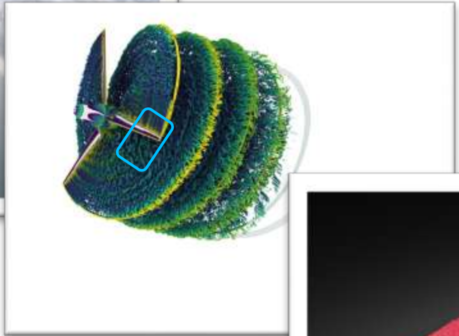

With friendly permission from
Pascal Weihing & Thorsten Lutz

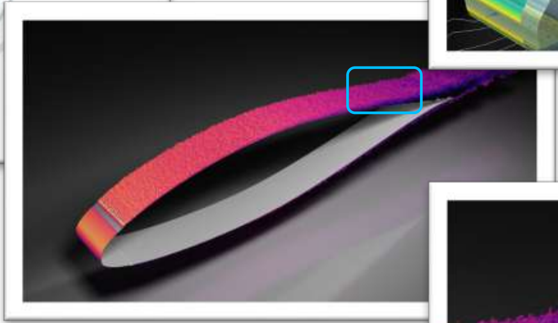

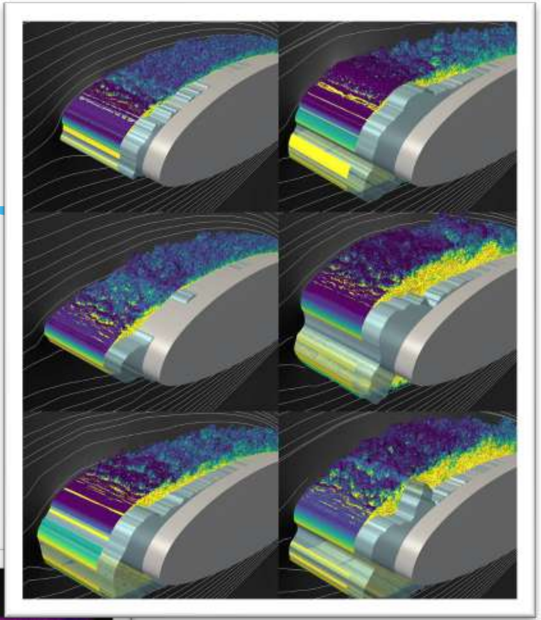

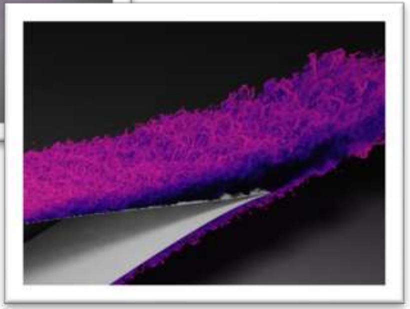

The Scale Gap

$$\underbrace{R\left(U_l, x_l, t_l\right) = 0}_{\text{Governing Equations at level l}}$$

$$\underbrace{R\left(U_L, x_L, t_L\right)}_{\text{Governing Equations at level L}} + \overbrace{M\left(U_l, U_L\right)}^{\text{Influence of level l on L}} = 0$$

$$\underbrace{R\left(U_\Delta, x_\Delta, t_\Delta\right)}_{\text{Governing Equations at level }\Delta} + \overbrace{M\left(U_L, U_\Delta\right)}^{\text{Influence of level L on }\Delta} = 0$$

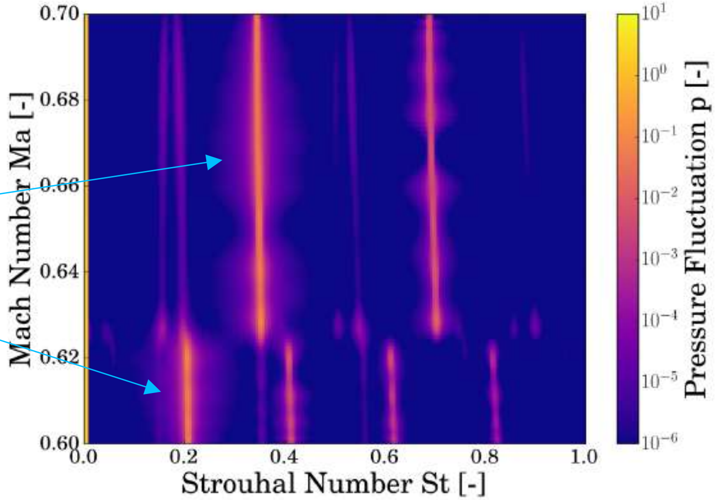

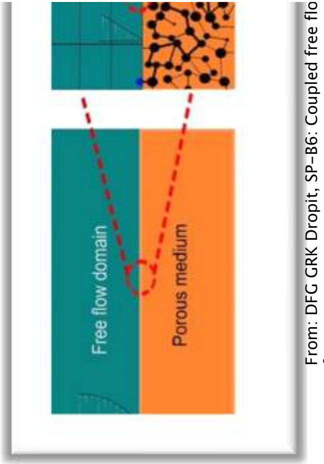

From: DFG GRK Dropit, SP-B6: Coupled free flo formation

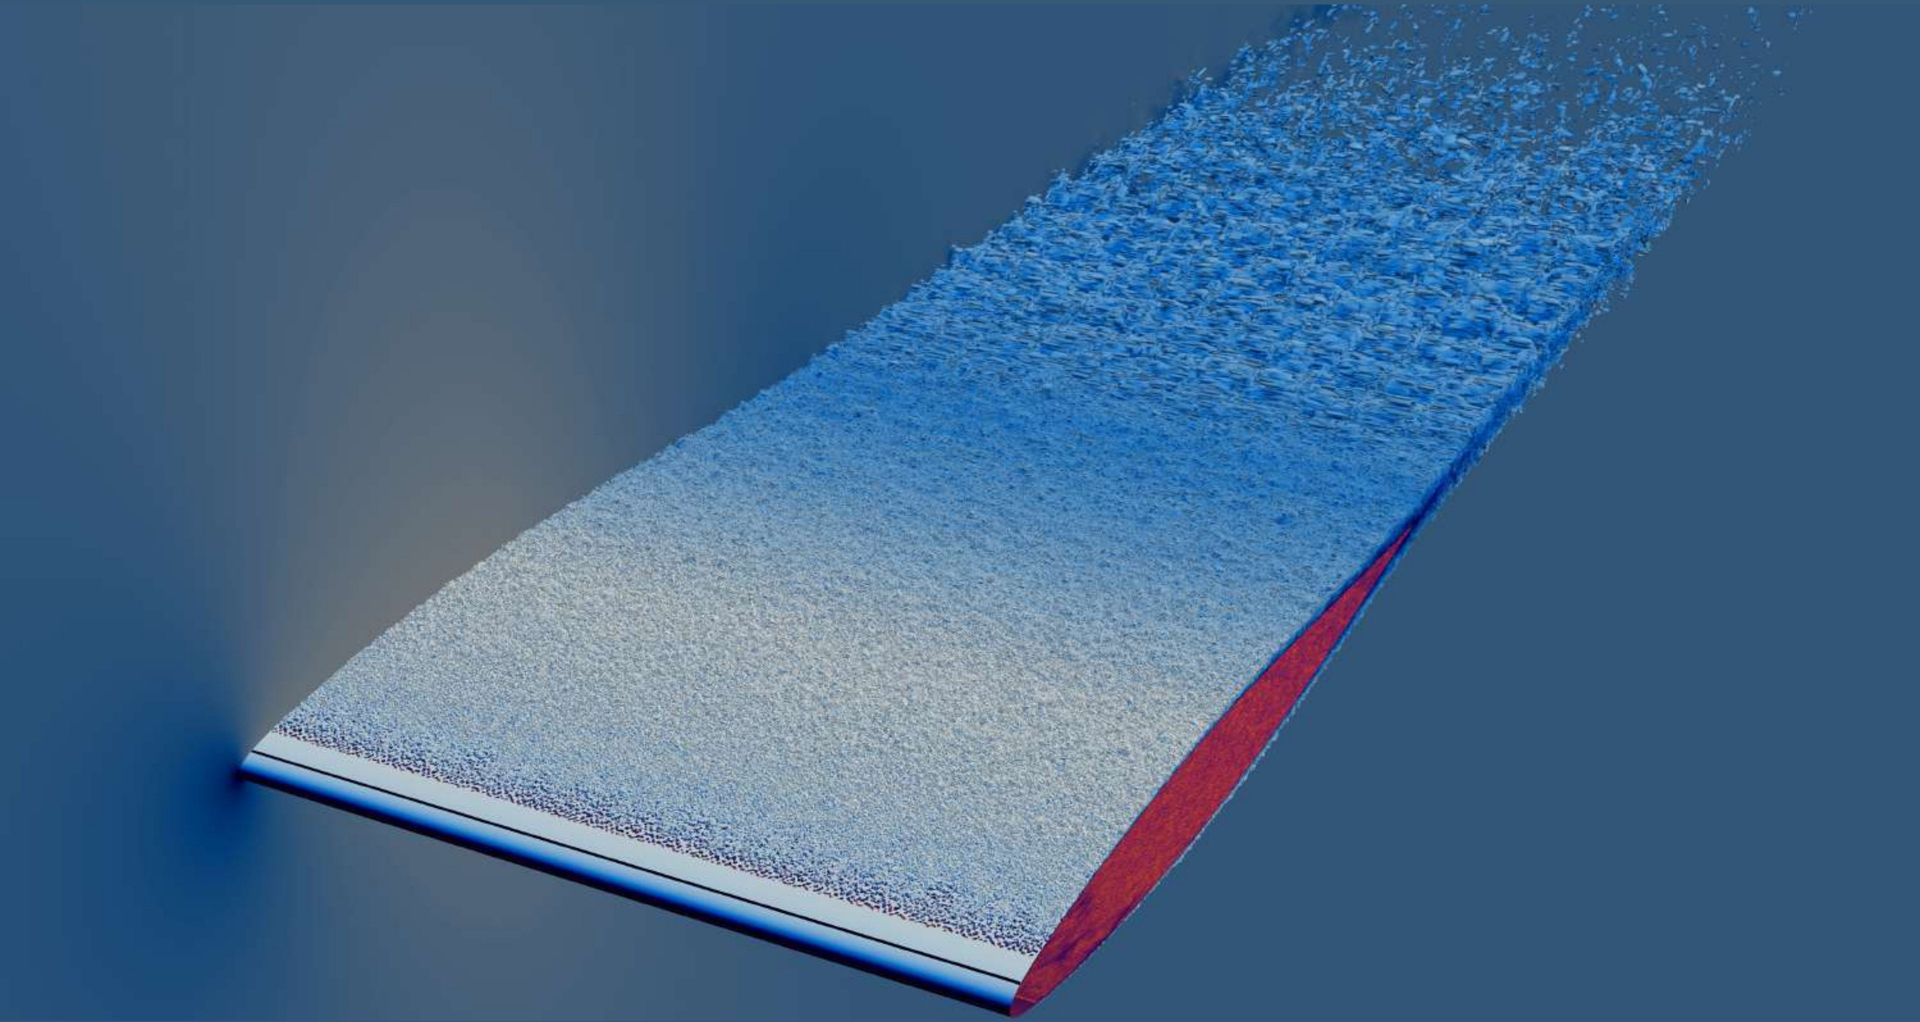

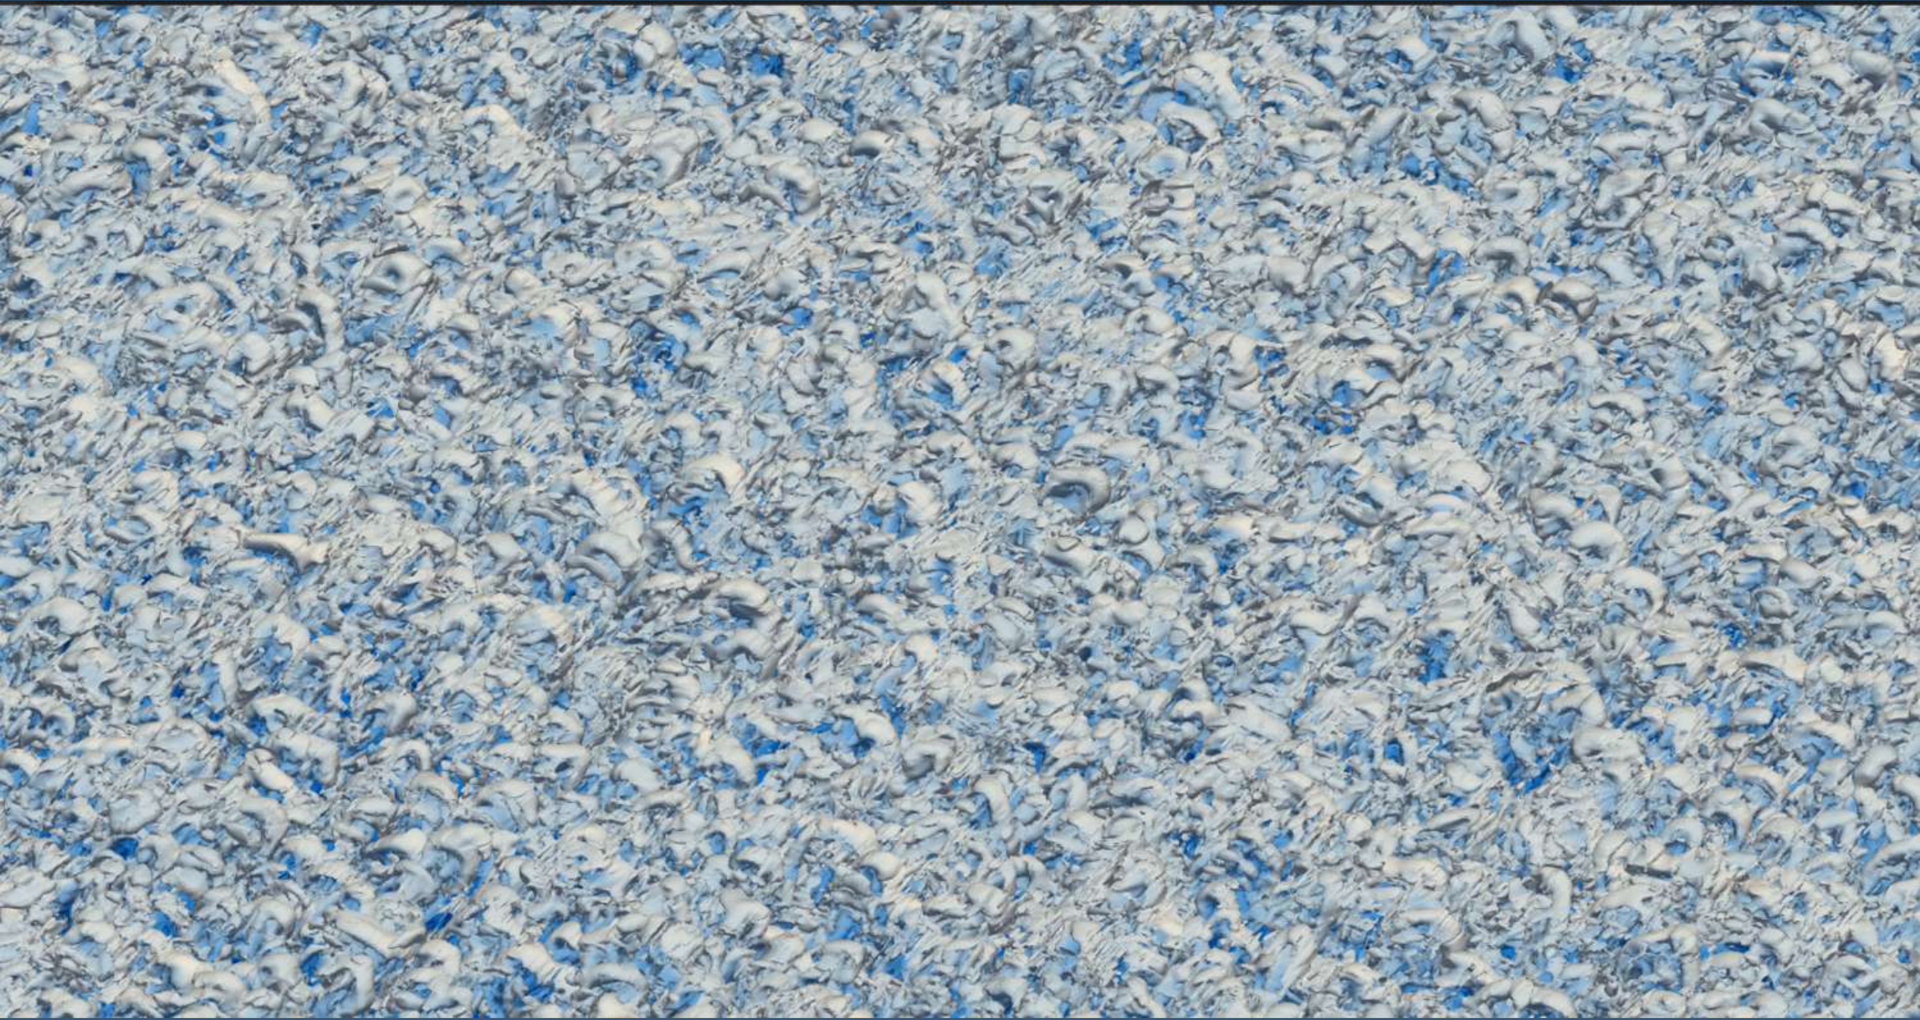

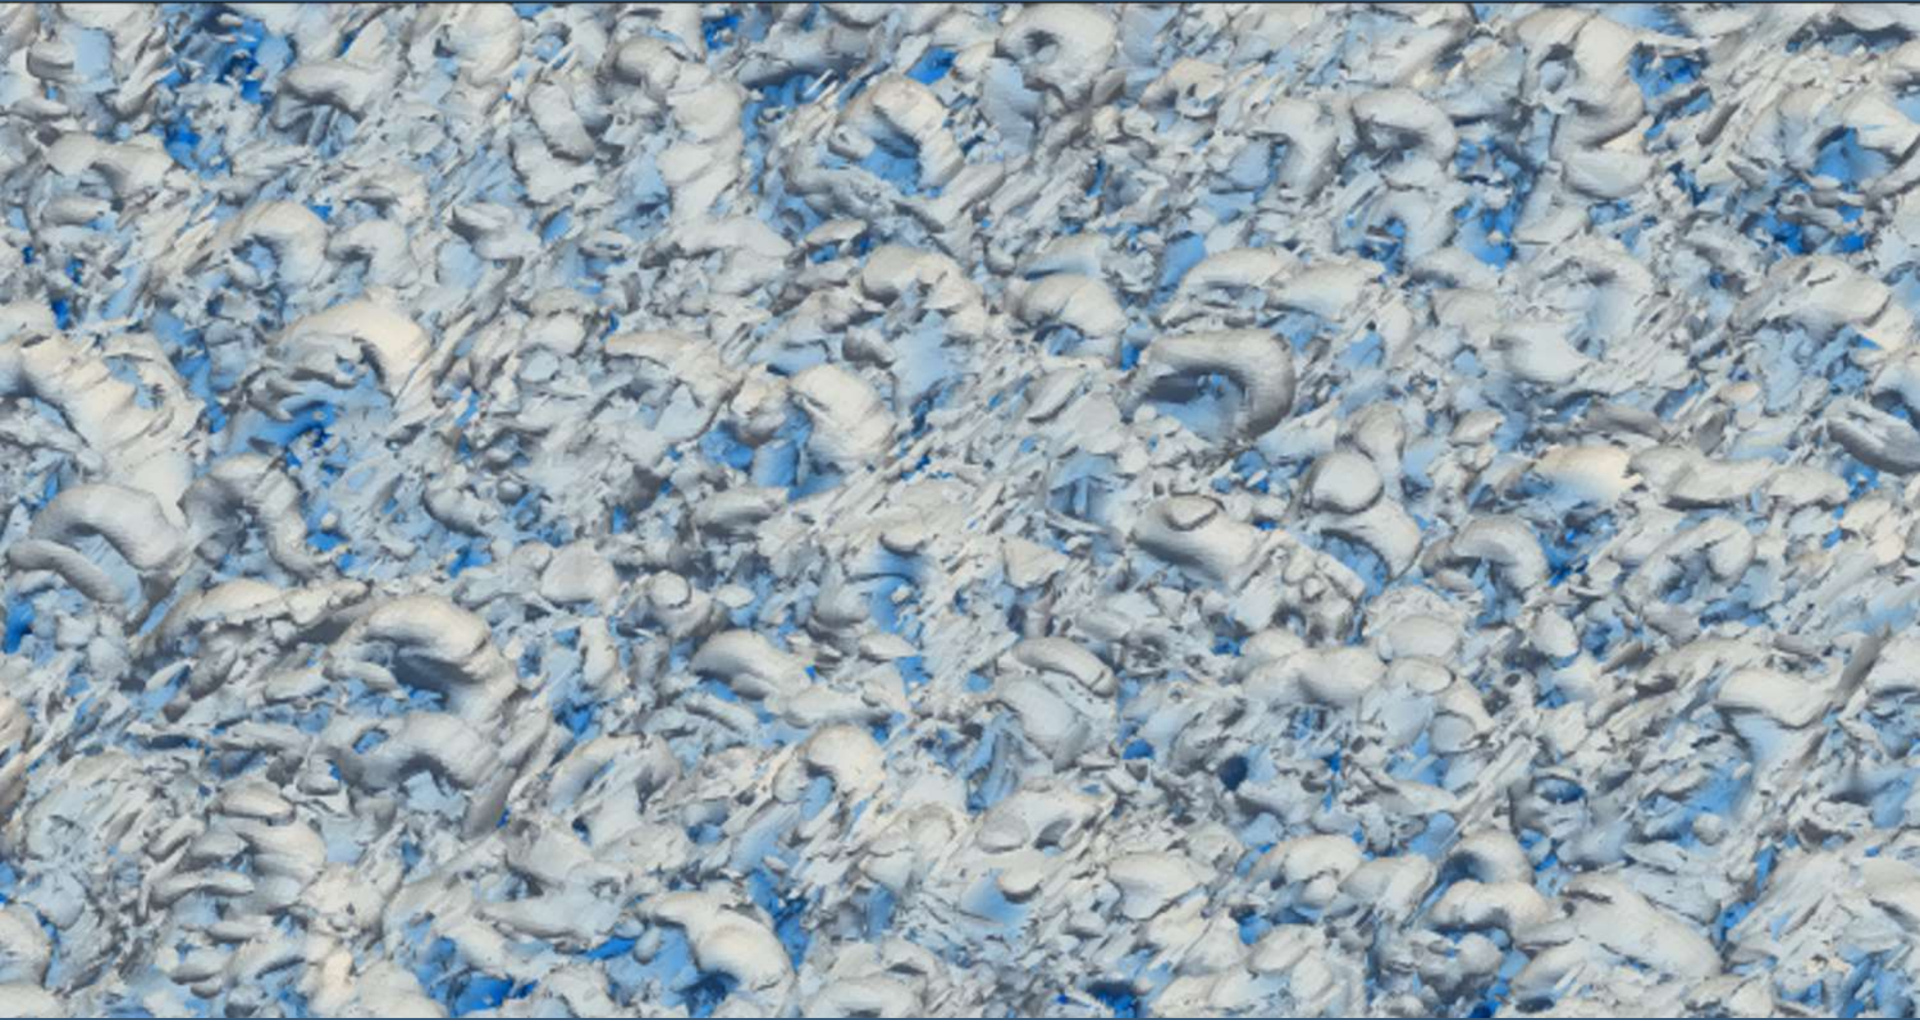

$$\int_E J(\vec{\xi}) U_t \phi \, d\vec{\xi} = \int_{-1}^{1}\int_{-1}^{1}\int_{-1}^{1} J(\vec{\xi}) \left(\frac{\partial}{\partial t} \sum_{r,s,t=0}^{N} \hat{U}_{rst}(t) \psi_{rst}^{N}(\vec{\xi}) \right) \psi_{ijk}^{N}(\vec{\xi}) d\vec{\xi^1} d\vec{\xi^2} d\vec{\xi^3}$$

$$= \sum_{\alpha,\beta,\gamma=0}^{N} J(\vec{\xi}_{\alpha\beta\gamma}) \left(\frac{\partial}{\partial t} \sum_{r,s,t=0}^{N} \hat{U}_{rst}(t) \underbrace{\ell_r^N(\xi_\alpha^1)}_{=\delta_{r\alpha}} \underbrace{\ell_s^N(\xi_\beta^2)}_{=\delta_{s\beta}} \underbrace{\ell_t^N(\xi_\gamma^3)}_{=\delta_{t\gamma}} \right) \psi_{ijk}^{N}(\vec{\xi}_{\alpha\beta\gamma}) \omega_\alpha \omega_\beta \omega_\gamma$$

$$= \sum_{\alpha,\beta,\gamma=0}^{N} J(\vec{\xi}_{\alpha\beta\gamma}) \frac{\partial}{\partial t} \hat{U}_{\alpha\beta\gamma}(t) \underbrace{\ell_i^N(\xi_\alpha^1)}_{=\delta_{i\alpha}} \underbrace{\ell_j^N(\xi_\beta^2)}_{=\delta_{j\beta}} \underbrace{\ell_k^N(\xi_\gamma^3)}_{=\delta_{k\gamma}} \omega_\alpha \omega_\beta \omega_\gamma$$

High Fidelity CFD: FLEXI

Multiscale Challenges to CFD

- Wide range of interacting scales: Non-linearity is the source of complexity and sensitivity
- For a smooth solution and a consistent scheme of order N, we have an error bound

$$\|u - u_h\|_{h,\Omega} \leq Ch^{N+1}$$

- Number of points per wavelength for a given error: Measure of information efficiency

DNS

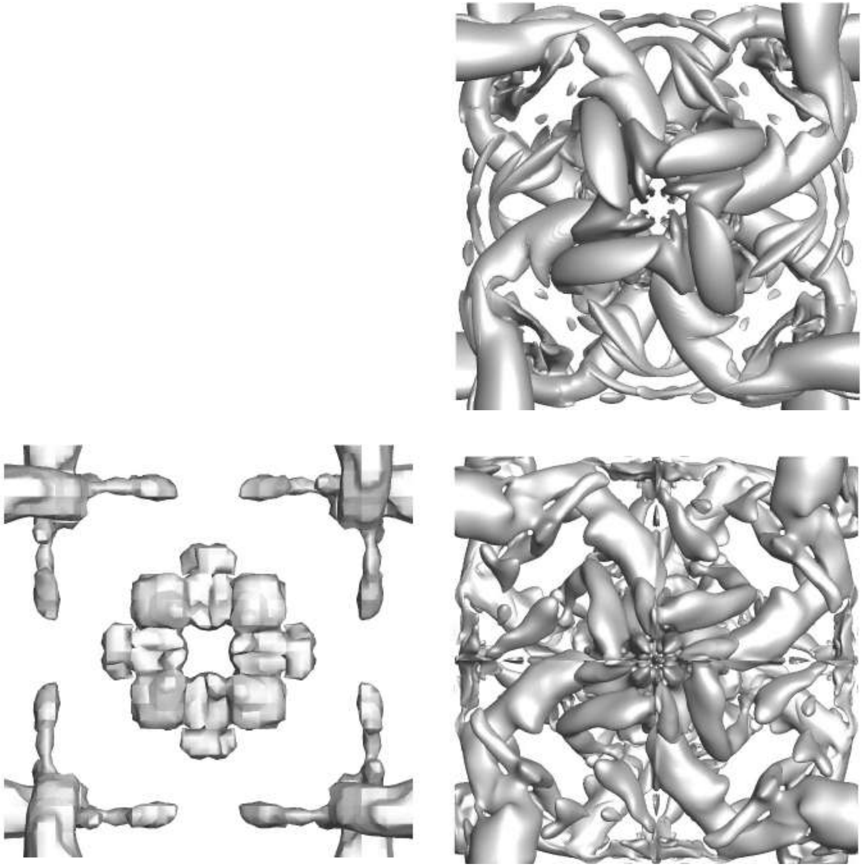

N=1, 64 DOF

N=15, 64 DOF

Discontinuous Galerkin Schemes

- Different Roads to High Order: Higher Derivatives, wider stencils: From local to global
- Discontinuous Galerkin schemes combine useful properties for multiscale problems
- Basic ideas:
  - High order polynomial basis with compact support
  - L_2 projection is optimal
  - Hybrid FE and FV scheme
- This gives flexibility, locality, conservation and stability (FV) and accuracy (FE)

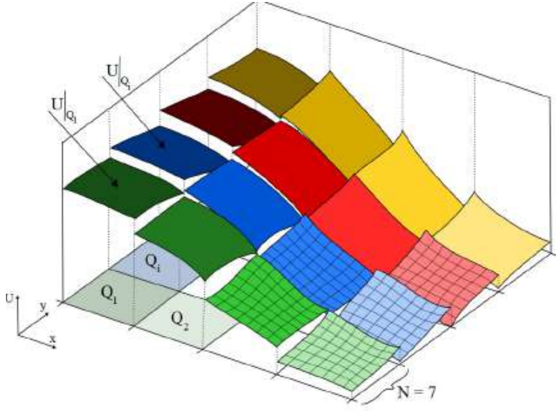

Simulation software: FLEXI
Introduction

[1]www.flexi-project.org

- High-order accurate open source solver[1] with excellent scaling behavior
- Discontinuous Galerkin spectral element method (DG-SEM)
- Written in modern Fortran and optimized for CPU based HPC systems
- Focus on DNS/LES of multiscale- and multi physics problems governed by the compressible Navier-Stokes equations
- Additional features
  - Lagrangian particle tracking (LES/DNS of particle laden flows)
  - Conservative sliding mesh interface for stator/rotor flow
  - Mesh deformation and mesh moving based on ALE formulation
  - hp-adaptivity
  - Intrusive and non-intrusive methods for uncertainty quantification
  - Management framework for optimal scheduling on HPC systems
  - A solver-in-the-loop framework for reinforcement learning

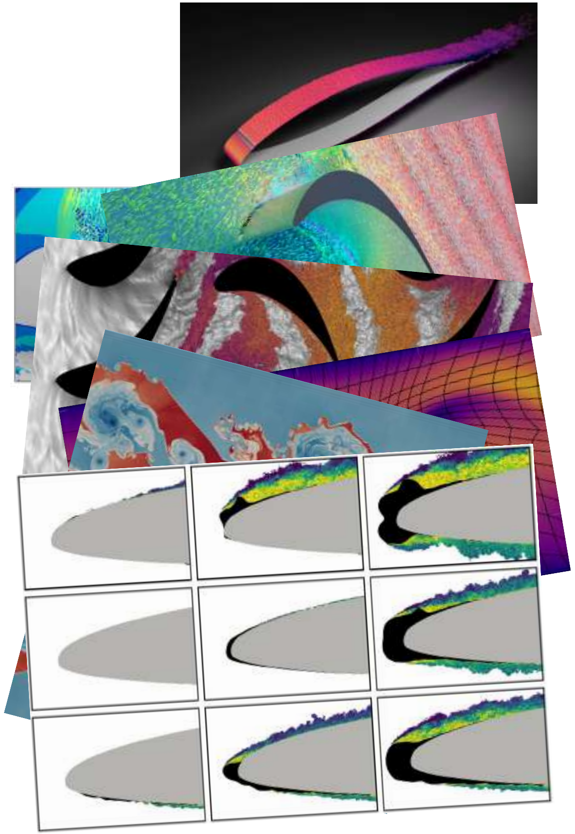

Simulation software: FLEXI

Discontinuous Galerkin Spectral Element Method (DG-SEM)

- DG-SEM:
  - Type of grid cells: Hexahedrons (curved elements, unstructured, hanging nodes)

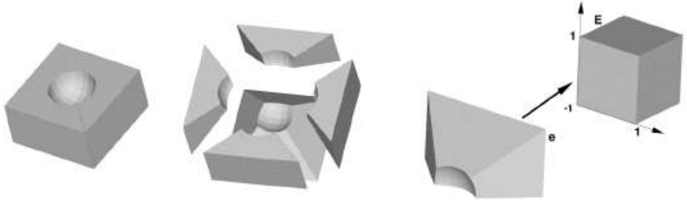

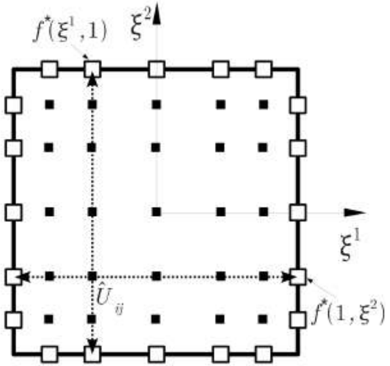

  - Set of basis functions: Tensor product, Lagrange polynomials at Gauß / Gauß-Lobatto points
    $U_h(\xi,t)\sum_{i,j=1}^{N}\hat{U}_{i,j}(t)\psi_i^N(\xi^1)\psi_j^N(\xi^2)$
  - Numerical integration: Collocation approach (SEM approach)
  - Time approximation: Explicit Runge-Kutta, IMEX
  - Numerical flux: Riemann solver, BR1/2
  - Stability: De-Aliasing, Split form (entropy / energy stable fluxes)
  - Shock-capturing: Finite volume sub-cells, h/p adaptivity

Simulation software: FLEXI
Shock-capturing

- Idea: Combination of DG and FV
  - DG in smooth parts of the flow and FV at shocks
- Oscillation, jump or ML indicators
- Troubled DG cell: $(N+1)^3$ equidistant FV sub-cells
- 2nd order TVD FV scheme on sub-grid
- Same number of degrees of freedom
- Concurrent calculation of FV and DG data
- Convex combination of FV and DG operator

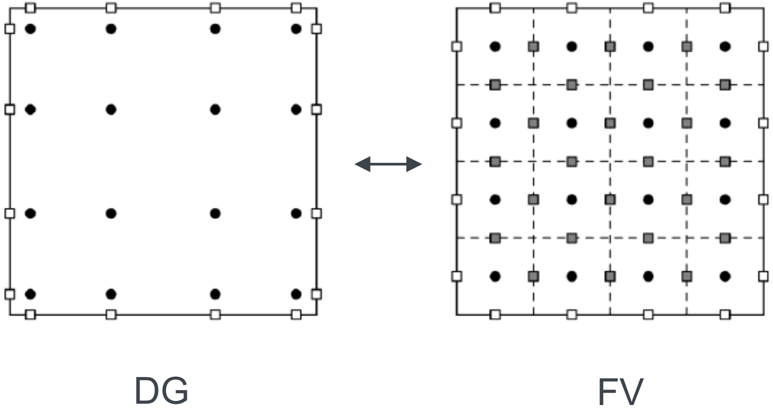

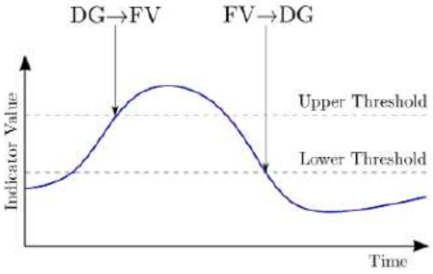

Simulation software: FLEXI

- Under development since 2010
- High order Discontinuous Galerkin SE framework with proven HPC capabilities
- Full framework: Preprocessor HOPR, FLEXI, Postprocessor POSTI and Paraview-Plugin, Blender Pipeline, HPC-UQ-Framework POUNCE, FLEXI-preCICE, FLEXI-OpenFOAM, FLEXI-TAU
- Reproduceability: Regression/unit testing, "compile from file", Development and Management via gitlab / on github


Computers & Mathematics with Applications
Volume 81, 1 January 2021, Pages 186-219

FLEXI: A high order discontinuous Galerkin framework for hyperbolic–parabolic conservation laws

Nico Krais [a], Andrea Beck [a], Thomas Bolemann [a], Hannes Frank [a], David Flad [a], Gregor Gassner [b], Florian Hindenlang [c], Malte Hoffmann [a], Thomas Kuhn [a], Matthias Sonntag [a], Claus-Dieter Munz [a]

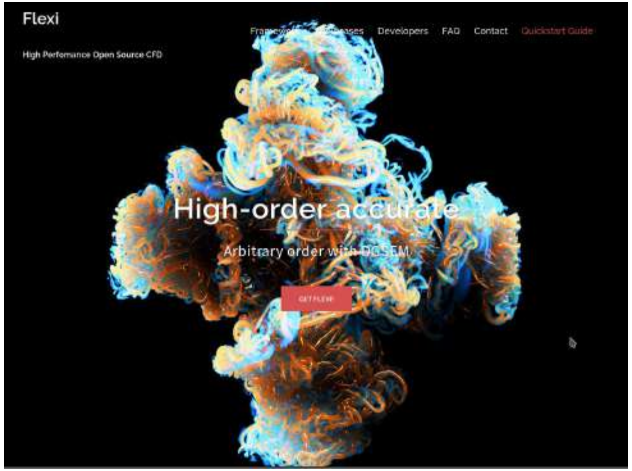

www.hopr-project.org
www.flexi-project.org
https://github.com/flexi-framework/flexi

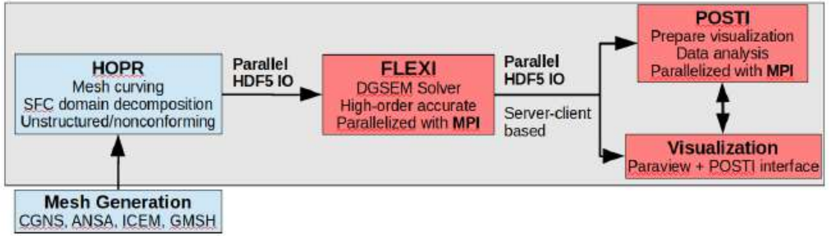

Simulation software: FLEXI

Adaptivity and error control – hp-refinement

- h-refinement[1]
  - Mortar interfaces with hanging nodes for optimal grids
  - Local grid-adaptation and grid-adaptation algorithm required
  - Up to 50% lower cost compared to human-generated grids
  - Residual estimation procedure
- p-refinement[2]
  - Local p-adaptivity in smooth solution regions
  - Dynamic Load Balancing (under development)

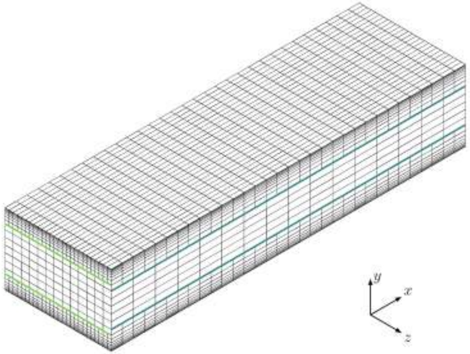

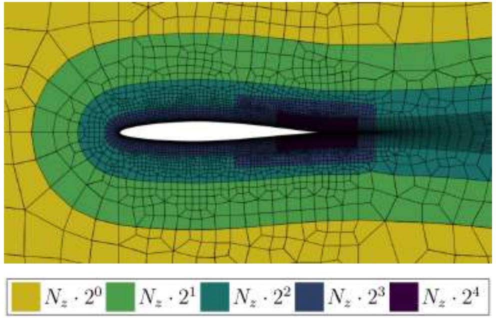

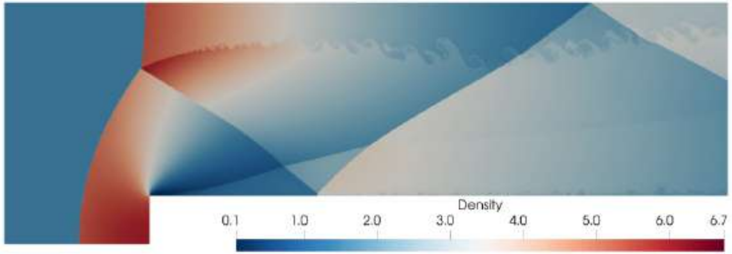

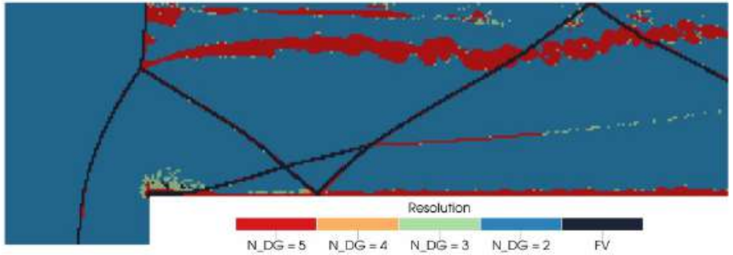

[1]Blind et al., Grid-Adaptation for Wall-Modeled Large Eddy Simulation Using Unstructured High-Order Methods. arXiv preprint
[2]Mossier et al., A p-Adaptive Discontinuous Galerkin Method with hp-Shock Capturing. Journal of Scientific Computing, 2022

Simulation software: FLEXI

hp-adaptive Multiphase Branch

- Sharp-Interface simulation of 3D water droplet – shock interaction at Ma = 2.4 and We = 100
- About 140 DOF / droplet diameter
- Hp-adaptive DGSEM / FV scheme with DLB
- 7200 CPUh on HAWK

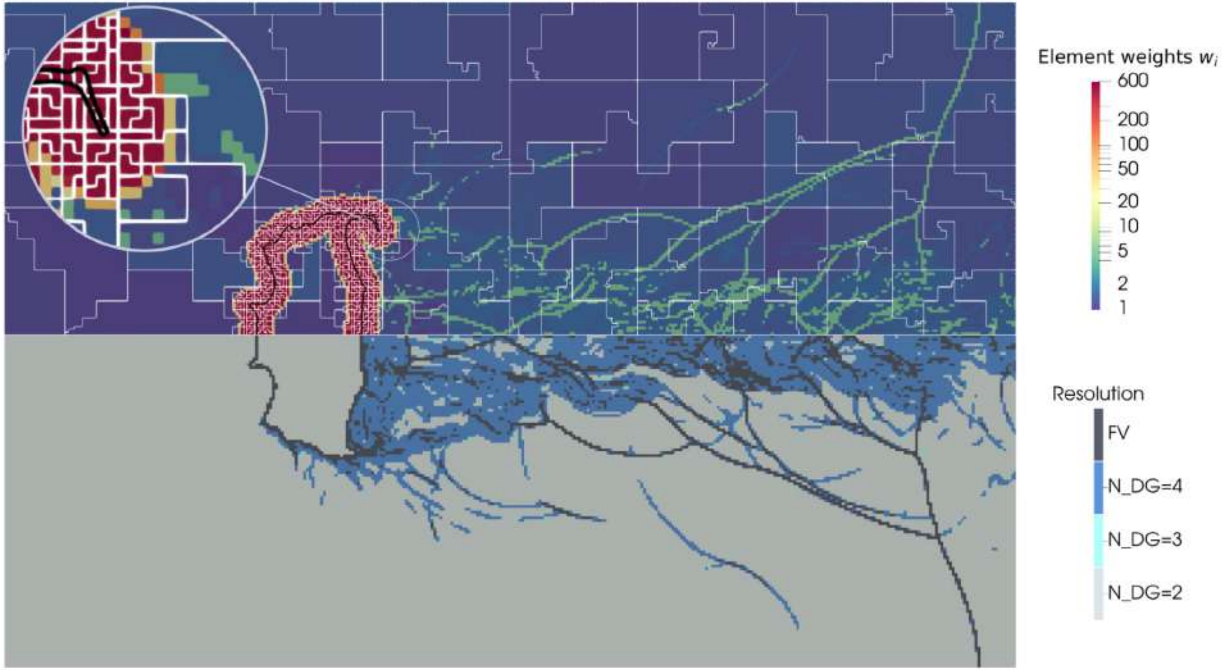

Mossier et al., An Efficient hp-Adaptive Strategy for a Level-Set Ghost Fluid Method, arxiv, 2023
Jöns et al., Riemann solvers for phase transition in a compressible sharp-interface method, JCP, 2023
Mossier et al., A p-Adaptive Discontinuous Galerkin Method with hp-Shock Capturing. Journal of Scientific Computing, 2022

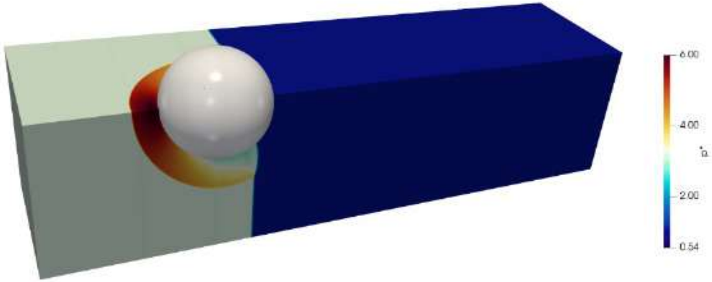

(a) $t^* = 0.8$

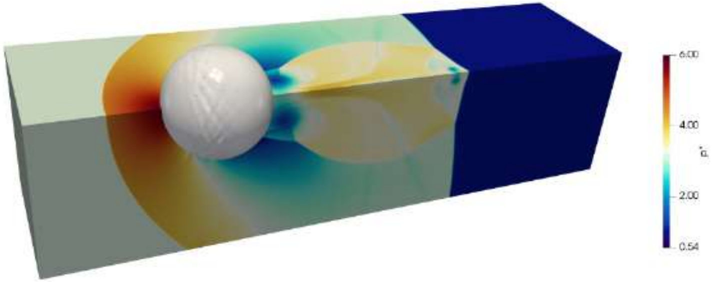

(b) $t^* = 2.4$

(c) $t^* = 4.7$

(d) $t^* = 7.0$

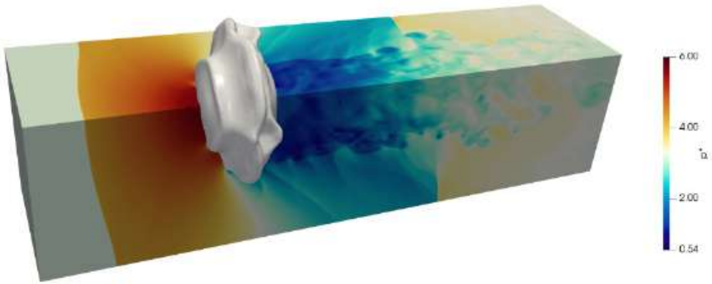

(e) $t^* = 9.4$

(f) $t^* = 11.8$

HPC-CFD: FLEXI

- Parallelization with the MPI paradigm
  - Domain decomposition using a space-filling curve
  - Communication latency hiding by local work
  - DGSEM operator requires only com. of surface fluxes
  - Small communication stencil
- Parallel I/O
- Small memory footprint
- Efficient cache usage at about 2000 DOF/core
- Excellent scaling up to over 10^6 cores
- Transition in HPC architecture requires support of accelerators (e.g. GPUs)

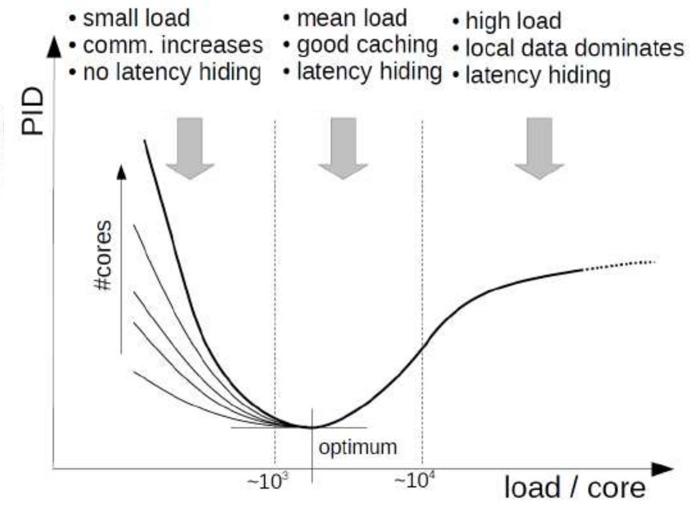

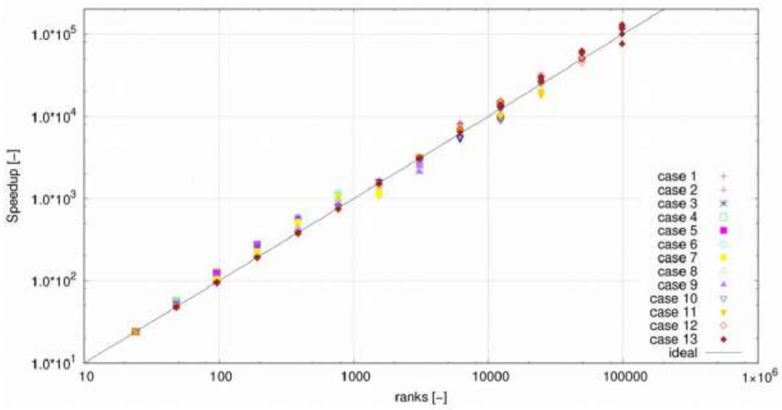

Simulation software: FLEXI
Postprocessing & visualization

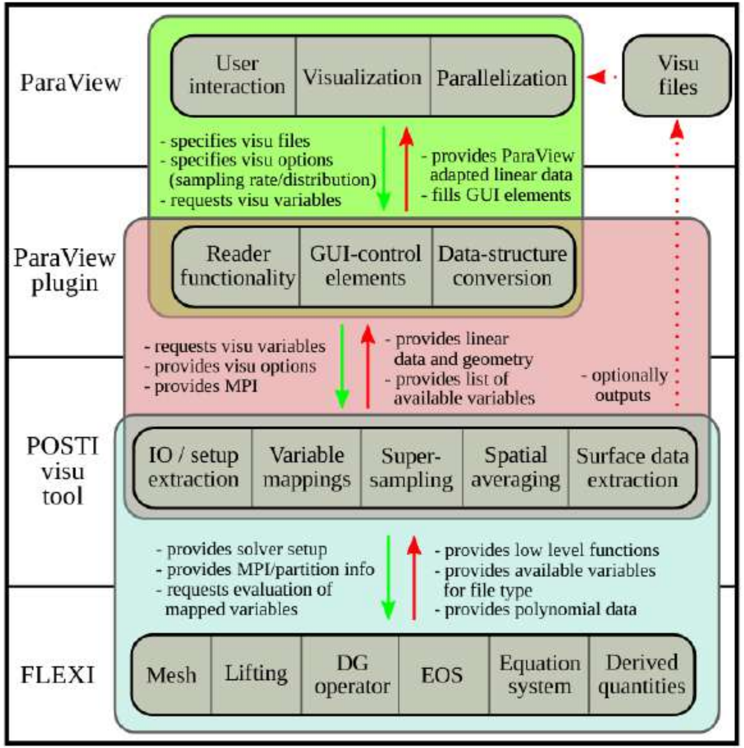

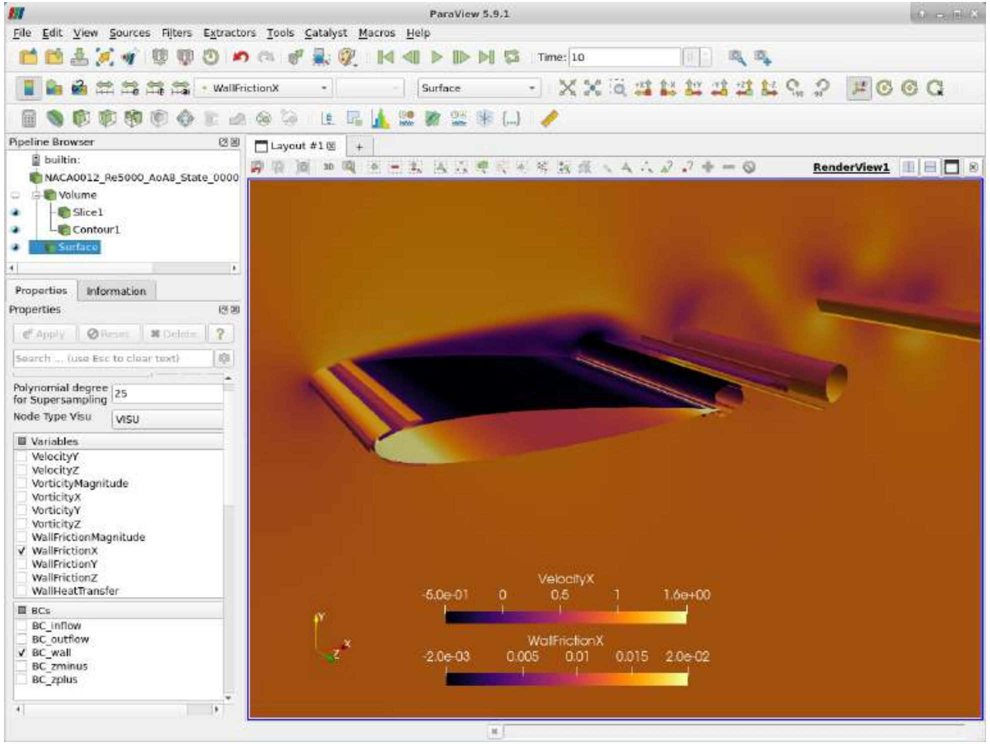

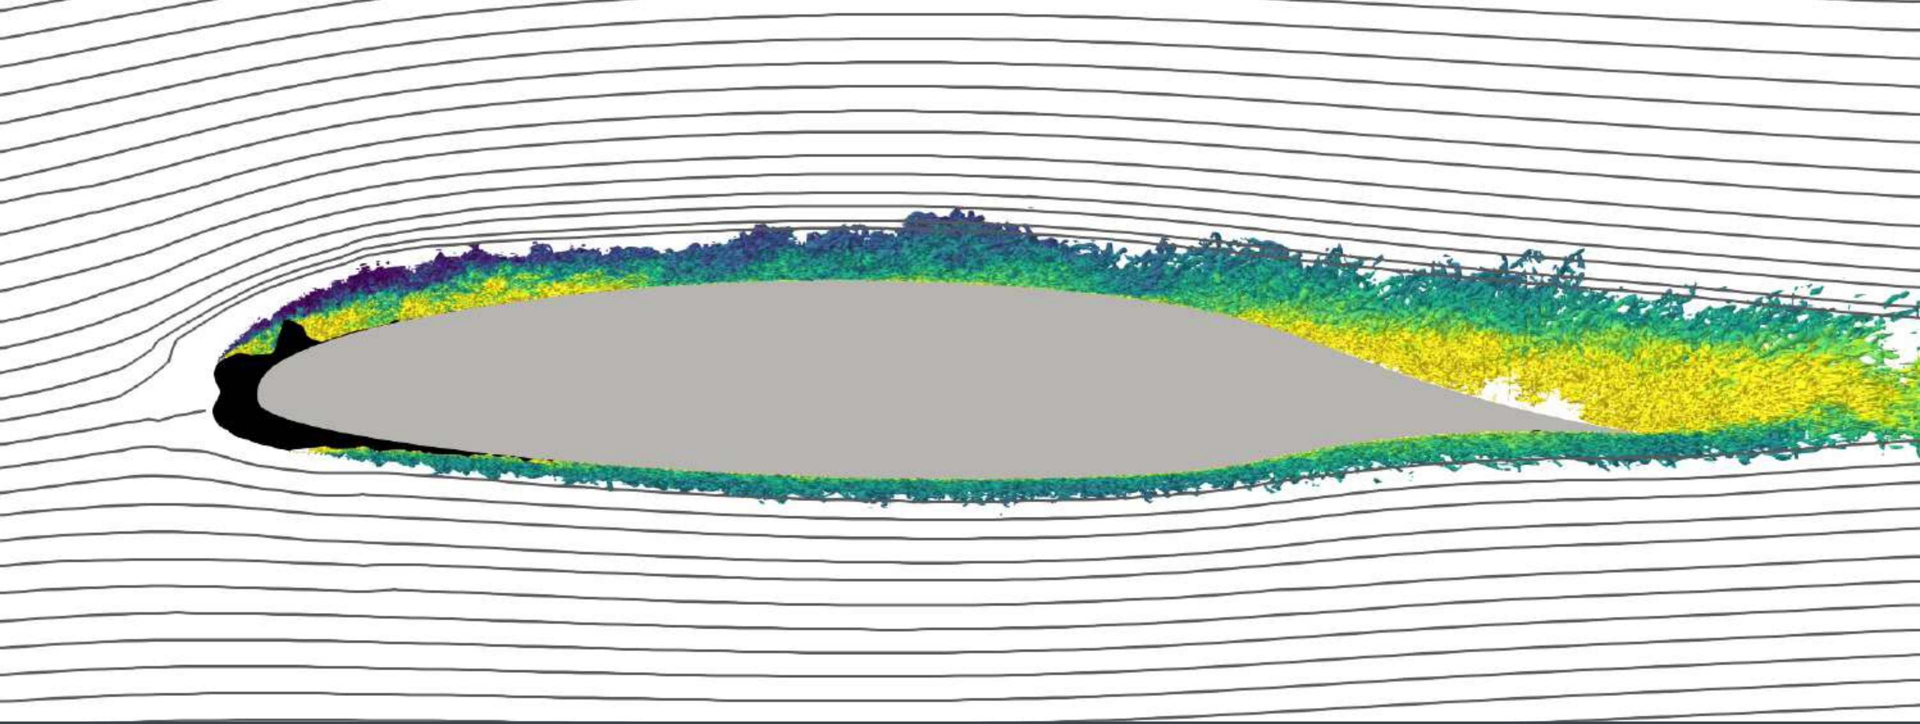

UQ

A UQ-HPC Framework for Icing

Input data
Wind tunnel ice geometries

Step 1
Set of continuous random variables

Goal
prediction of airfoil lift & drag

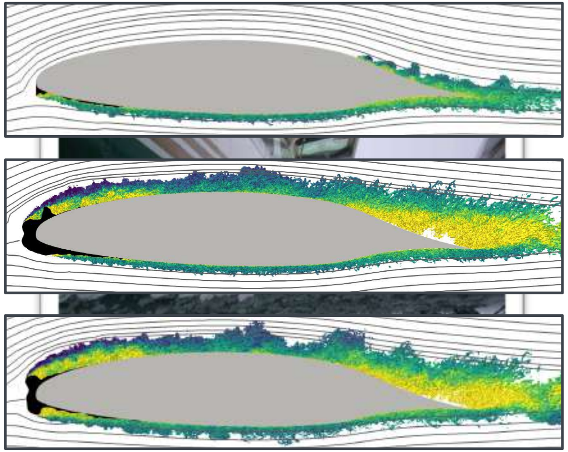

Principal component analysis

Uncertainty propagation: NISP, MLMC, MFMC

Model management (POUNCE)

Simulation (FLEXI)

Meshing (LaPlus)

Intrusive and Non-Intrusive UQ

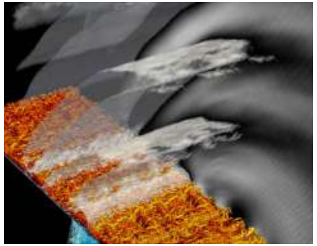

Non-Intrusive / NISP, MLMC

SIAM J. Sci. Comput., 42(4), B1067–B1091. (25 pages)

Computational Methods in Science and Engineering

hp-Multilevel Monte Carlo Methods for Uncertainty Quantification of Compressible Navier--Stokes Equations

Andrea Beck, Jakob Dürrwächter, Thomas Kuhn, Fabian Meyer, Claus-Dieter Munz, and Christian Rohde

Related Databases

Web of Science

Journal of Theoretical and Computational Acoustics | Vol. 27, No. 01, 1850044 (2019) No Access

Uncertainty Quantification for Direct Aeroacoustic Simulations of Cavity Flows

Thomas Kuhn, Jakob Dürrwächter, Fabian Meyer, Andrea Beck, Christian Rohde and Claus-Dieter Munz

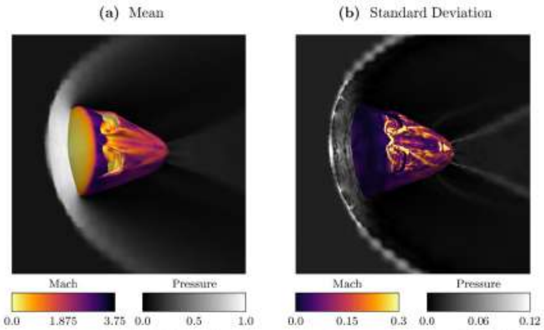

Fig. 6. Spacecraft: mean and standard deviation; Mach number on spacecraft surface, pressure in slice through flow field. Example 5.4.

Intrusive / SG

Contents lists available at ScienceDirect

Computers and Fluids

journal homepage: www.elsevier.com/locate/compfluid

A high-order stochastic Galerkin code for the compressible Euler and Navier-Stokes equations

Jakob Dürrwächter[a,*], Fabian Meyer[b], Thomas Kuhn[a], Andrea Beck[a], Claus-Dieter Munz[a], Christian Rohde[b]

[a] Institute of Aerodynamics and Gas Dynamics, University of Stuttgart, Pfaffenwaldring 21, 70569 Stuttgart, Germany
[b] Institute of Applied Analysis and Numerical Simulation, University of Stuttgart, Pfaffenwaldring 57, 70569 Stuttgart, Germany

Accelerating Monte Carlo

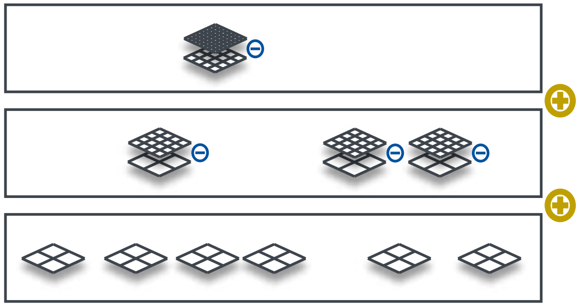

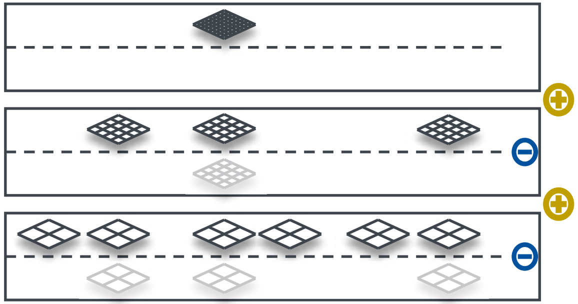

Multilevel MC:

- Cheap models: coarse mesh resolutions
- Driving factor: Mesh convergence

Multifidelity MC:

- Cheap models: arbitrary correlated models
- Driving factor: Model covariance

- Wall-resolved LES of the clean geometry: an HPC problem
- Optimal number of samples is solution dependent!
- HPC challenge: Extreme simulation cost variations
- We need millions of computations

POUNCE (Propagation Of Uncertainties)

PoUnce: A framework for automatized uncertainty quantification simulations on high-performance clusters

Jakob Duerrwaechter[1¶], Thomas Kuhn[1], Fabian Meyer[2], Andrea Beck[1,3], and Claus-Dieter Munz[1]

1 Institute of Aerodynamics and Gas Dynamics, University of Stuttgart, Germany 2 Institute of Applied Analysis and Numerical Simulation, University of Stuttgart, Germany 3 The Laboratory of Fluid Dynamics and Technical Flows, Otto von Guericke University Magdeburg, Germany ¶ Corresponding author

DOI: 10.21105/joss.04683

- Fully automated model management framework for UQ on HPC systems (OpenSource)
  - runs multiple iterations with different models: Optimal stacking according to length and size
  - Common MPI Communicator & common IO
  - each model includes pre- and post-processing
  - focus on efficient machine use: different machines via ssh
- Object-oriented Python, modular Design
- NISP, MLMC, MFMC – Hazel Hen, HAWK, local cluster

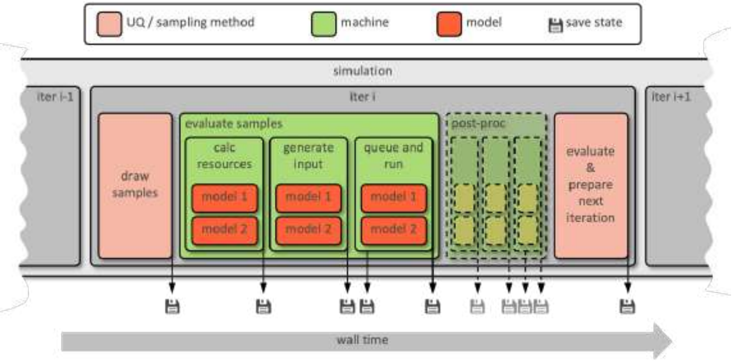

Structured Grid Algorithm

Laplacian Smoothing (Average by Neighbours)

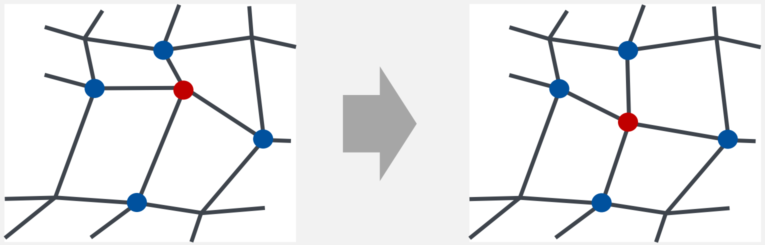

Additional Terms for Grid Quality

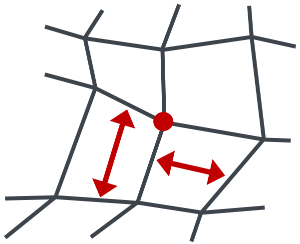

Minimum spacing

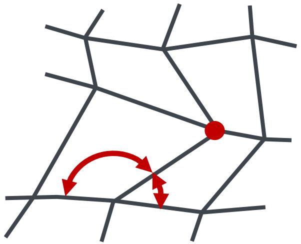

Maximum skewness

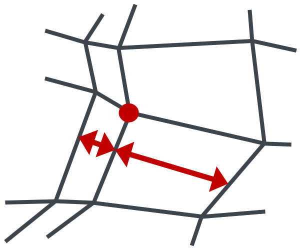

Max stretching ratio

Grid: Results

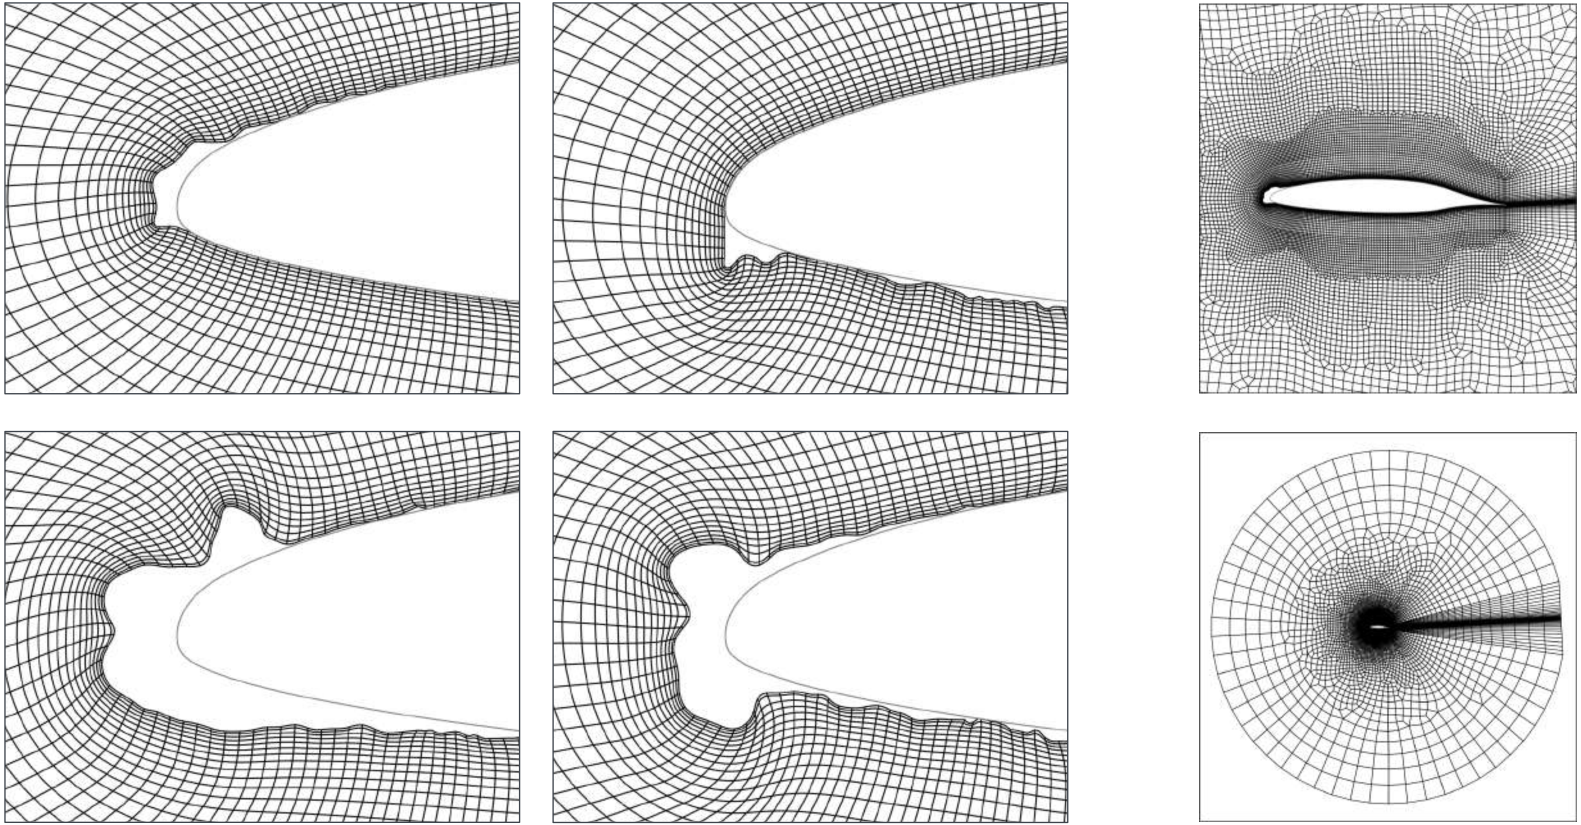

Simulation Setup

Baseline Simulation

- Laminar flow airfoil
- 3D wall-resolved LES
- Low Reynolds number
- Minimal required resolution

Parameters	
Re	500.000
AoA	3°
$\Delta y^+_{\mathrm{wall}}$	≈ 4
n_{elems}	190.000
N_{DG}	5
t_{end}	$10\ d/u_\infty$
computation time	$\approx$ 40.000 CPUh

Uncertainty Quantification

- **NISP**
  - Two PCA modes, $N_{stoch} = 4$
  - 25 LES
  - 0.5 M CPUh
- **MLMC**
  - 3D $N_{DG} = 5, 3, 1$
  - 16 // 63 // 406 simulations
  - 1 M CPUh
- **MFMC**
  - 3D $N_{DG} = 5$ // 2D $N_{DG} = 5, 3, 1$
  - 23 // 174 // 587 // 3676 simul.
  - 1 M CPUh

Iced Airfoil Performance

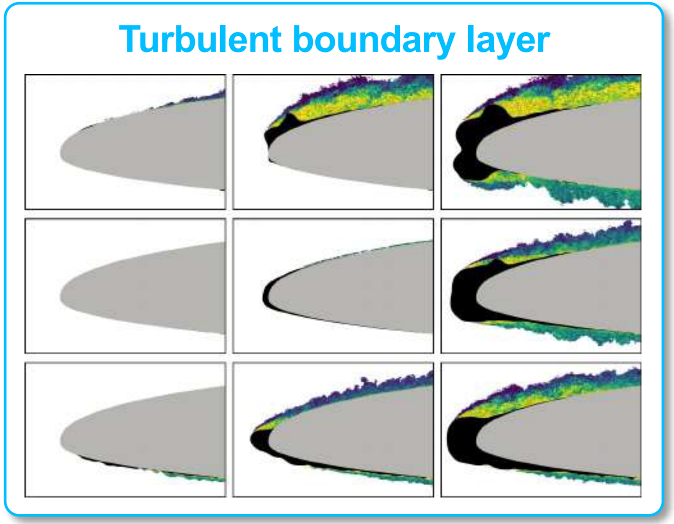

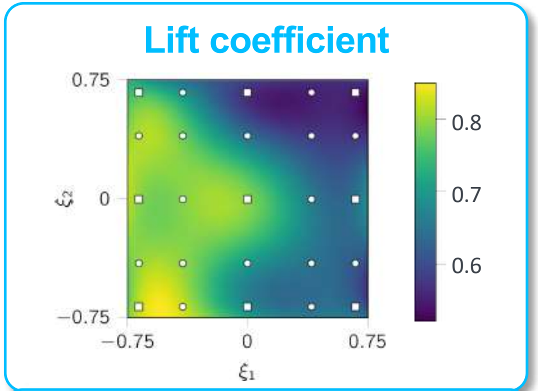

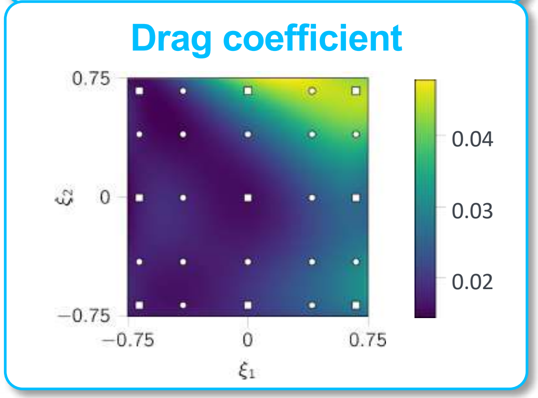

FLEXICE: towards Multi-X/UQ

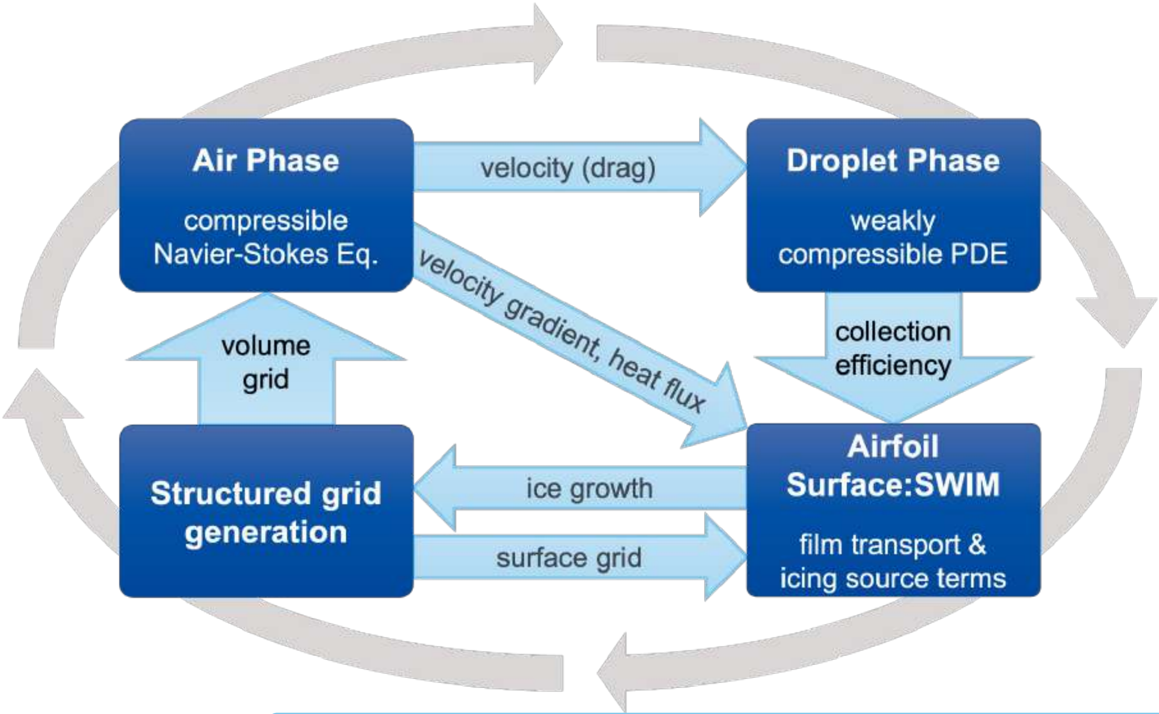

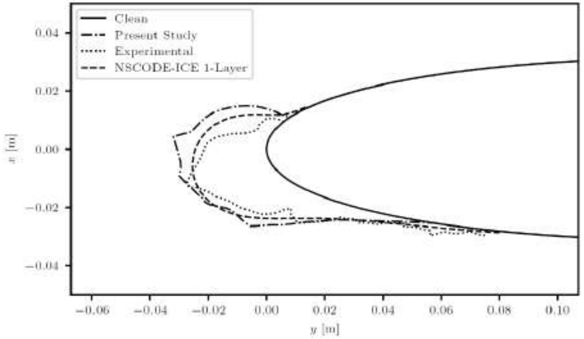

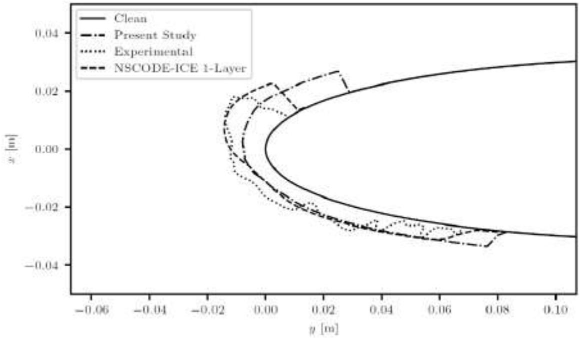

Closing the Scale Gap (a bit) through UQ

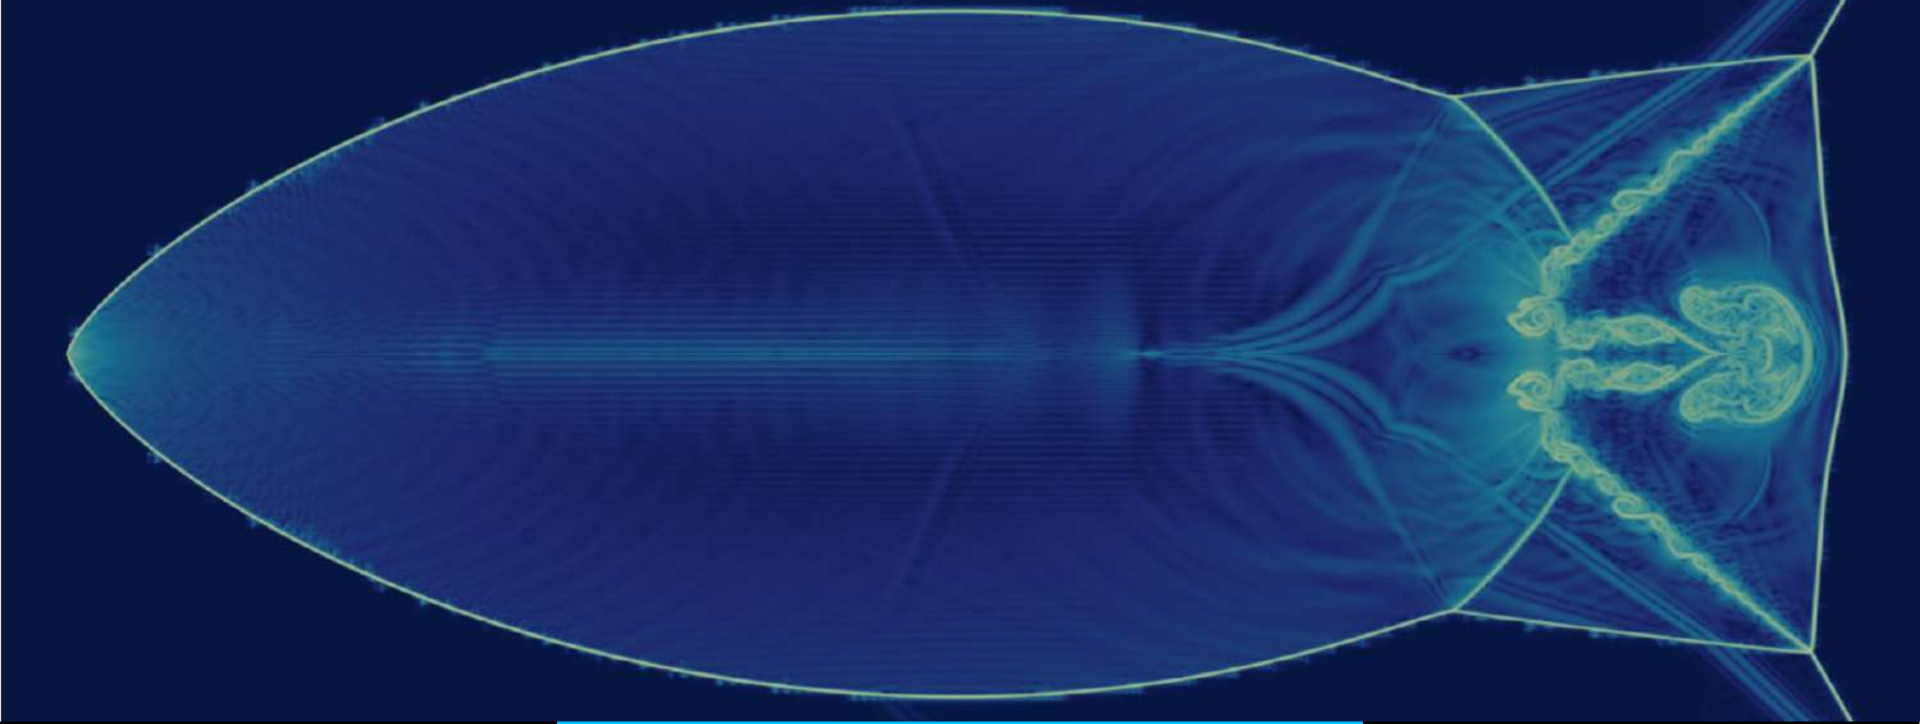

ML: SL + FLEXI

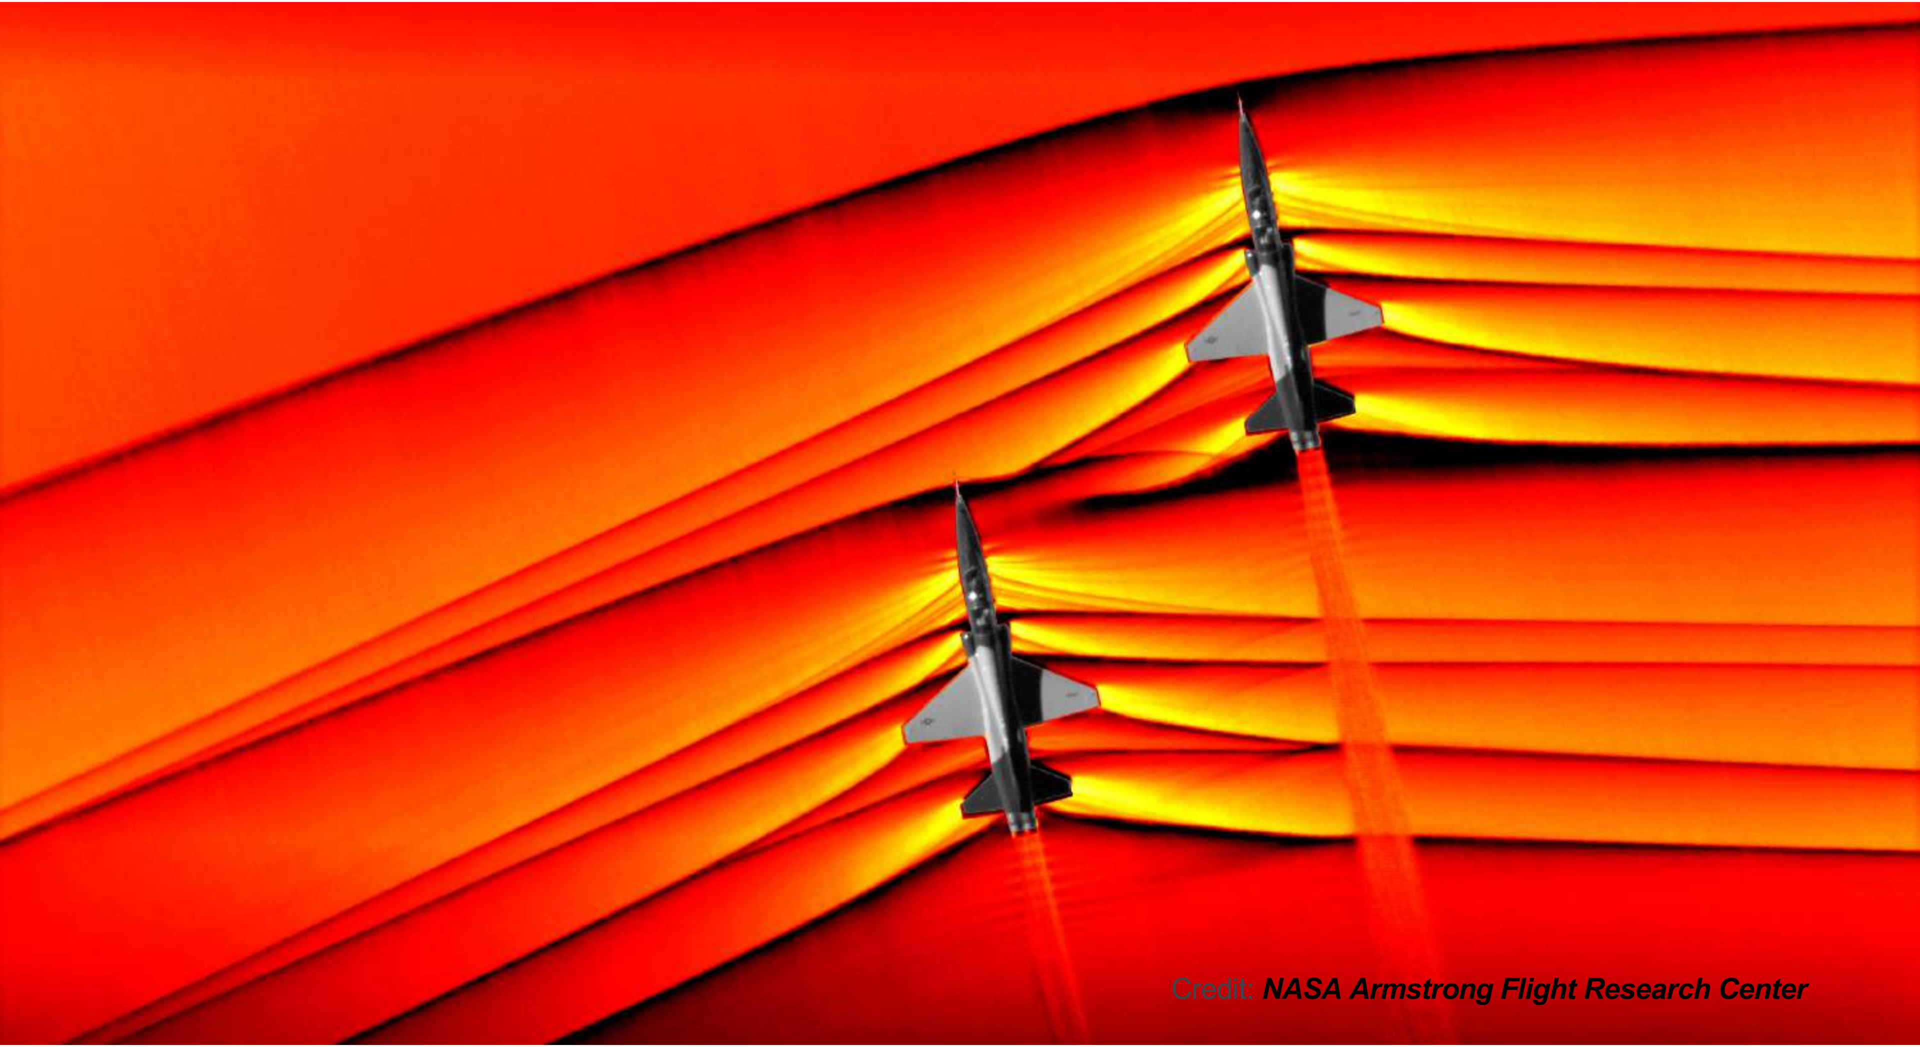

Credit: ***NASA Armstrong Flight Research Center***

Data-informed Shock Capturing for High Order Methods

- Stable numerical approximation through Shock Capturing: improves stability, decreases accuracy: use sparingly!
- Detecting the occurrence of shocks: non-trivial, empiricism, many parameters
- For HO methods: Just detecting a “troubled cell” is not good enough: We need localization on the element subscale

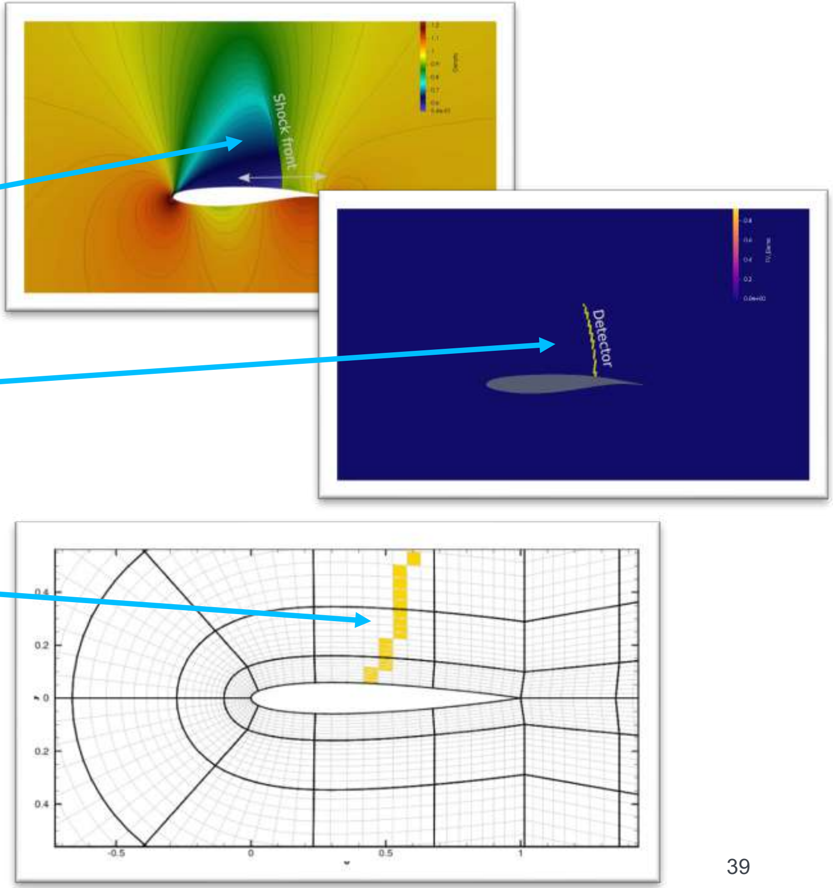

Data-informed Shock Capturing for High Order Methods

- Supervised learning of a classifier from analytical smooth and non-smooth data
- Convolutional neural networks for spatial correlations
- About 100,000 samples per class, classes balanced
- Train/validate/test split
- Cross-entropy loss, ADAM optimizer, minibatch GD
- Tensorflow 1.6, coupled to FLEXI
- Resulting F1 score > 0.96

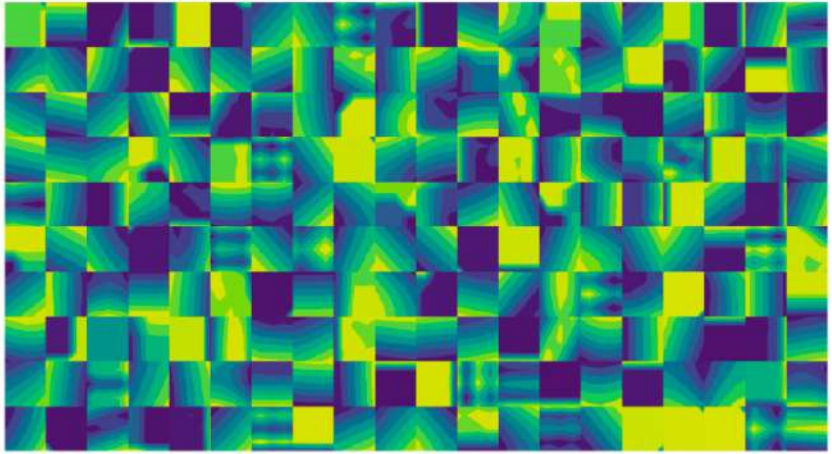

Data-informed Shock Capturing for High Order Methods

- Multiscale-CNNs for edge detection
- Consistent subscale localization, contiguous shock fronts: On different grids, for different problems (same model)
- On ”bad but practical” grids: stable & accurate

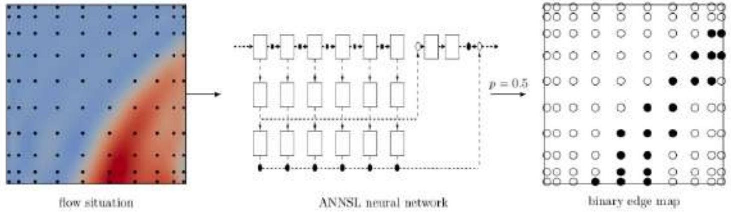

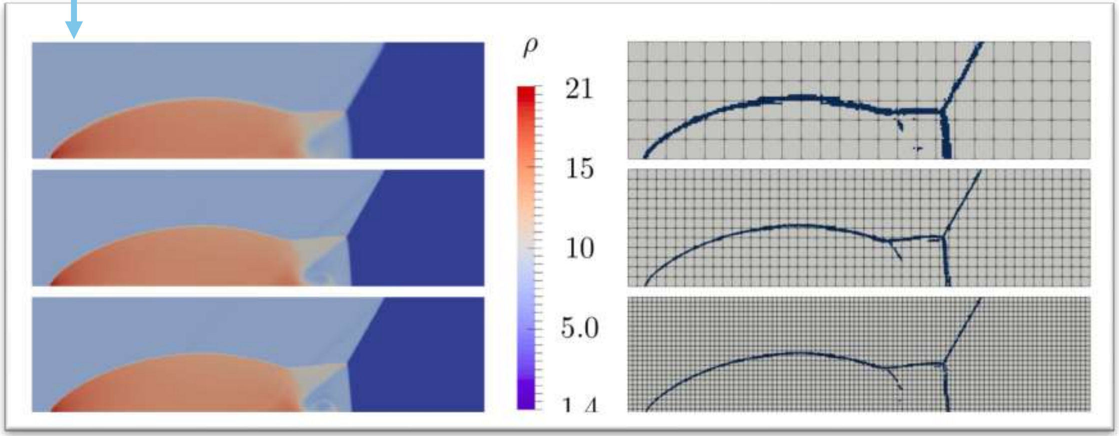

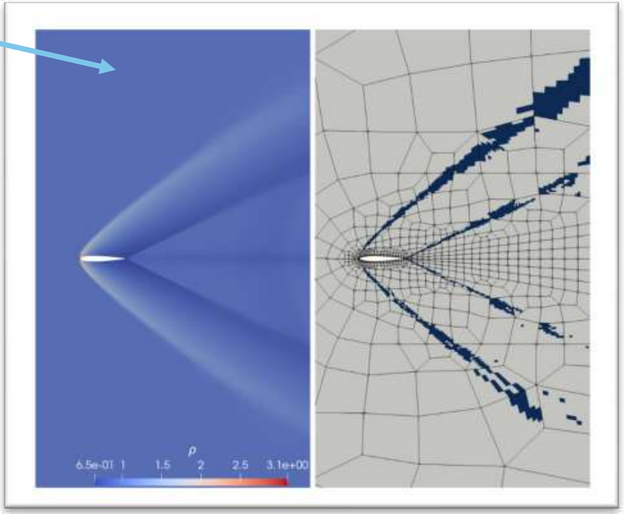

Localized AV for shock capturing

- Use prediction of ”shocked nodes” (binary edge map) and smooth with high order Radial Basis functions (RBF) interpolation

$$\mu_a(\vec{x}) = \mu_{a\text{scale}} \sum_{i=1}^{n_s} \alpha_i \phi_r \left(\|\vec{x} - \vec{x}_{s_i}\|_2 \right) ,$$

- With ϕ being the chosen RBF, α_i as interpolation coefficients and the spatial position of the inner-element solution nodes labeled “shocks” as $\vec{x}_{s_i}$

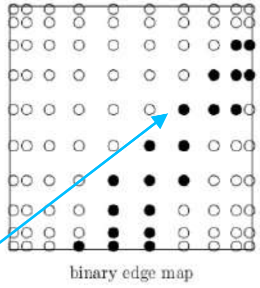

- This leads to a global, but weakly coupled Vandermonde matrix.
- Solve linear system with PETSc or approximate as local problems (compact support)

Data-informed Shock Capturing for High Order Methods

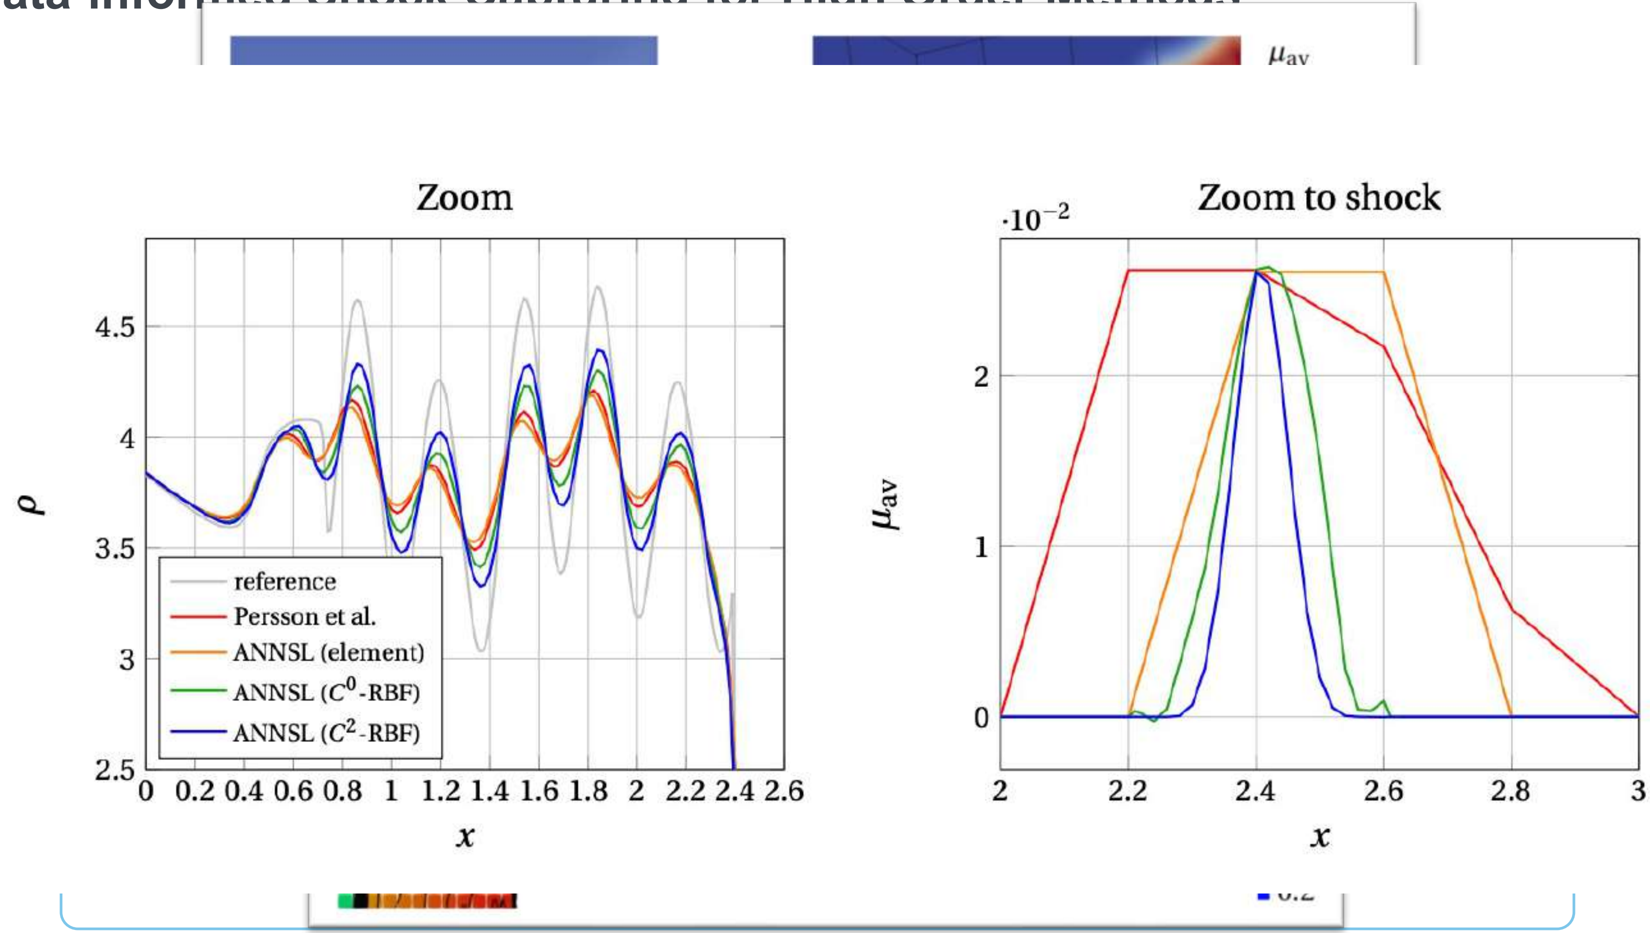

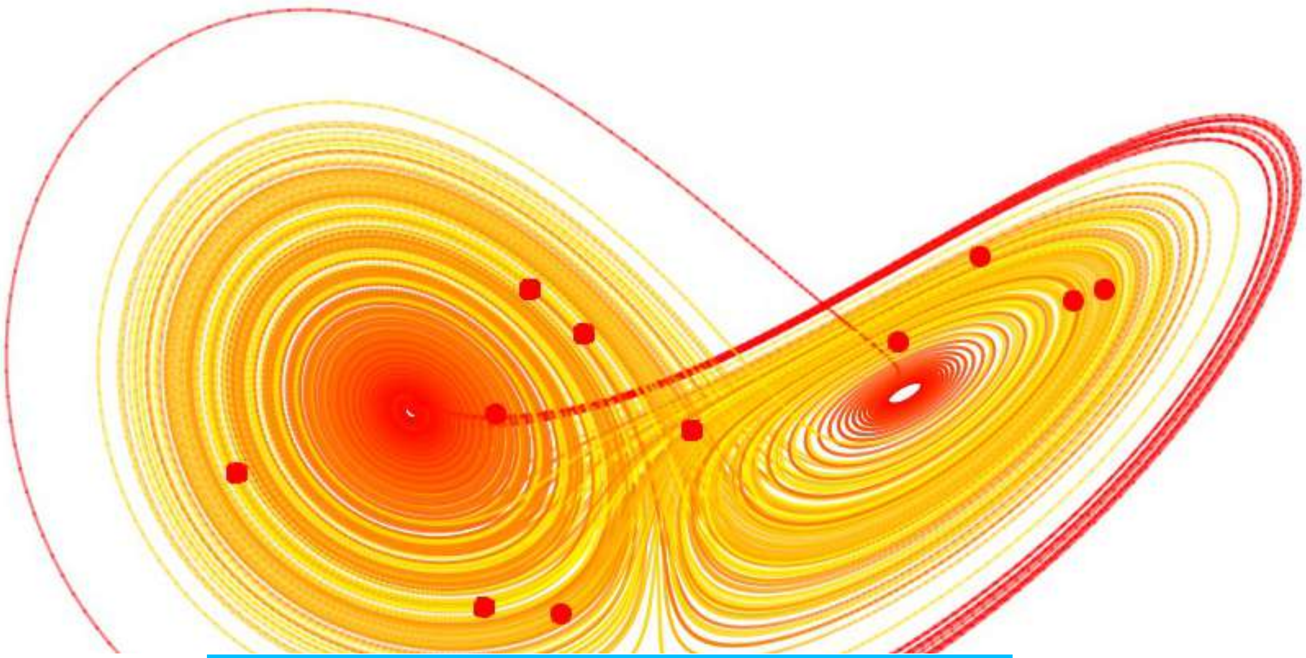

ML: RL + FLEXI

The Scale Gap in Turbulence

$$\underbrace{R\left(U_l, x_l, t_l\right) = 0}_{\text{Governing Equations at level l}}$$

$$\underbrace{R\left(U_L, x_L, t_L\right)}_{\text{Governing Equations at level L}} + \overbrace{M\left(U_l, U_L\right)}^{\text{Influence of level l on L}} = 0$$

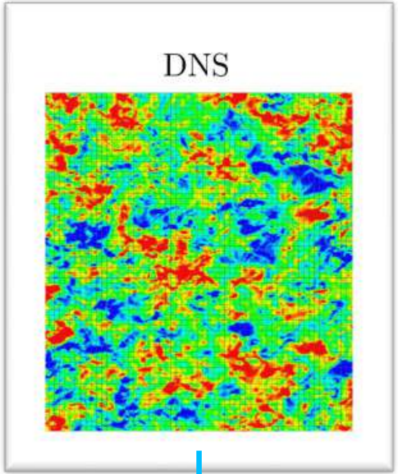

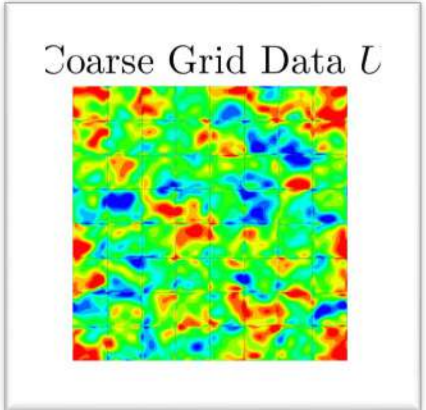

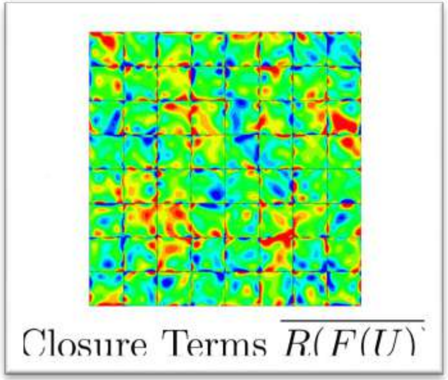

The Reynolds stresses:
The footprint of the fine scale turbulence

Data-driven models

Large Eddy Simulation - Definition

$$u = \overline{u}(x,t) + u(x,t)' \Rightarrow \overline{u}(x,t) = \int_{\Delta r} u(r,t)G(x-r)dr$$

- Separate Δr and h
- Explicit filtering: $h \to 0$
- Discretization scheme not relevant
- Caveat: homogeneity and isotropy, boundary conditions, realizability, commutation…

0.5% on google scholar

- Joined Δr and h
- Implicit filtering: $h \not\to 0$
- Discretization scheme defines filter kernel
- $\Leftarrow$ plus discretization parameters and errors

99.5% on google scholar, 100% for “industrial” LES

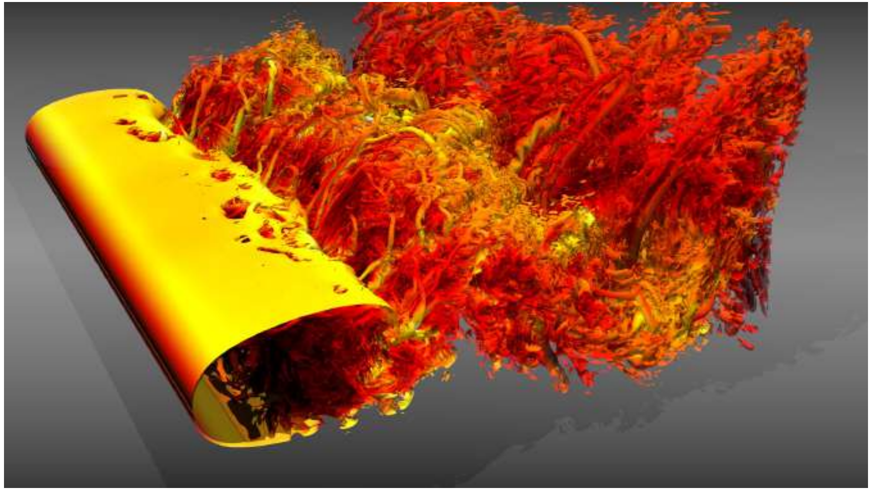

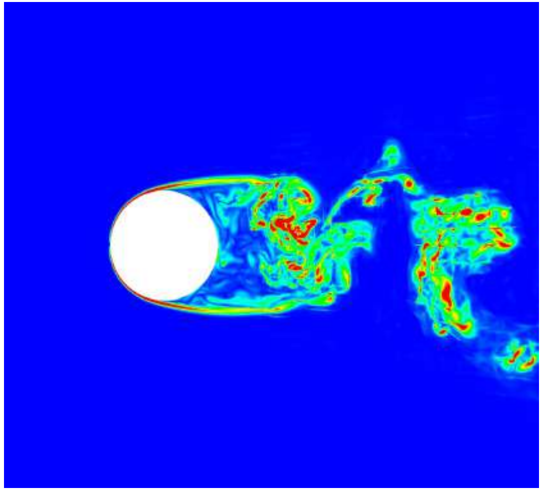

"Same equations, domains, models…"

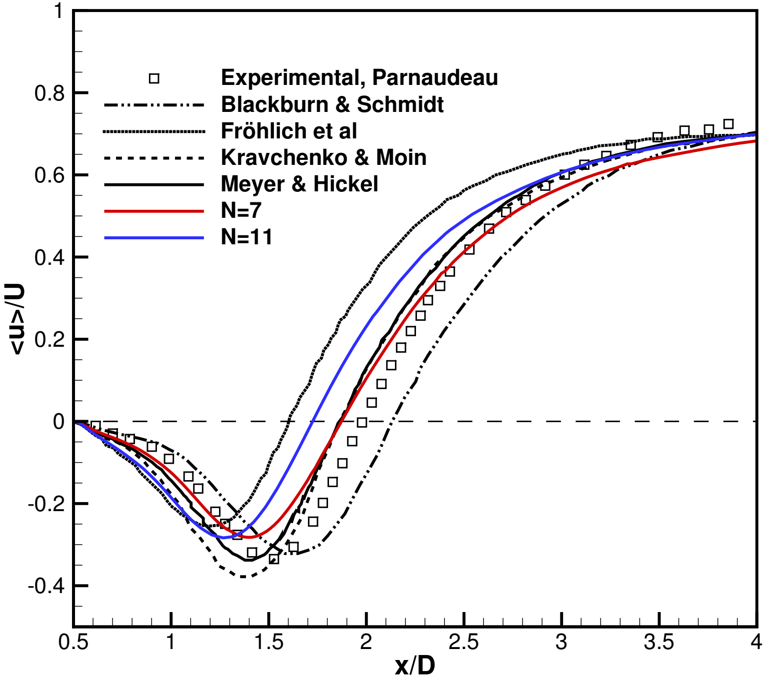

The LES dilemma

$$\overline{u(x)} = \int_{-\Delta x}^{\Delta x} u(r)G(x-r)dr$$

- Commutation introduces errors: IFF isotropy, homogeneity, linearity of the filter are given and grid and filter are completely separated, then the closure term is:

$$\overline{u}\,\overline{u} - \overline{uu}$$

- Otherwise, it is

$$D\left(\tilde{u}\,\tilde{u}\right) - \frac{\partial}{\partial x}\overline{uu}$$

Which should be modelled / closed?

Inhomogenous, non-linear

The filter shape is not known a priori

What is the coarse scale LES solution?

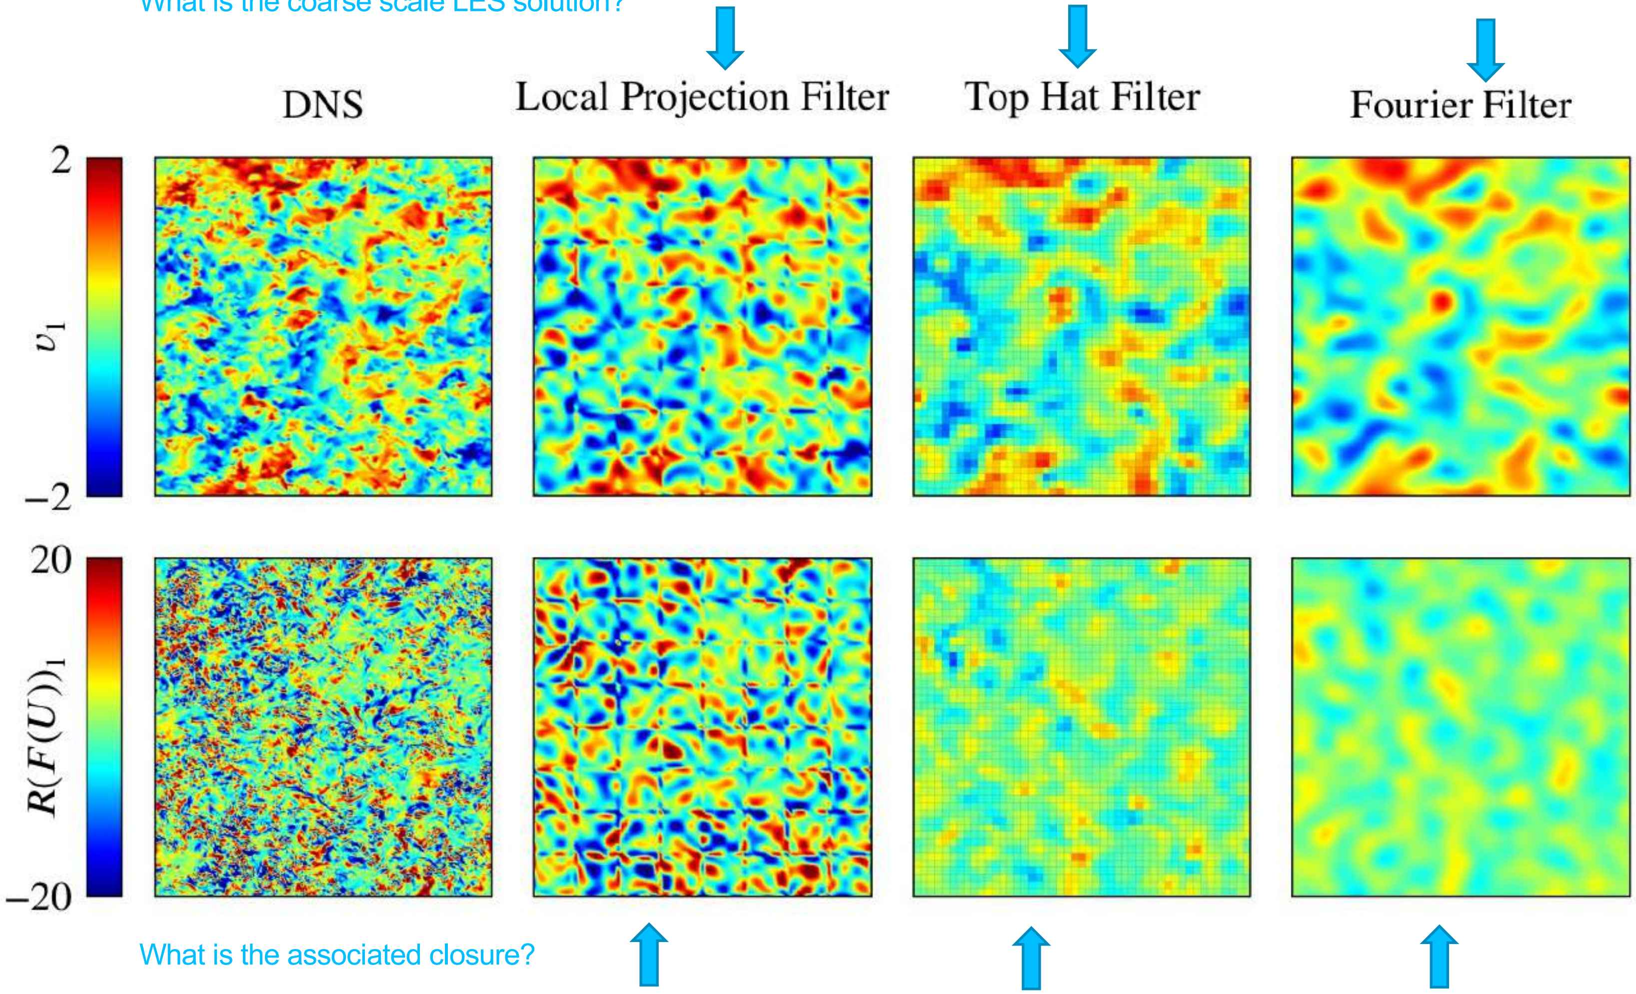

What is the associated closure?

LES Closure Terms: Supervised learning

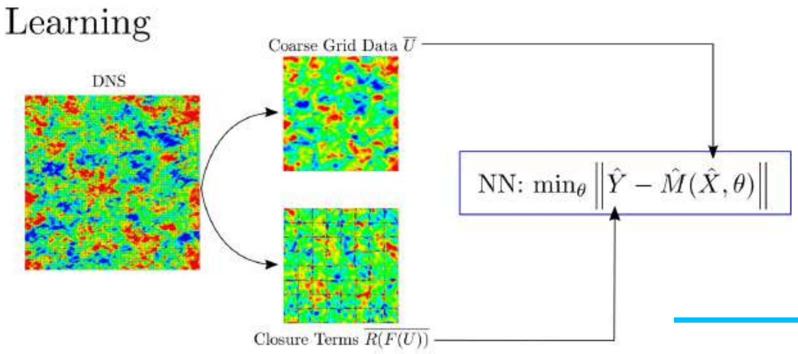

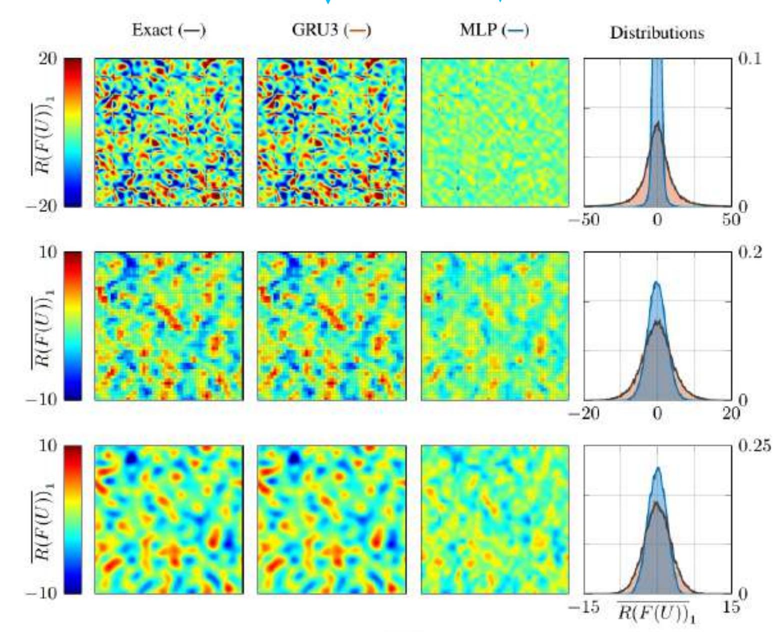

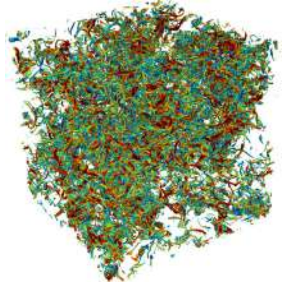

- Turbulence is a non-local phenomena: Pointwise data only not sufficient
- LES closure term prediction from spatial (CNN) or temporal (GRU/LSTM) data
- 99.99% CC, error < 0.01%

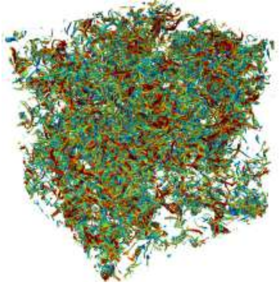

LES Closure Terms: Supervised learning

- Directly close LES equations with ML-learned Reynolds forces
- Implicitly filtered LES: Expected stability problems
- Incorporating physical constraints and ML stabilization for uncertainty estimation helps

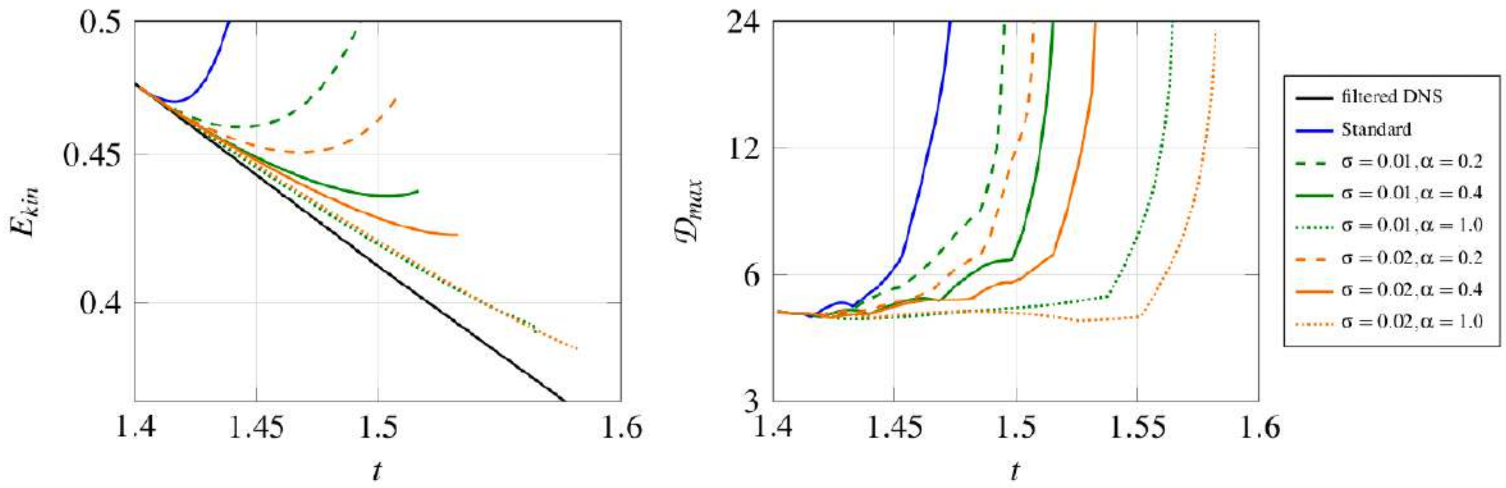

LES Closure Terms: Supervised learning

- Turn the ML prediction into a useful model:
  - include dissipation constraint from TKE in cost function
  - project forces prediction on stable basis
- Eddy viscosity approach with adaptive viscosity in space and time

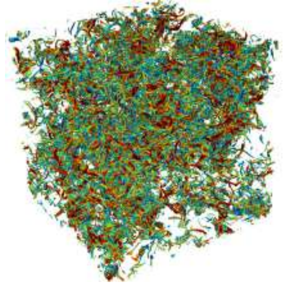

$$\widetilde{R}(F(\overline{U^i})) - \overline{R(F(\overline{U^i}))} \approx \mu_{ANN}\, \widetilde{R}(F^{visc}(\overline{U^i}, \nabla\overline{U^i}))$$

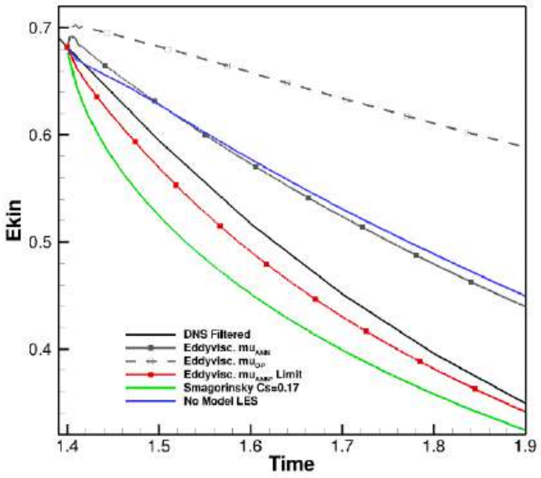

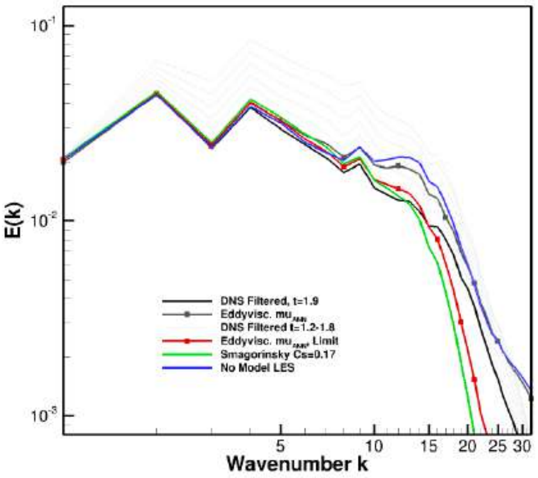

Control Problems

- How to model decisions in dynamical systems under uncertainty?

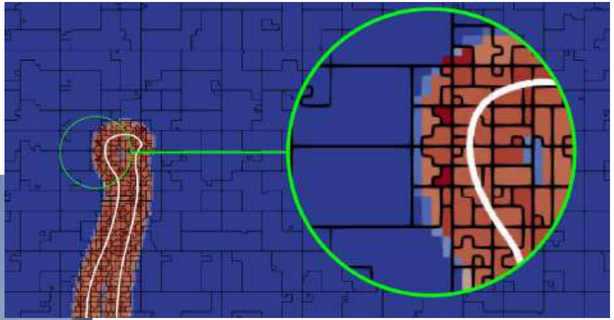

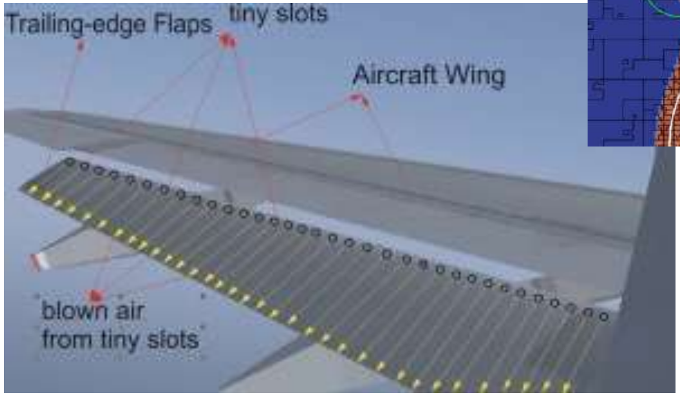

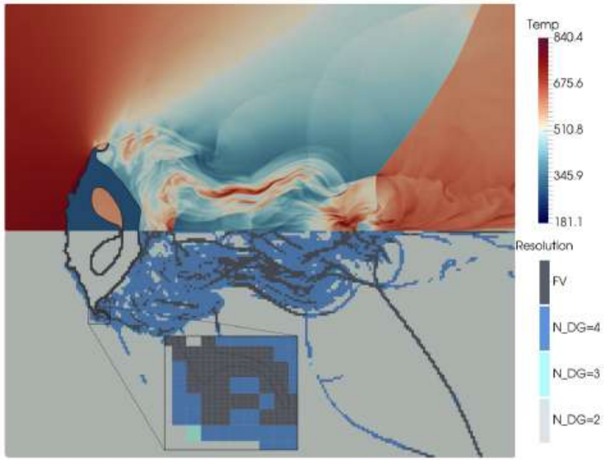

Credit: Department Safety Critical Systems & Systems Engineering, DLR

Learning for dynamical systems

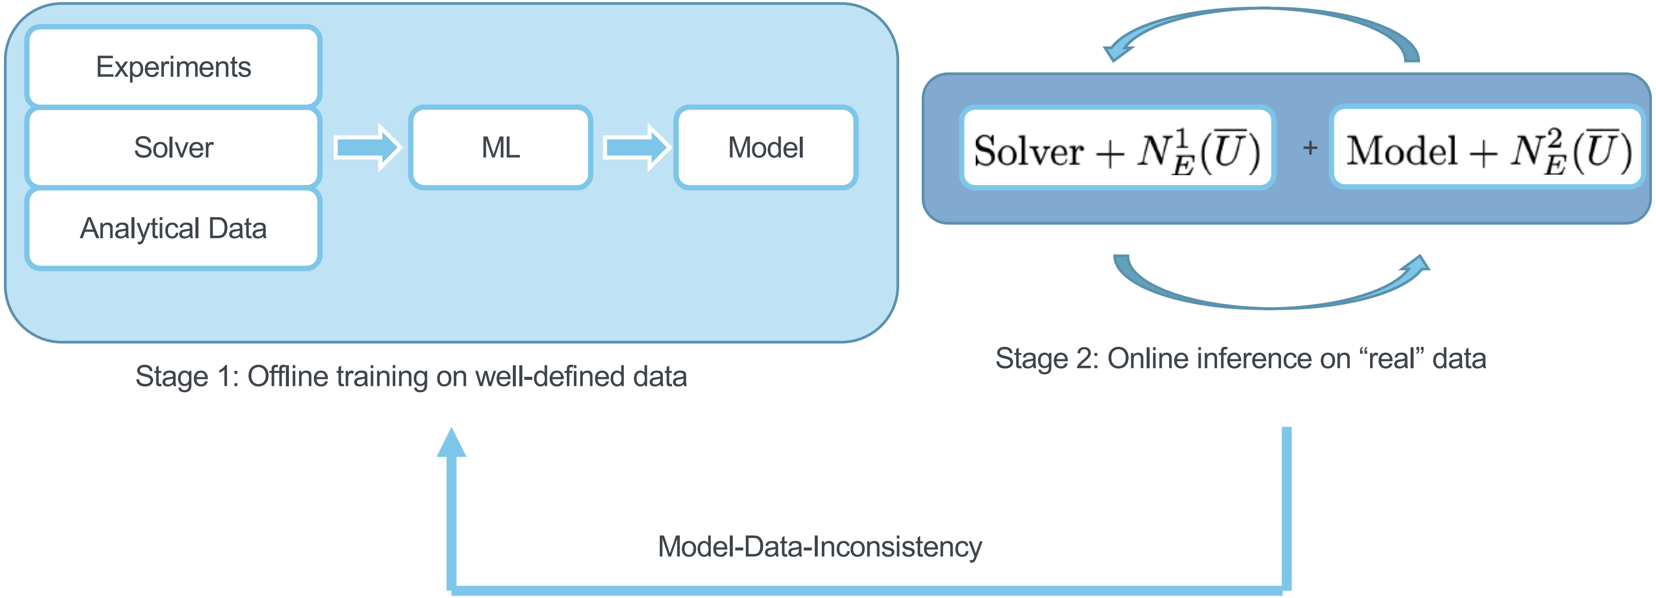

Learning for dynamical systems

- Simulations are (discrete) dynamical systems
- RL is a different paradigm from ML: Learning to take optimal actions in a dynamical system from experience (on the coarse level)
- Supervised learning has two problems: How do we define a lot of true data and make models robust?
- Reinforcement Learning is the method behind the recent successes: Self-Driving Cars, Autonomous Robots, AlphaGo, Starcraft, …

$$\widetilde{R}\,(U_L, x_L, t_L) = -\widetilde{M}\,(U_L)$$

ML with Solver-in-the-loop: discretization-awareness

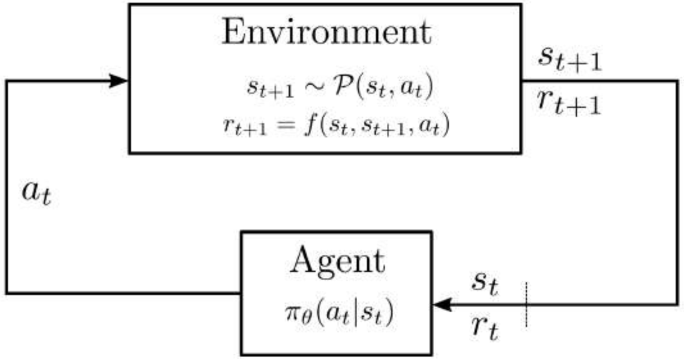

An Example: RL for LES closures

- Turbulence is a prime example of the scale gap
- We need models to augment the coarse-grained equations (LES & RANS)
- We can formulate this as an RL problem: Find a strategy for choosing the best model

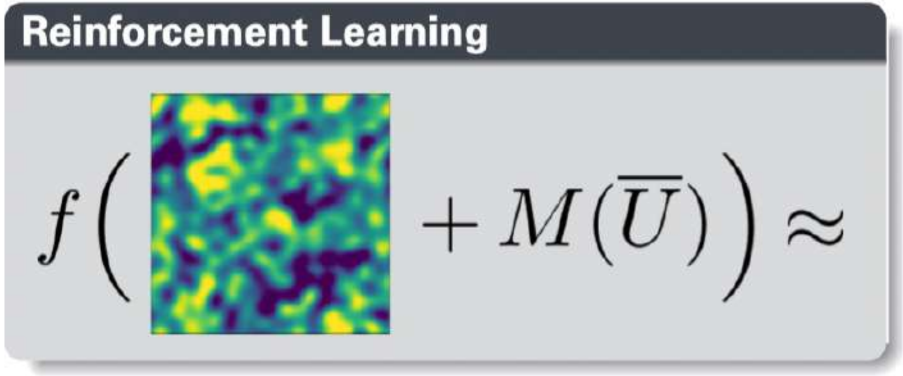

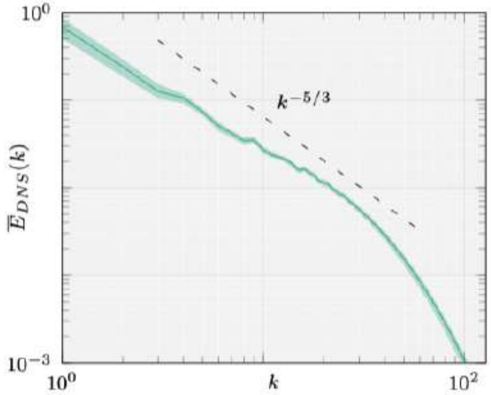

Formulation

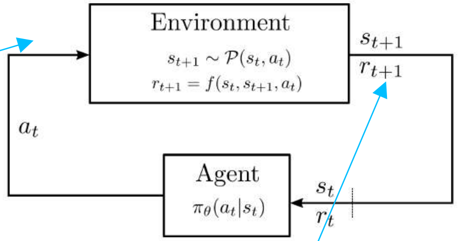

Environment

- **Implicitly filtered LES** with High-Order DG scheme.
- Homogeneous Isotropic Turbulence ("Turbulence-in-a-box"), $Re_\lambda = 180$
- Periodic boundaries
- Forcing for **statistically stationary** flow

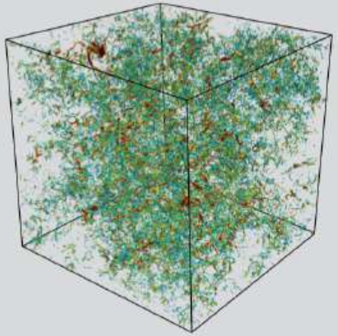

Reward

- Reward based on **error in spectrum** of turbulent kinetic energy
- Spectrum of precomputed DNS as target
- Reward scaled to $\mathbf{r_t} \in [-1, 1]$ with exponential function

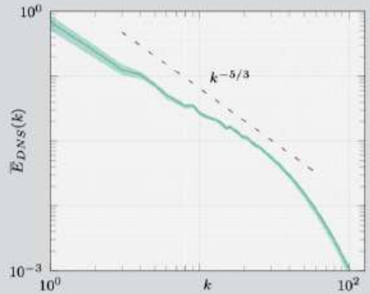

Formulation

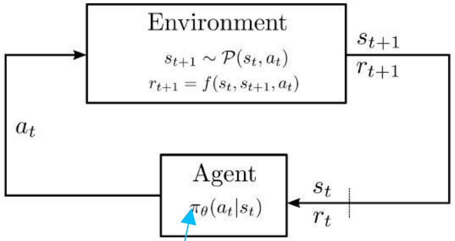

Actions

Smagorinsky's model:

$$\nu_t = (C_s\Delta)^2 \sqrt{2\overline{S}_{ij}\overline{S}_{ij}}$$

- C_S: Model coefficient
- Δ: Filter width
- $\overline{S}$: Rate-of-strain tensor

- Adapt model parameter dynamically in **space and time:** $C_s = C_s(x, t)$
- First step: **elementwise** constant C_s

Policy

- Elementwise **convolutional** architecture

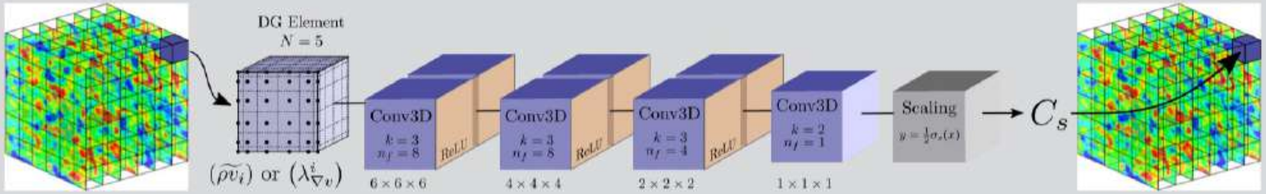

Simulation software: FLEXI
Reinforcement learning framework – ReLeXI[1]

- Distribution on hybrid HPC systems via the SmartSim Library[2]
- Dedicated GPU node („Head“) for training and model evaluation
- FLEXI instances interactively distributed across multiple CPU nodes („Workers“)
- Communication via in-memory database with the Redis library
- Easily extendable to other codes

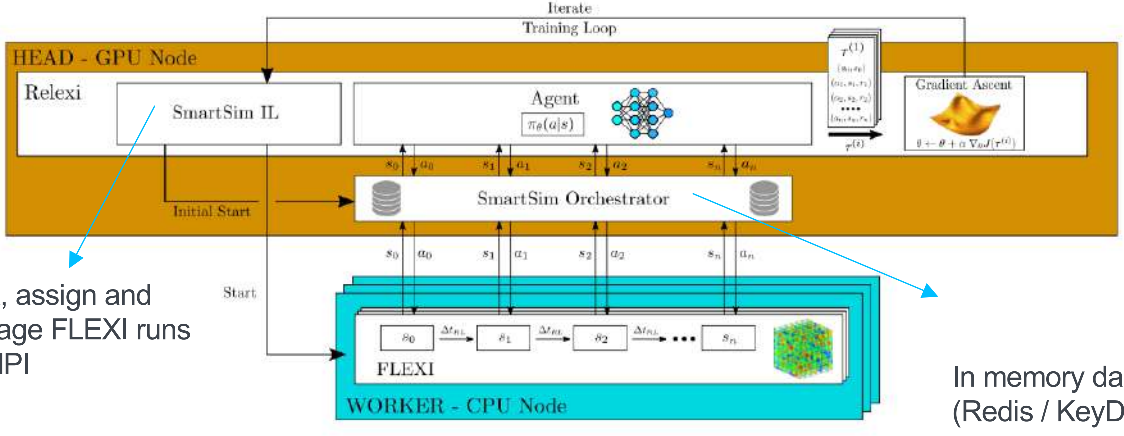

Start, assign and manage FLEXI runs on MPI

In memory database (Redis / KeyDB)

[1]Kurz et al., Relexi—A scalable open source reinforcement learning framework for high-performance computing. Software Impacts, 2022
[2]https://github.com/CrayLabs/SmartSim

Relexi Framework

- The framework scales well over many parallel runs
- We can efficiently evaluate each policy to get reliable gradients
- We can run many samples **in parallel**

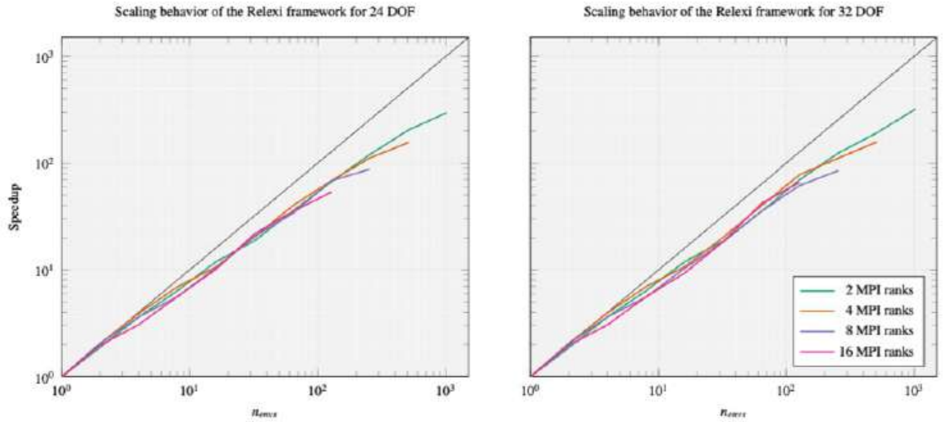

Figure 3: Scaling behavior of the Relexi framework on up to 16 Hawk compute nodes (2048 MPI ranks) and one Hawk-AI node for the HIT test case with 24 DOF and 32 DOF for 2, 4, 8 and 16 MPI ranks per FLEXI instance. The black line indicates perfect scaling.

Results

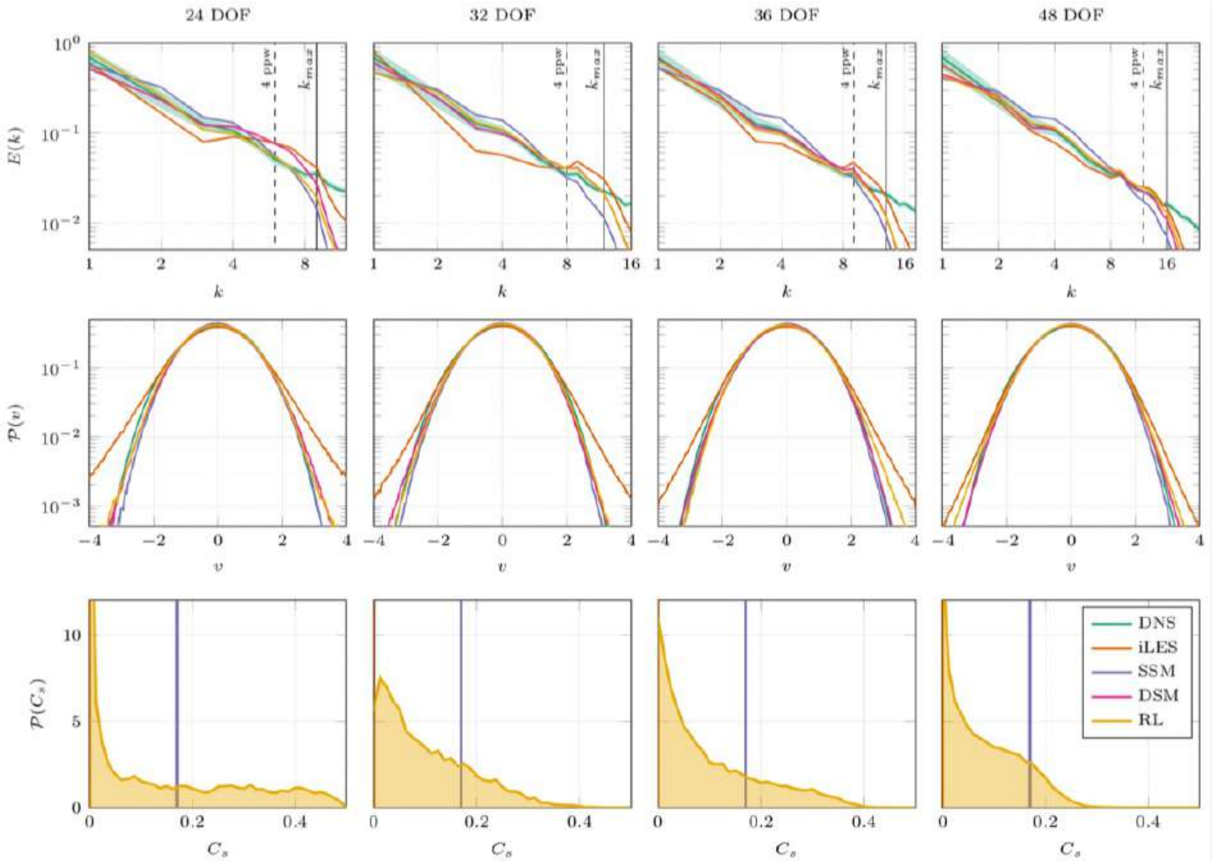
24 DOF
32 DOF
36 DOF
48 DOF
$E(k)$
k
4 ppw
k_{max}
$\mathcal{P}(v)$
v
$\mathcal{P}(C_s)$
C_s
DNS
iLES
SSM
DSM
RL

Results

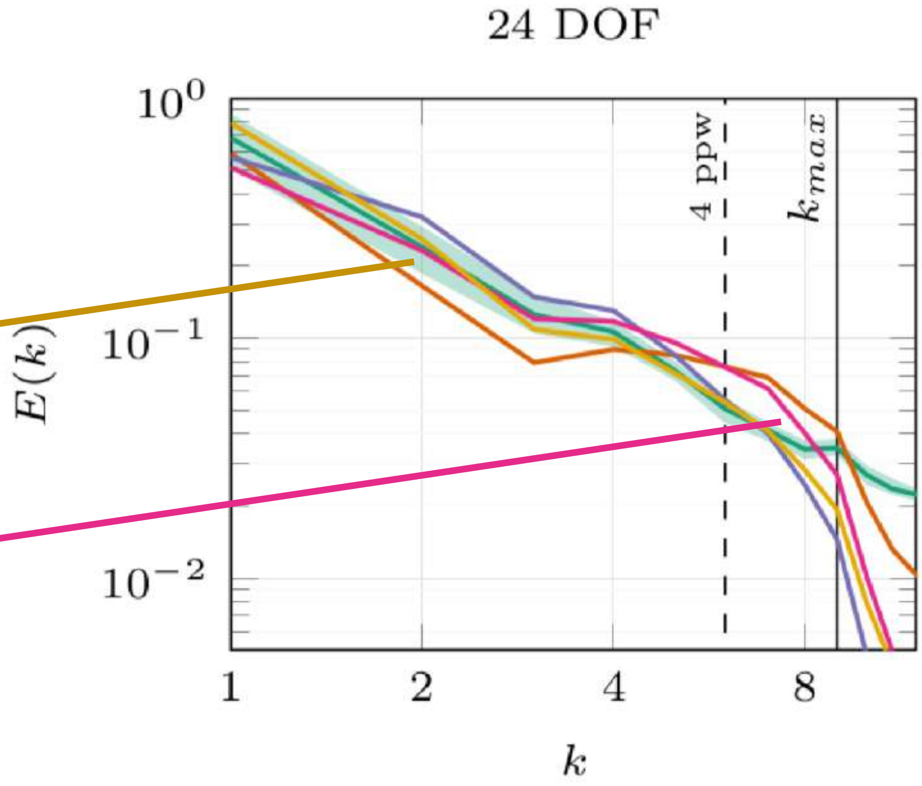

- RL Agent
- Dyn. Smag.

Results

- Optimal model for different discretizations

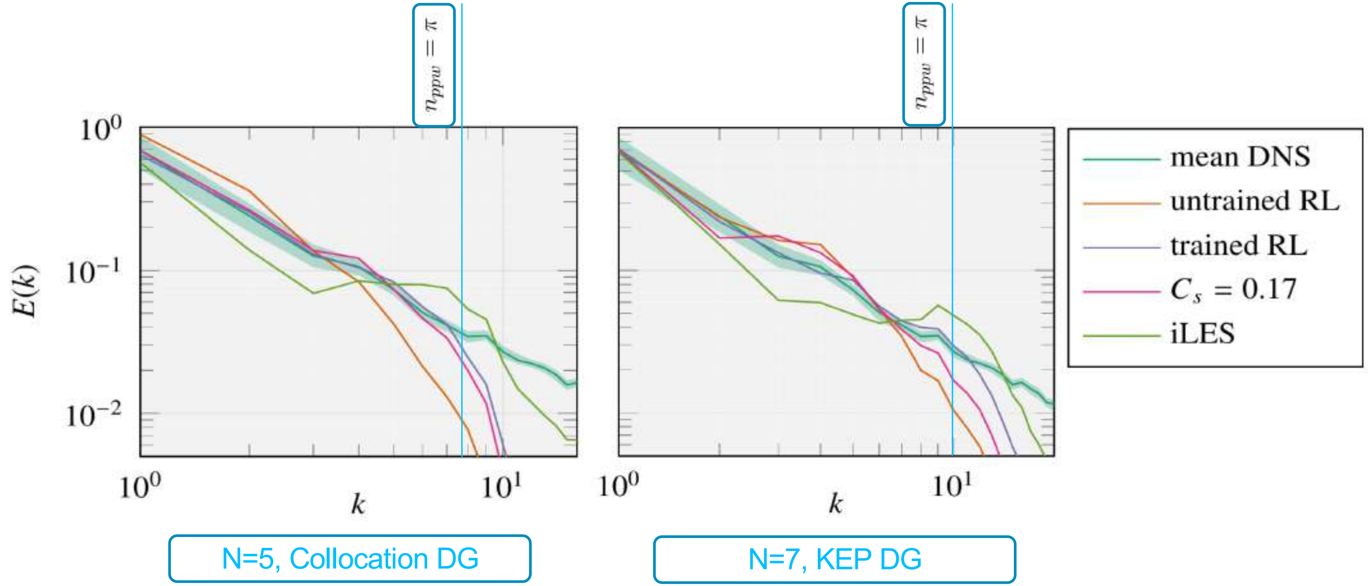

Summary

- Supervised learning of models with guarantees can successfully augment CFD codes
- Examples: Data-driven shock capturing for DG, turbulence closures
- Learning tasks that can be framed as MDPs: Reinforcement learning
- RL can learn from uncertainty, non-linear environments but needs lots of runs to gather experience: HPC
- We have developed a framework for coupling PDE solvers with RL on HPC systems in a plug-and-play style
- With this, we can derive optimal, discretization-specific strategies for turbulence closure
- Generally: an optimizable PDE solver

More Information

- UQ
  - Duerrwaechter, Jakob, et al. “Data-integrated uncertainty quantification for iced airfoil performance prediction”, arXiv:2302.10294v1
  - Duerrwaechter, Jakob, et al. "PoUnce: A framework for automatized uncertainty quantification simulations on high-performance clusters." *Journal of Open Source Software* 8.82 (2023): 4683.
  - Dürrwächter, Jakob, et al. "A high-order stochastic Galerkin code for the compressible Euler and Navier-Stokes equations." *Computers & Fluids* 228 (2021): 105039.
  - Beck, Andrea, et al. "hp-Multilevel Monte Carlo Methods for Uncertainty Quantification of Compressible Navier--Stokes Equations." *SIAM Journal on Scientific Computing* 42.4 (2020): B1067-B1091.
- Reinforcement Learning
  - Kurz, Marius, Philipp Offenhäuser, and Andrea Beck. "Deep reinforcement learning for turbulence modeling in large eddy simulations." *International Journal of Heat and Fluid Flow* 99 (2023): 109094.
  - Kurz, Marius, et al. "Relexi—A scalable open source reinforcement learning framework for high-performance computing." *Software Impacts* 14 (2022): 100422.
  - Schwarz, Anna, et al. “A Reinforcement Learning Based Slope Limiter for Second-Order Finite-Volume Schemes”, 10.13140/RG.2.2.24411.64806
  - Kurz, Marius, et al. "Deep reinforcement learning for computational fluid dynamics on HPC systems." *Journal of Computational Science* 65 (2022): 101884.
- Supervised Learning
  - Schwarz, Anna, et al. "A neural network based framework to model particle rebound and fracture." *Wear* 508 (2022): 204476.
  - Zeifang, Jonas, and Andrea Beck. "A data-driven high order sub-cell artificial viscosity for the discontinuous Galerkin spectral element method." *Journal of Computational Physics* 441 (2021): 110475.
  - Beck, Andrea, and Marius Kurz. "A perspective on machine learning methods in turbulence modeling." *GAMM-Mitteilungen*44.1 (2021): e202100002.
  - Beck, Andrea D., et al. "A neural network based shock detection and localization approach for discontinuous Galerkin methods." *Journal of Computational Physics* 423 (2020): 109824.

Thank you!