Towards a Virtual Wind Tunnel - Fluid Simulations in the SeRC Exascale Flagship

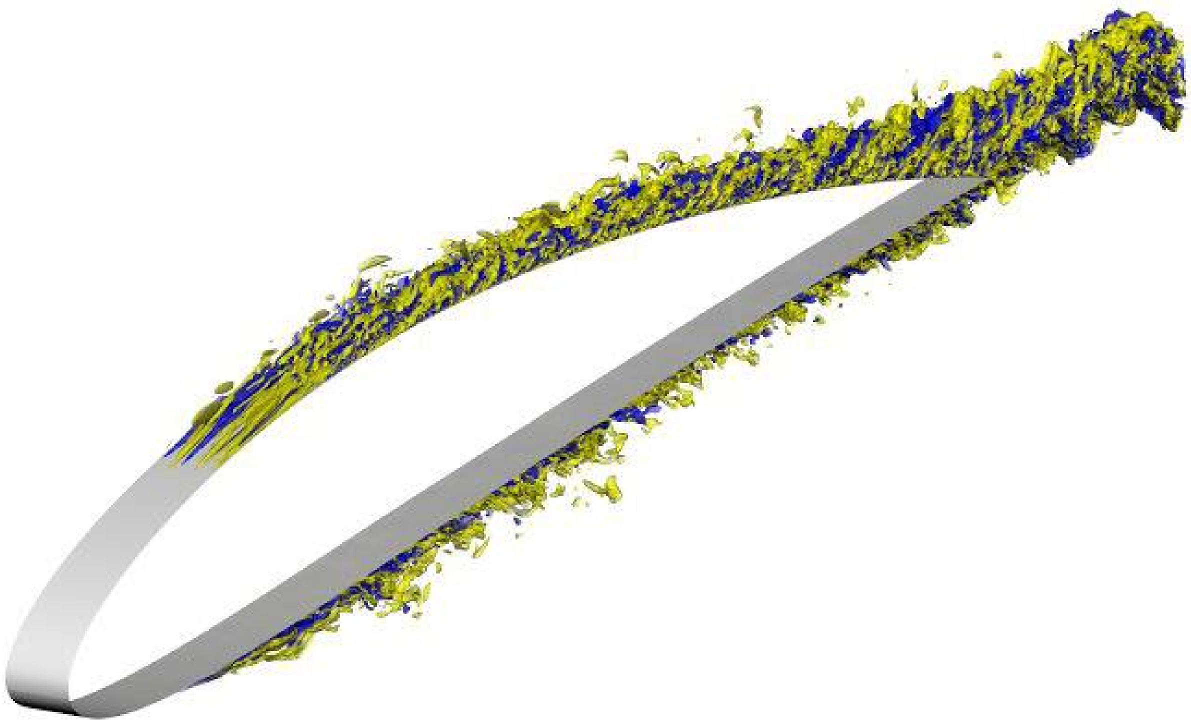

Matthew de Stadler
Dept. of Mechanics, Linné Flow Centre, KTH

Collaborators: D. Henningson, P. Schlatter, A. Hanifi, I. Bouya, S. M. Hosseini, J. Gong, A. Peplinski, P. Fischer

What is Exascale computing?

Exaflop: 10^{18} floating point operations / second

30 times faster than today's fastest machine

The fastest computer in the world today (As of November 2013)

- Tianhe-2 in China
- 33.8 Petaflops / second (Peta = 10^{15})
- 3.12 million cores
- Power 17.8 MW

Tianhe-2 Image from TOP500 list: http://www.top500.org/system/177999

Can we scale up today's systems?

- 3.12 million cores → 93.6 million cores
- 17.8 MW → 534 MW → supercomputer needs its own power plant

Why do we need exascale computing?

Jumbo jet in straight and level flight

~10 quadrillions (10^{16}) grid points needed to simulate the near surface turbulence with reasonable detail

- Moin and Kim. Scientific American 1996

Want better, faster, cheaper solutions to complex problems

- Industrial applications: airplanes, ships, cars
- Accurate weather prediction

Flow physics understanding

- More complex situations previously studied mainly through experiments

Better understanding → better models → better design tools

Why do we care about turbulence?

Turbulence is the rule rather than the exception

- Modeling turbulence is extremely challenging

Drag reduction: 10% of world energy use spent to overcome turbulent friction (PRACE Scientific Case)

- Laminar flow control
  - Suppressing turbulence decreases drag
  - 15% fuel reduction!

EU New Aircraft Concepts Research Project

Goal: be able to predict and subsequently control turbulence

- Take advantage of its good properties (higher mixing)
- Avoid its bad properties (more drag)

FLOW
LINNÉ FLOW CENTRE

KTH
VETENSKAP OCH KONST

Does turbulence scale to exascale?

We can always make our problems bigger

- Reynolds number: separation between large and small scales

Increasing Re → smaller scales → more gridpoints

Re_D= 5,000 Re_D= 10,000 Re_D= 20,000 Re_D= 40,000

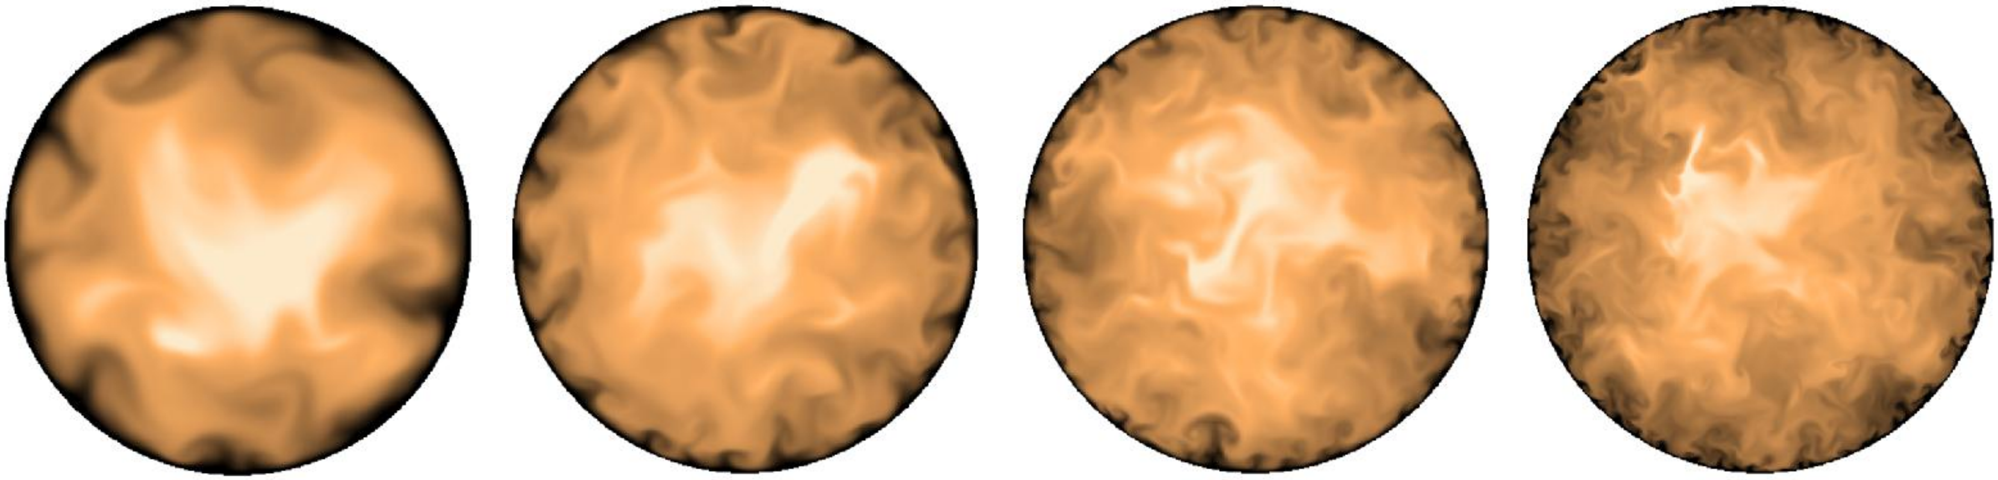

Higher Re → New physics: Complicated flow phenomena

- Greater understanding needed for turbulence models

Hardware: Good news and bad news

Good news: supercomputers keep getting faster and more efficient

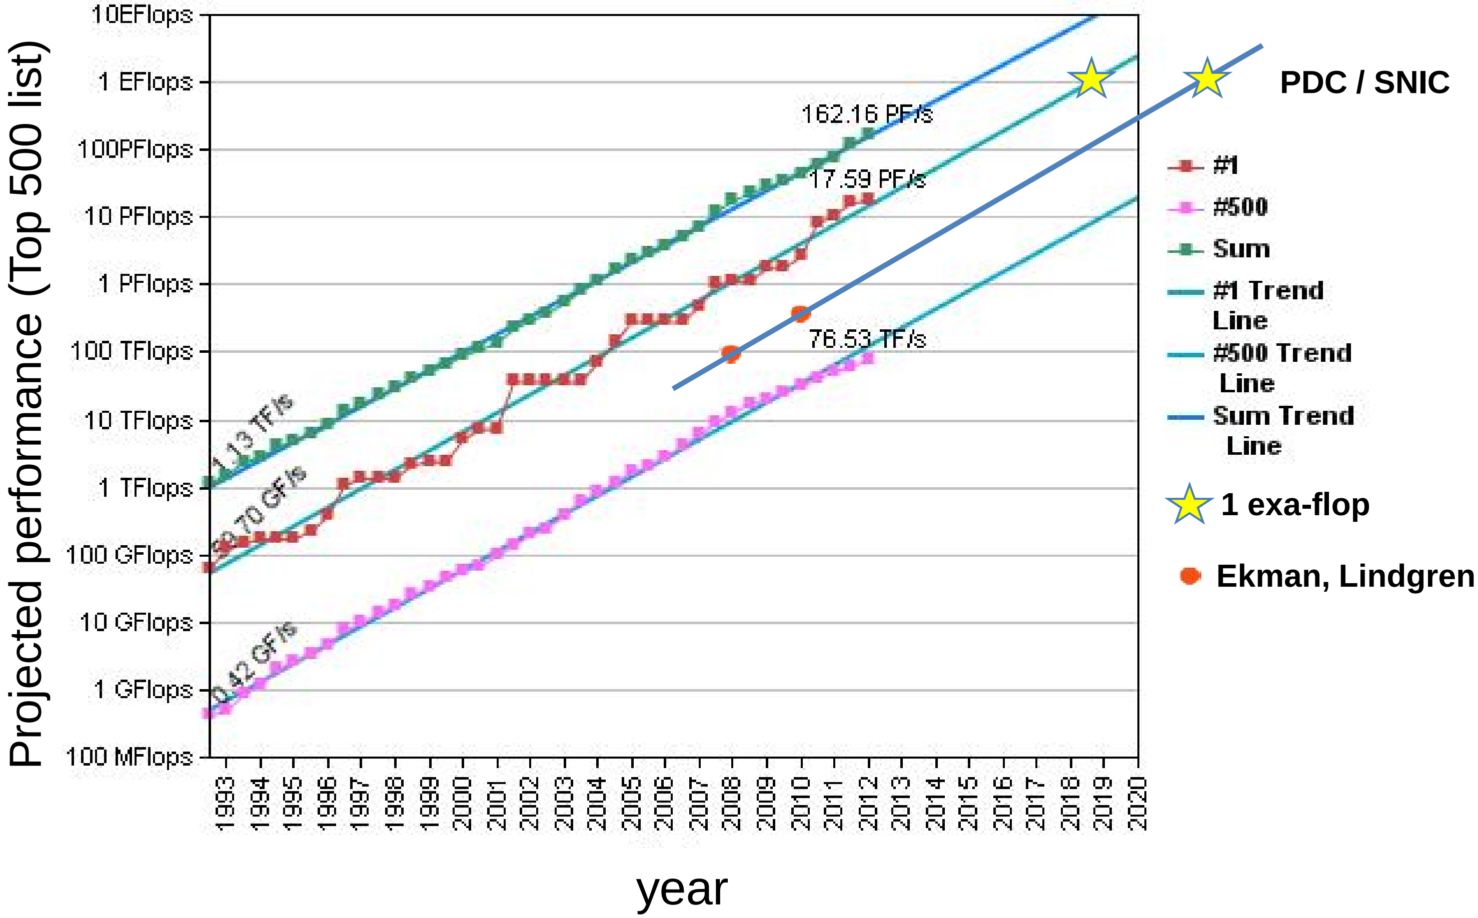

Bad news: Supercomputers are becoming more...

- Heterogeneous — - Harder to program (disruptive technologies)
- Parallel — - Power costs are soaring

FLOW
LINNÉ FLOW CENTRE

KTH VETENSKAP OCH KONST

How fast does a ’real program’ run

Speed rankings of supercomputers from the LINPACK benchmark

- Not indicative of performance of a production code
  - Best application codes — 30% of peak efficiency
  - Typical application code — less

Software lags several years behind hardware

- Collaboration needed between computing specialists and domain scientists to ensure good performance

Open question: How to design the next generation of efficient software?

- GPUs?
- MPI?
- New algorithms needed?

LINNÉ FLOW CENTRE

SESSI: SeRC Exascale Simulation Software Initiative

A group of about 15 experts in various areas of high performance computing

- Similar in spirit to EU CRESTA project

3 main application areas

1. Computational fluid dynamics

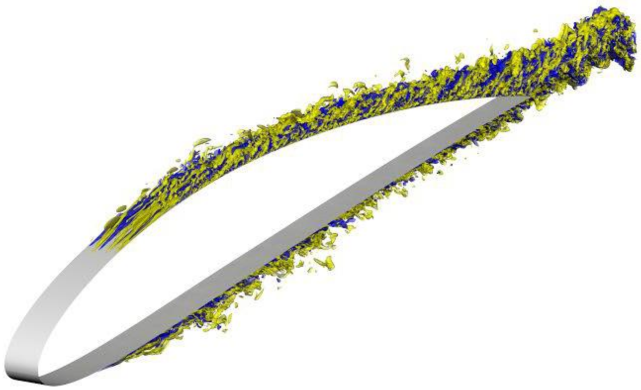

2. Molecular dynamics (GROMACS)

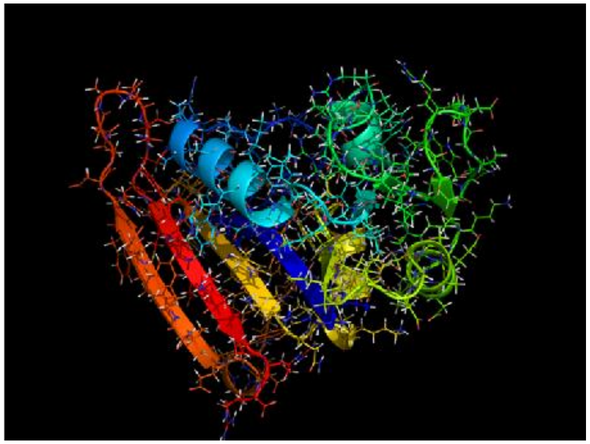

3. Computing specialists (PDC)

Goals of the group:

- Short term: Faster wing simulations
- Long term: To combine knowledge and experience to create efficient software for running on exascale platforms

FLOW
LINNÉ FLOW CENTRE

KTH
VETENSKAP OCH KONST

Computational fluid dynamics group in SESSI

General interest:

(1) Study transition to turbulence, turbulent flow

(2) Aerodynamics applications

(3) Massively parallel highly accurate simulation

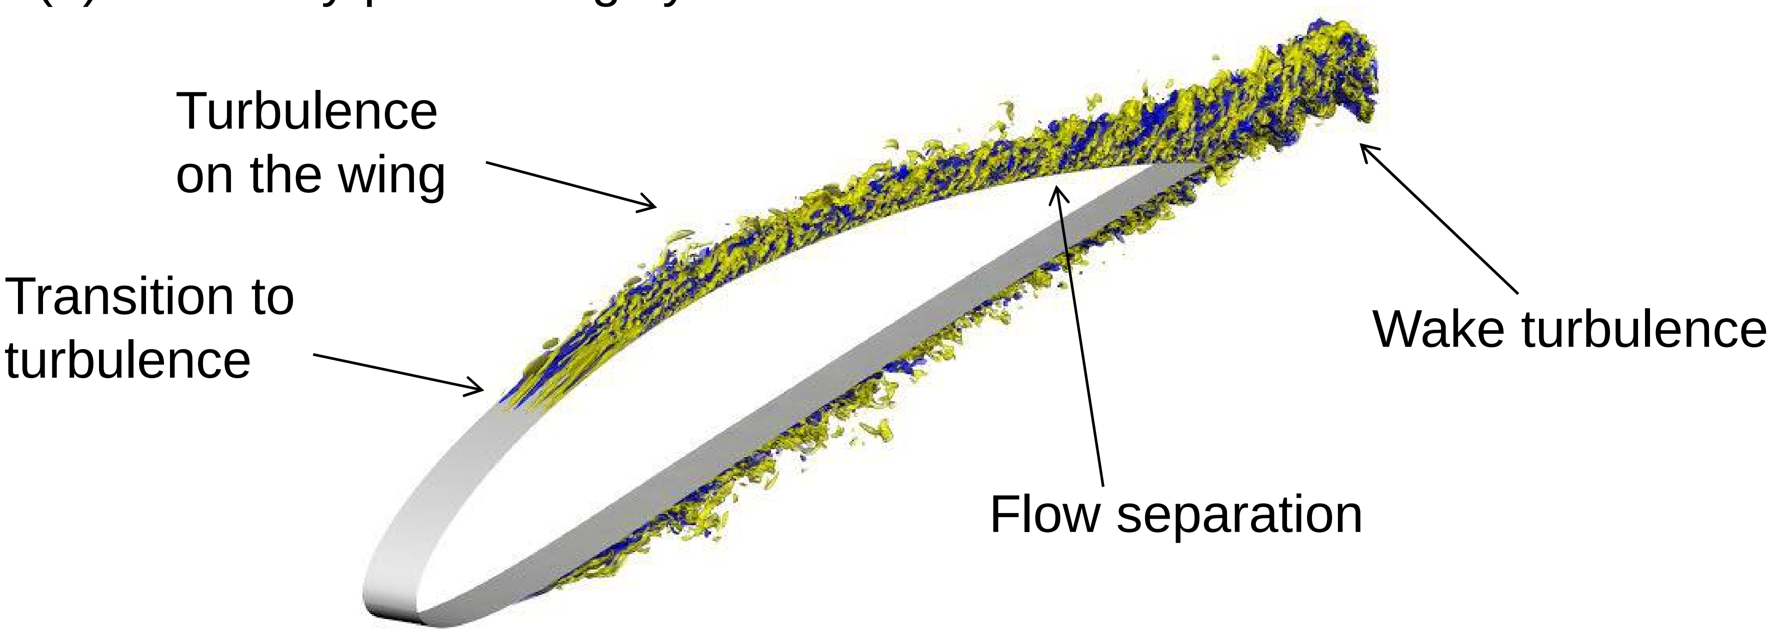

Study fundamental problems → understand turbulent phenomena

Everything is coupled: study interaction

SEM: Combines flexibility with accuracy

Finite element (FE) **Spectral** **Spectral element (SE)**

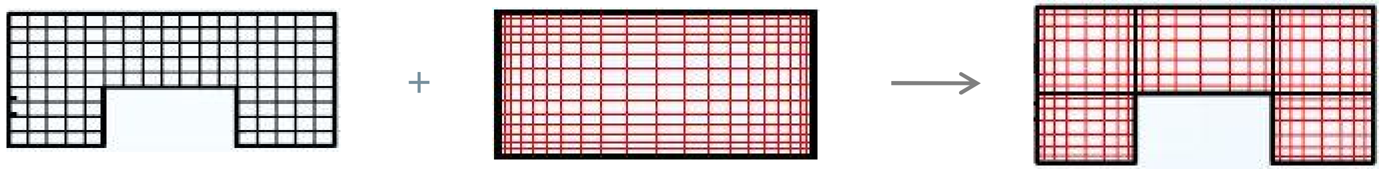

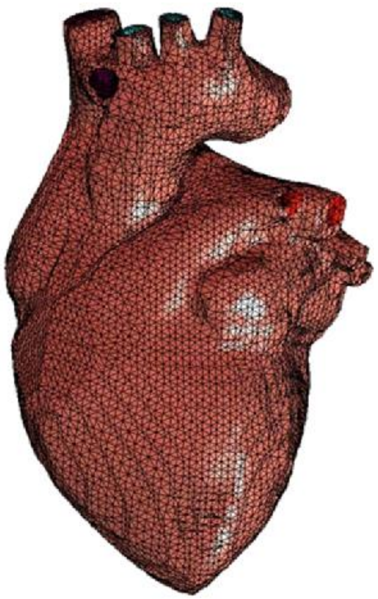

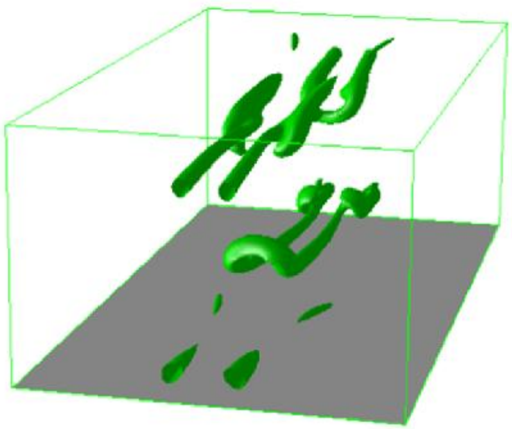

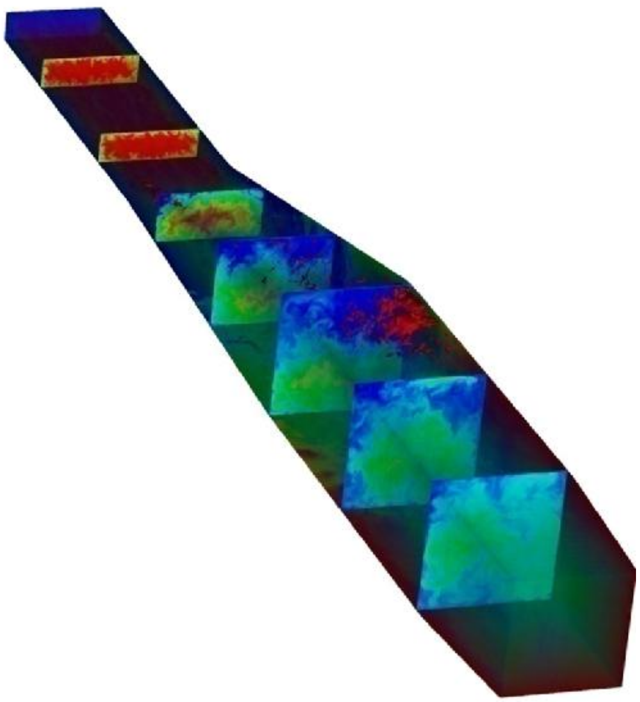

Nek5000 code by **Paul F. Fischer**, Argonne National Lab, USA
Open source: nek5000.mcs.anl.gov

FLOW
LINNÉ FLOW CENTRE

Over 120 users worldwide and cited in 200 journal papers

KTH
VETENSKAP OCH KONST

Strengths of Nek5000

High accuracy at low cost

- Rapid (exponential) convergence in space
- 3rd-order accuracy in time
- Optimized matrix-matrix multiplication kernels

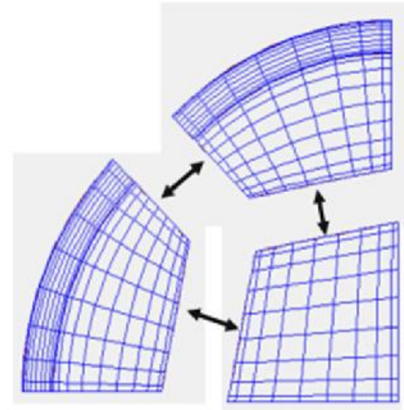

Highly scalable

- Standard domain decomposition + message-passing based parallelism
- Loosely coupled elements (C_0 continuity between elements)
- Efficient crystal router communications library
- Fast scalable coarse-grid solvers
- Iterative solvers with dense local work

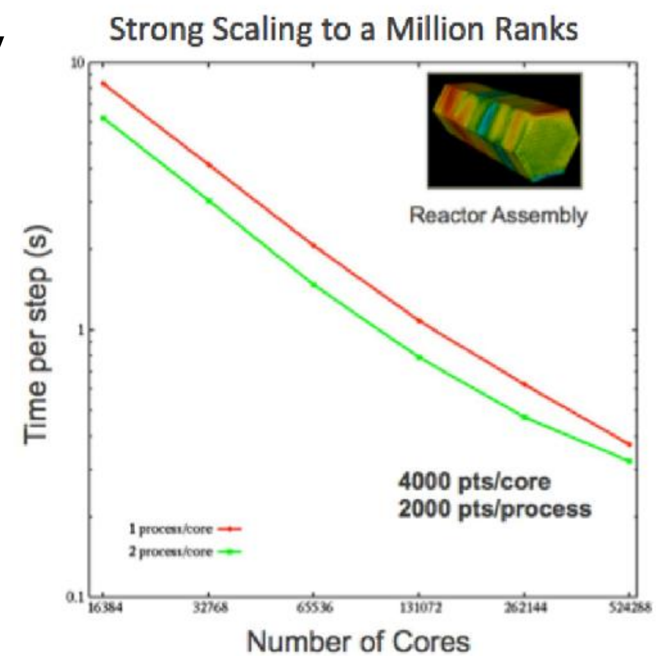

Scales to > 1,000,000 ranks

Can we go to exascale with Nek5000?

Key to good performance

- Minimize global communication
- Local work has to outweigh cost for communication
- Number of grid points *N* per processor *P* important
  - For good scaling: $(N/P) \sim 10{,}000\text{-}50{,}000$ sufficient

Grey: local work
Blue: communication

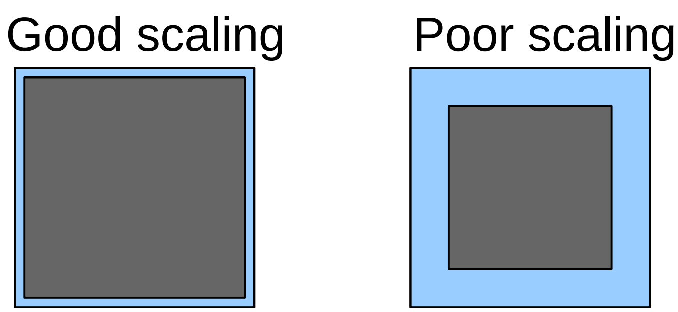

Must increase problem size for efficiency at exascale

- No problem for higher Reynolds numbers: $N \sim Re^{2.25}$
- An exascale machine will have 10^8 processors
  - Minimum of $N \sim 10^{12}$ gridpoints to scale to $P = 10^8$

FLOW LINNÉ FLOW CENTRE

KTH VETENSKAP OCH KONST

SESSI collaboration to date

Production version of Nek5000

- Designed for portability, architecture independent
- Optimized for use on many CPUs

With the PDC people

- Optimization of communications routines
- Development of a GPU version

512 nodes (8192 cores) (seconds/step)	512 GPU (seconds/step)	Speed-up (GPU/Node)
7.02s	4.41s	1.59

With the Gromacs people

- Hand tune most computationally intensive routines
    - First tests show a 15% speed up

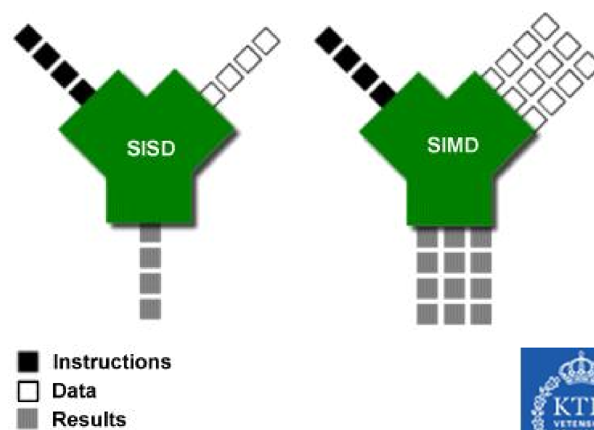

FLOW
LINNÉ FLOW CENTRE

KTH
VETENSKAP OCH KONST

Turbulent flow close to solid walls

Turbulent flow close to solid walls

Simulation results

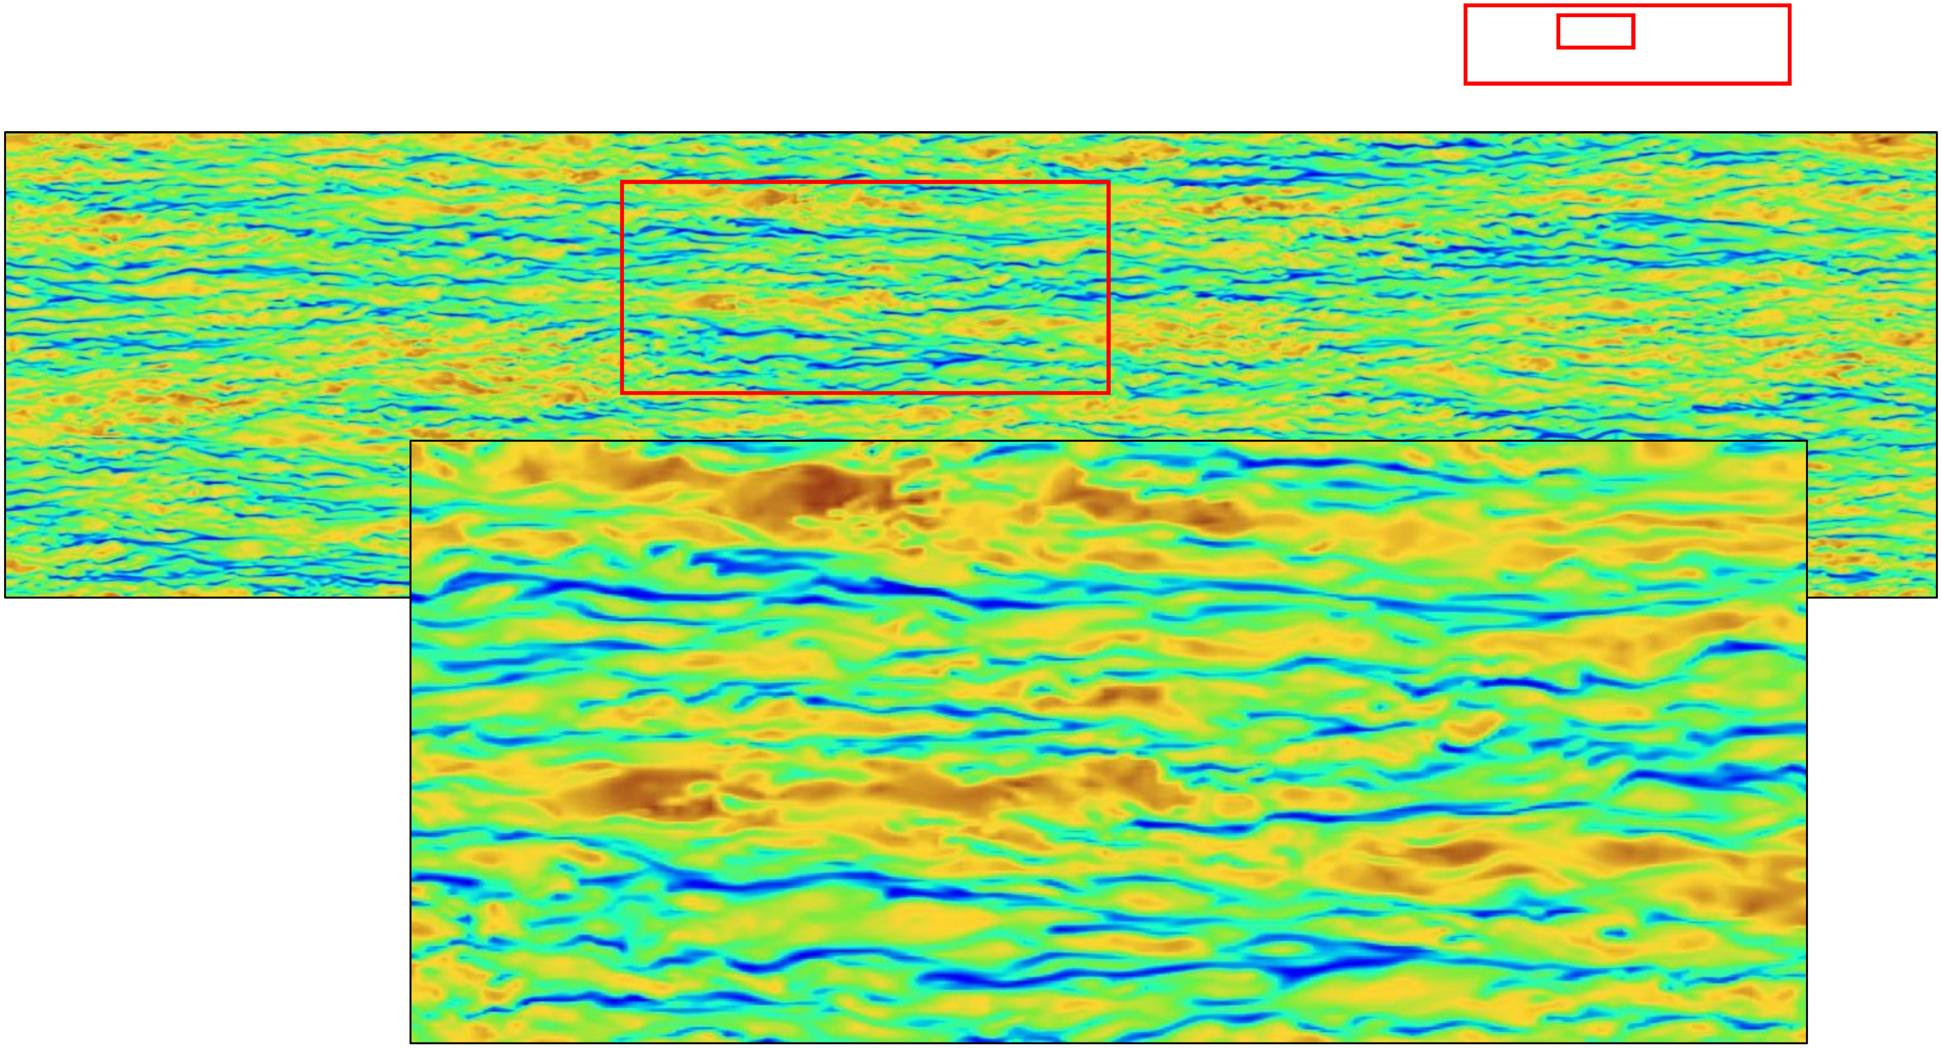

Zoom in on turbulent flow close to solid walls

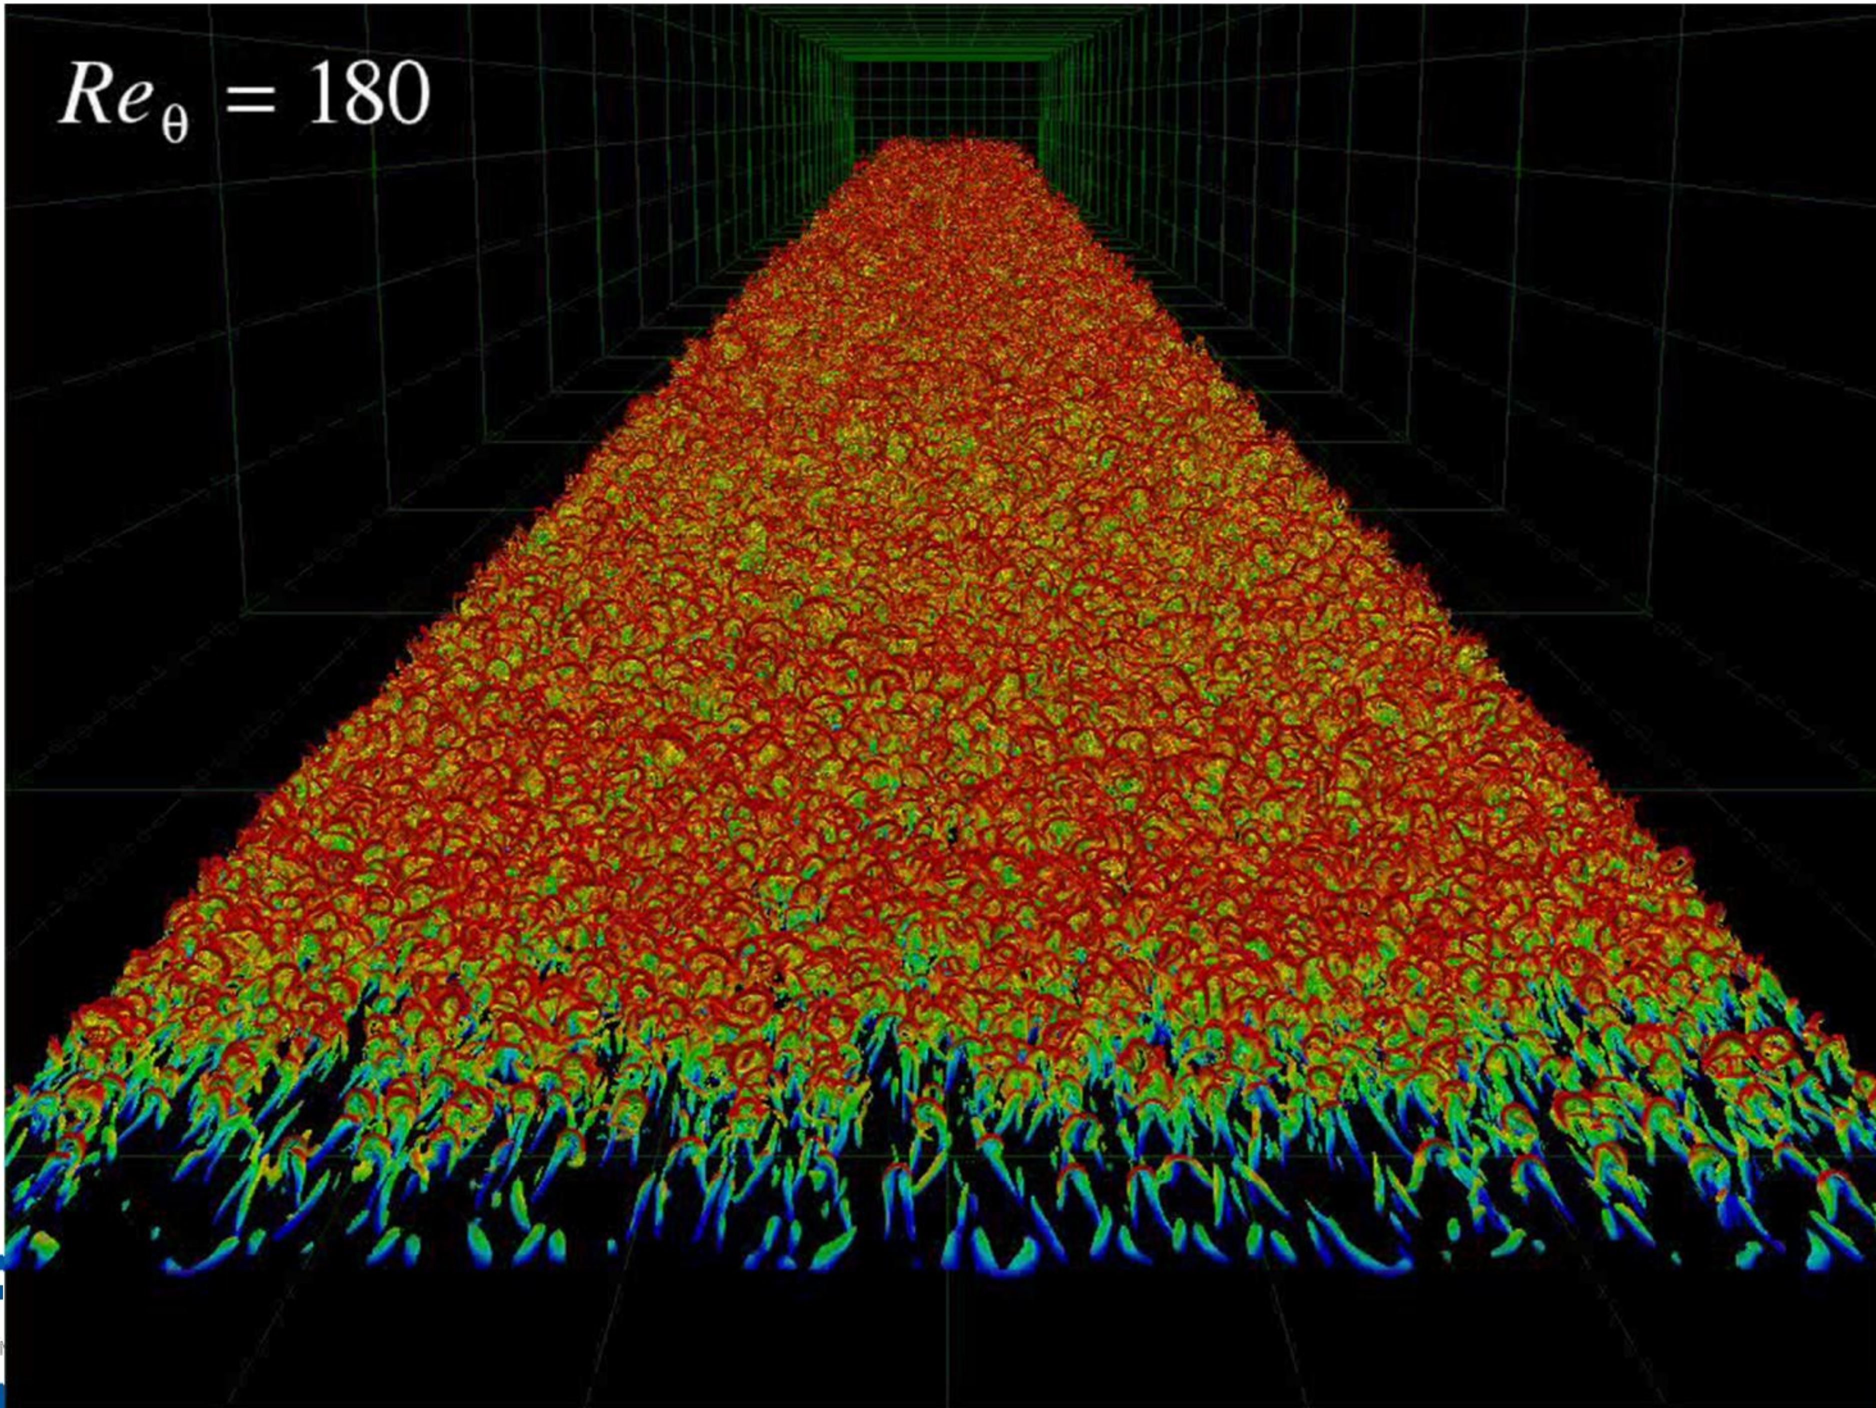

Flow past a wing in a virtual wind tunnel

Project goal: Move beyond traditional wind tunnel testing to a virtual wind tunnel

- Simpler, more accurate and less expensive wing testing
- Provide access to previously inaccessible flow data
  - Study interaction of turbulent phenomena

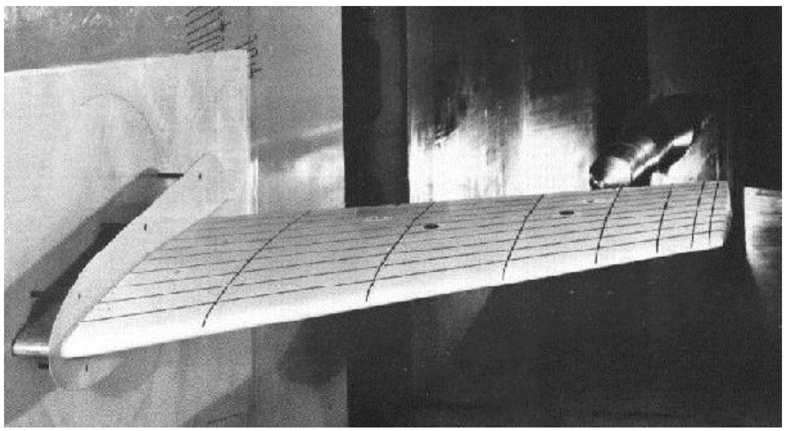

Project objective: Demonstrate proof of concept that high resolution methods can be used to calculate flows of practical engineering interest using massively parallel computing resources

- KAW Academy Fellow Project

Numerical wind tunnel

EU-project RECEPT
KTH Mechanics

Laminar Flow Control Experiment:
Re = 1 million

DNS of a typical wind tunnel experiment requires

- 10 billion grid points
- 1.5 million timesteps
- 60 TB of memory required
- 100 terabytes of data
- Run time: 64 days on 32,768 cores
- On an Exascale machine: 2 hours

Where are we now?

- PRACE TIER-0 application submitted
- Benchmark case of flow past a NACA4412 wing at Re=1.52 million
  - Experimental data for validation
- 85 million CPU hours requested

FLOW LINNÉ FLOW CENTRE

KTH VETENSKAP OCH KONST

Conclusions

We are preparing for exascale computing

- Short term: speed up wing calculations
- Long term: GPU version and optimal communication patterns

Upcoming challenges with exascale computing

- Unprecedented computational power → huge amount of data
  - I/O issues
  - How to analyze the Petabytes of data output? Storage?
- Redesign of algorithms for increased parallelism?
- Meshing and mesh quality

Wing simulations are pushing the bounds of what is possible

- Virtual wind tunnel: unprecedented level of detail
  - Study flow interaction in previously impossible way

Acknowledgments

CREST

VisIt

ParaView

KTH VETENSKAP OCH KONST
ROYAL INSTITUTE OF TECHNOLOGY

SeRC
Swedish e-Science Research Centre

FLOW
LINNÉ FLOW CENTRE

PDC HPC

NSC
National Supercomputer Centre at Linköpings universitet

PRACE

Knut och Alice Wallenbergs Stiftelse

Vetenskapsrådet

SNIC
Swedish National Infrastructure for Computing