Turbulence on large computers

Dardel Inauguration

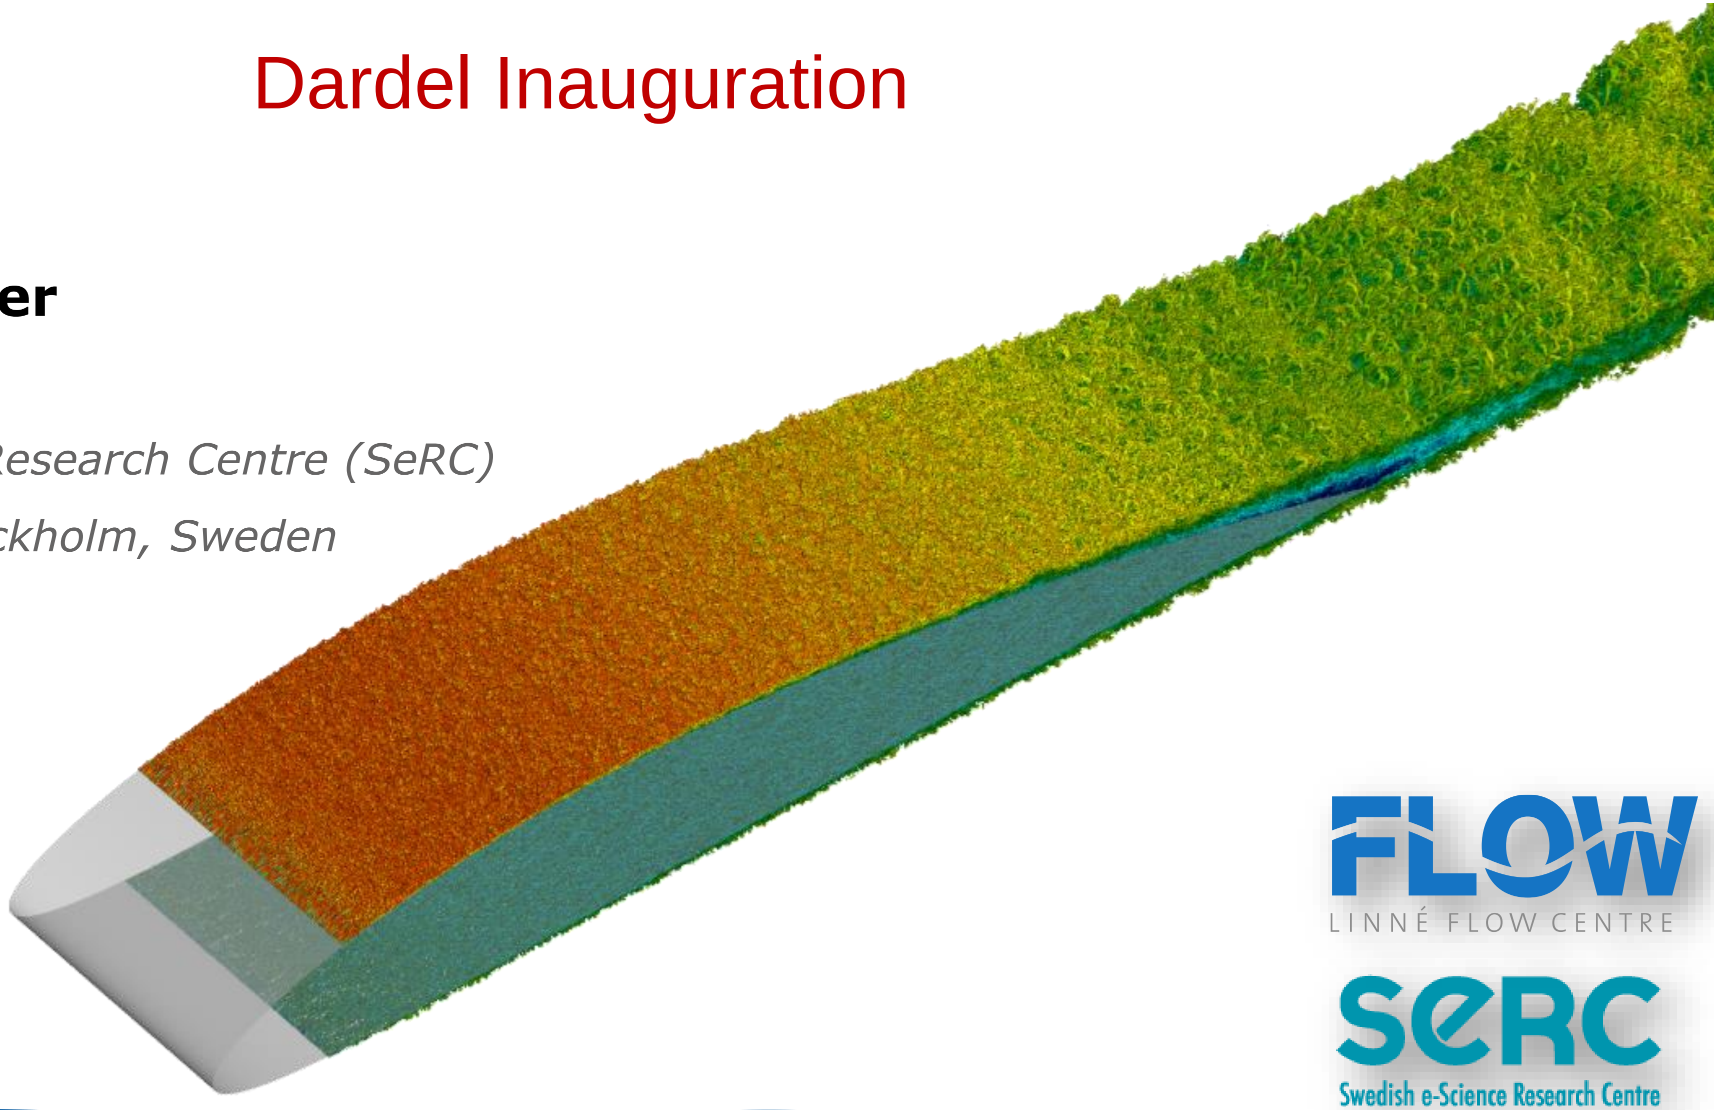

Philipp Schlatter

FLOW Centre and

Swedish e-Science Research Centre (SeRC)

KTH Mechanics, Stockholm, Sweden

PDC, January 25, 2022

Why turbulence?

Skin friction/drag reduction is the key for economically and ecologically more efficient transport

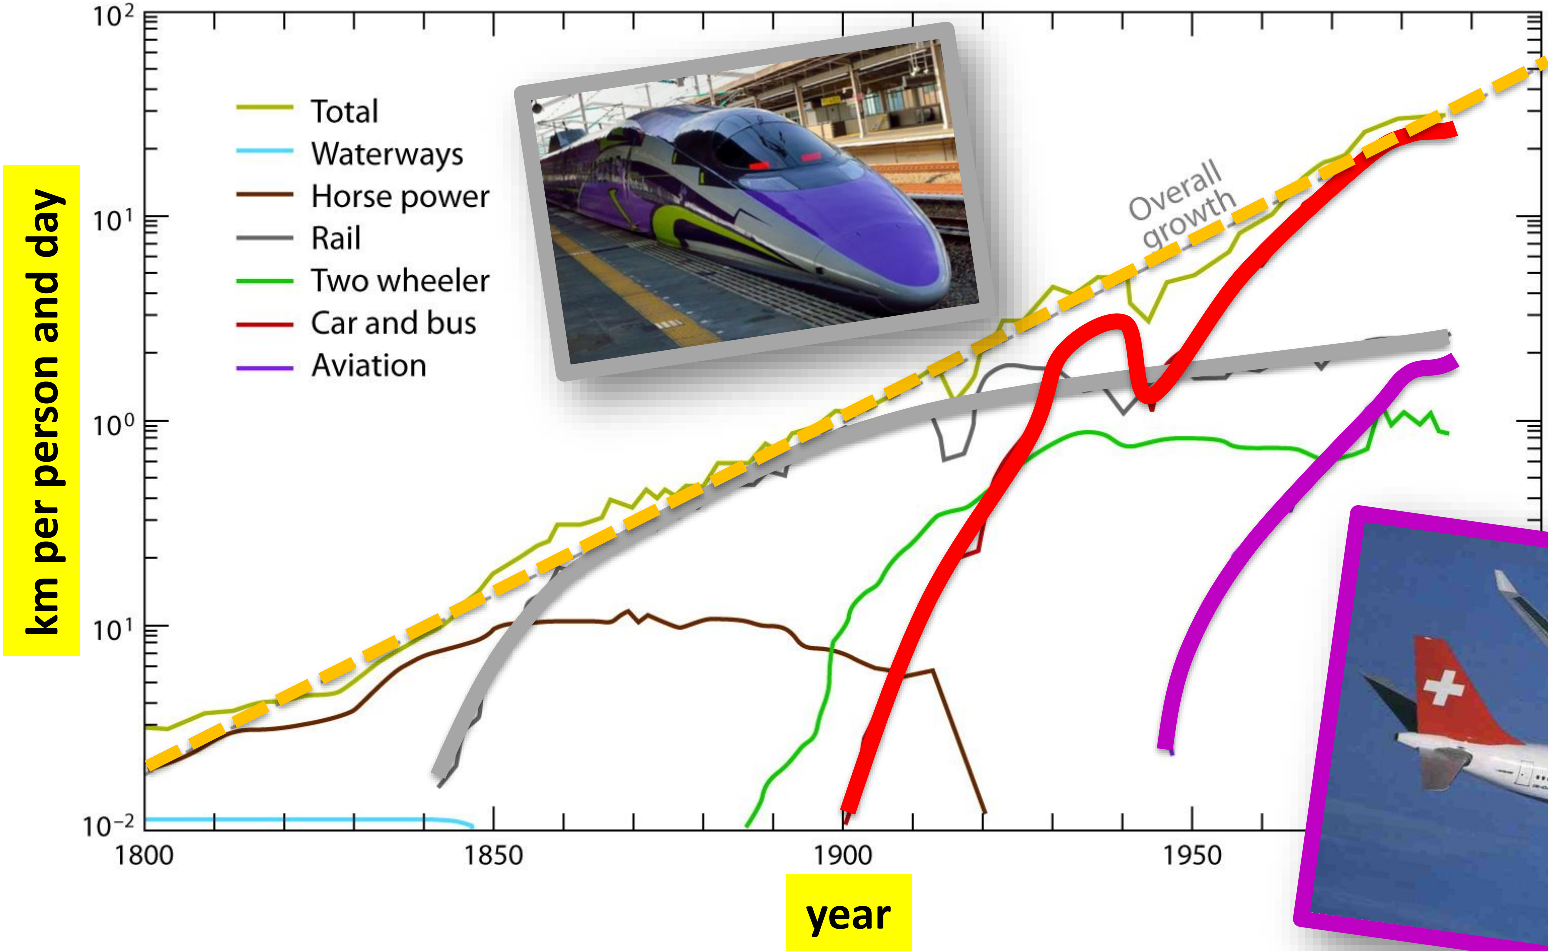

Banister D, et al. 2011.
Annu Rev. Environ. Resour. 36:247–70

Why are we here…?

- “When a sufficiently **advanced computer** becomes available, we believe it will **replace the wind tunnel** as the principal facility for providing aerodynamic flow simulations”
- “If past trends continue, such computer performance should **be available in the mid-1980s**…”

Chapman, D. R., Mark, H., Pirtle, M. W., “Computers vs. wind tunnels for aerodynamic flow simulations”, Astronautics & Aeronautics **13**(4):22-30, 1975 (NASA Ames)

The Airbus Challenge, David Hills, 2008

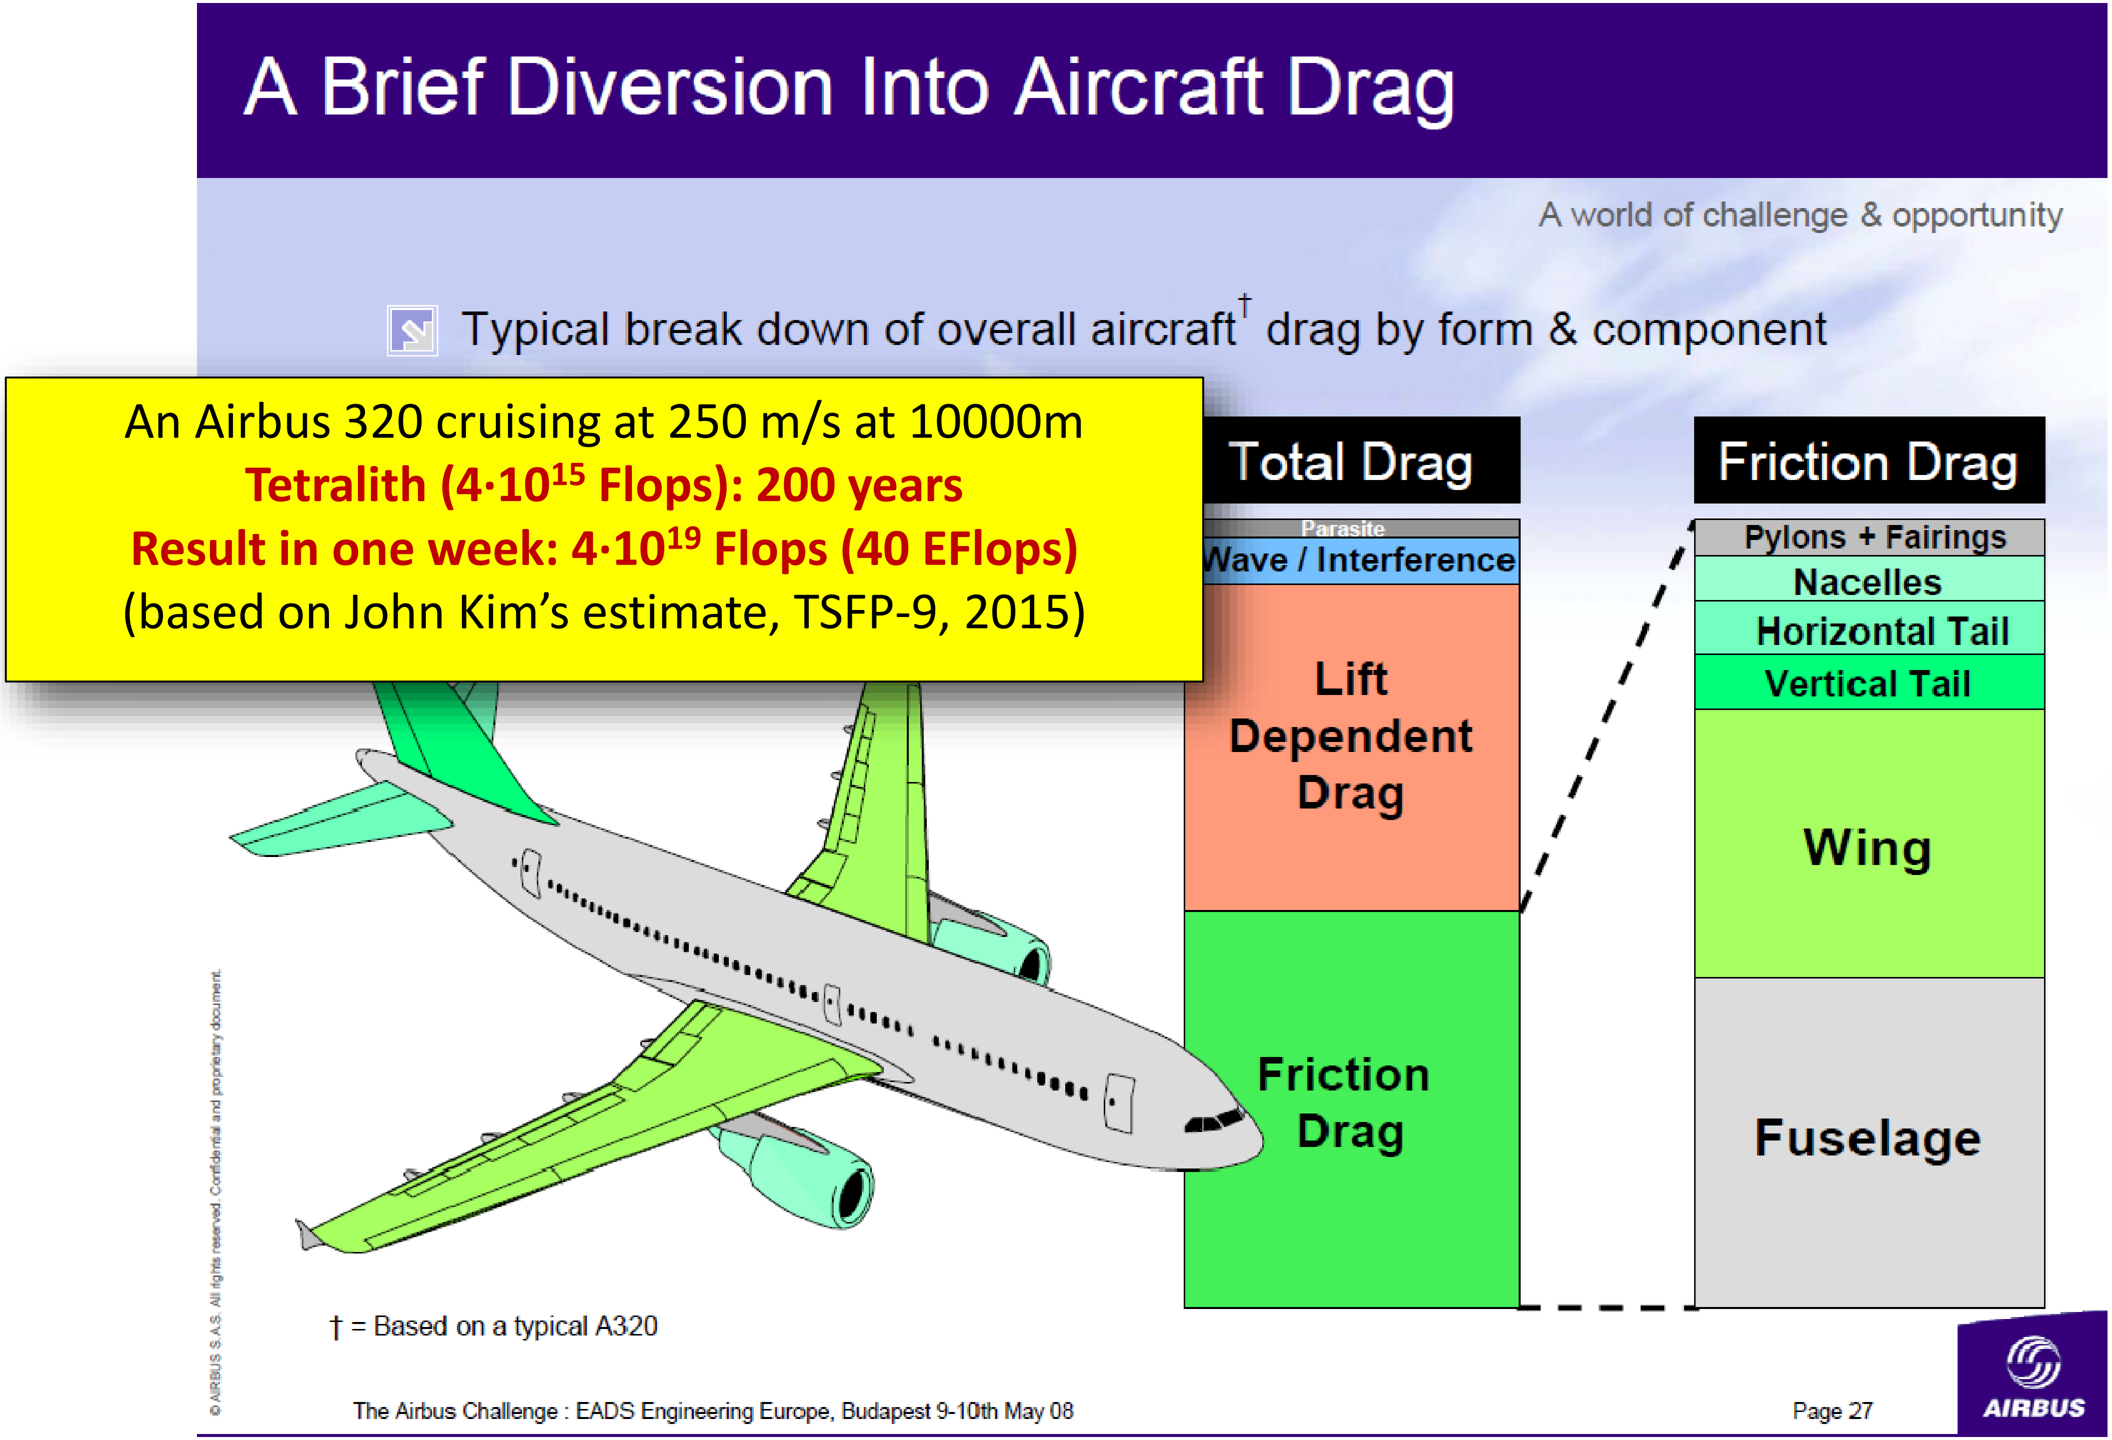

Humble beginnings 100 years ago…

- Lewis Fry Richardson (1881-1953)

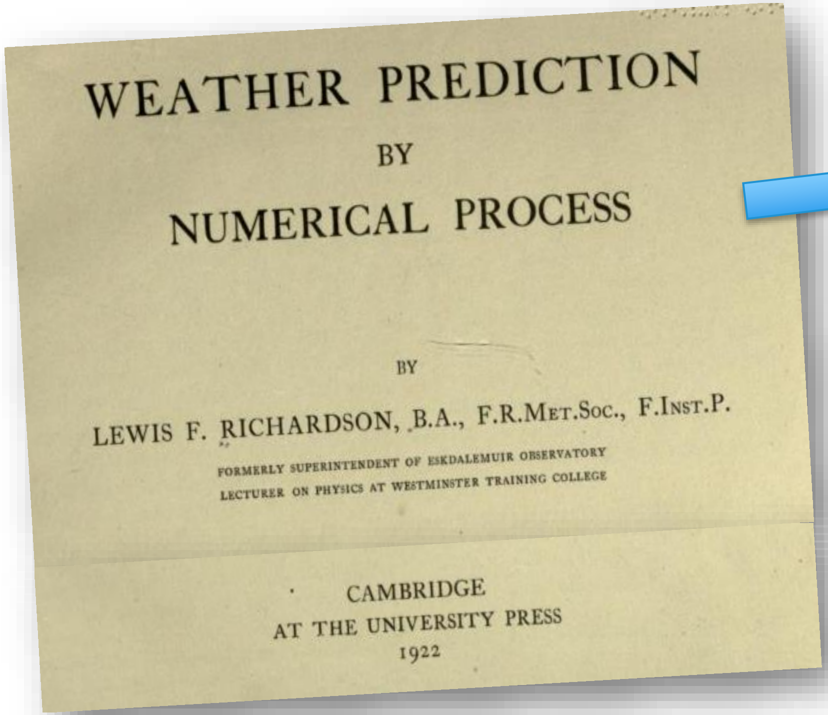
WEATHER PREDICTION
BY
NUMERICAL PROCESS

BY

LEWIS F. RICHARDSON, B.A., F.R.MET.SOC., F.INST.P.
FORMERLY SUPERINTENDENT OF ESKDALEMUIR OBSERVATORY
LECTURER ON PHYSICS AT WESTMINSTER TRAINING COLLEGE

CAMBRIDGE
AT THE UNIVERSITY PRESS
1922

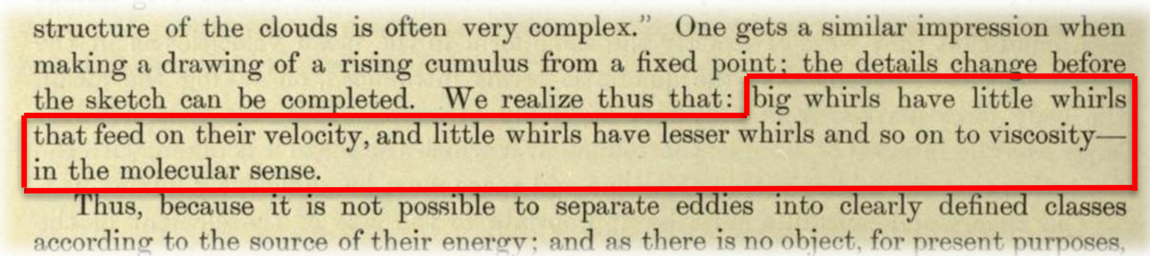
structure of the clouds is often very complex." One gets a similar impression when making a drawing of a rising cumulus from a fixed point; the details change before the sketch can be completed. We realize thus that: big whirls have little whirls that feed on their velocity, and little whirls have lesser whirls and so on to viscosity—in the molecular sense.

Thus, because it is not possible to separate eddies into clearly defined classes according to the source of their energy; and as there is no object, for present purposes,

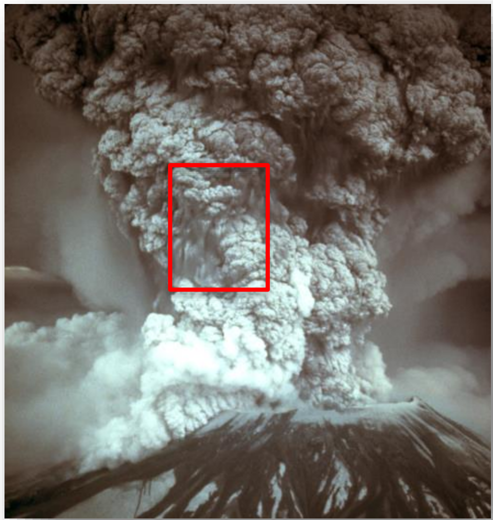

play on Augustus de Morgan's famous paraphrasing[3] of Jonathan Swift

Humble beginnings 100 years ago…

- Lewis Fry Richardson (1881-1953)

WEATHER PREDICTION

BY

NUMERICAL PROCESS

BY

LEWIS F. RICHARDSON, B.A., F.R.MET.SOC., F.INST.P.

FORMERLY SUPERINTENDENT OF ESKDALEMUIR OBSERVATORY

LECTURER ON PHYSICS AT WESTMINSTER TRAINING COLLEGE

CAMBRIDGE

AT THE UNIVERSITY PRESS

1922

structure of the clouds is often very complex." One gets a similar impression when making a drawing of a rising cumulus from a fixed point; the details change before the sketch can be completed. We realize thus that: big whirls have little whirls that feed on their velocity, and little whirls have lesser whirls and so on to viscosity—in the molecular sense.

Thus, because it is not possible to separate eddies into clearly defined classes according to the source of their energy; and as there is no object, for present purposes,

"First simulations" 1920: Eight hours weather prediction in 6 weeks, using **2000 *human* computers**

→ "Forecast-Factory"

The Weather Forecasting Factory by Stephen Conlin (1986)

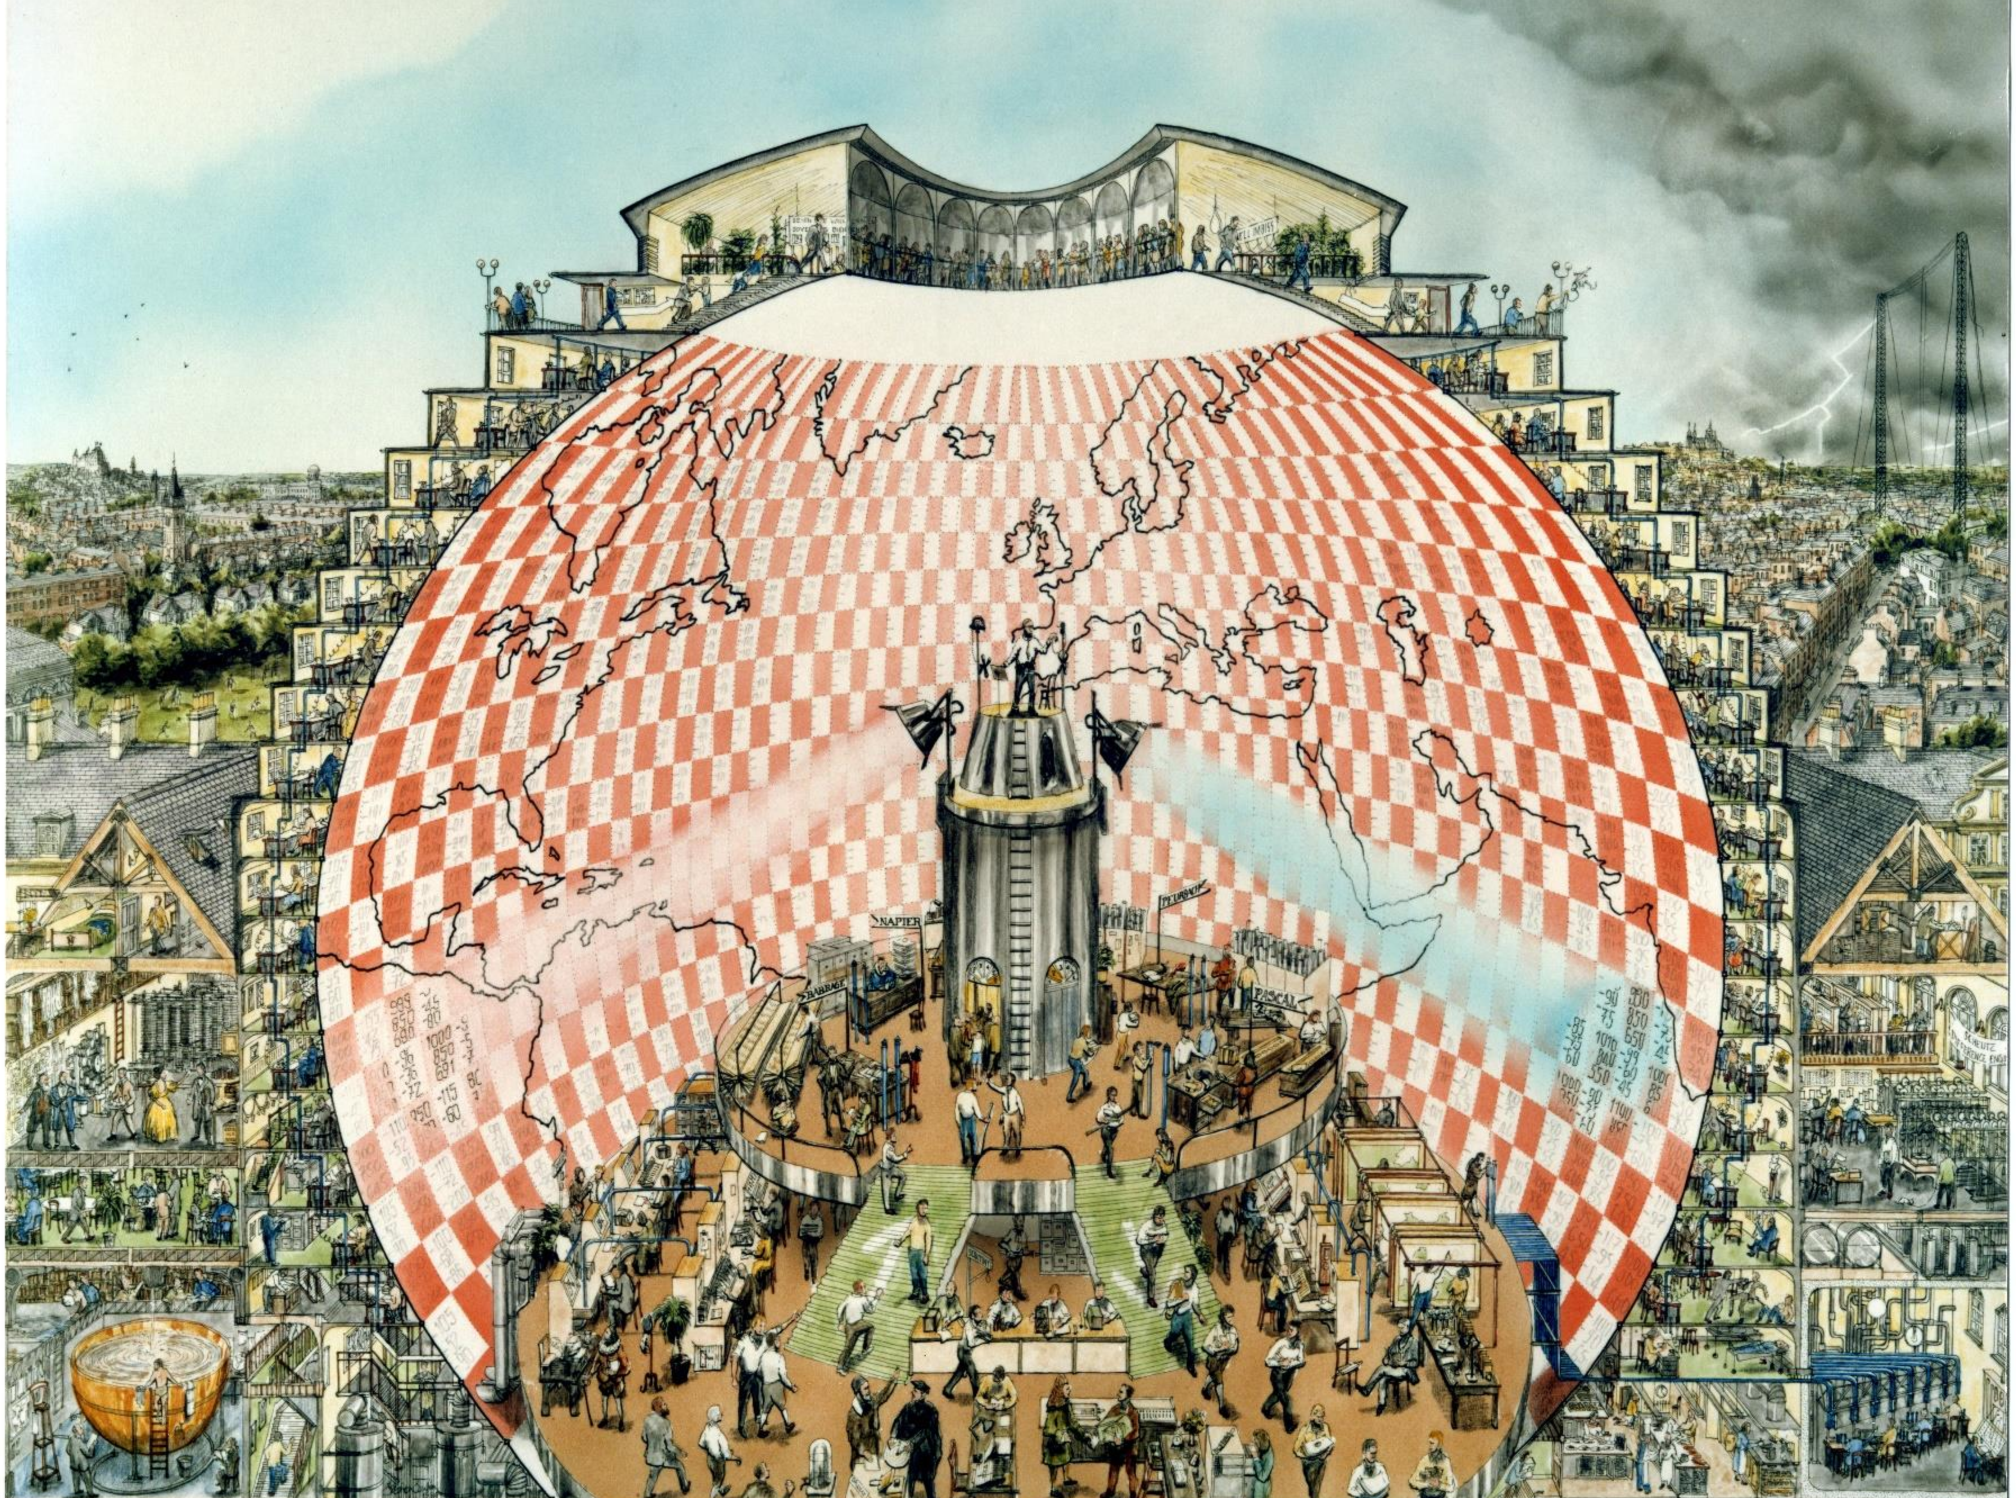

Humble beginnings…

Von Kármán vortex street @ KTH TR8

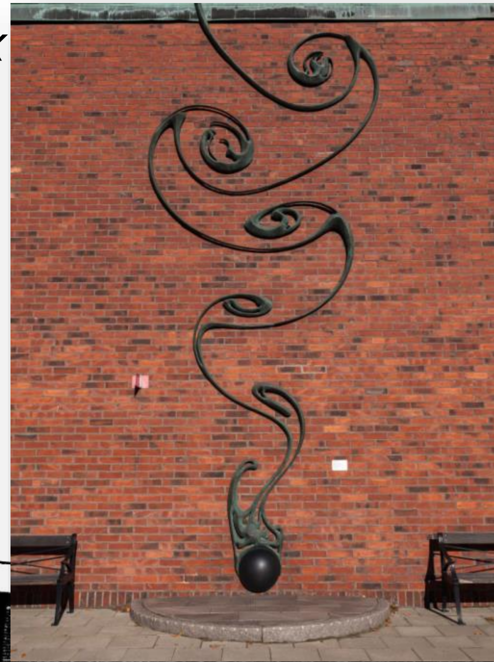

- Development of numerics (1940ies, 1950ies)
- Low-*Re* cylinder wakes by Thom (1933), Kawaguti (1953) and Fromm & Harlow (1963), Los Alomos

THE PHYSICS OF FLUIDS VOLUME 6, NUMBER 7 JULY 1963

Numerical Solution of the Problem of Vortex Street Development

JACOB E. FROMM AND FRANCIS H. HARLOW
University of California, Los Alamos Scientific Laboratory, Los Alamos, New Mexico
(Received 6 February 1963)

A method is described for the solution of time-dependent problems concerning the flow of viscous incompressible fluids in several space dimensions. The method is numerical, using a high-speed computer for the solution of a finite-difference approximation to the partial differential equations of motion. The application described here is to a study of the development of a vortex street behind a plate which has impulsively accelerated to constant speed in a channel of finite width; the Reynolds-number range investigated was $15 \leq R \leq 6000$. Particular attention was given to those features for which comparison could be made with experiments, namely, critical Reynolds number for vortex shedding, drag coefficient, Strouhal number, vortex configuration, and channel-wall effects. The nature of the early stages of flow-pattern development was also investigated.

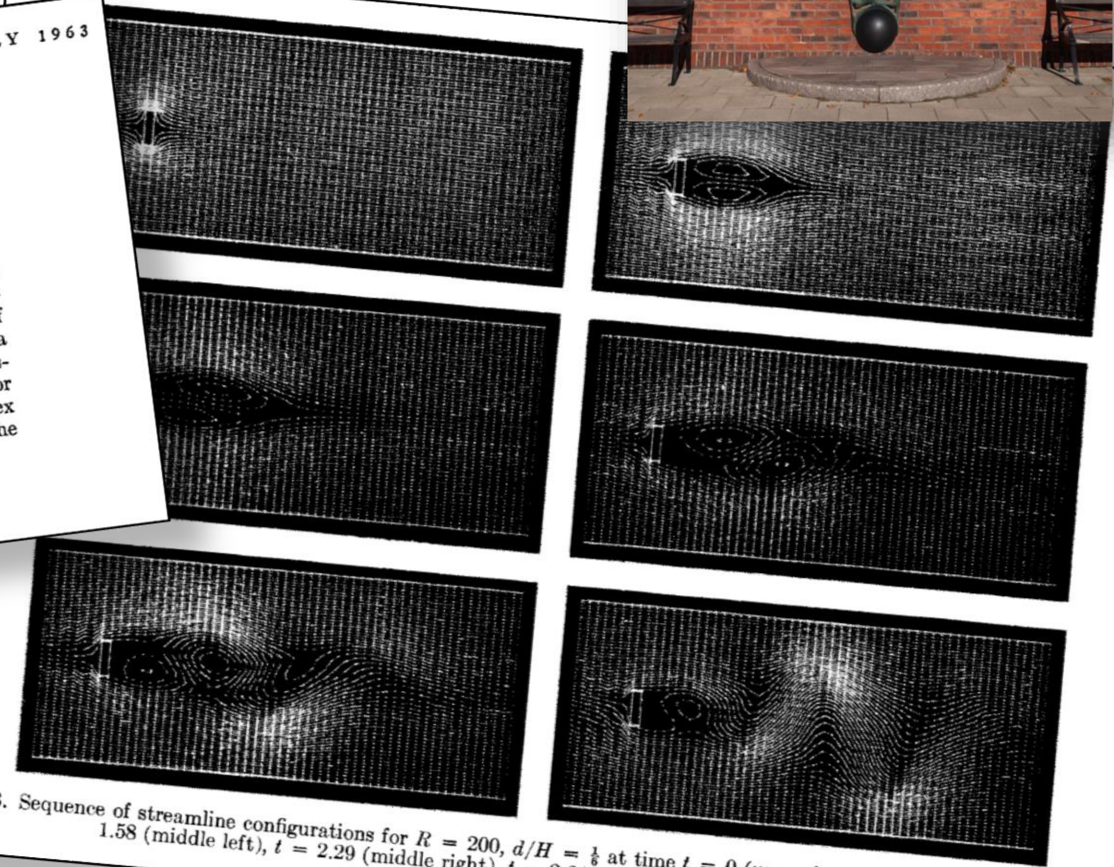

FIG. 3. Sequence of streamline configurations for $R = 200$, $d/H = \frac{1}{6}$ at time $t = 0$ (upper left), $t = 0.70$ (upper right), $t = 1.58$ (middle left), $t = 2.29$ (middle right), $t = 2.64$ (lower left), and $t = 3.16$ (lower right).

1440 grid points!

FLOW
LINNÉ FLOW CENTRE

Turbulent flow close to solid walls...

Ekman 2008-2013

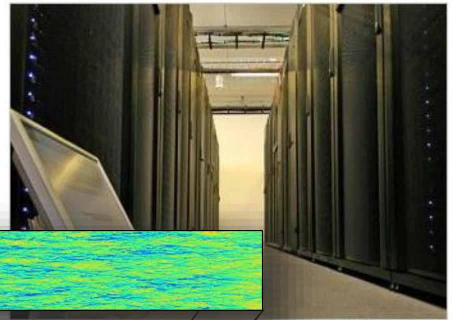

Turbulent flow close to solid walls...

Ekman 2008-2013

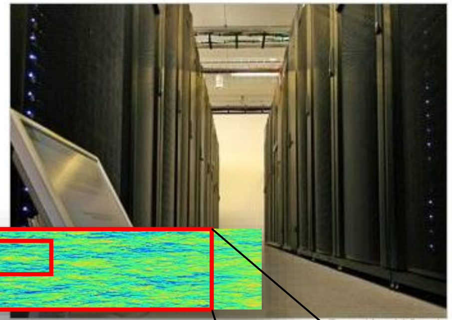

simulation result

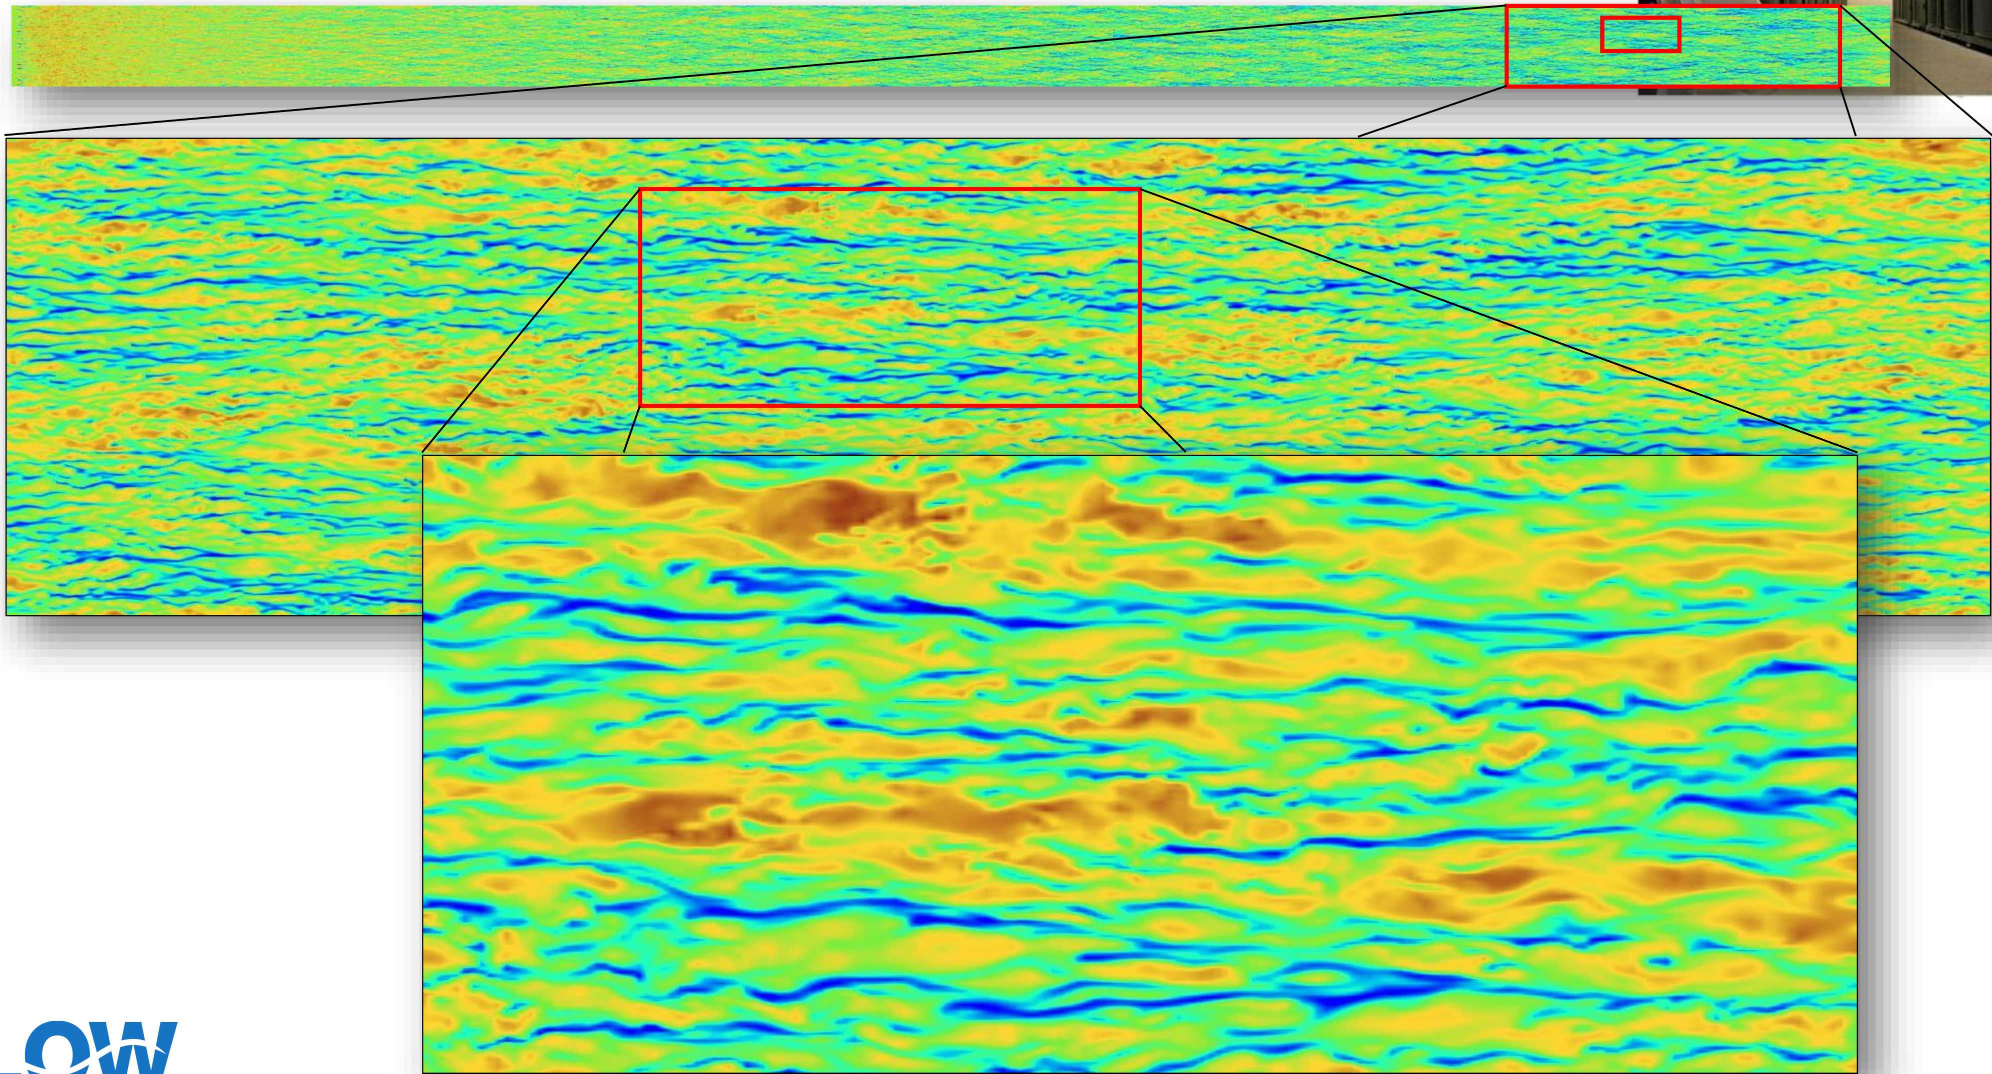

Turbulent flow close to s

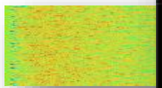

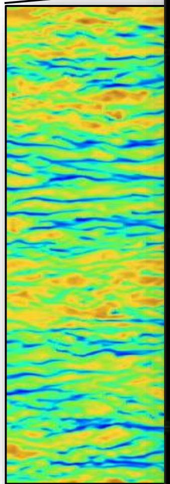

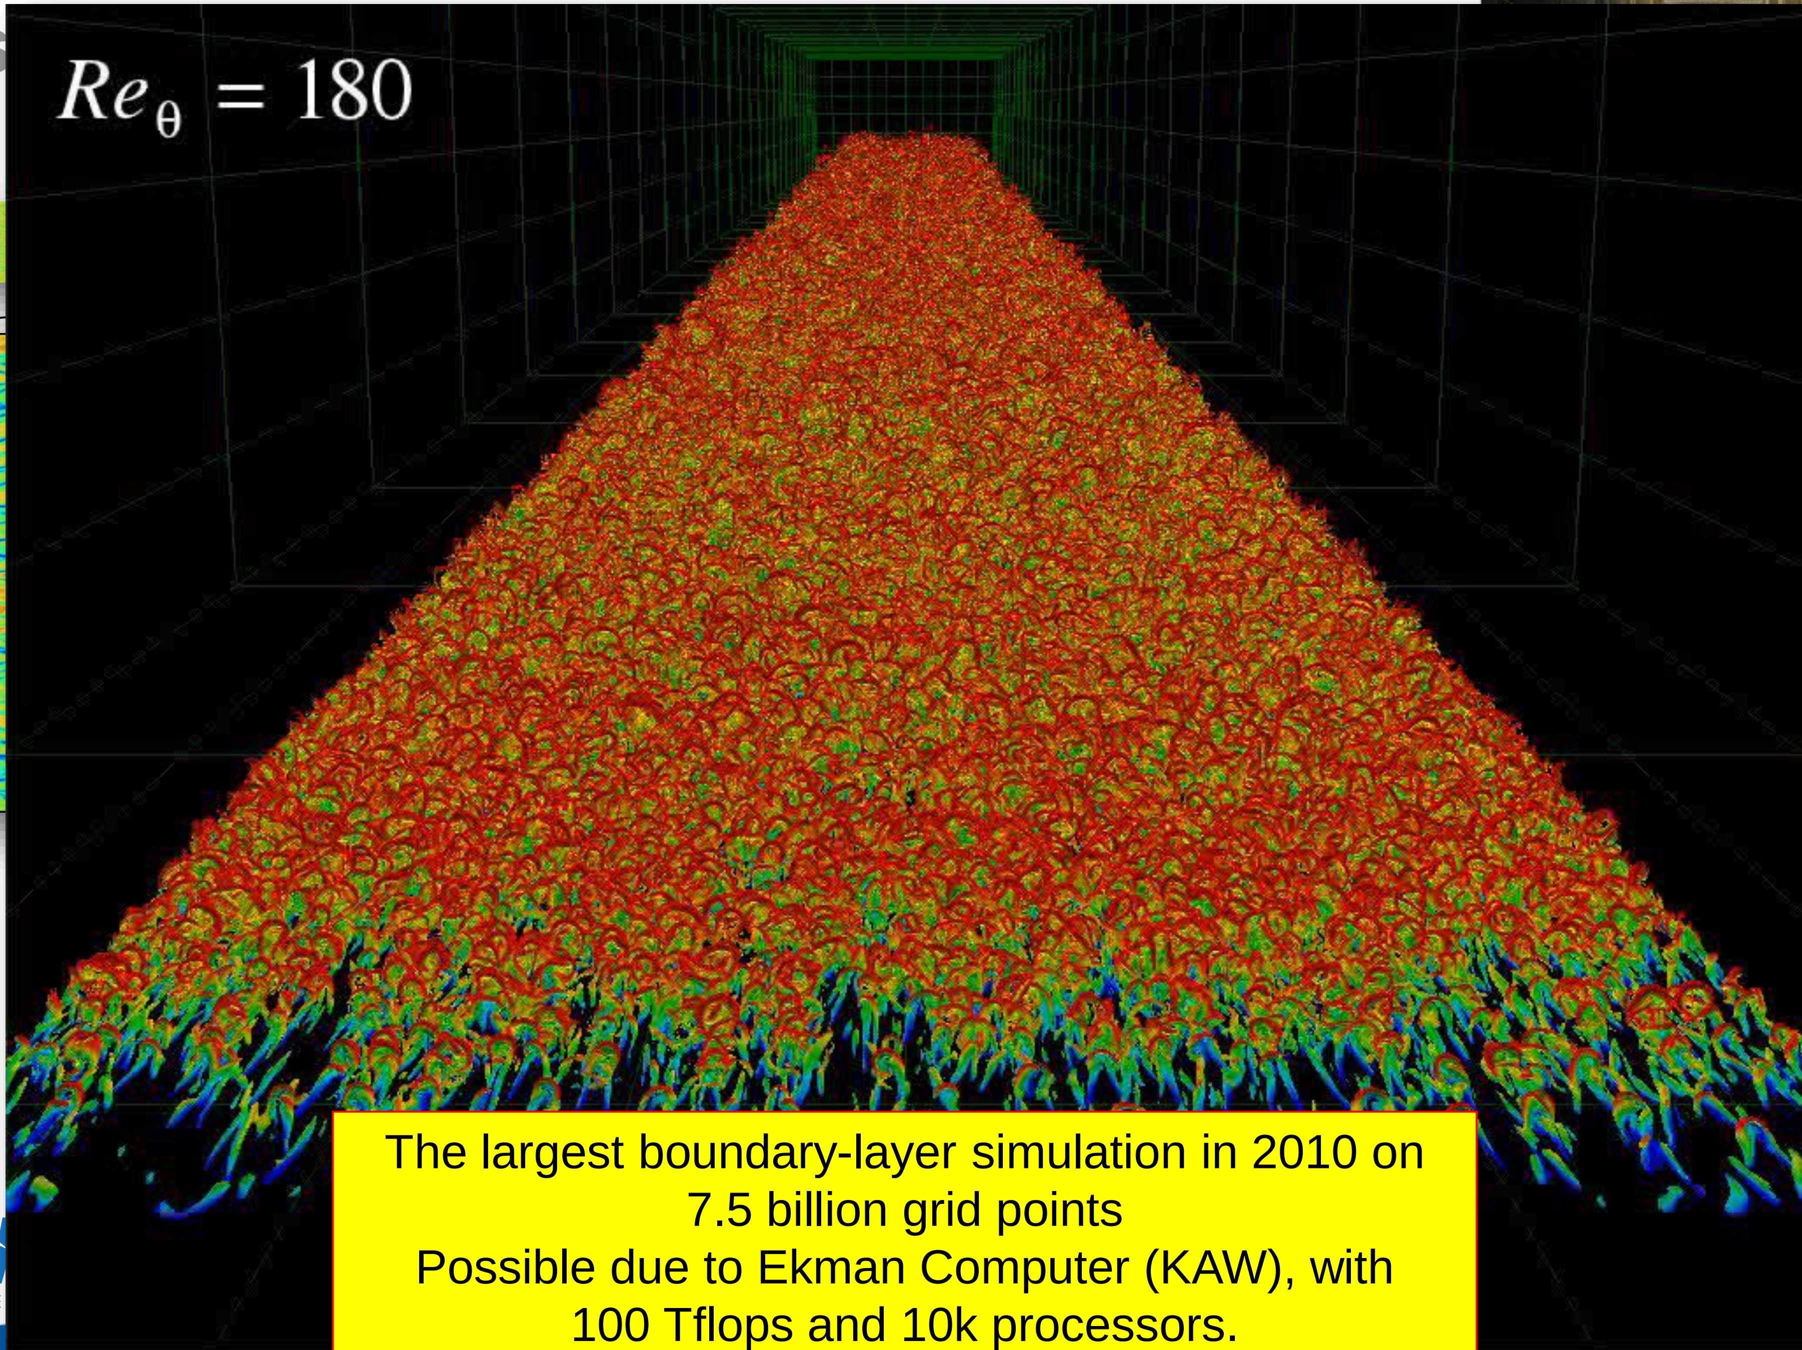

The largest boundary-layer simulation in 2010 on
7.5 billion grid points
Possible due to Ekman Computer (KAW), with
100 Tflops and 10k processors.

Ongoing: Turbulent Pipe Flow

Lindgren 2011-2015

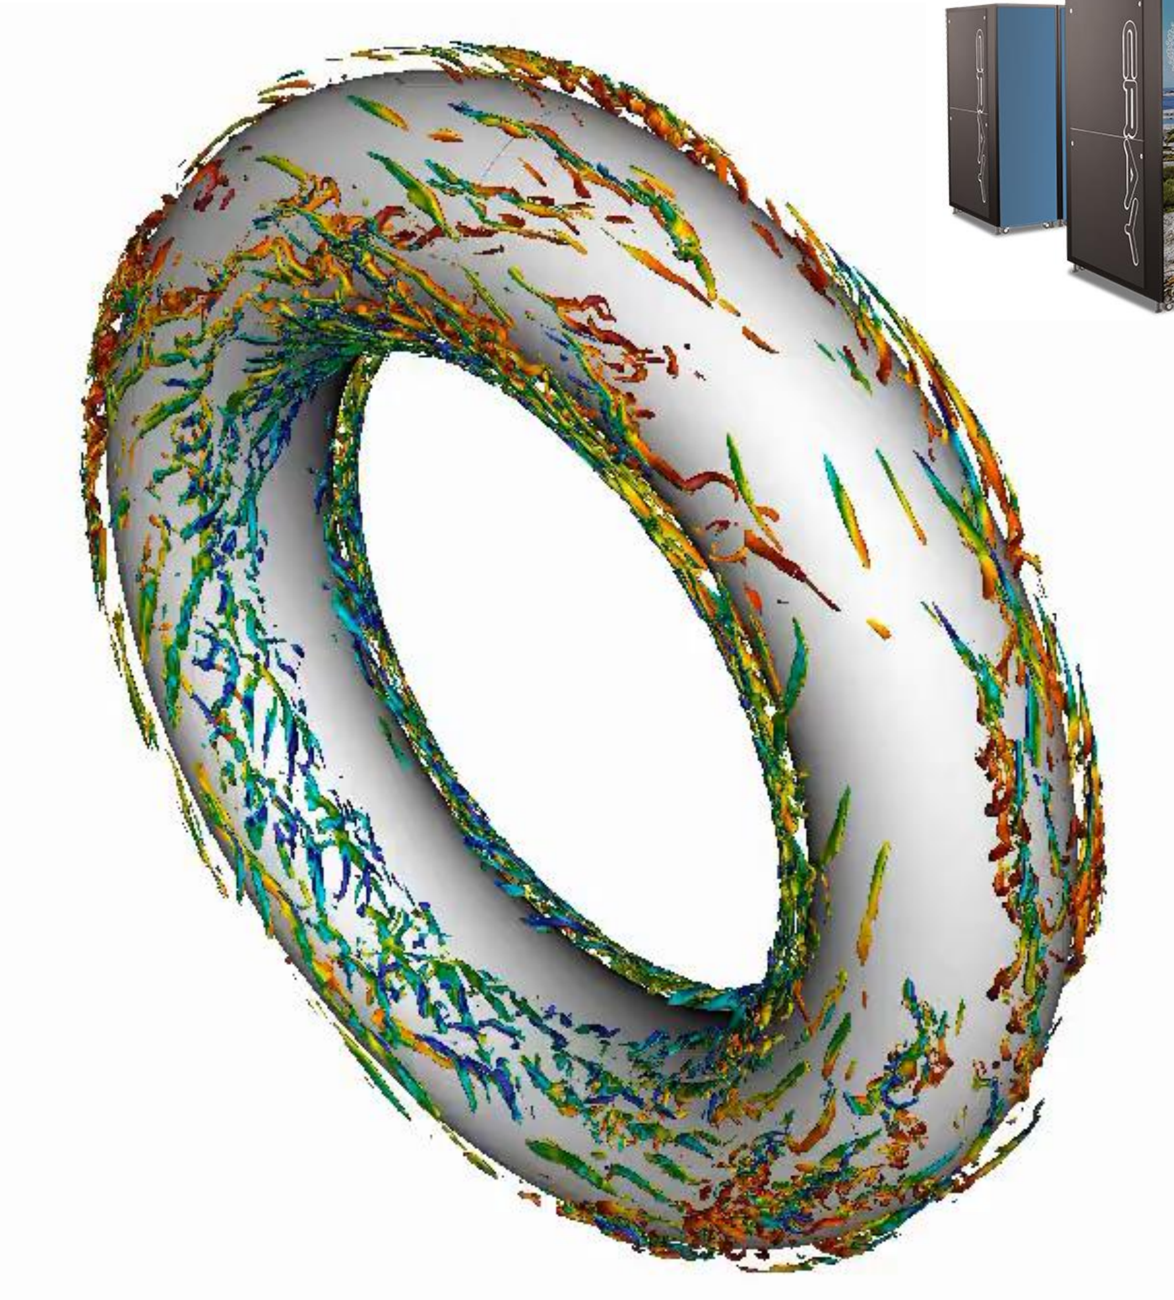

$$Re_\tau = 360$$

$$\kappa = 0.3$$

Ref. Noorani *et al.* 2015

DNS of flow around a NACA4412 wing section; Re_c=400 000 and AoA=5°

Beskow 2015-2021

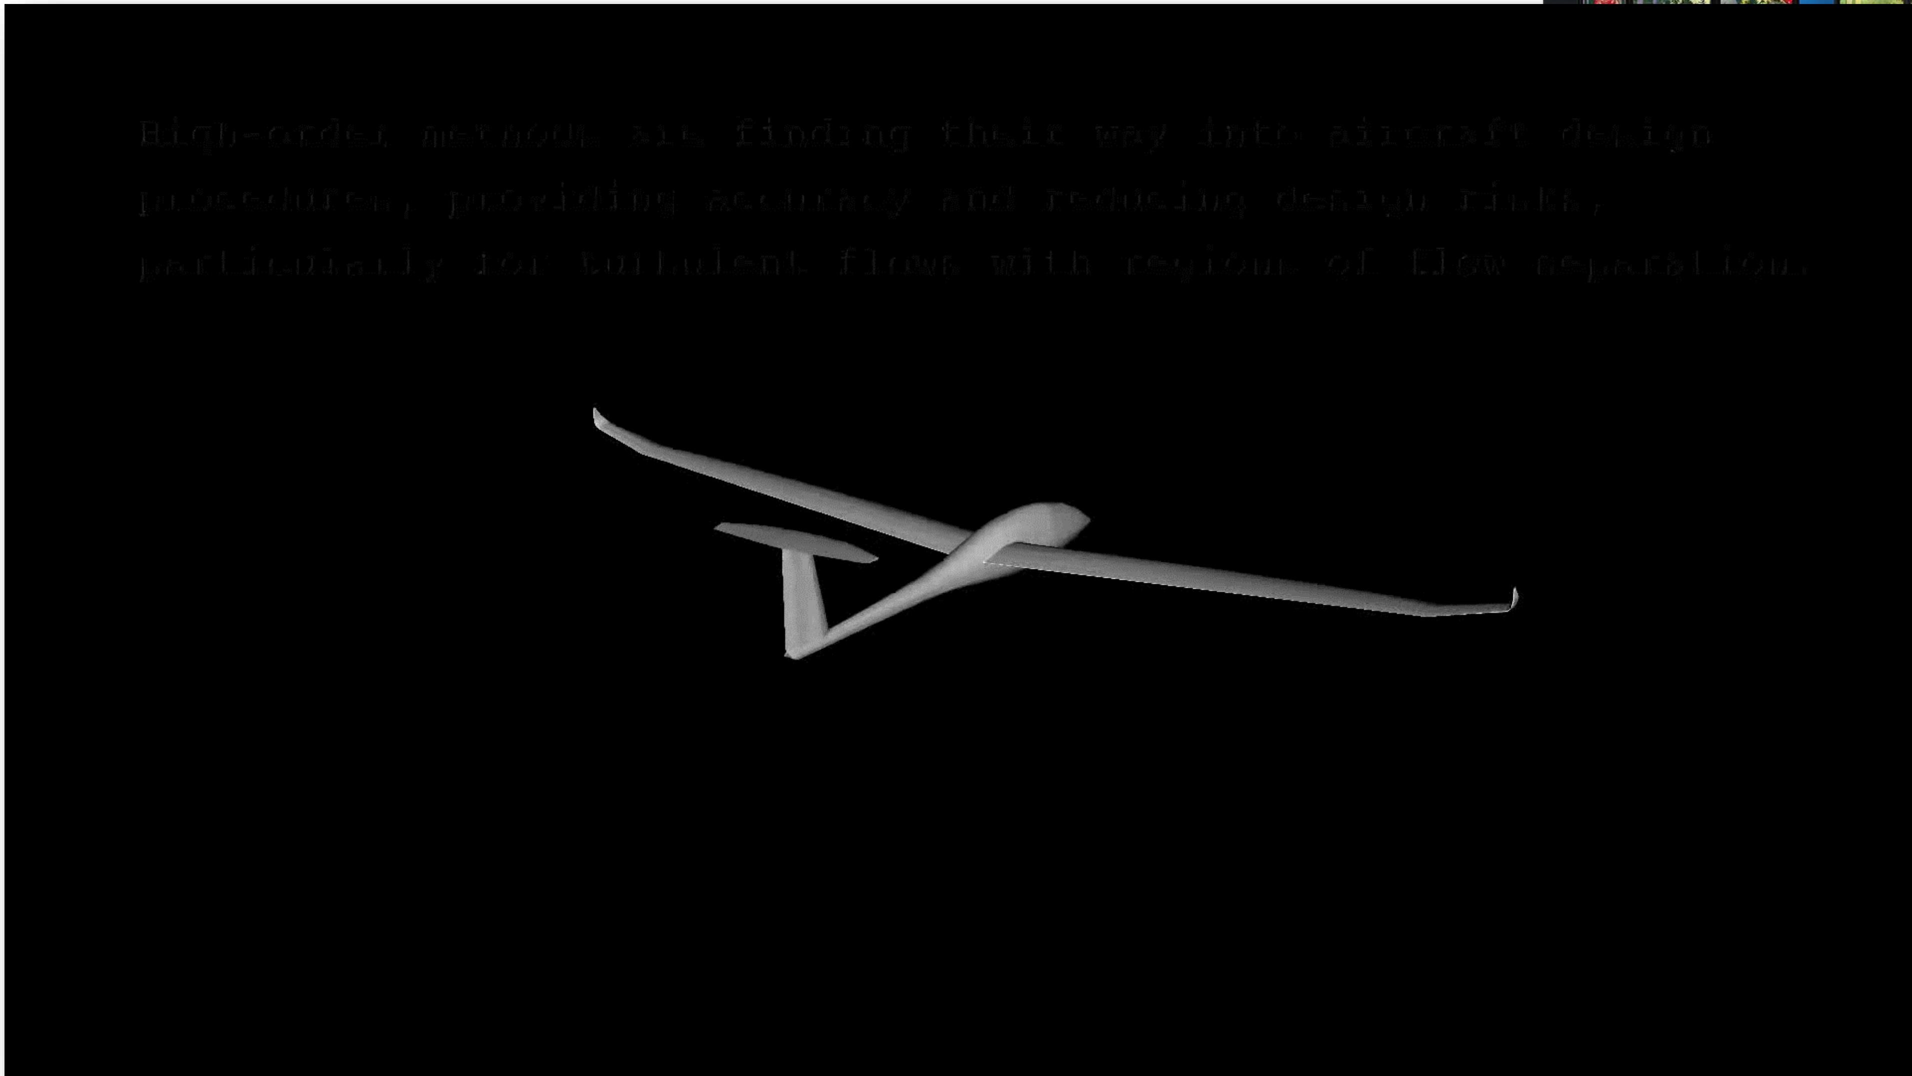

DNS Rotating Propeller

- Physical details:
  - NACA0012, 5 degrees angle of attack, rounded wingtips
  - Re=10k, 100k
    (based on chord and tip speed)
- Simulation details:
  - 1.1 million elements
    polynomial order 5
  - 300 million grid points
- One complete rotation:
  250k core hours
  - 10 days on 1024 cores

- Collaboration with CINECA via EXCELLERAT
- Video by Antonio Memmolo

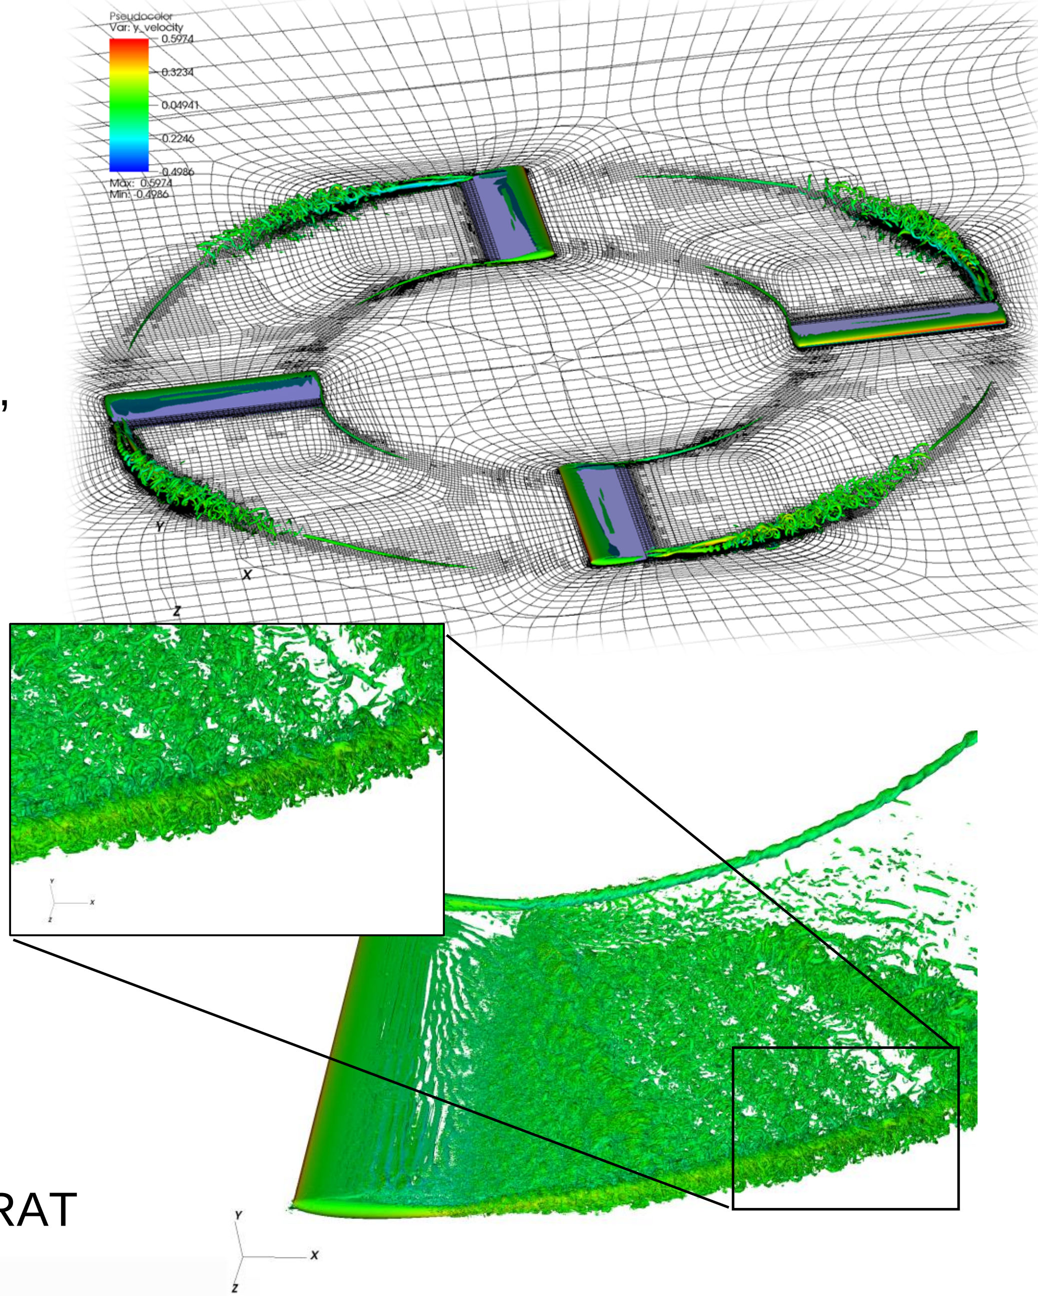

Re=100k

DNS Rotating Propeller (Re=10k)

Dardel 2021

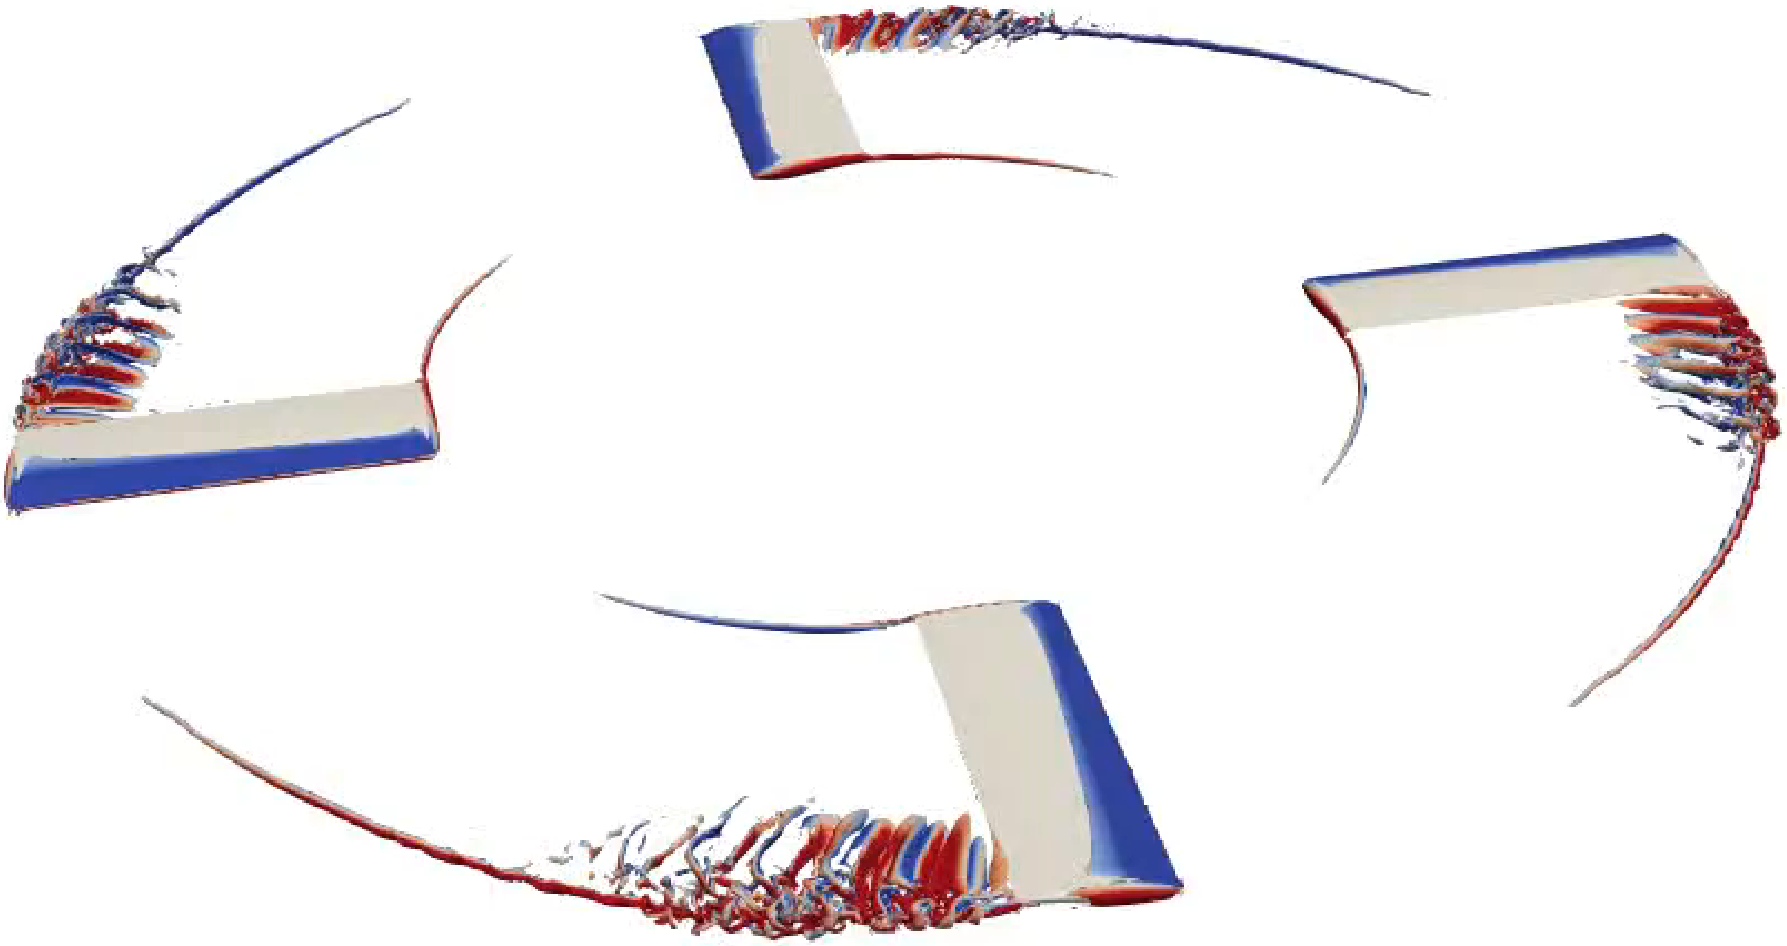

Nek5000 – Spectral Elements

- SEM code by **Paul F. Fischer**, Argonne National Lab, USA
  Open source: **`nek5000.mcs.anl.gov`**
- 80 000 lines of **Fortran 77** (some C for I/O), MPI (no hybrid)
- **Gordon Bell Prize 1999** for algorithmic quality and performance
- **KISS** ("Keep it simple, stupid") – world's most powerful computers have very weak operating systems

- **EU Projects on algorithms** (CRESTA, ExaFLOW, Exellerat, Admire…): adaptive meshing, **GPUs**, compression…

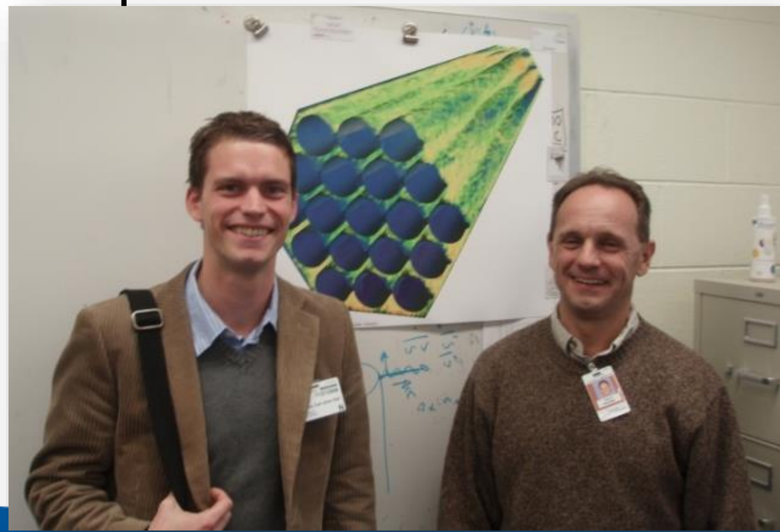

- Good scaling up to >**1,000,000** ranks

Why Spectral Elements?

- **Higher order p** (vs. **smaller grid spacing h**) means more work per core/communication: **”convecting cone”** [Gottlieb & Orszag 1977]

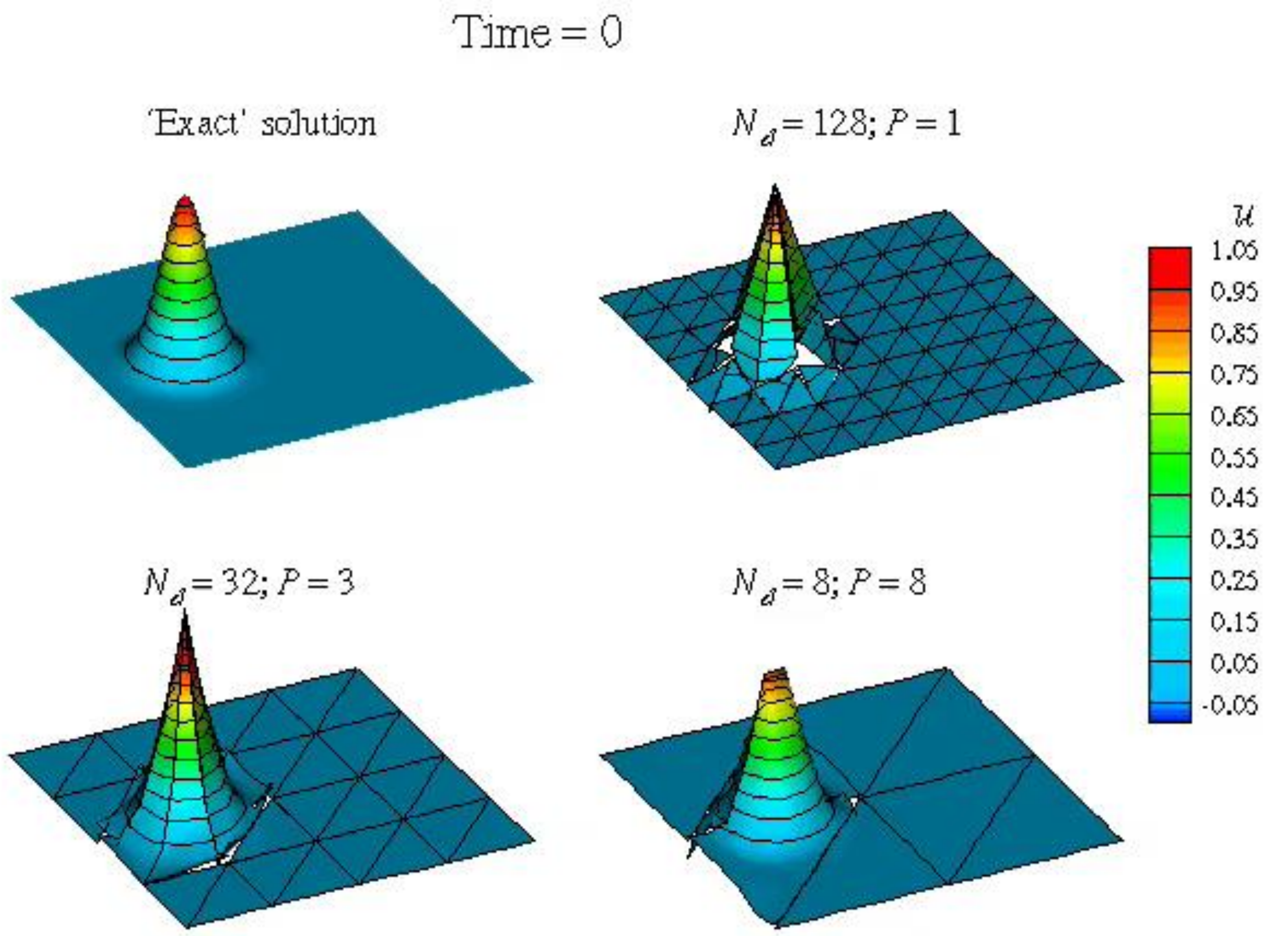

From David Moxey, Univ. Exeter

Before we go any further; can we quantify accuracy?

- Let's look at standard channel flow, Re_τ=300, DNS
- Two solvers: Nek5000 and OpenFOAM
- Consider varying $\Delta x^+, \Delta z^+$, fixed $\Delta y_w^+ = 0.5$ (fine)

- Quantities of interest (QoIs) for UQ problem: $\langle u_\tau \rangle, \langle u \rangle, \langle u_i' u_j' \rangle$

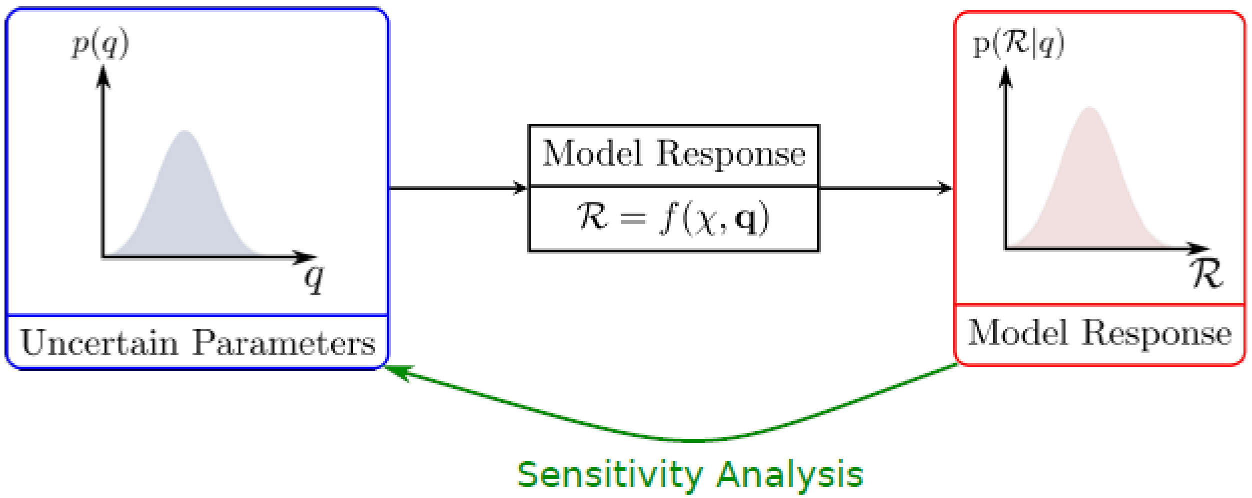

Ref. Smith 2013, Ghanem *et al.* 2017, Rezaeiravesh *et al.* 2021

UQ: Robustness Channel Flow

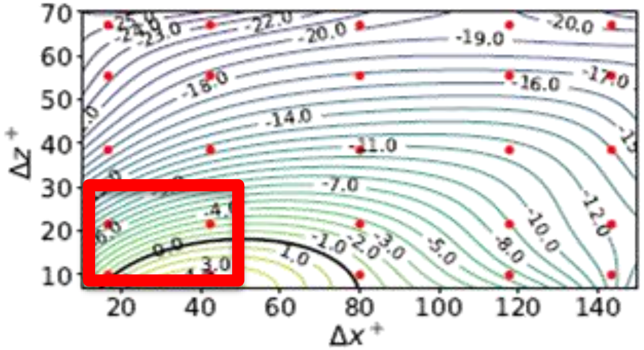

- How robust is a QoI with respect to variations of the numerical parameters? [Santner *et al.* 2003, Smith 2013]
- Uncertainty propagation by polynomial chaos expansion [Xiu&Karniadakis 2002] $\Delta x^+ \sim \mathcal{U}[10, 50]$ and $\Delta z^+ \sim \mathcal{U}[7, 30]$

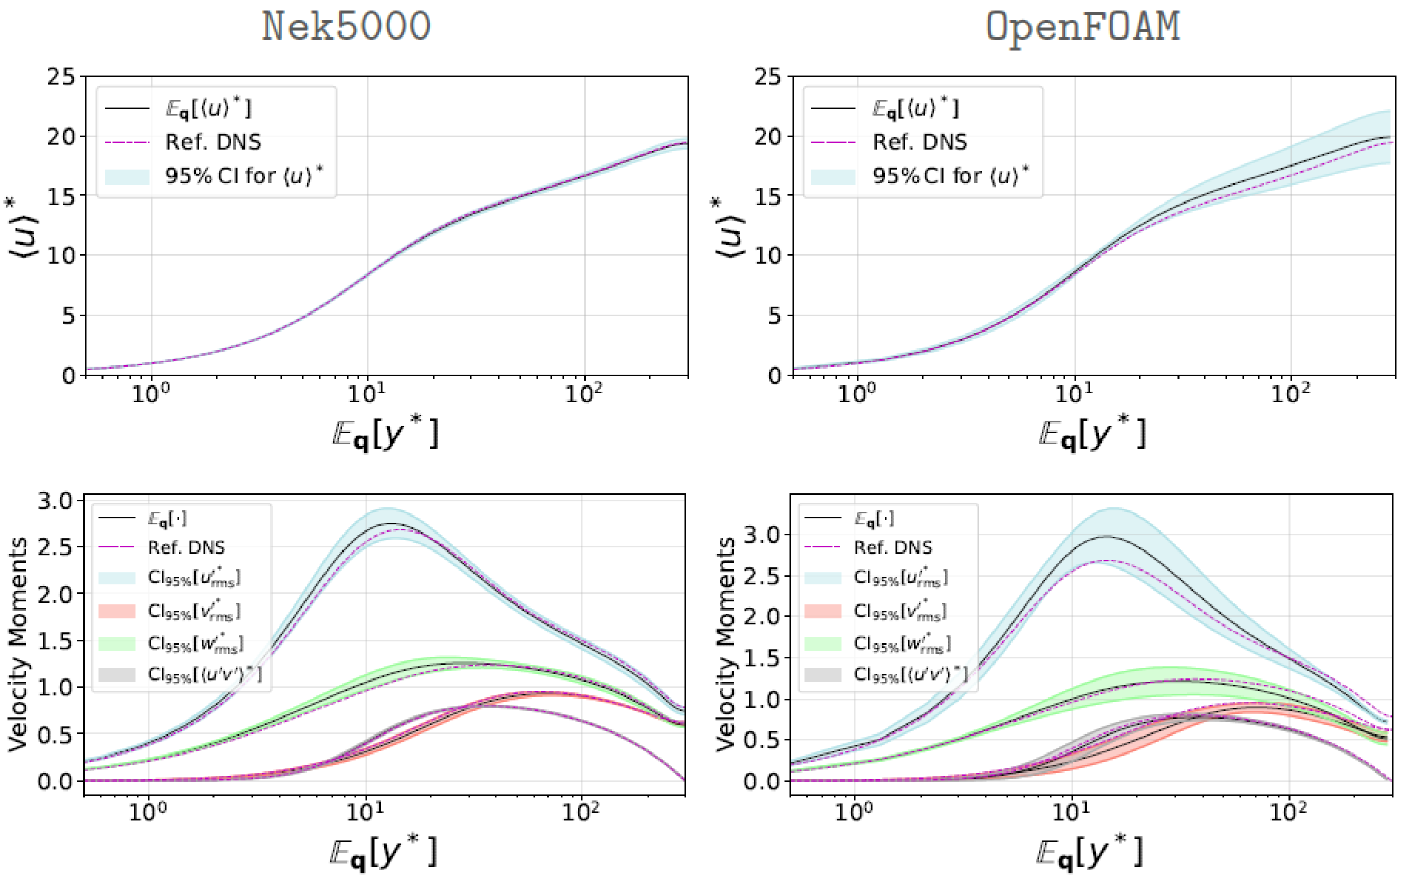

Some scaling data
Dardel pilot phase

- Three-dimensional wing tip case, non-conformal meshes, 600k elements (500M grid points)

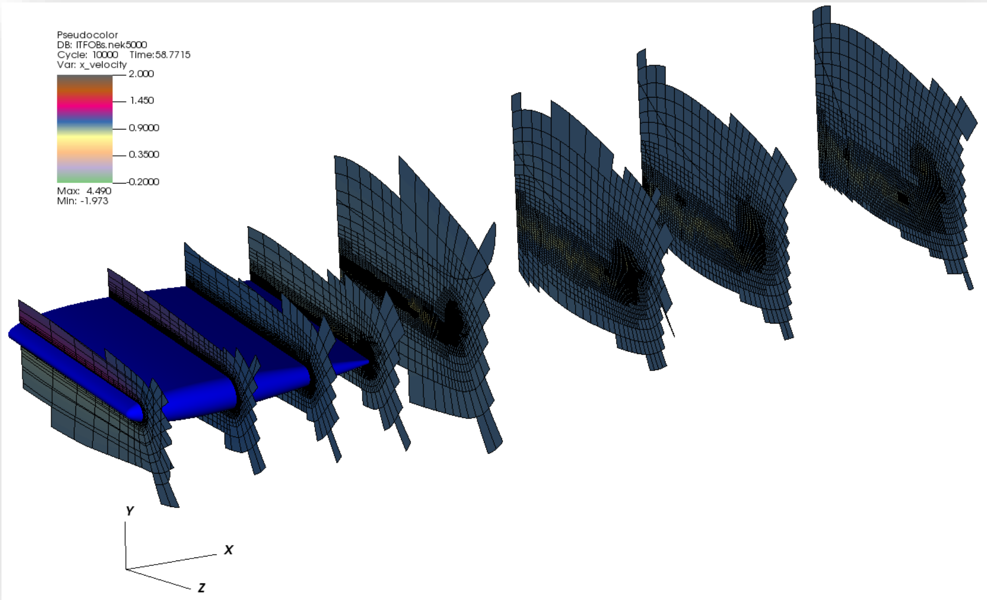

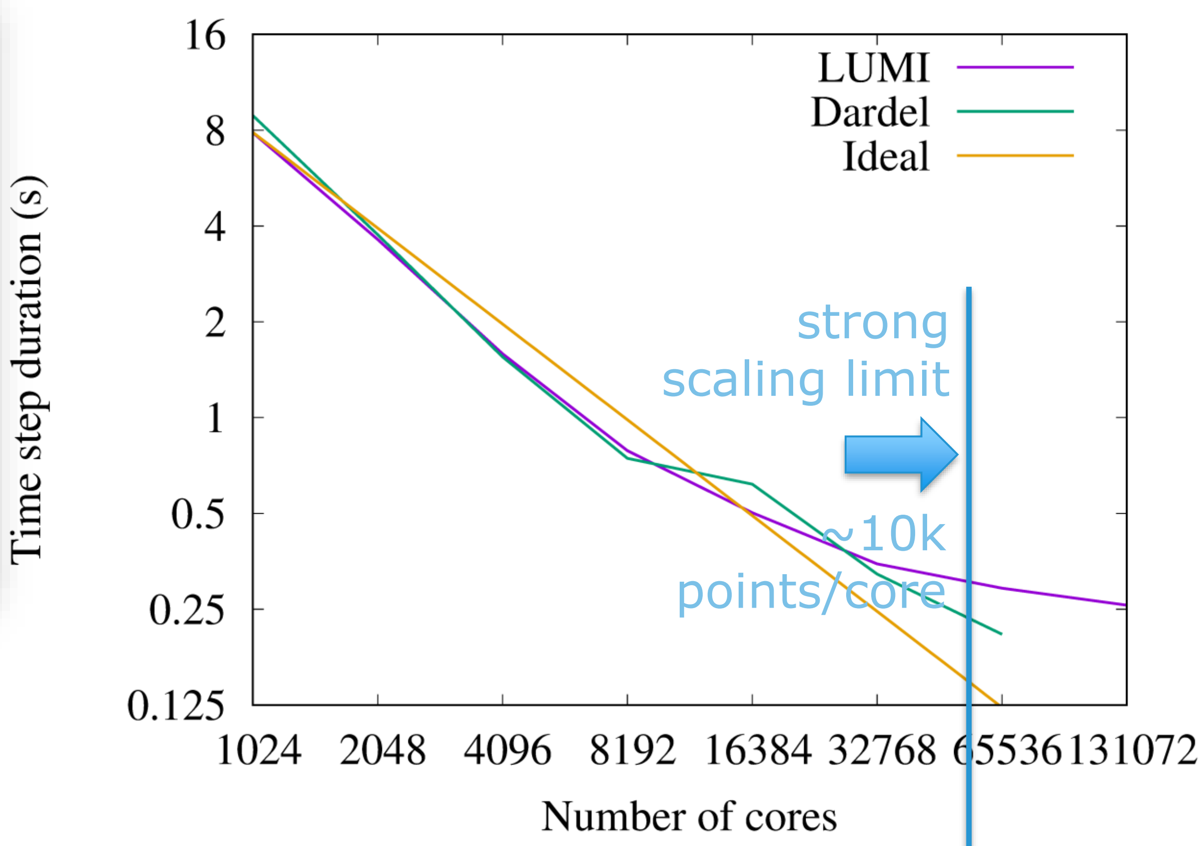

Main challenge: DATA

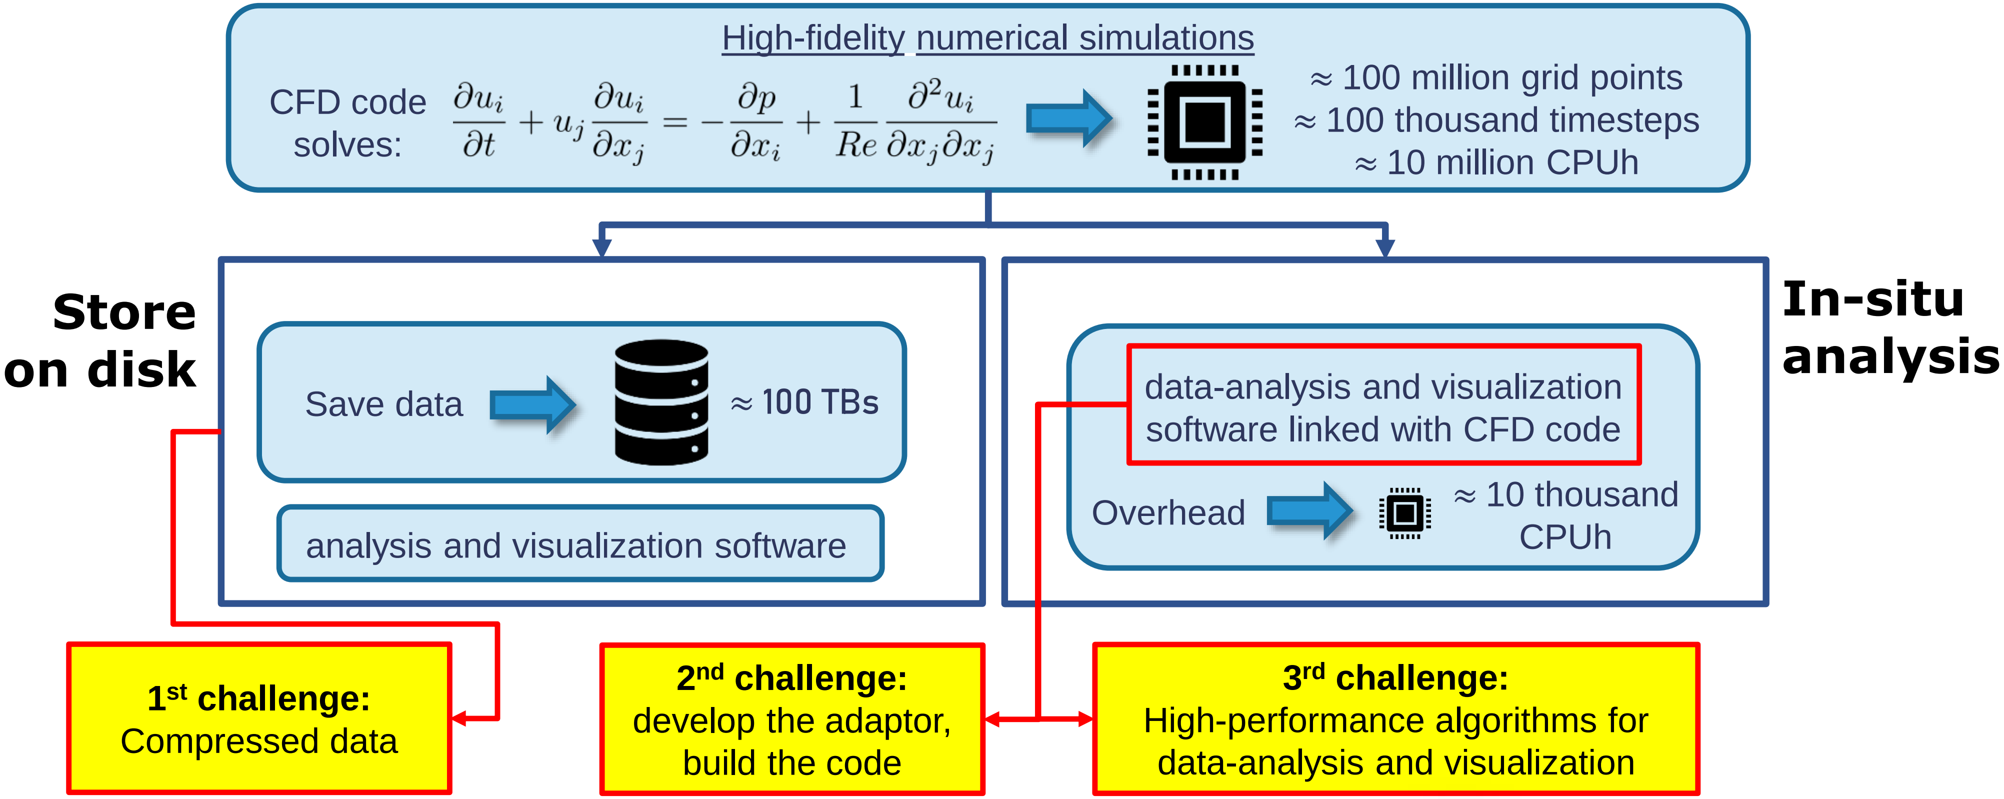

Multiple frameworks exist:

Paraview / Catalyst
ADIOS2
VisIt / Libsim
SENSEI

EXCELLERAT
FlowZoom project
ADMIRE malleable data solutions for HPC

Time-averaging Uncertainty

Reliable Estimation of the uncertainty due to the finite time-averaging

- Autoregressive model
- Modeled Autocorrelation Function

$$\hat{\mu} \sim \mathcal{N}(\mu, \sigma^2(\mu))$$

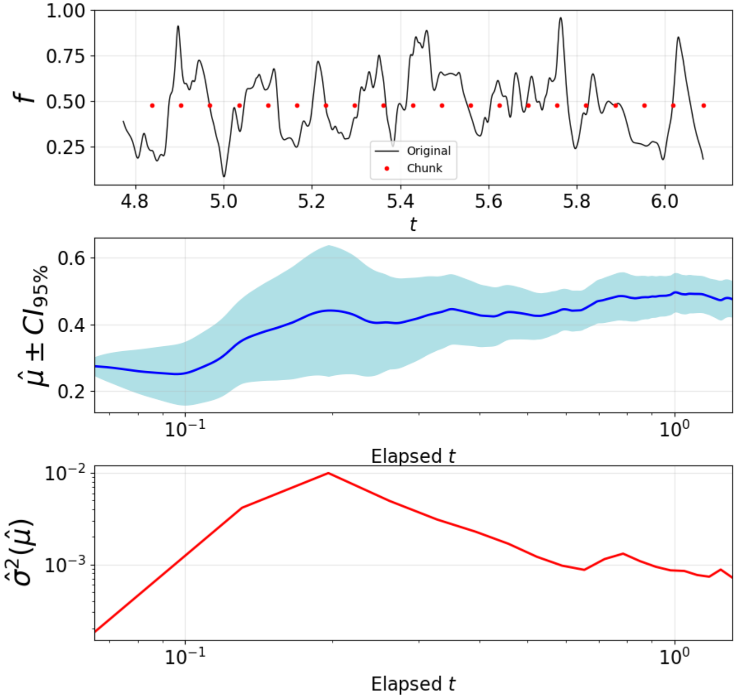

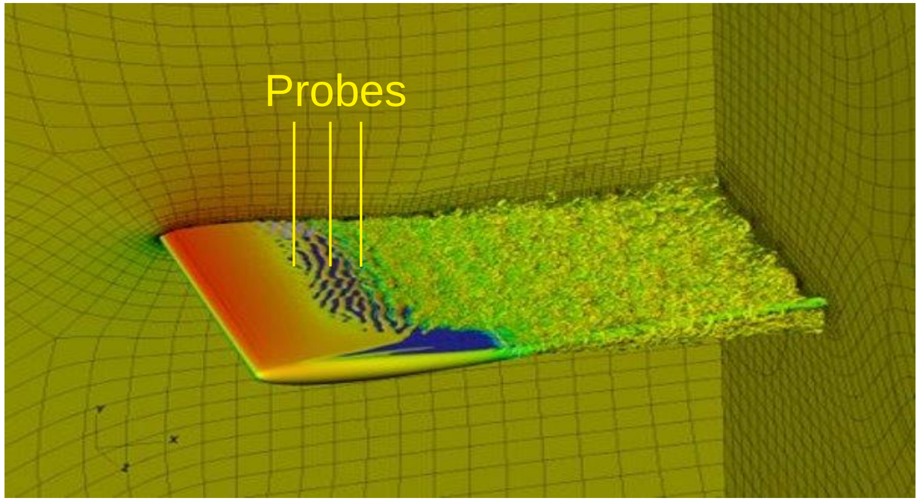

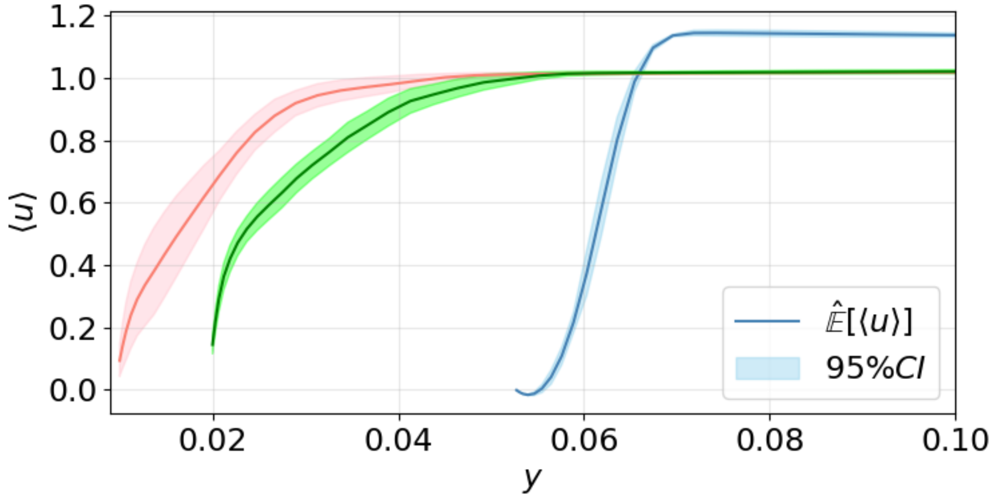

Mean velocity profiles and associated uncertainty

Time-averaging Uncertainty

- **Updating algorithms** for sample-mean estimator, batch means, variance, and autocorrelations
- No need to dump the time-series samples!
- Collaboration within EXCELLERAT (KTH and Fraunhofer-SCAI)

Nek5000 → ParaView Catalyst → UQit

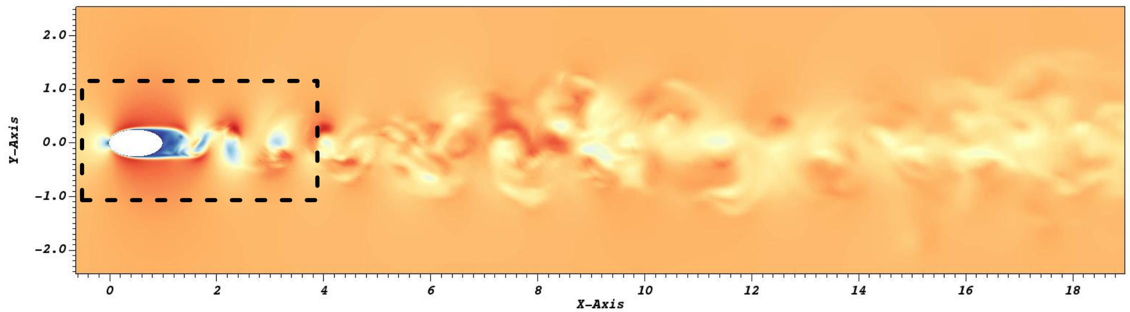

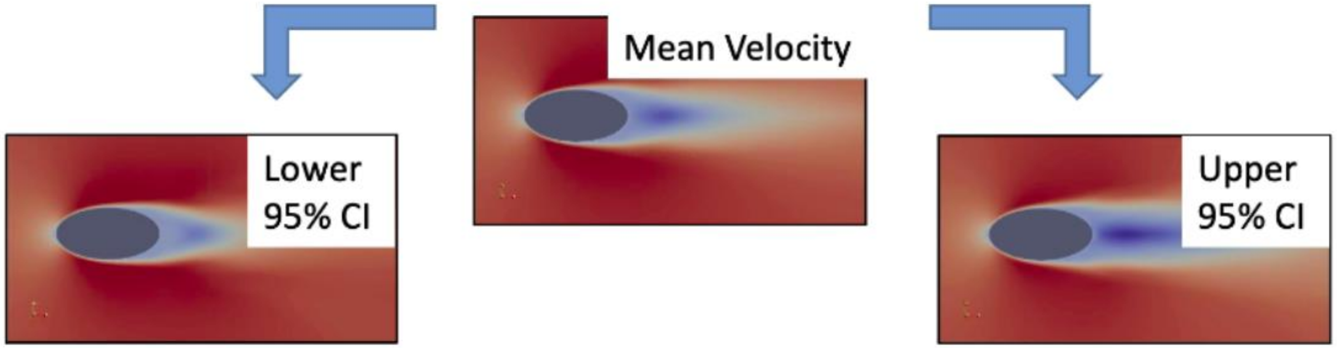

Lossy compression of turbulence fields

- A lot of entropy in turbulence, loss-less compression does not help much
- JPG-inspired algorithm in spectral space (Otero *et al.* 2018)

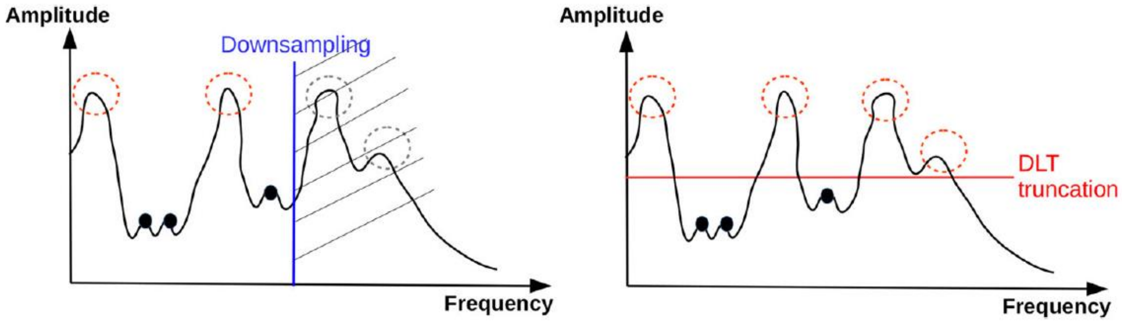

- Runtime encoding using ADIOS2 library in collaboration with Max Planck Computing and Data Facility

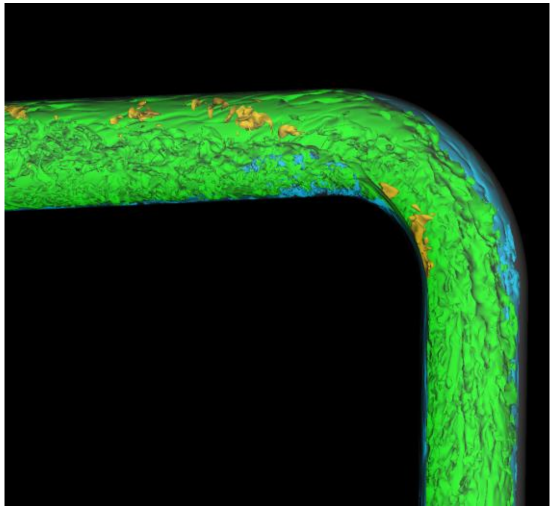

In-situ compression

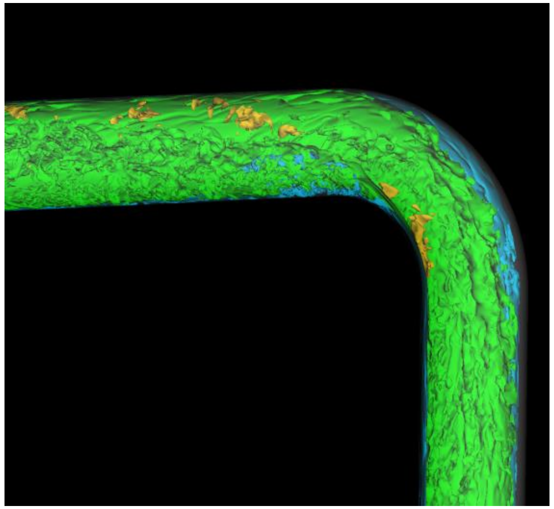

11% of data, error 0.01%

FLOW LINNÉ FLOW CENTRE

ADMIRE malleable data solutions for HPC

Nek5000 results – CPU vs GPU scaling

- $Re_\tau = 550$, max polynomial order 9
- $Re_\tau = 360$, max polynomial order 9

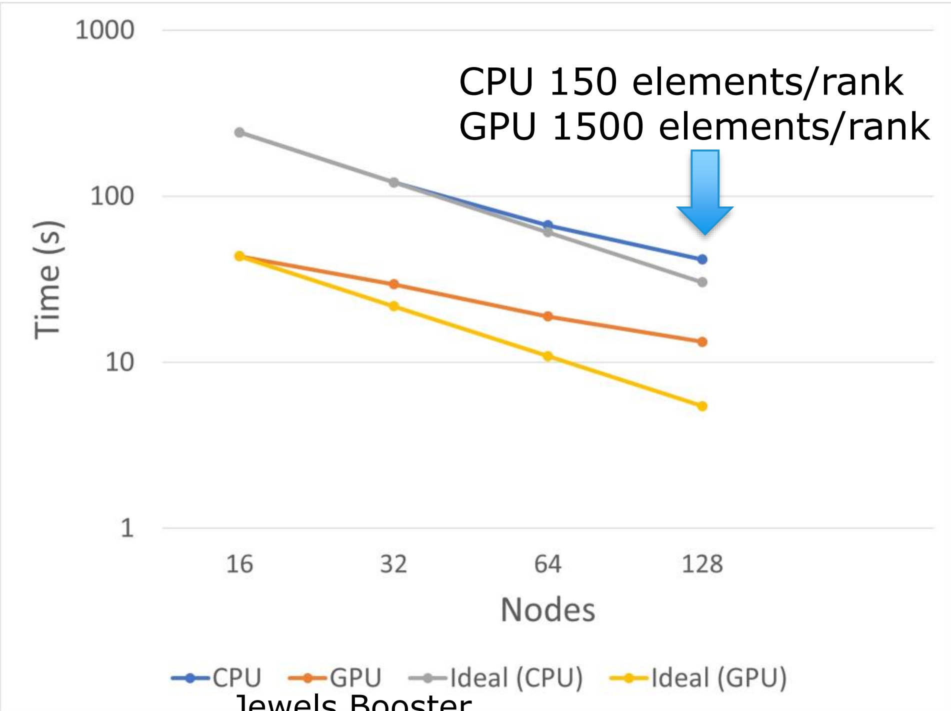

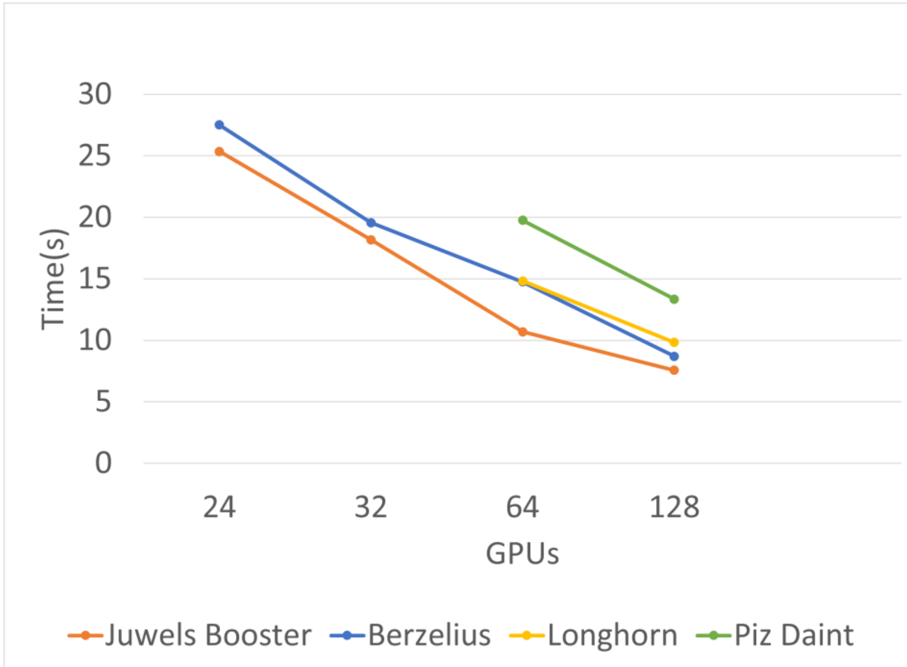

Vincent *et al.* HPC Asia 2022, https://arxiv.org/abs/2109.03592

Neko: Portable Spectral Element Framework

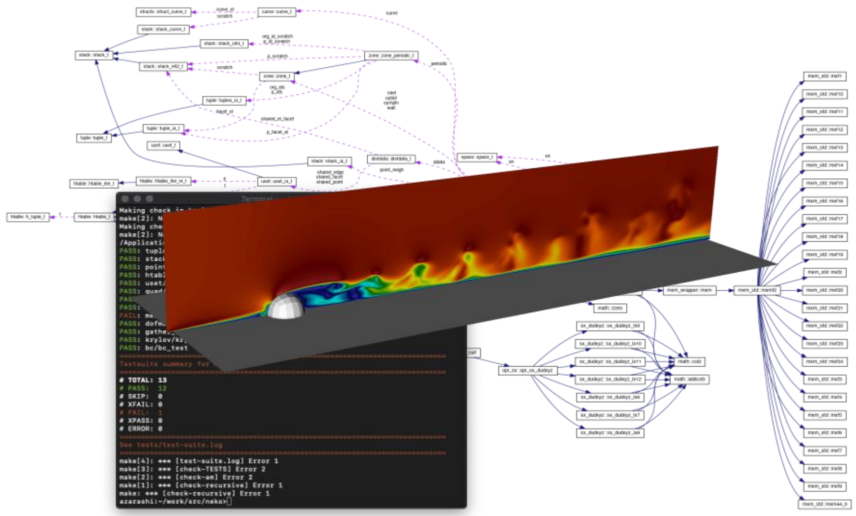

- High-order spectral element flow solver
  - Matrix-free formulation, small tensor products
  - Modern object-oriented approach
    - > Based on the abstract problem, $a(u, v) = L(v) \quad \forall v \in V$
    - > Modern Fortran, F2008/F2018 (coarrays!)

```
! Base type for a matrix-vector product providing Ax
type, abstract :: ax_t
 contains
   procedure(ax_compute), nopass, deferred :: compute
end type ax_t

! Abstract interface for computing Ax
abstract interface
   subroutine ax_compute(w, u, coef, msh, Xh)
     implicit none
     type(space_t), intent(inout) :: Xh
     type(mesh_t),  intent(inout) :: msh
     type(coef_t),  intent(inout) :: coef
     real(kind=dp), intent(inout) :: w(:,:,:,:)
     real(kind=dp), intent(inout) :: u(:,:,:,:)
   end subroutine ax_compute
end interface
```

- Various hardware-backends
  - CPUs, GPUs down to exotic vector processors and FPGAs
    - > Regularly tested on Intel, AMD, A64FX and SX-Aurora
  - Multi-tier abstractions of the solver stack
    - > HW specific tuning e.g. kernel merging, redundant computation

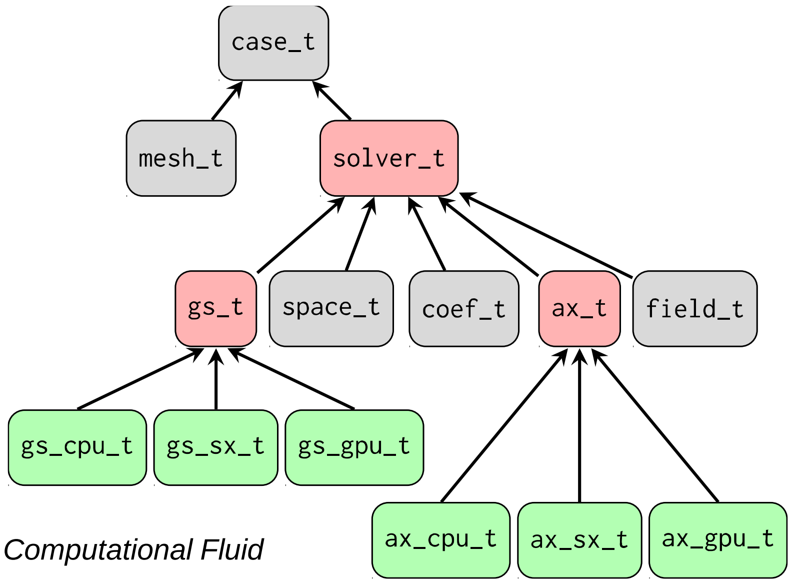

N. Jansson et al. *Neko: A Modern, Portable, and Scalable Framework for High-Fidelity Computational Fluid Dynamics,* arXiv preprint arXiv:2107.01243, 2021.

Neko: Performance Evaluation

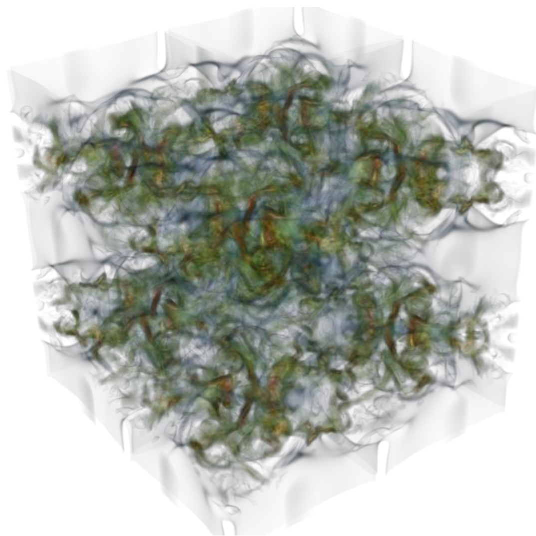

Strong scalability test of Neko compared to NekRS/Nek5000

- Taylor-Green vortex at $Re = 5000$
  - 261,144 hexahedral elements
  - 9^{th} order polynomials (10 GLL points)
- **Beskow**, Cray XC40 at PDC
  - Two 16 cores Intel E5-2698v3
  - Tested on 16 up to 1024 nodes
- **Vulcan**, NEC A300-8 nodes at HLRS
  - Eight SX-Aurora Type 10B per node
  - Tested on 2 up to 8 nodes (entire machine)
- **Dardel**, HPE Cray EX at PDC
  - Two 64 cores AMD EPYC 7742
  - Tested on 32 up to 512 nodes

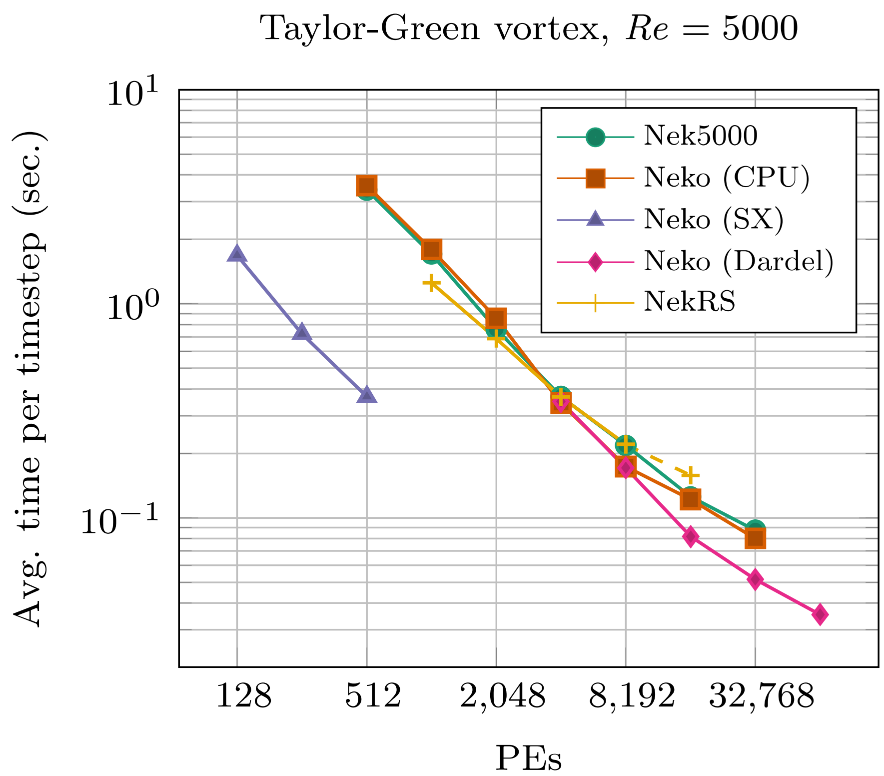

https://github.com/ExtremeFLOW/test-suite

To conclude…

- Let me put on a different hat…

- SNAC (Swedish National Allocation Committee)
- **Dardel phase 1 gives 47 Mh per month**
- For SNIC LARGE projects → 34 Mh/month
- 20 different large-scale projects, 7 universities in Sweden, 8 different areas of science
- National machine to the benefit of all HPC users!

Acknowledments

KTH Engineering Mechanics

- Adam Peplinski
- Saleh Rezaeiravesh
- Adalberto Perez

PDC

- Niclas Jansson
- Jonathan Vincent

ENCCS

- Jing Gong