Co-design in CFD - Initial Experiences from the ExaFLOW Project

Philipp Schlatter, KTH, Stockholm, SWEDEN

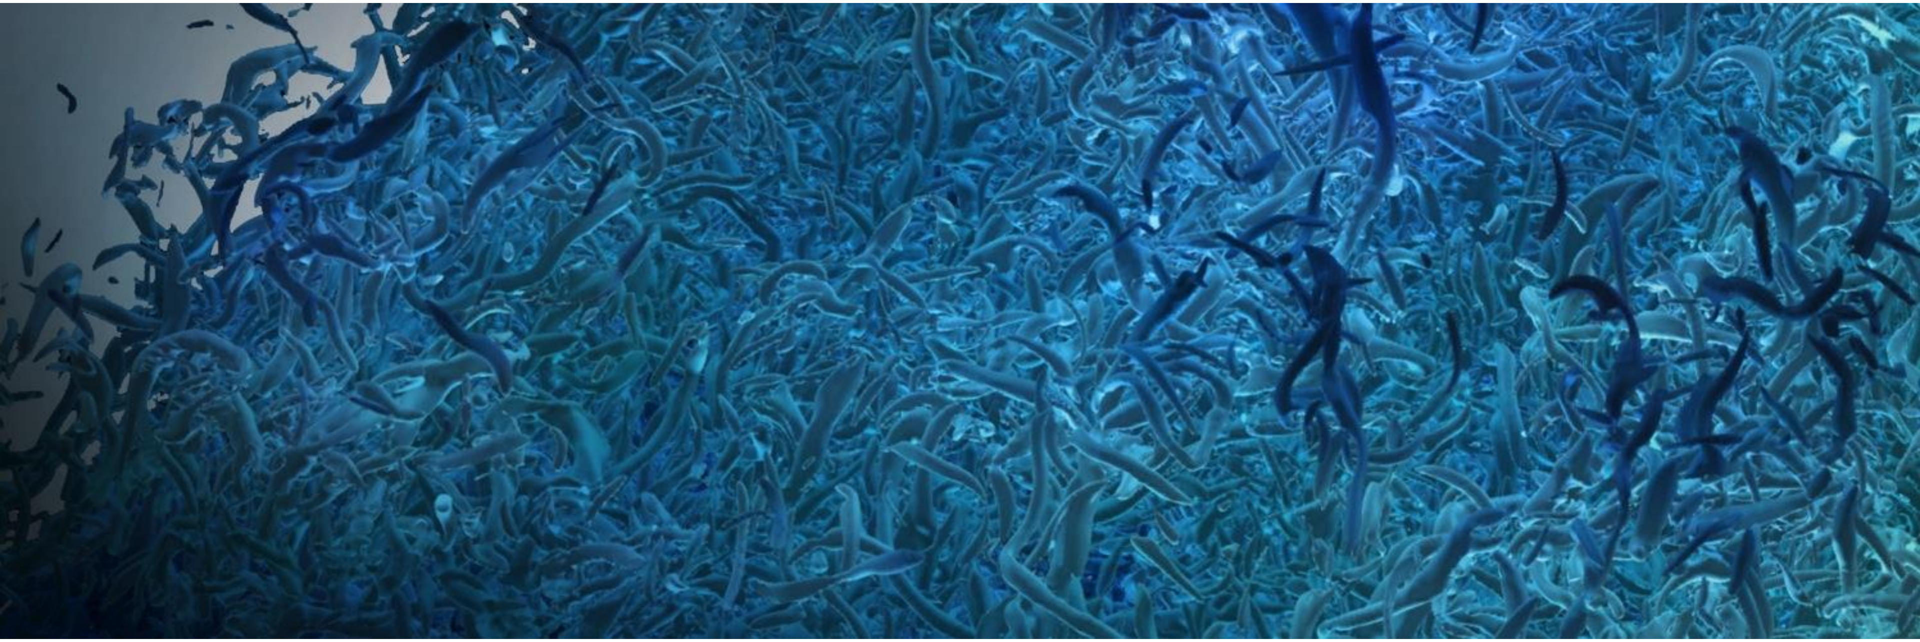

→ Address current algorithmic bottlenecks to enable the use of accurate CFD codes for problems of practical engineering interest

ExaFLOW

→ Address current **algorithmic** bottlenecks to enable the use of **accurate CFD codes** for problems of practical **engineering interest**

The Airbus Challenge, David Hills, 2008

A Brief Diversion Into Aircraft Drag

A world of challenge & opportunity

- Typical break down of overall aircraft† drag by form & component

An Airbus 310 cruising at 250 m/s at 10000m
Teraflops machine (10^{12} Flops): $8 \cdot 10^{5}$ years
Result in one week: $4 \cdot 10^{19}$ flops machine (40 EFlops)
(based on John Kim's estimate, TSFP-9, 2015)

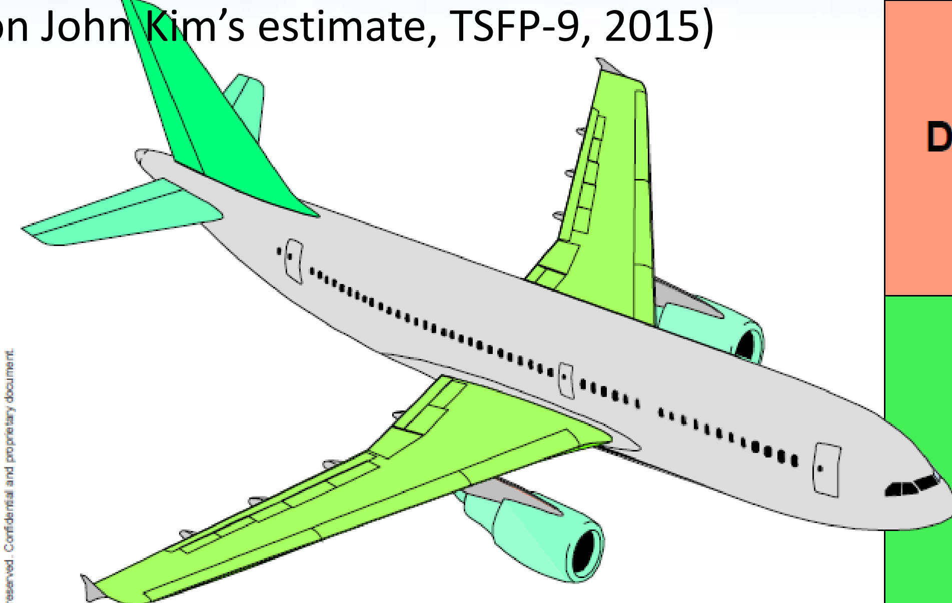

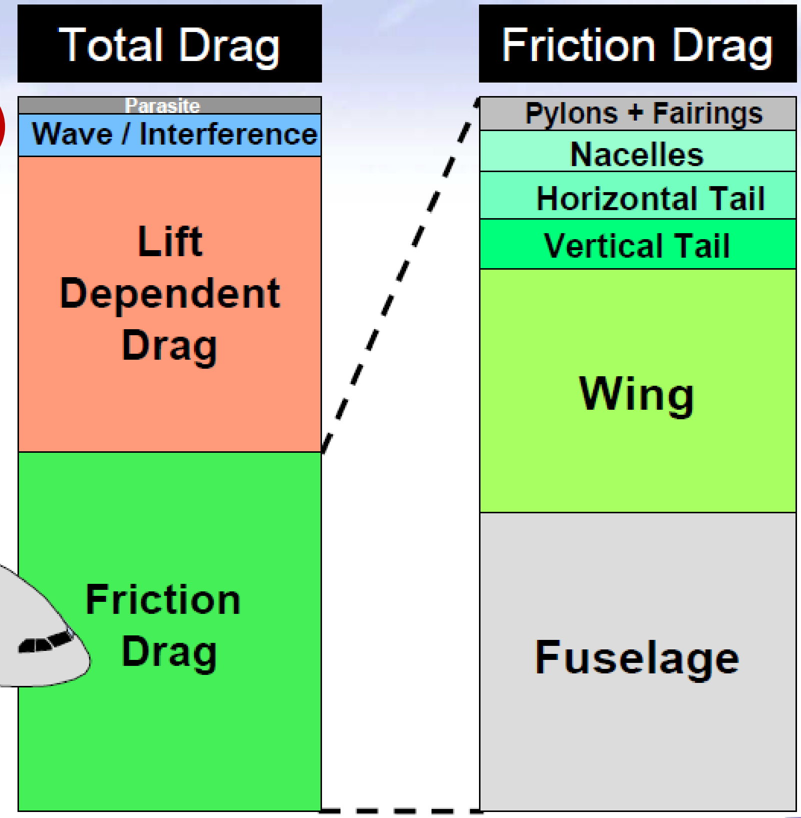

† = Based on a typical A320

What can we do today?

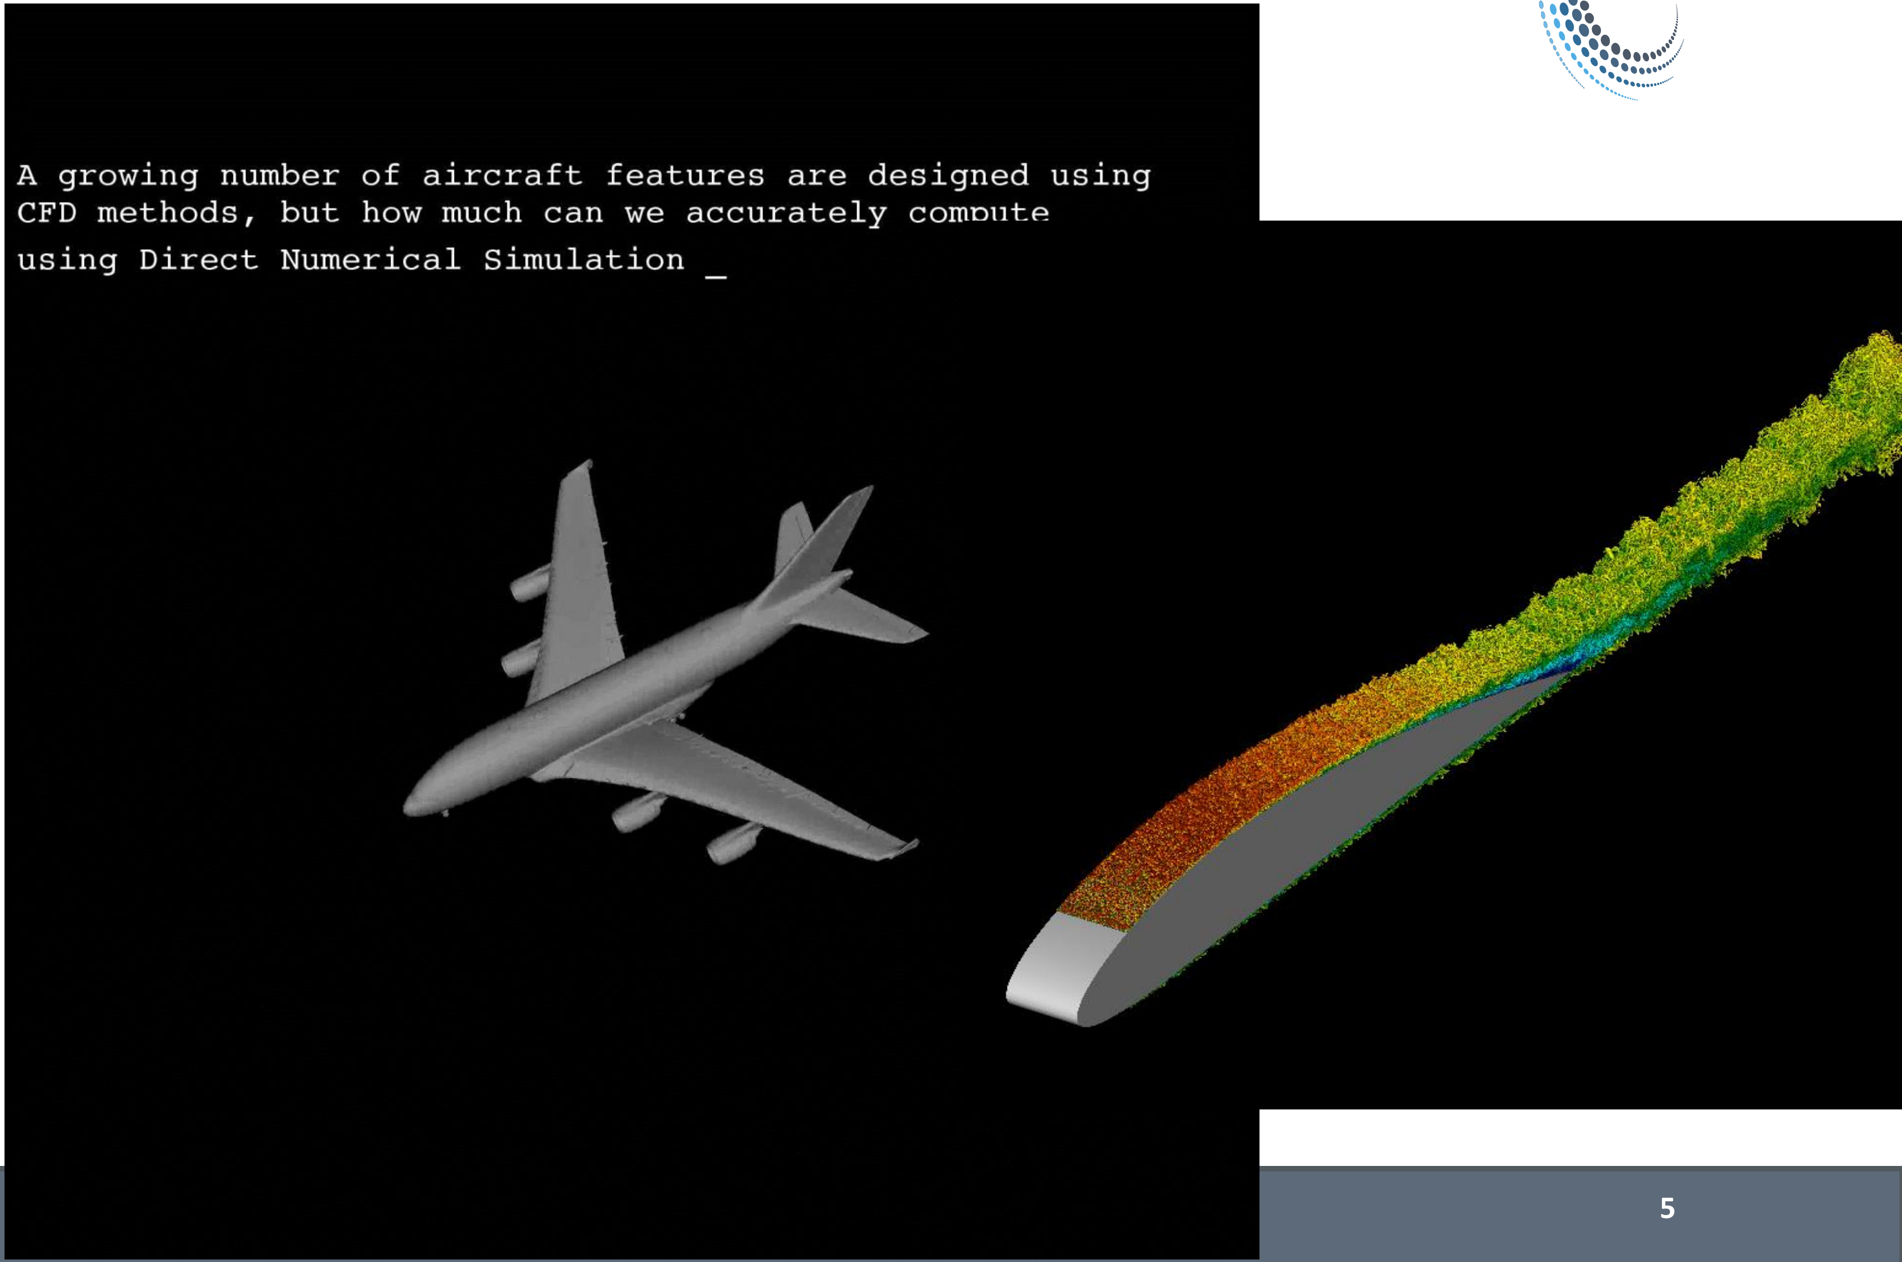

Navier – Stokes equations

ExaFLOW

Data from Mira (2013), million core hours

• Engineering/CFD	525	19%
• Subsurface flow & reactive transport	80	3%
• Combustion	100	4%
• Climate	280	10%
• Astrophysics	28	1%
• Supernovae	105	4%
	1118	**40%**

(fraction of Navier-Stokes based simulations on current supercomputers)

ExaFLOW Overview

The main goal of the project is to address current algorithmic bottlenecks to enable the use of accurate CFD codes for problems of practical engineering interest. The focus will be on different simulation aspects including:

- accurate **error control** and **adaptive mesh refinement** in complex computational domains,
- **solver efficiency** via mixed discontinuous and continuous Galerkin methods and appropriate optimised preconditioners,
- strategies to ensure **fault tolerance** and resilience,
- **heterogeneous modelling** to allow for different solution algorithms in different domain zones,
- parallel **input/output** for extreme data, employing novel data reduction algorithms (feature-based in-situ analysis),
- **energy awareness** of high-order methods,
- 4 different high-order codes.

ExaFLOW Partners

- KTH Stockholm, PDC and Mechanics (Coordinator)
- Imperial College, London, CFD
- University of Southampton, Aerodynamics
- University of Edinburgh, EPCC
- University of Stuttgart, HLRS and Aerodynamics
- EPF Lausanne, Mathematics

- McLaren Racing, UK
- Automotive Simulation Center Stuttgart

Different Codes

High-order accurate discretisations

- Potential for exascale applications (work vs communication, ”turbulence problem”)
- Clear advantages in the prediction of the physics
- 2x finite difference methods (compressible)
- 2x spectral element methods (incompressible)

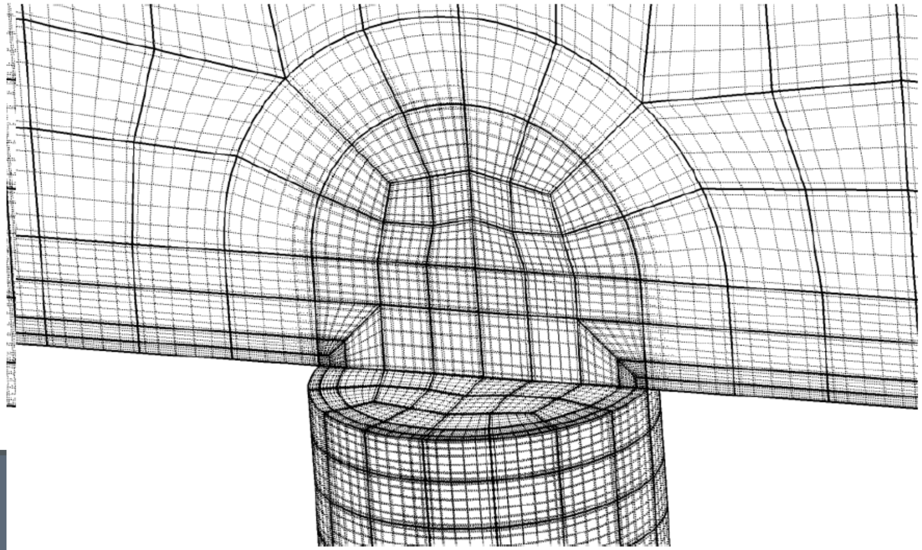

ExaFLOW Partners

- KTH Stockholm, PDC and Mechanics
- Imperial College, London, CFD
- University of Southampton, Aerodynamics
- University of Edinburgh, EPCC
- University of Stuttgart, HLRS and Aerodynamics
- EPF Lausanne, Mathematics

- McLaren Racing, UK
- Automotive Simulation Center Stuttgart

ExaFLOW Partners

- KTH Stockholm, PDC and Mechanics: **Nek5000**
- Imperial College, London, CFD: **Nektar++**
- University of Southampton, Aerodynamics: **OpenSBLI**
- University of Edinburgh, EPCC
- University of Stuttgart, HLRS and Aerodynamics: **NS3D**
- EPF Lausanne, Mathematics

- McLaren Racing, UK
- Automotive Simulation Center Stuttgart

Algorithmic Development

ExaFLOW

- 3 main objectives to develop the **mathematical and algorithmic knowledge** needed to tackle key objectives and enable exascale-level CFD software:
  - **Objective 1:** error control through adaption, heterogeneous modelling and resilience
  - **Objective 2:** strong scaling at exascale
  - **Objective 3:** techniques for I/O at exascale

WP1: 6 month progress highlights

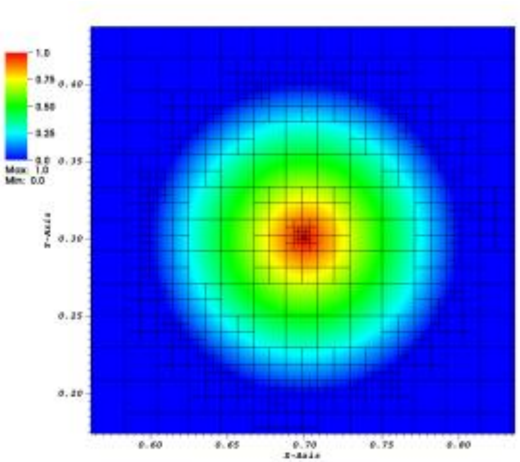

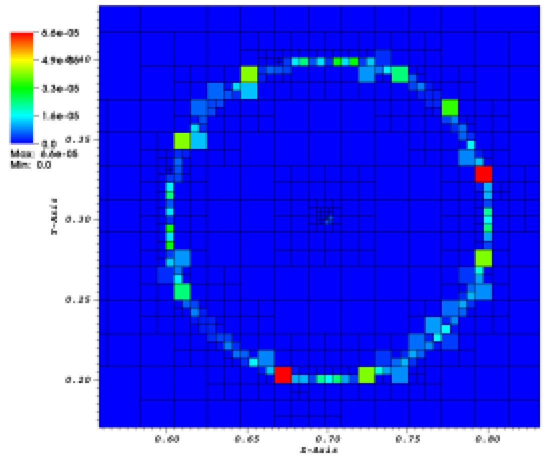

Working implementation of a spectral error estimator to drive a *p* and *h*-adaptive process for Objective 1 **(KTH)**

Initial formulation of hybrid CG-HDG on one node undergoing initial testing for Objective 2, aim to reduce communication bottleneck **(ICL)**

Preprint of investigation into resilience measures for hard/soft errors during runtime for Objective 1 **(EPFL)**

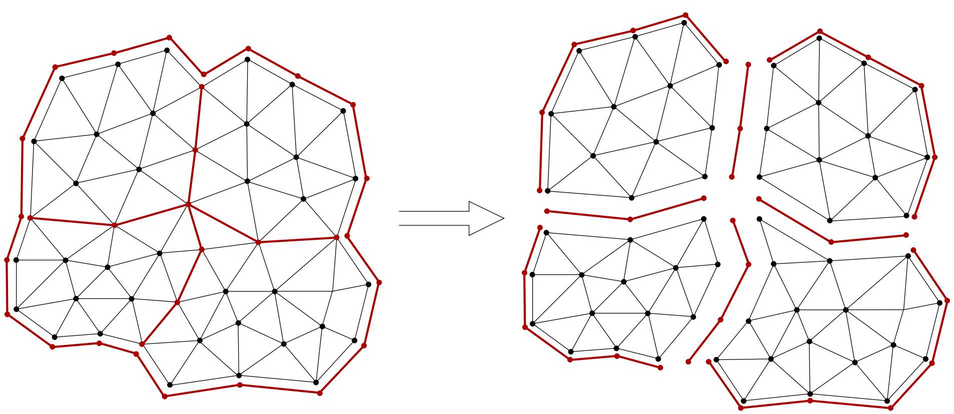

WP3: Initial Scaling – Nek5000

Characteristics

	Syst. arch.	Core arch.	# cores	Cores/node	Topology
Mira	IBM BG/Q	PowerPC A2	786,432	16	3D Torus
Titan	Cray XK7	AMD Opteron	299,008	16	5D Torus
Beskow	Cray XC40	Intel Haswell	53,632	32	DragonFly

Latency and inverse bandwidth

	α* (μs)	β* (μs/wd)	t_a (μs)	α	β
Mira	4	5×10^{-3}	1.1×10^{-3}	3600	5
Titan	2.25	1.42×10^{-3}	0.65×10^{-3}	3500	2.2
Beskow	2.55	0.825×10^{-3}	0.15×10^{-3}	17000	5.5

Linear model for communication (m = number of 64 bits words) :

$$t_c = (\alpha + \beta\, m)\, t_a$$

Hardware and instrumentation

- **Ping-pong test :**
  - Mira : very predictable,
  - Titan : random peaks,
  - Beskow : high noise,
  - Latency is the limiting parameter.

- **Code instrumentation :**
  - Hardware Performance Monitor for Mira (BG),
  - CrayPat sampling for Beskow (Cray) and Titan.

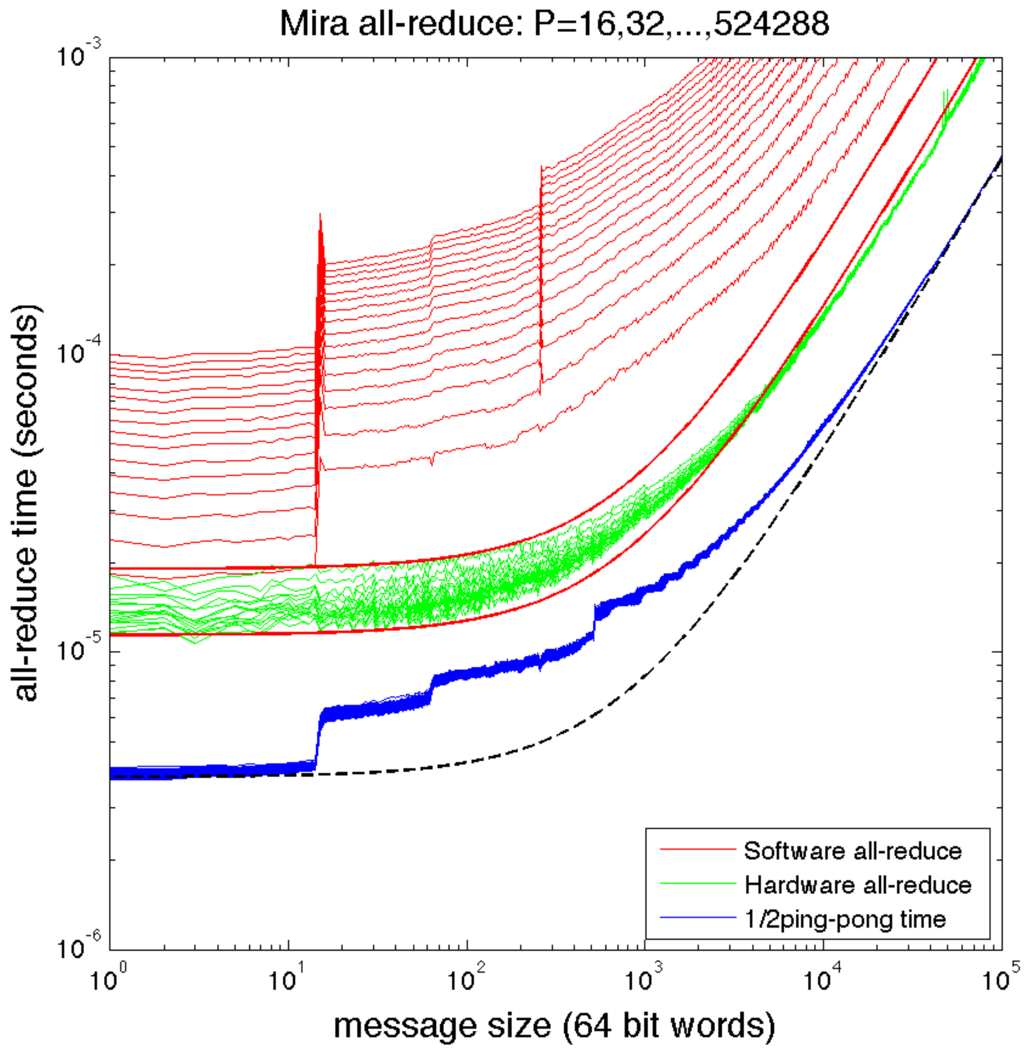

Scaling on Mira - Re_τ = 550

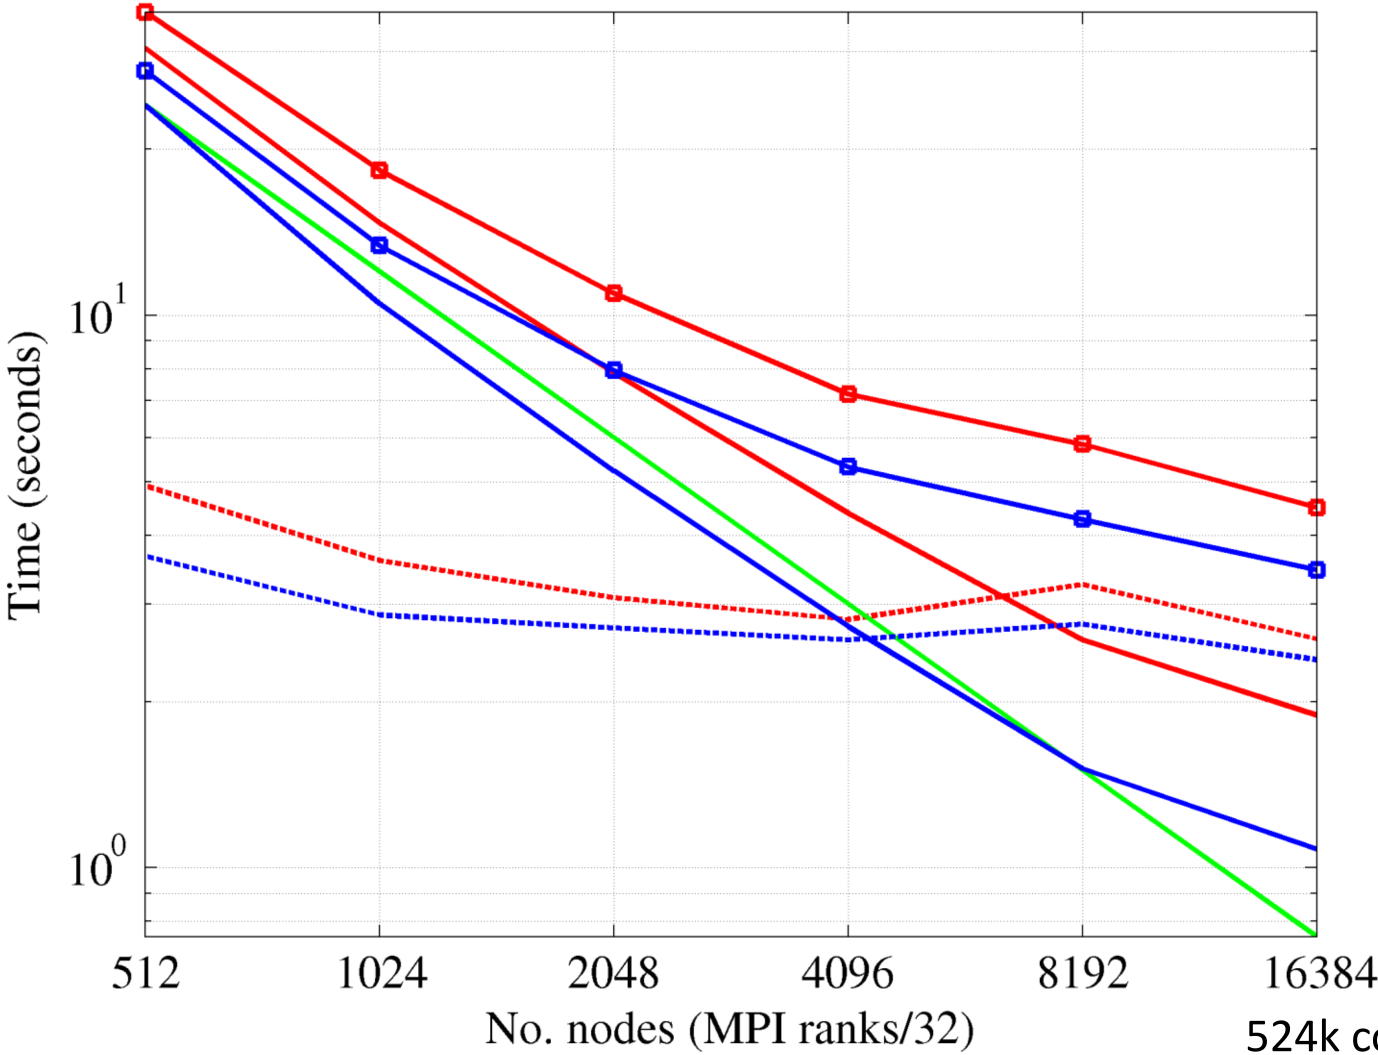

524k cores

AMG
XXT
Dashed: comm.
Solid: comp.
Square: total
Linear scaling

Adaptive Grid Refinement

Convected cone problem

- Introduced by Gottlieb and Orszag
- Passive scalar transport problem in which unit-height cone with a base-radius of 0.1 subjected to plane rotation
- 3-dimensional adaptation
  - sphere-shape cone - strong scaling; 33864 elements
  - cylinder-shape cone - weak scaling; 117192 elements at 2048 cores
- Cone evolution according to energy equation in Nek5000
- We follow advected features in the flow (the cone); refinement every 50 steps

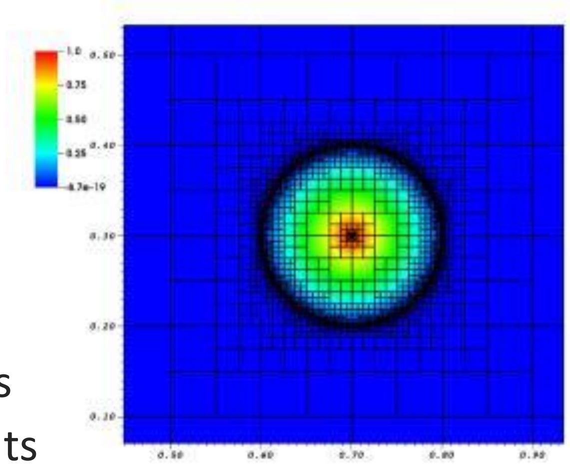

Passive scalar profile

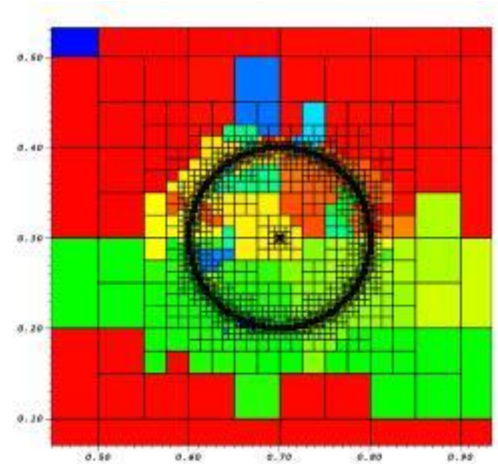

Grid partitioning

Power and energy profiling to build baselines

- Building power and energy profiles of Nektar++ and Nek5000 to establish baseline upon which to measure improvements.
- Using full test cases to exercise complete code.
- Energy-to-solution and time-to-solution as a function of processor frequency (left); looking for optimal trade-offs.
- Variation seen between otherwise identical runs of the same code and test-case (right); looking to quantify uncertainty in measurements.

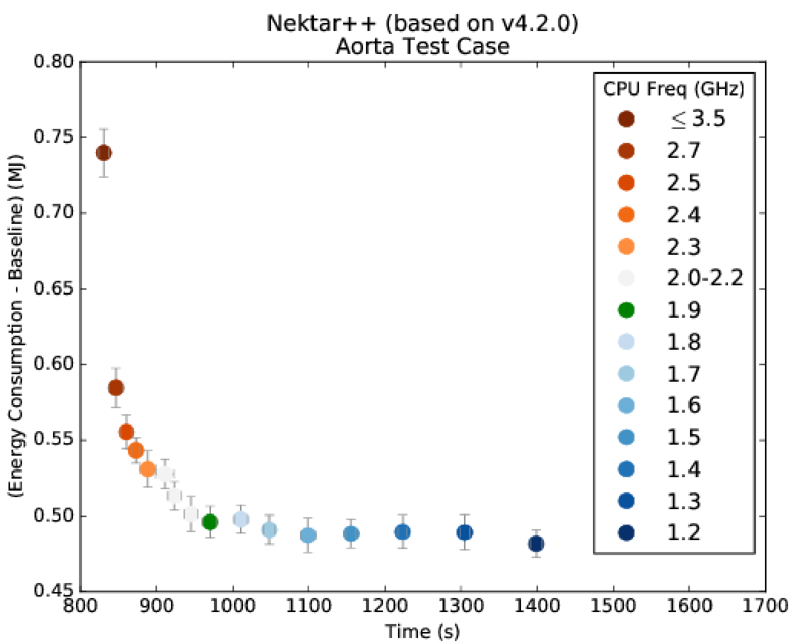

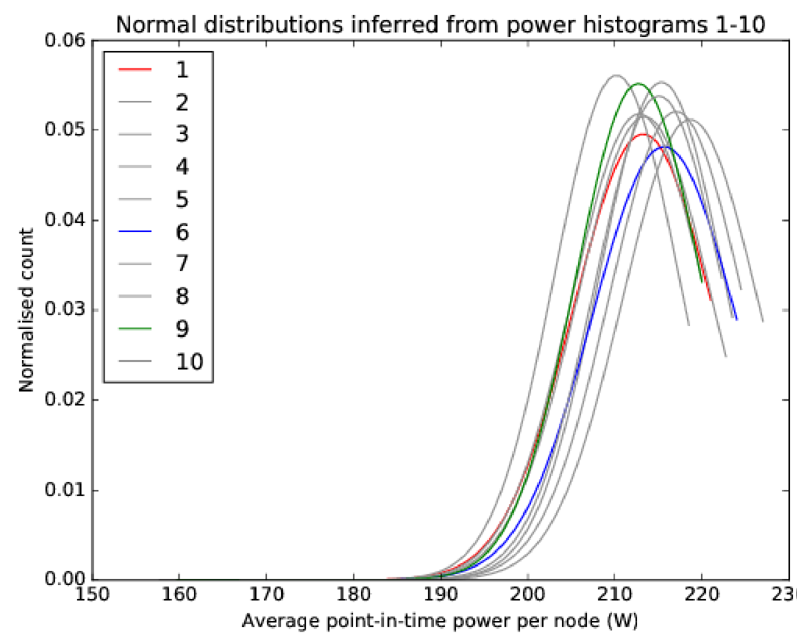

Adding modern parallel I/O

ExaFLOW

- Code bottleneck found in I/O.
  - Slows checkpoint / restart and final results writes.
- Previously multiple files, per process writing.
- Now using HDF5 atop MPI-IO for parallel I/O to single file.
  - Non-trivial mapping of elements on each process to file.
  - Allows checkpoint on X nodes, restart on Z nodes; X!=Z.
  - Fit problem to available resources, even when solution requires >1 jobs of different lengths/widths.
- Builds on top of work done in prior project, now ready to merge into trunk.
- Provides implementation useful to other project codes doing similar I/O with non-trivial mappings.
- Runtime compression to reduce data (POD, DMD etc.)

Pilot Cases

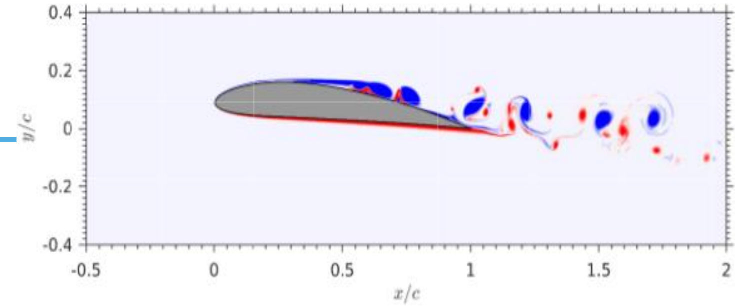

Five computationally-demanding use cases suitable for demonstrating the need for exascale capabilities have been created (Deliverable 3.1)

- **NACA4412 (compressible)** - Southampton
- **NACA4412 (incompressible)** - KTH
- **Jet in crossflow** - KTH/Stuttgart
- **Automotive/flow past a car** - ASCS
- **Imperial Front wing** – Imperial/McLaren

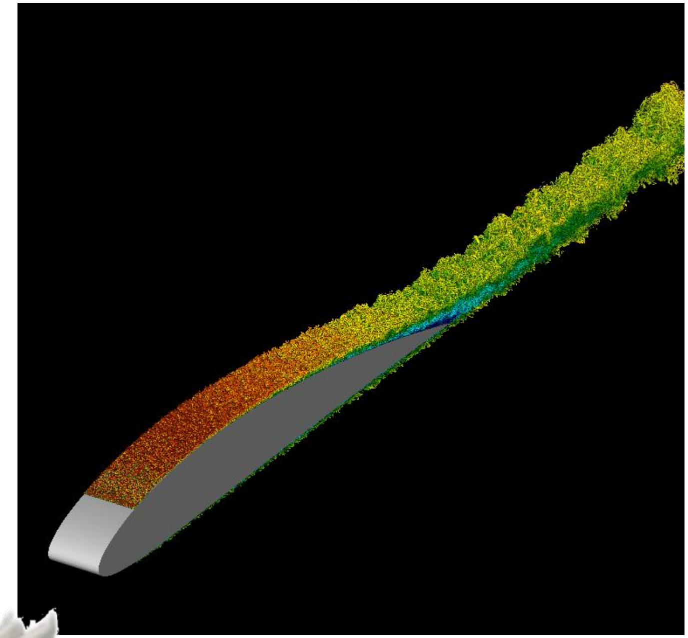

Quantitative measures to ensure correct flow physics is reproduced after code optimisations have been defined.

Computational requirements have been investigated and internal evaluation of use cases is under way.

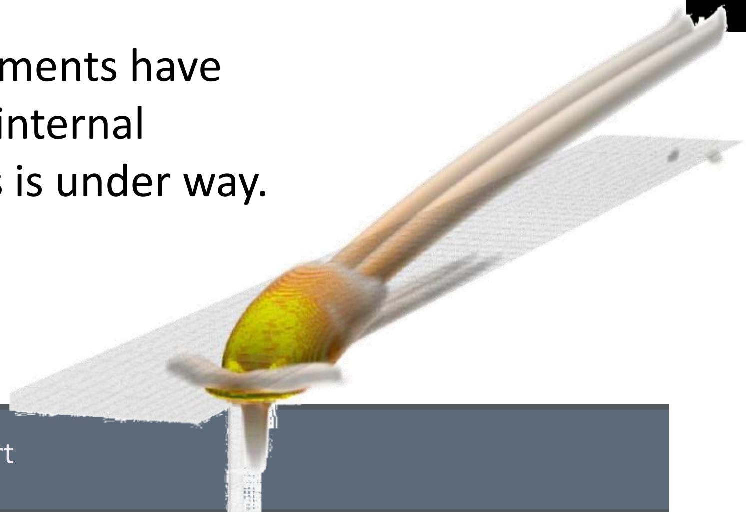

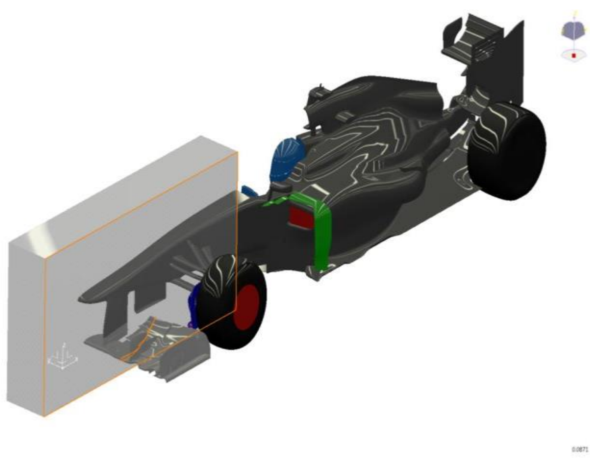

Conclusions and Outlook

- Fluid mechanics is a **prime example** for exascale
- ExaFLOW will address some of the issues when it comes to **practial applications**
  - error control and adaptive meshing for larger and more complex simulation domains; capable of dynamic remeshing if necessary.
  - Heterogeneous modelling
  - Resilience & fault tolerance
  - data handling, complex feature extraction (in-situ) and sharing of simulation data.
- Enhancing **community codes** (Nek5000, Nektar++, SBLI); Open-source development of all components