Overcoming barriers to adoption of high-fidelity CFD

Presenter:
Charles Mockett
Managing Director
charles.mockett@upstream-cfd.com

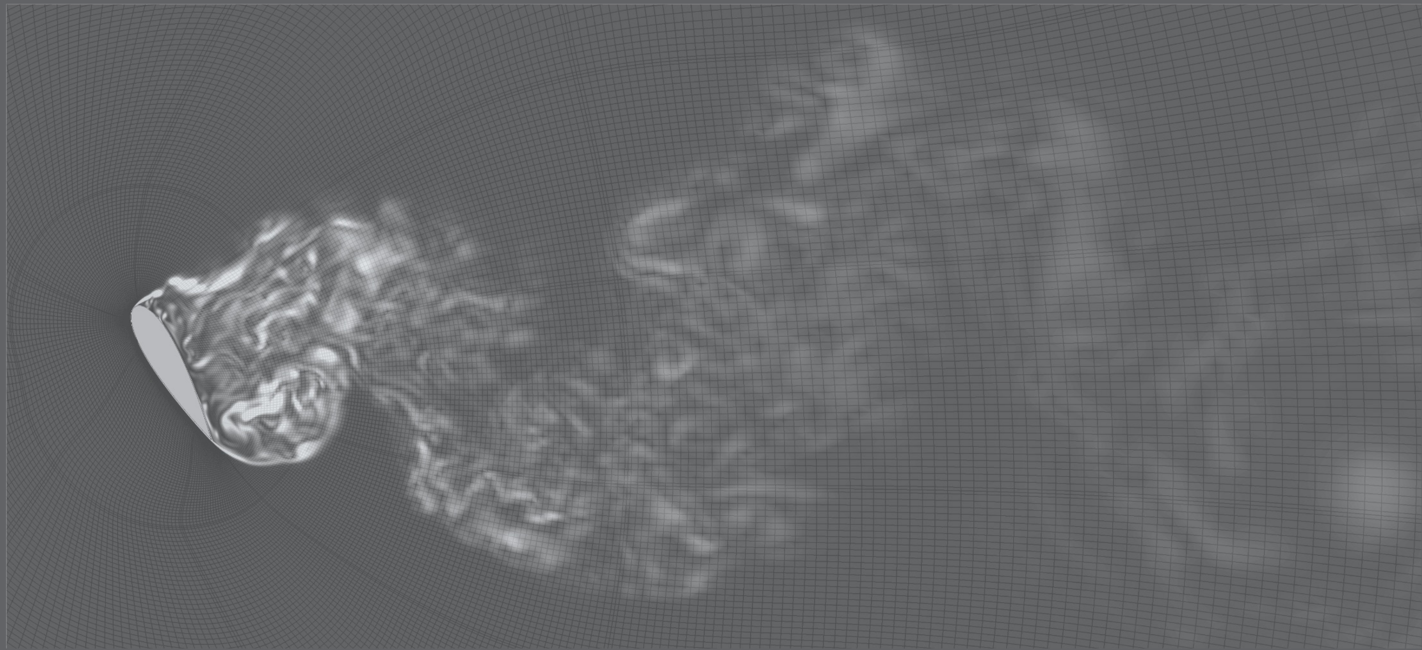

6th German OpenFOAM User Network (GOFUN)
23rd March 2022

High-fidelity CFD

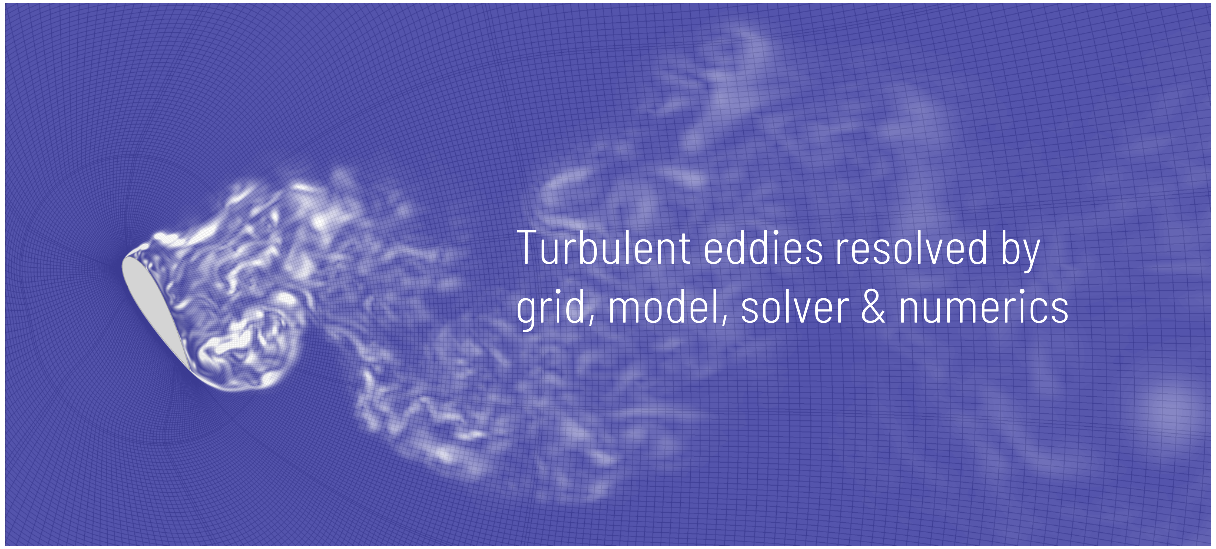

"High-fidelity" a.k.a. turbulence-resolving methods:

- By definition 3D and time-resolved
- Different levels of turbulence resolution

RANS	Hybrid RANS-LES	LES	DNS
Reynolds-averaged Navier-Stokes	e.g. DES, SAS, WMLES, SBES…	Large-eddy simulation	Direct Numerical Simulation
All turbulence modelled		≈ 80% of turbulence resolved	All turbulence resolved
High empiricism		Low empiricism	No empiricism
Ready: 1990		Ready: 2070	Ready: 2080

Increasing computational cost and method accuracy

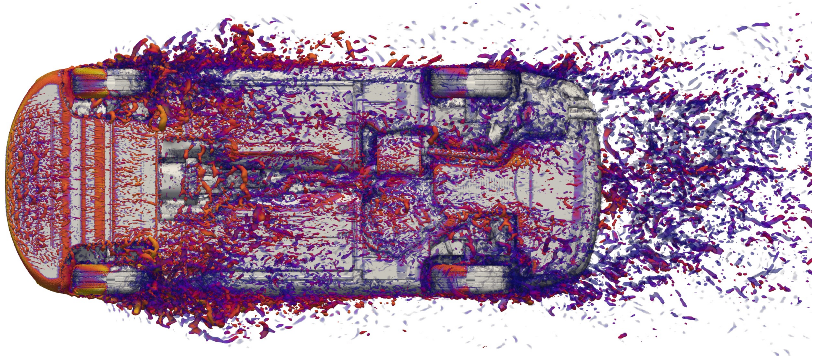

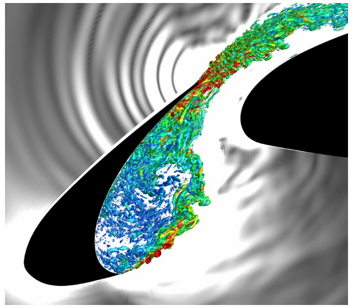

- Advantages of high-fidelity methods over RANS:
  - More accurate prediction of flows with large-scale separation
  - Resolving turbulent unsteadiness necessary for some applications, e.g. broadband aeroacoustics

Upstream CFD & today's presentation

upstream**CFD**

- Experts in high-fidelity CFD, aerodynamics & aeroacoustics
- Innovative SME based in Berlin, Germany
- Mission:
  - To enable the deployment of high-fidelity CFD in industrial design processes

- Today's talk:
  - Present some of the barriers to adoption of high-fidelity CFD
  - ...and some of the steps we are taking to overome them
  - More of a flyover than a deep dive
    - Hyperlinks to more details (slides will be published)
  - Focus on OpenFOAM-relevant content for the GOFUN audience

Components for effective CFD in industrial environments:

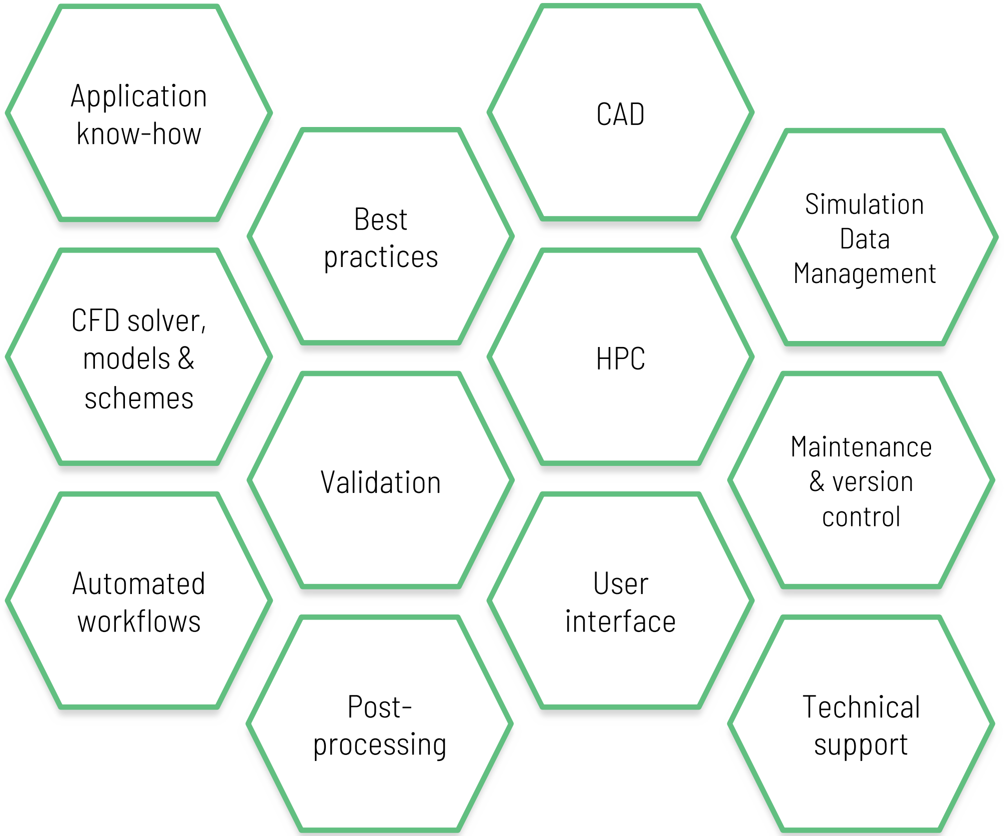

Barriers to adoption of high-fidelity CFD

Barriers:	Efforts to overcome them:
High computational cost, long turnaround times Software license costs HPC investment & maintenance	High-performance computing (HPC) in the cloud
	Best practices tuned for performance
	Evaluating latest developments from the R&D community (exaFOAM)
	Accelerated solvers optimised for new HPC architectures
Method limitations (e.g. turbulence models, numerics)	Overcoming the "grey area" problem in hybrid RANS-LES methods
	Initial activities to remedy the boundary layer shielding problem in DES
	Robust, low-dissipation solvers and schemes
Complicated to use and error-prone	Validated and automated simulation workflows
	Graphical user interfaces
	Automatic and adaptive meshing
	Smart simulation control (averaging start time and sample size)
Standard CFD methods are fine for many applications!	Recognise cases where RANS is the best choice

Outline of talk

upstream CFD

Barriers:	Efforts to overcome them:
High computational cost, long turnaround times	High-performance computing (HPC) in the cloud
	Best practices tuned for performance
	Evaluating latest developments from the R&D community (exaFOAM)
	Accelerated solvers optimised for new HPC architectures
Method limitations (e.g. turbulence models, numerics)	Overcoming the "grey area" problem in hybrid RANS-LES methods
	Initial activities to remedy the boundary layer shielding problem in DES
	Robust, low-dissipation solvers and schemes
Complicated to use and error-prone	Validated and automated simulation workflows
	Graphical user interfaces
	Automatic and adaptive meshing
	Smart simulation control (averaging start time and sample size)
Standard CFD methods are fine for many applications!	Recognise cases where RANS is the best choice

Outline of talk

Barriers:	Efforts to overcome them:
High computational cost, long turnaround times	High-performance computing (HPC) in the cloud
	Best practices tuned for performance
	Evaluating latest developments from the R&D community (exaFOAM)
	Accelerated solvers optimised for new HPC architectures
Method limitations (e.g. turbulence models, numerics)	Overcoming the "grey area" problem in hybrid RANS-LES methods
	Initial activities to remedy the boundary layer shielding problem in DES
	Robust, low-dissipation solvers and schemes
Complicated to use and error-prone	Validated and automated simulation workflows
	Graphical user interfaces
	Automatic and adaptive meshing
	Smart simulation control (averaging start time and sample size)
Standard CFD methods are fine for many applications!	Recognise cases where RANS is the best choice

HPC in the cloud with settings tuned for performance

- Improvements to in-house DES best practices:
  - Factor 5.6 speedup (raw computational cost) compared to previous settings
- Impressive scalability to thousands of cores seen on Amazon EC2 cloud
  - AWS ParallelCluster with c5n.18xlarge instance type and Elastic Fabric Adapter high-performance interconnect
  - Turnaround time reduction: Large DES cases possible in a few hours

powered by aws

OpenFOAM DES results for DrivAer notchback case on grid with 121M cells

Computation costs and turn-around times for optimised setup:

nCores	72	144	648	1296	2448	7200
Core-hrs	37.77k	37.13k	35.98k	34.60k	42.26k	55.10k
Parallel efficiency	100%	102%	105%	109%	89%	69%
Turnaround time / hrs	525	263	58	29	15	5

(Assuming 121M Cells, 3s simulated time, 1e-4s time step)

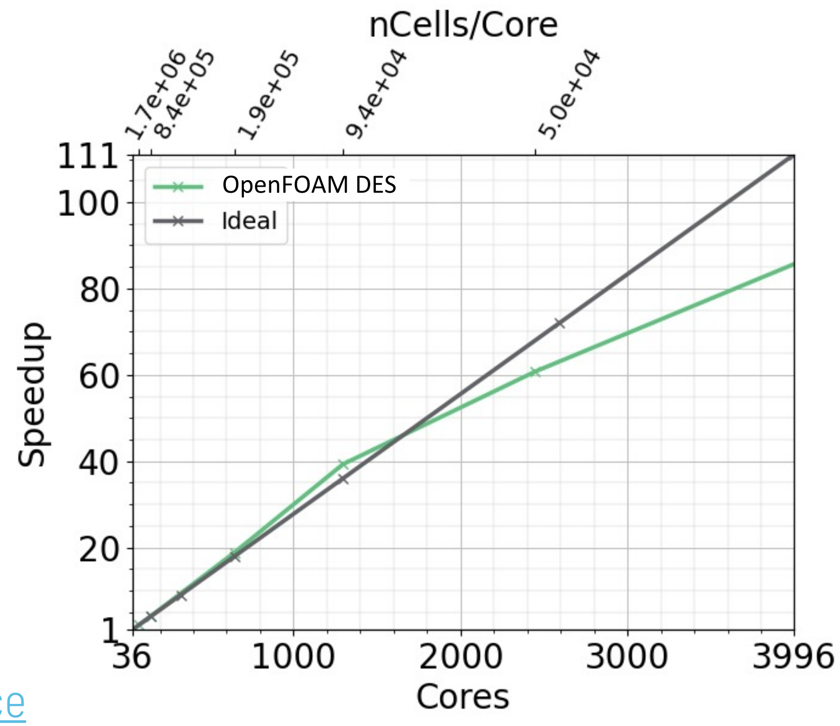

For more details, see Hetmann et al. presentation at the 2021 OpenFOAM Conference

Outline of talk

Barriers:	Efforts to overcome them:
High computational cost, long turnaround times	High-performance computing (HPC) in the cloud
	Best practices tuned for performance
	Evaluating latest developments from the R&D community (exaFOAM)
	Accelerated solvers optimised for new HPC architectures
Method limitations (e.g. turbulence models, numerics)	Overcoming the "grey area" problem in hybrid RANS-LES methods
	Initial activities to remedy the boundary layer shielding problem in DES
	Robust, low-dissipation solvers and schemes
Complicated to use and error-prone	Validated and automated simulation workflows
	Graphical user interfaces
	Automatic and adaptive meshing
	Smart simulation control (averaging start time and sample size)
Standard CFD methods are fine for many applications!	Recognise cases where RANS is the best choice

Open-Source CFD on exascale systems: exaFOAM

- R&D project with joint EU & member state funding
  - Objective: Reduction of bottlenecks to performance scalability of CFD on massively-parallel HPC systems
- Maximum exploitation of project results through publication in the open source CFD software OpenFOAM
- Project started in April 2021, 36 months duration
- Total budget approx. 5.4 M€
- 12 partners from 7 European countries and 17 industrial associates / stakeholders
- For more information: https://exafoam.eu

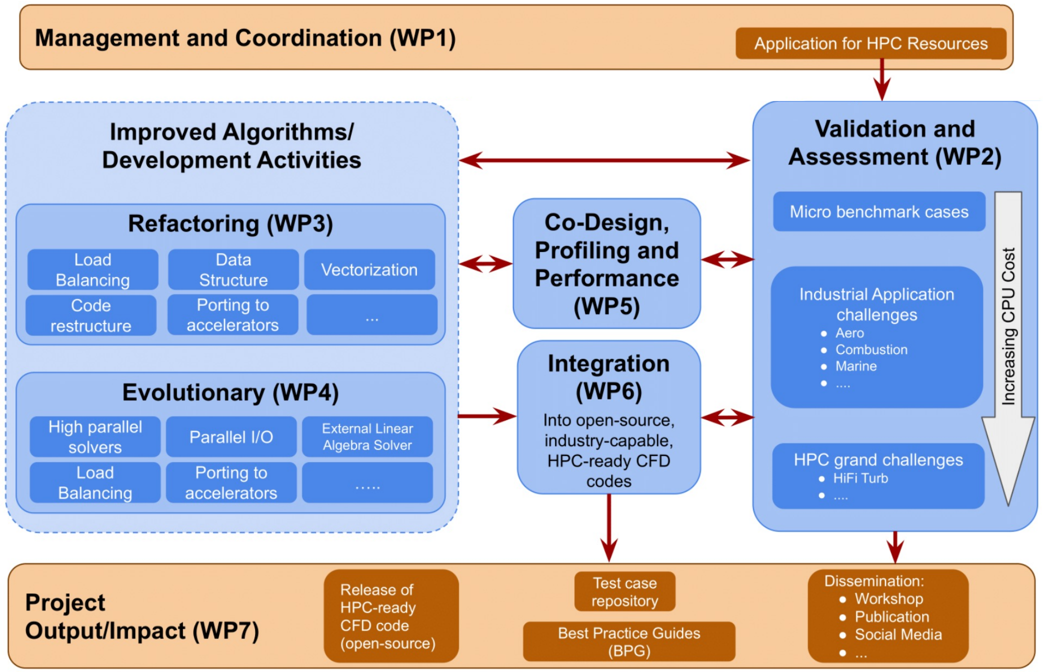

GEFÖRDERT VOM

Bundesministerium für Bildung und Forschung

This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956416. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, United Kingdom, Germany, Italy, Croatia, Spain, Greece, Portugal

Industrial Application Challenge: DrivAer with rotating wheels

- Ford Open Cooling DrivAer (OCDA) model:
  - Full scale
  - Based on original ("closed cooling") DrivAer model (TU-Munich), but with some minor differences in underbody geometry
  - Measurements available
- Specific variant selected for exaFOAM:
  - 140kph, 0° yaw
  - Notchback geometry (considered most challenging to predict)
  - Detailed underbody
  - Closed cooling (despite the name…)
  - Meshing with snappyHexMesh ~ 250M cells

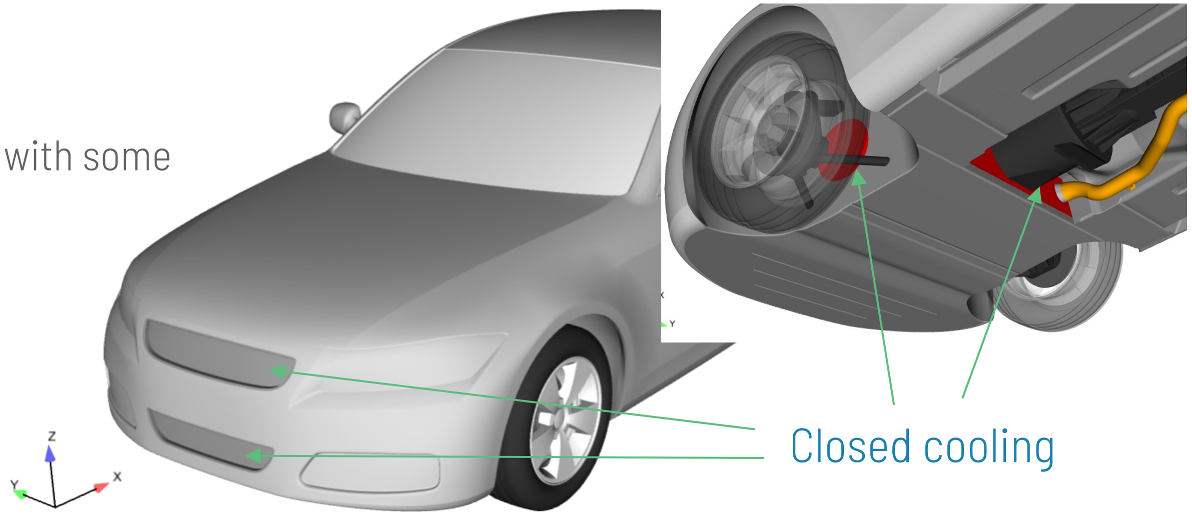

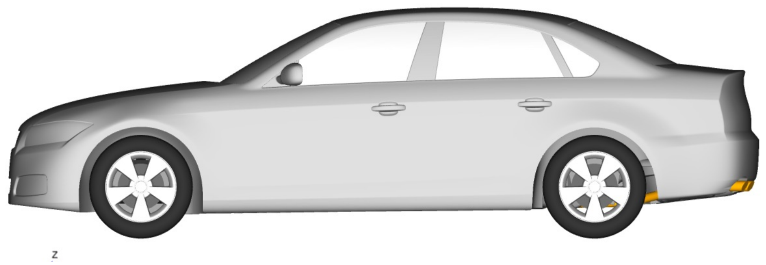

- Microbenchmark: Isolated rotating wheel
  - Isolated from full scale car
  - Rotating mesh
  - Meshing with snappyHexMesh ~ 20M cells

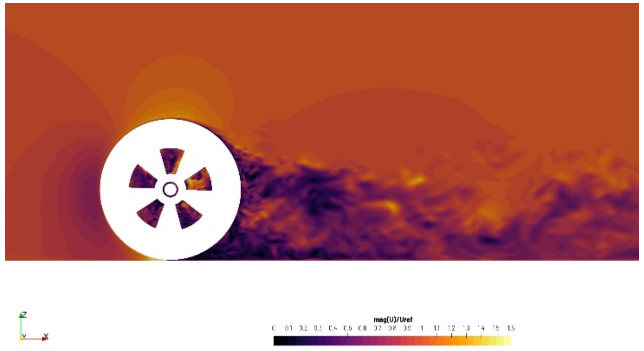

The project leading to this application has received funding from the European High-Performance Computing Joint Undertaking Joint Undertaking (JU) under grant agreement No 956416.
The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, United Kingdom, Germany, Italy, Croatia, Spain, Greece, Portugal

HPC Grand Challenge: WMLES of NASA CRM aircraft model

- Test case at AIAA CFD High Lift Prediction Workshops
- Geometry and meshes from different sources publicly available
- $Re_c = 5{\times}10^6$, $M = 0.2$, $\alpha = 16°$
- Wall-modelled LES (WMLES), potentially zonal RANS/WMLES
- Targeting 20 billion cells on 28.000 cores
- Meshing with snappyHexMesh

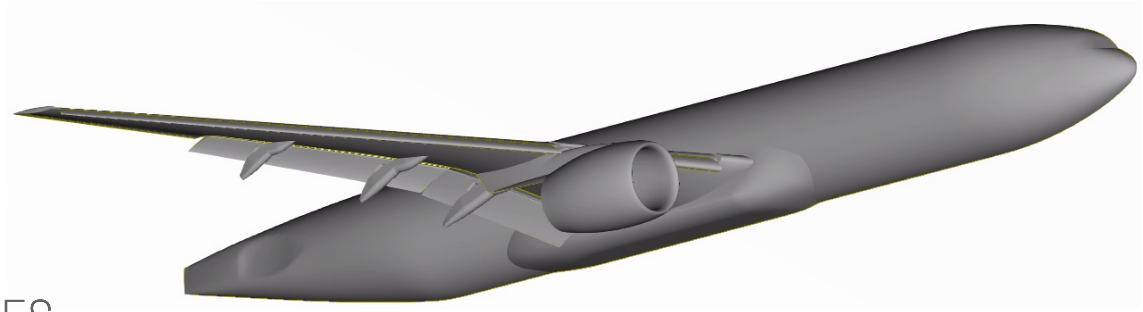

- Microbenchmark:
  - Extruded 30P30N
  - Flight conditions as above
  - Reference data exists

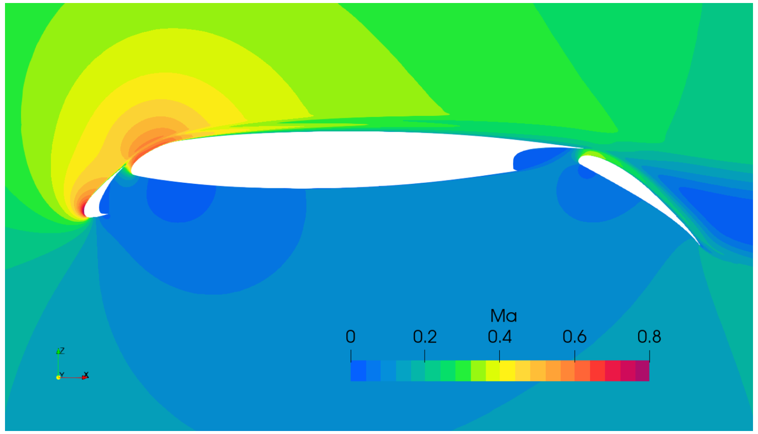

The project leading to this application has received funding from the European High-Performance Computing Joint Undertaking Joint Undertaking (JU) under grant agreement No 956416.
The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, United Kingdom, Germany, Italy, Croatia, Spain, Greece, Portugal

Outline of talk

Barriers:	Efforts to overcome them:
High computational cost, long turnaround times	High-performance computing (HPC) in the cloud
	Best practices tuned for performance
	Evaluating latest developments from the R&D community (exaFOAM)
	Accelerated solvers optimised for new HPC architectures
Method limitations (e.g. turbulence models, numerics)	Overcoming the "grey area" problem in hybrid RANS-LES methods
	Initial activities to remedy the boundary layer shielding problem in DES
	Robust, low-dissipation solvers and schemes
Complicated to use and error-prone	Validated and automated simulation workflows
	Graphical user interfaces
	Automatic and adaptive meshing
	Smart simulation control (averaging start time and sample size)
Standard CFD methods are fine for many applications!	Recognise cases where RANS is the best choice

The Grey Area problem and its mitigation

- "Grey Area": Delayed transition from RANS to LES in free shear layer following BL separation
- Many practical flows are affected, e.g. shallow separation, vortices, jets…

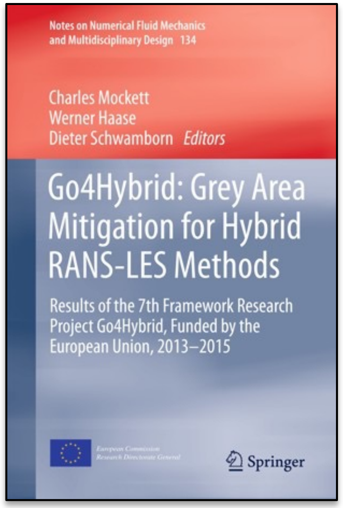

EU-funded Go4Hybrid project (2013-2015)

- We proposed a modified DES version named σ-DDES:
  - First publication: Mockett et al. (2015)
  - An extension to DDES, maintaining all key features of the original model
  - Strong reduction of Grey Area for a wide range of fundamental and complex flows: Fuchs et al. (2020)
- Results of σ-DDES better than or equivalent to standard DDES for all cases tested…
- …exception: boundary layer shielding issues
  - A separate and important problem that we hope to tackle soon

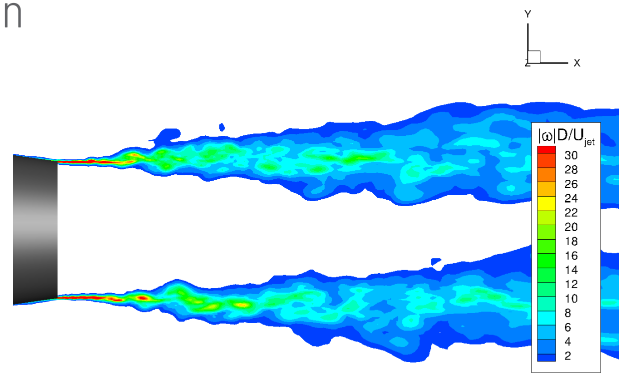

What we want…

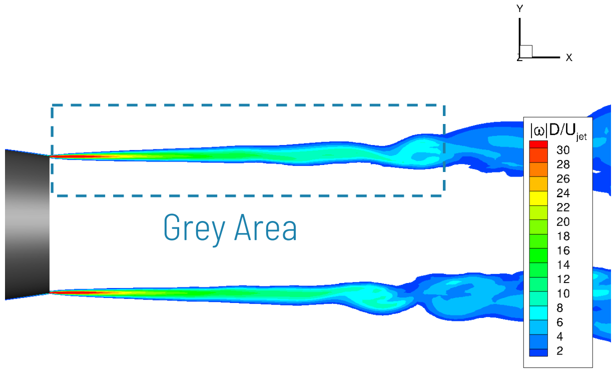

…what we get

OpenFOAM®

σ-DDES is currently being integrated into the central OpenFOAM release – target is v2206

Example results with σ-DDES

Direct comparisons with standard DDES (identical grid, numerics etc.)

	Delta wing	2D hump	Round jet
Std. DDES	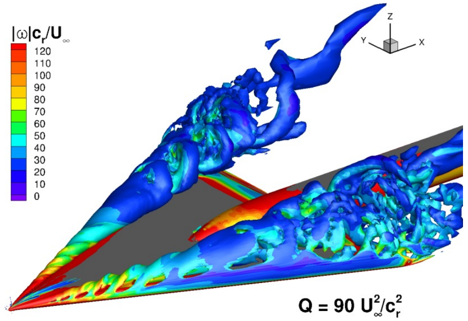	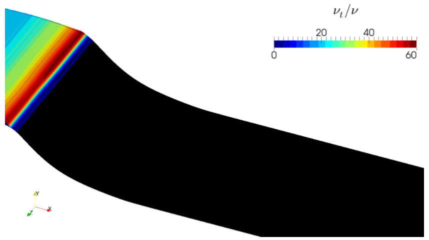	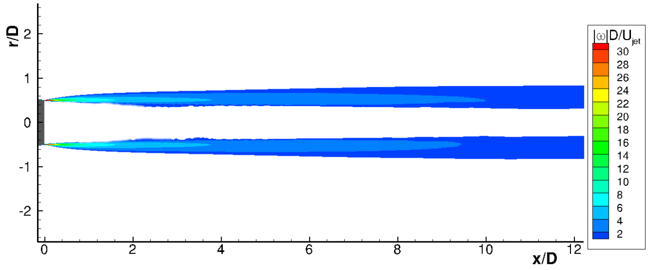
σ-DDES	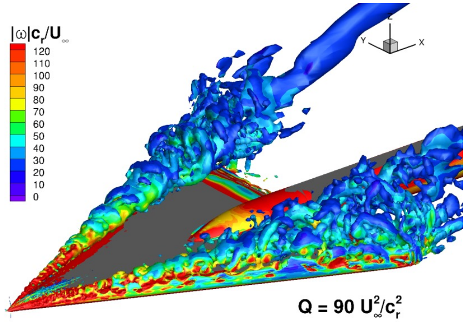	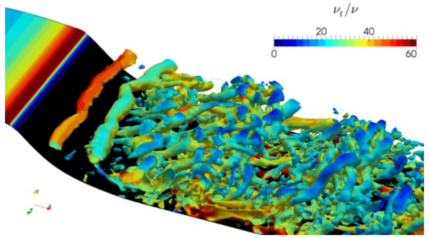	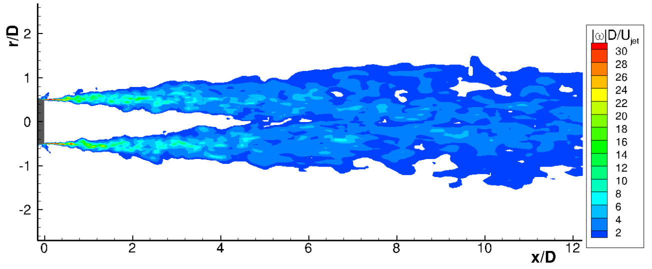

For more information (and quantitative results) see e.g.:

- Fuchs et al. (2015) – Delta wing case
- Fuchs et al. (2020) – Review of formulation, 2D hump, Ahmed body, rudimentary landing gear cases

Outline of talk

Barriers:	Efforts to overcome them:
High computational cost, long turnaround times	High-performance computing (HPC) in the cloud
	Best practices tuned for performance
	Evaluating latest developments from the R&D community (exaFOAM)
	Accelerated solvers optimised for new HPC architectures
Method limitations (e.g. turbulence models, numerics)	Overcoming the "grey area" problem in hybrid RANS-LES methods
	Initial activities to remedy the boundary layer shielding problem in DES
	Robust, low-dissipation solvers and schemes
Complicated to use and error-prone	Validated and automated simulation workflows
	Graphical user interfaces
	Automatic and adaptive meshing
	Smart simulation control (averaging start time and sample size)
Standard CFD methods are fine for many applications!	Recognise cases where RANS is the best choice

Robust, low-dissipation solver: WRLES of pipe flow

- Fully-developed pipe flow, $Re_\tau = 550$, constant driving pressure gradient
- Wall-resolved LES with σ SGS model (Nicoud et al.)
- Structured grid, $\Delta x^+ = 40, \Delta\theta^+ = 10$
- Direct comparison of incompressible OpenFOAM solvers
- Advantages of improved solver:
  - More robust thanks to consistent pressure-velocity coupling
  - Lower numerical dissipation
- 50% reduction in bulk velocity error with improved solver

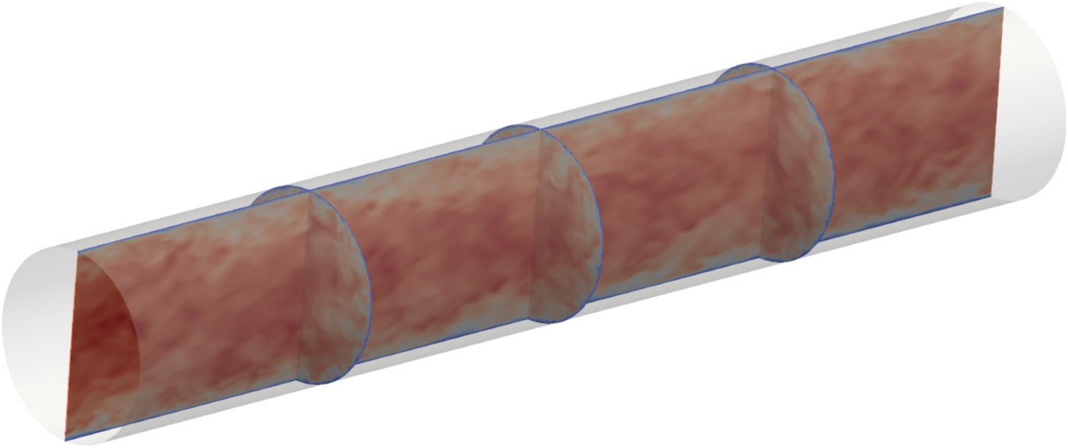

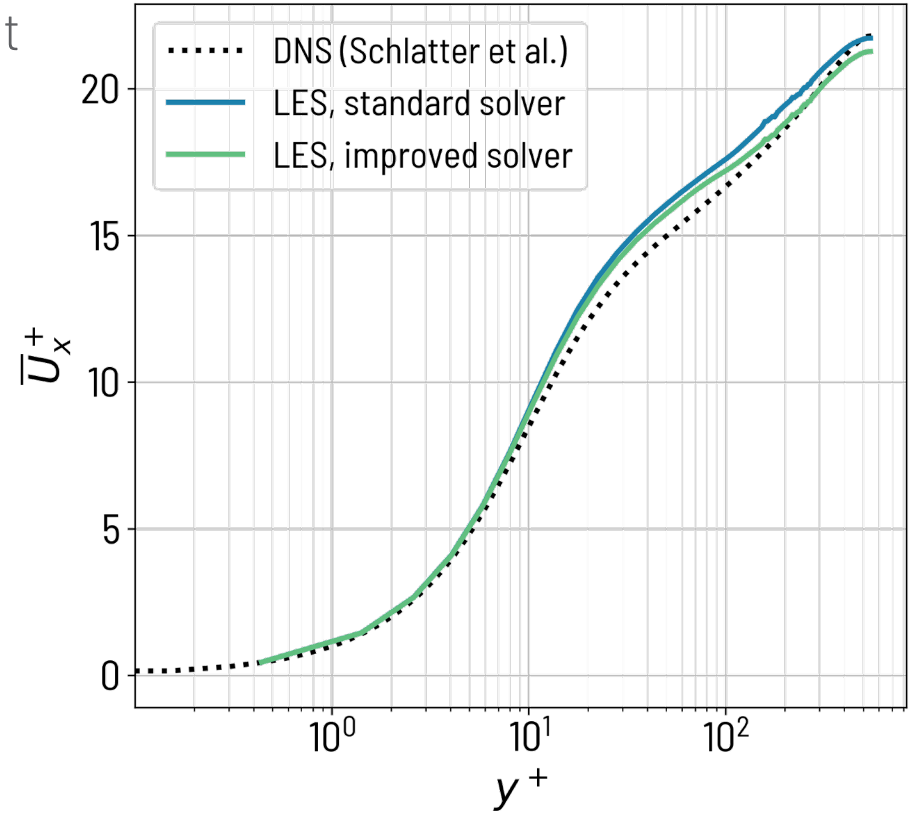

We plan to publish the solver details in the OpenFOAM Journal and contribute the code to the central release (hopefully this year)

Complete DES methodology for DrivAer Notchback configuration

- DrivAer Notchback case from 2nd Automotive CFD Prediction Workshop
- Computed on committee grid (128M cells, ANSA, wall functions)
- All components of advanced DES methodology demonstrated for complex case:
  - σ-DDES (grey area improvement)
  - "Enhanced protection" BL shielding function of Deck & Renard (2019)
  - Improved solver and advanced hybrid scheme to minimise dissipation and maximise robustness

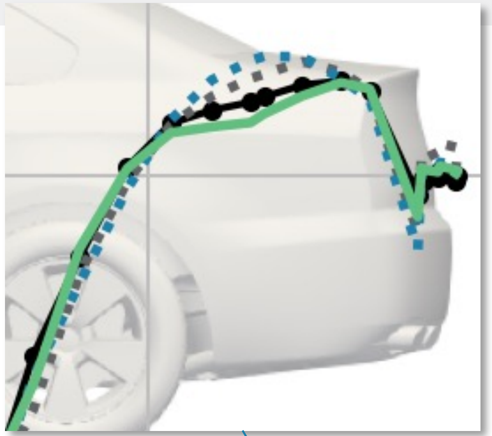

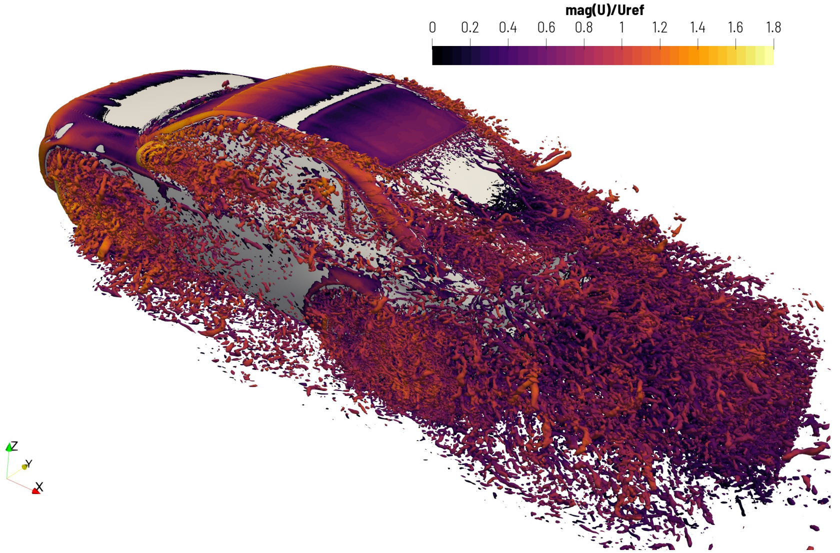

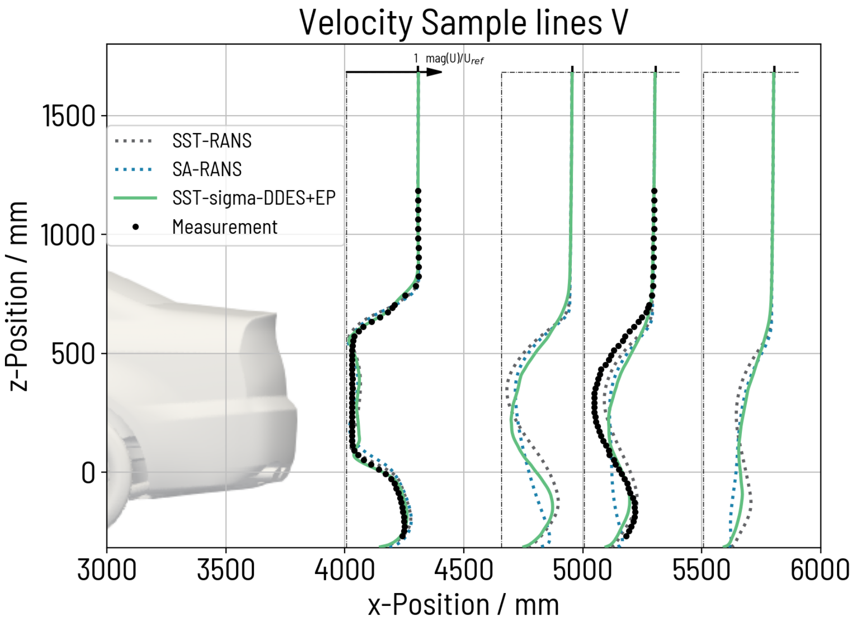

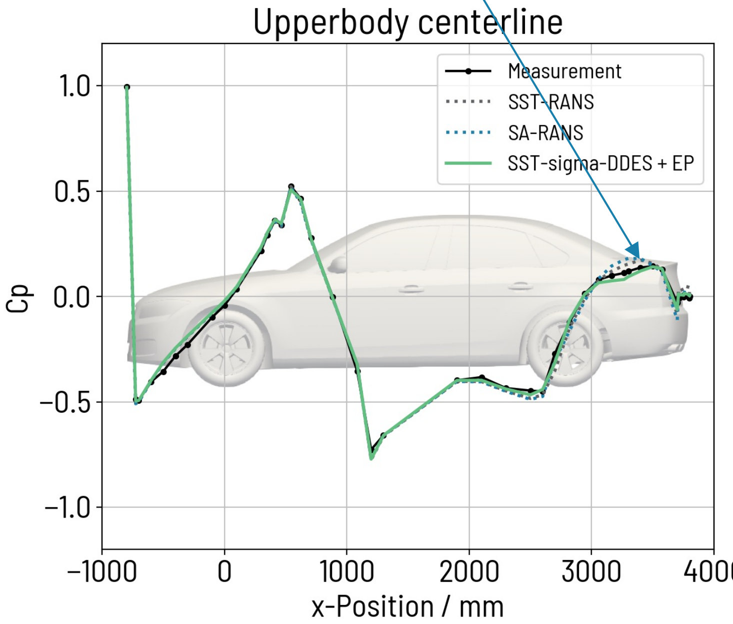

Outline of talk

Barriers:	Efforts to overcome them:
High computational cost, long turnaround times	High-performance computing (HPC) in the cloud
	Best practices tuned for performance
	Evaluating latest developments from the R&D community (exaFOAM)
	Accelerated solvers optimised for new HPC architectures
Method limitations (e.g. turbulence models, numerics)	Overcoming the "grey area" problem in hybrid RANS-LES methods
	Initial activities to remedy the boundary layer shielding problem in DES
	Robust, low-dissipation solvers and schemes
Complicated to use and error-prone	Validated and automated simulation workflows
	Graphical user interfaces
	Automatic and adaptive meshing
	Smart simulation control (averaging start time and sample size)
Standard CFD methods are fine for many applications!	Recognise cases where RANS is the best choice

Smart control of unsteady simulations

- In-house statistical processing tool *Meancalc* for analysis of time signals:
  1. Automatic detection and removal of initial transient
  2. Quantification of the statistical error (e.g. 95% confidence interval) on mean and standard deviation

Example for automotive CFD force coefficient:

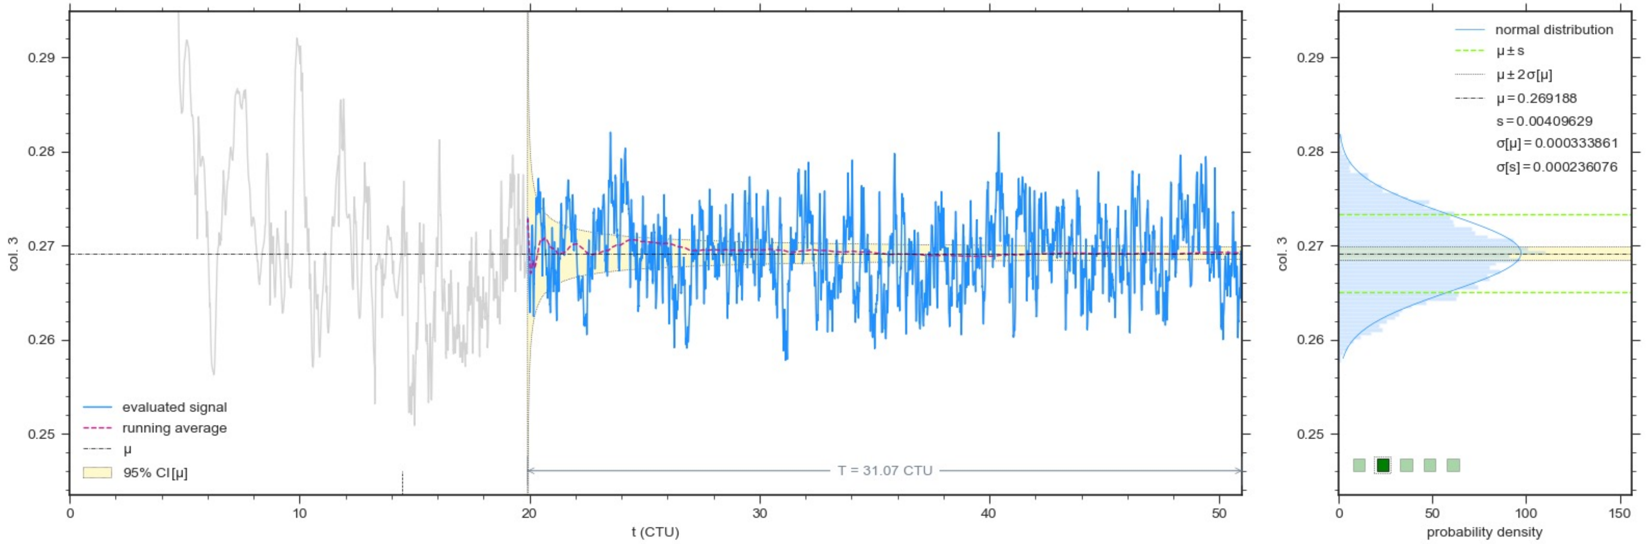

coming soon
to a cloud near you!

Benefits:

- Optimise use of HPC resources
- Quality Assurance: Guarantee that accuracy targets are met
- Insight: Separate real differences in results from statistical noise

- New components:
  – Meancalc–OpenFOAM interface for run-time control of simulations
  – Module to automatically generate insightful plots
- Currently being deployed for pilot customers from the automotive industry

Closing remarks

- Despite significant promise, numerous barriers to the adoption of high-fidelity CFD remain
- Upstream CFD are working hard, together with our clients, partners and the OpenFOAM Community, to overcome these, e.g.:
  - Contribution of improvements to OpenFOAM via the Governance framework (e.g. Turbulence Technical Committee)
  - Assessing HPC developments in the exaFOAM project
  - Harnessing the power of HPC in the cloud, and helping customers/partners to do the same
  - Software modules for smart, automated and efficient simulation workflows
- Our **vision** for the future:
  - Automated, application-specific workflows consisting of linked SaaS modules in the cloud

- High-fidelity methods are great, but RANS will remain valuable for decades to come
  - Be cautious of CFD offerings without steady-state RANS capabilities!

- Please get in touch if you are facing any challenges in the adoption and deployment of high-fidelity CFD!

upstreamCFD

Thank you for your attention