HPC et Mécanique des Fluides

A. Cadiou, Ch. Pera
Marc Buffat, Lionel Le Penven

Laboratoire de Mécanique des Fluides et d'Acoustique
CNRS, Université Lyon I, École Centrale de Lyon, INSA de Lyon

École calcul intensif, Lyon, Septembre 2013

avec

Fluid Mechanics

Navier-Stokes equations

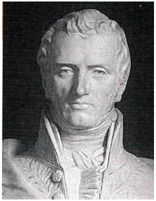

Claude Louis NAVIER

(1785-1836)

For incompressible flow, homogeneous fluid

velocity $\boldsymbol{U}$, pressure p

$$\partial_t \boldsymbol{U} + \boldsymbol{U}.\nabla \boldsymbol{U} = -1/\rho \ \nabla p + \nu \Delta \boldsymbol{U}$$
$$\nabla.\boldsymbol{U} = 0$$

density ρ, viscosity ν

+ initial and boundary conditions

George Gabriel STOKES

(1819-1903)

Well-known equations, but open mathematical problems ....

Fluid Mechanics

Historical milestones

- NEWTON, 1687, Publishes *Principia Mathematica*,
- NAVIER, 1823, Momentum equation with frictional resistance,
- HAGEN, 1839, Noted (but could not explain) two regimes of flow in pipes,
- STOKES, 1845, Coefficient of viscosity and mathematical analysis of the Navier-Stokes equations,
- BOUSSINESQ, 1877, Eddy-viscosity concept,
- REYNOLDS, 1883, Laminar and turbulent flow regimes in pipes determined by the Reynolds number,
- REYNOLDS, 1895, Reynolds-averaged equations for turbulent flow,

O. Reynolds' pipe flow experiment (1883)

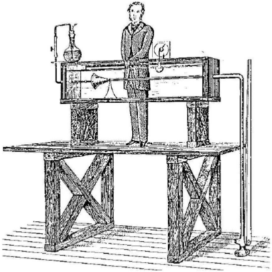

The general results were as follows :—

(1) When the velocities were sufficiently low, the streak of colour extended in a beautiful straight line through the tube, Fig. 3.

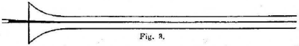

Fig. 3.

(2) If the water in the tank had not quite settled to rest, at sufficiently low velocities, the streak would shift about the tube, but there was no appearance of sinuosity.

(3) As the velocity was increased by small stages, at some point in the tube, always at a considerable distance from the trumpet or intake, the

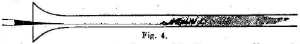

Fig. 4.

colour band would all at once mix up with the surrounding water, and fill the rest of the tube with a mass of coloured water, as in Fig. 4.

Any increase in the velocity caused the point of break down to approach the trumpet, but with no velocities that were tried did it reach this.

On viewing the tube by the light of an electric spark, the mass of colour resolved itself into a mass of more or less distinct curls, showing eddies, as in Fig. 5.

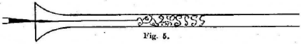

Fig. 5.

Laminar/turbulent

Laminar jet $R_e = 2000$

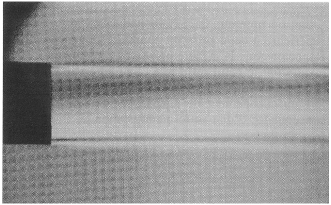

Turbulent jet $R_e = 50000$

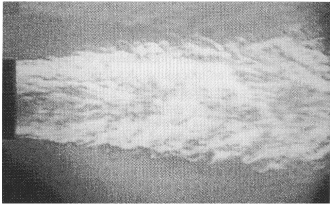

P. Chassaing, IMFT

$\Downarrow$

Observation :

In the laminar jet, the fluid layers are rectilinear.

When increasing the velocity, a transition is observed :

layers oscillate, scramble and have a disorderly appearence

The turbulent state is characterized by an agitated motion

Practical consequences of turbulence

- Promote energy transfer to smaller scales and increase dissipation.

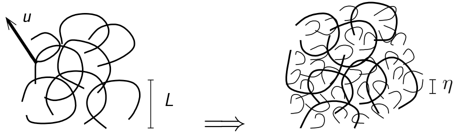

$\frac{L}{\eta} \approx Re_t^{3/4}$

$Re_t = \frac{uL}{\nu}$

- Increase spatial transfer for mass, momentum, energy.
  - momentum transfer on solid wall $\Rightarrow$ force,

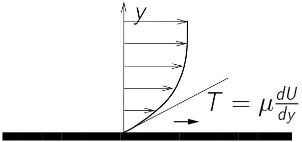

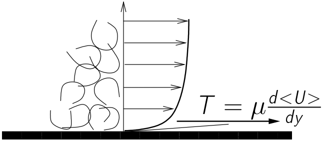

- Consequences may be favorable or not, depending on applications.

Example

Flow around a sphere

$R_e = 15000$

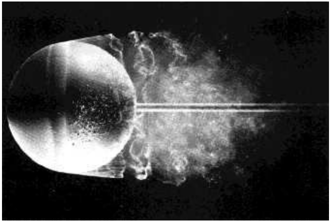

$R_e = 30000$

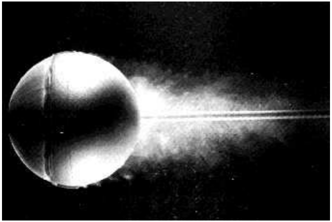

H. Werlé, ONERA

Reduction of global drag
when boundary layer transition is triggered
with a small wire located ustream the equator.

Example

Flow around a delta wing

Laminar

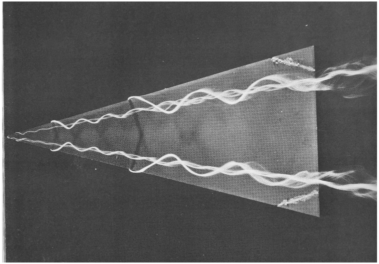

Turbulent

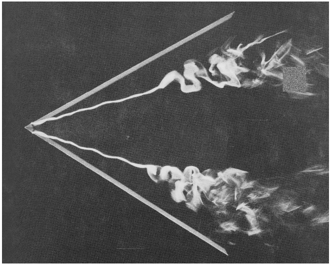

H. Werlé, ONERA

Generation of longitudinal vortices
Breakup in turbulent flow followed by a lift decay

Near wall behaviour

Development of boundary layers

Laminar

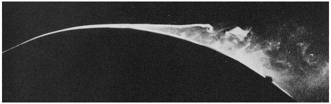

Turbulent

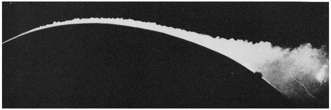

M. Van Dyke

Modification of turbulent properties in the wall vicinity

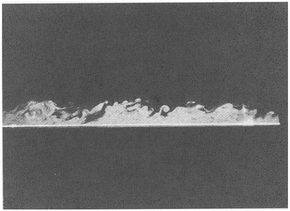

F. Laadhari, LMFA

Boundary layer : from analytical to numerical solution

Historical milestones

- PRANDTL, 1904 Boundary-layer concept,
- BLASIUS, 1908 Solution for laminar flat-plate boundary layer,
- VON KÁRMÁN 1921 Momentum integral theory for flat-plate boundary layer,
- PRANDTL, 1925, Mixing length concept,
- TOLLMIEN, 1929, Theoretical critical Reynolds number for instability of the zero-pressure-gradient flat-plate boundary layer,
- NIKURADSE, 1933, Friction factor for artificially-roughened pipes,
- COLEBROOK and WHITE, 1937, transition formula for commercial pipe friction,
- CLAUSER, 1954, Turbulent boundary layers with pressure gradient,
- KLEBANOFF, 1955, Detailed experimental measurements in a flat-plate boundary layer,
- COLES, 1956, Law of the wake in the outer layer,
- 1968, Stanford Conference: Computing the turbulent boundary layer,
- SPALART, 1988, Direct numerical simulation of a turbulent flat-plate boundary layer up to $R_e = 1400$.

D. Apsley, University of Manchester

L.F. Richardson's numerical solution (1910)

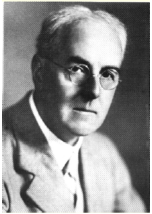

Lewis Fry RICHARDSON

(1881-1953)

Weather forecast calculation (V. Bjerknes, 1904)

First numerical solution
to differential equation, by hand, 1910
⇒ predict at a single point location
a pressure change in weather over a six-hour period
⇒ took him six weeks to calculate
and the prediction turned out to be
completely unrealistic

"Weather Forecast Factory",
with 64000 people (human computer)

Computational Fluid Dynamics

Milestones

- RICHARDSON, 1910, hand calculation, 2000 operations per week
- Relaxation methods (1920's-50's), solve potential linearized equations
- COURANT, FRIEDRICHS and LEWY, 1928, Landmark paper for hyperbolic equations
- VON NEUMANN, 1950, Stability criteria for parabolic problems
- RICHARDSON, 1950, Climate prediction on ENIAC, iterative method
- HARLOW and FROMM, 1963, computed unsteady vortex street using a digital computer
- PATANKAR and SPALDING, 1972, Solution techniques for incompressible flows (SIMPLE)
- JAMESON, 1981, Compute Euler flow over complete aircraft

Computational challenge

Mainly associated to the resolution of small-scale turbulence phenomena

Wide range of spatial and temporal scales

- Large scale L, scale of most kinetic energy containing eddies
- Kolmogorov scale η, smallest eddies

$$\frac{L}{\eta} = Re_t^{3/4}$$

where $Re_t = uL/\nu$ turbulent Reynolds number,
u turbulent velocity scale

$$\eta = \left(\frac{\nu^3}{\varepsilon}\right)^{1/4}$$

where ε dissipation of turbulent kinetic energy
$K = 2/3\, u^2$

- Taylor scale λ, most dissipative eddies $\varepsilon = 15\, \nu\, \frac{u^2}{\lambda^2}$

Scale speparation increases with $R_\lambda \sim Re_t^{1/2}$

Atmospheric flows $Re \sim 10^8$

Turbulence and wavenumber

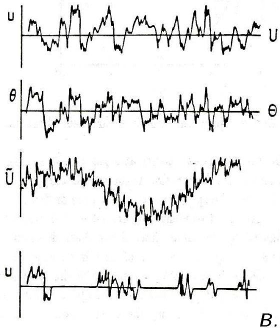

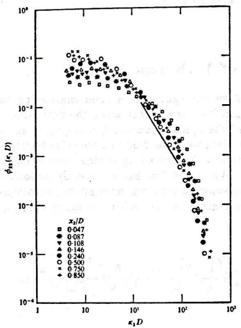

B. Launder

Spatial resolution, $N = L/\eta \;\Rightarrow N^3 = Re_t^{9/4}$ Time frequency $\tau = Re_t^{11/4}$

For $Re_t = 1000$, at least one million of modes are needed.

Blood in aorta	$Re \sim 1000$
Aircraft	$Re \sim 10^6$
Large ship	$Re \sim 10^9$

Kolmogorov hypothesis, 1941

$$K = \int E(k)dk \quad \text{et} \quad \varepsilon = 2\,\nu \int k^2 E(k)dk$$

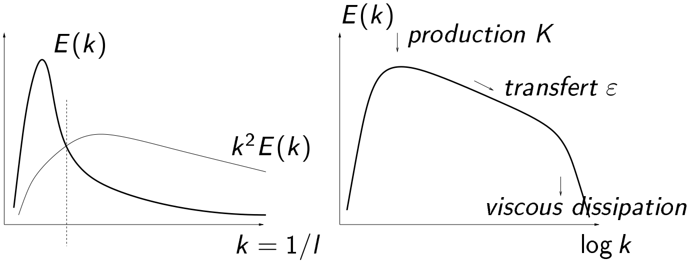

Cascade of energy from large to small scales

Numerical simulation of turbulence

Historical milestones

- Orszag 1969-1971: spectral and pseudo-spectral methods
- Riley and Patterson, 1972: particle tracking (32^3), $R_\lambda \sim 35$
- Leonard, 1974: large-eddy simulation
- Rogallo, 1981: homogeneous turbulence (128^3)
- Kim, Moin and Moser, 1987: channel flow (Chebyshev)
- 512^3 simulations, $R_\lambda \sim 200$
- Kaneda, Ishihara, Yokokawa, Itakura and Uno, 2002: 4096^3 on Earth Simulator, Japan, $R_\lambda \sim 1200$

Wide range of numerical methods and strategies for CFD

Eulerian methods

- Finite Difference Method (FDM),
- Finite Element Method (FEM),
- Finite Volume Method (FVM),
- Spectral Methods

Particle methods

- lattice gas automata (LGA),
- lattice Boltzmann equation (LBE),
- discrete velocity methods (DVM),
- dissipative particle dynamics (DPD),
- smoothed-particle hydrodynamics (SPH),
- direct simulation Monte Carlo (DSMC),
- stochastic rotation dynamics (SRD),
- molecular dynamics (MD)

Hybrid methods

- Coupled Eulerian/Lagrangian methods

Lagrangien/Eulerian

Example on dispersion of polluant

Lagrangian

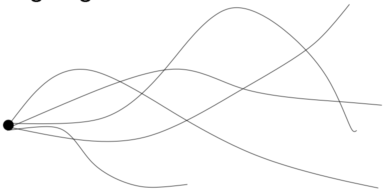

Eulerian

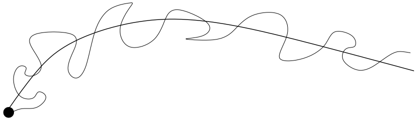

Turbulence :
Dispersion, mixing, transport, transfert, dissipation

Common characteristics of computational fluid dynamics

CFD is associated with computers with

- large memory size
- high processing speed

CFD need

- data storage
- interconnection systems
- data processing and vizualisation

CFD process

Three basic steps

- preprocessing
  - selection of the method
  - mesh generation
- simulation
  - solver convergence
  - vectorization/parallel efficiency
- post-processing
  - data transfert
  - analysis
  - vizualisation

Fluid Mechanics and GENCI resources

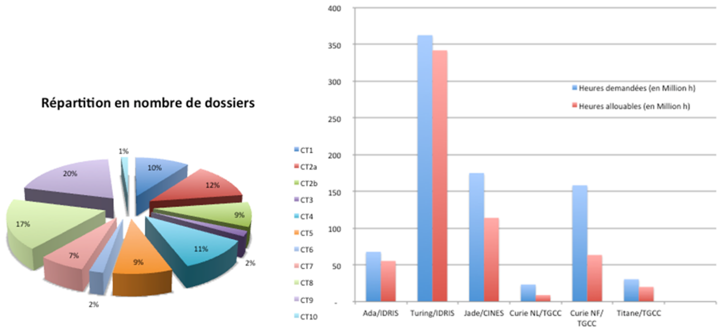

Increasing number of project requesting more than 1 million of core hours.
Largest projects in

- CT2 (mécanique des fluides, fluides réactifs, fluides complexes),
- CT4 (astrophysique et géophysique),
- CT5 (physique théorique et physique des plasmas),
- ...

What is expected from HPC resources ?

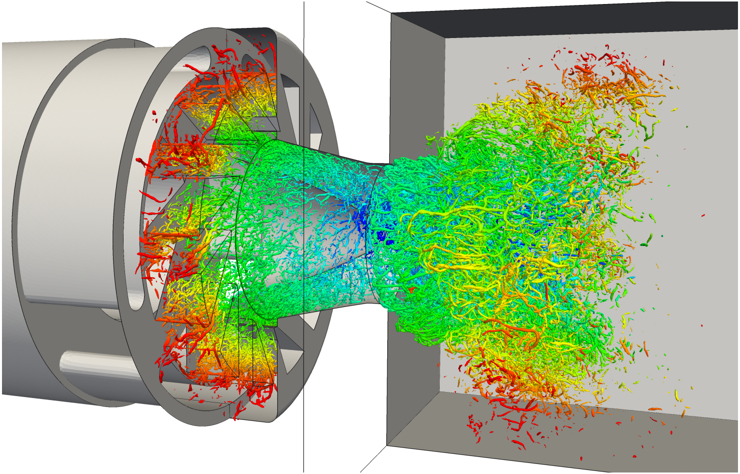

V. Moureau, CORIA

Complex geometry and coupled physics.

FV solver for massively parallel platforms. 2.6 billions of tetrahedrons.

… for industry applications

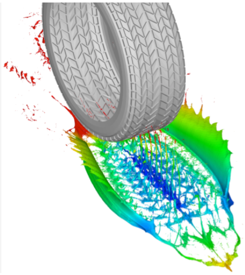

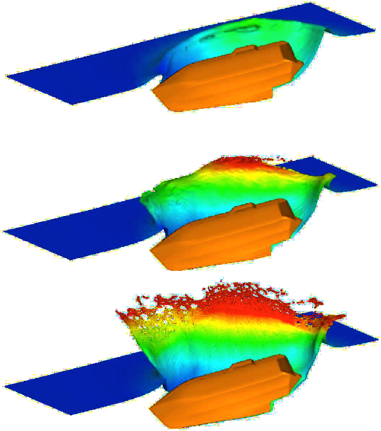

G. Oger, ECN, D. Guibert, HydrOcean

Fluid-structure interaction.

SPH highly scalable. CPU time compatible with industry constraints.

… for academic research

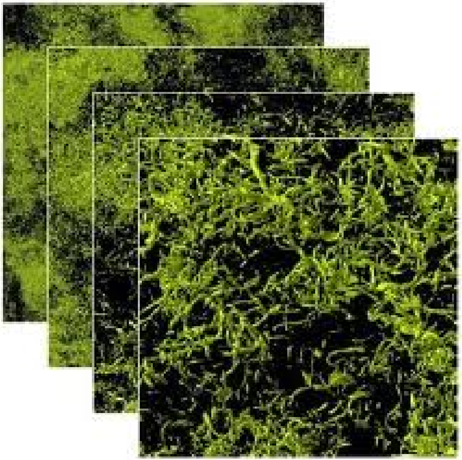

Y. Kaneda, Nagoya University

Homogeneous isotropic turbulence.

Spectral approximation. 69 billions of modes (4096^3).

What did we learn ?

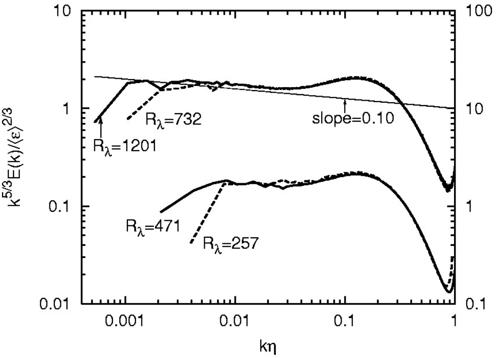

Y. Kaneda et al., 2003

Catch the tail of universality by scaling

Difficulty : small scale activity dominated by dissipation and enstrophy

Intermittency : strong localized fluctuations in space and time at high R_λ

What is expected from larger resolution ?

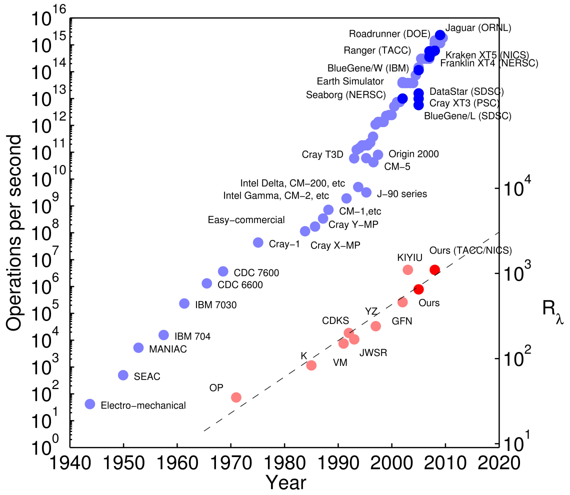

D. Donzis

Le Laboratoire de Mécanique des Fluides et d'Acoustique

Domaines de recherche

- Physique et modélisation de la turbulence,
- instabilités hydrodynamiques,
- écoulements diphasiques,
- mécanique des fluides environnementale,
- aérodynamique interne,
- phénomènes thermiques couplés,
- aéroacoustique,
- propagation acoustique,
- méthodes de résolution numérique des équations de Navier-Stokes,
- contrôle actif ou passif des écoulements,
- microfluidique.

Le Laboratoire de Mécanique des Fluides et d'Acoustique

Secteurs d'application et partenaires industriels

- **Aéronautique et Spatial :** SAFRAN/SNECMA, CNES, ONERA, Turbomeca, EADS, Dassault
- **Automobile** Renault, Volvo, Valeo, Delphi, IAV, CNRT, PO, Pôles LUTB, Moveo, ID4car
- **Environnement - Bruit** DGA, CNES, ONERA, SNCF, Eurocopter, EADS, PSA, CEA DAM
- **Environnement - atmosphérique, hydrologique** EDF, INERIS, CEA, CEMAGREF, Total, Londres, Turin, Air Parif...
- **Génie des procédés, Energie** CEA, Aventis, Geoservices, Andritz...

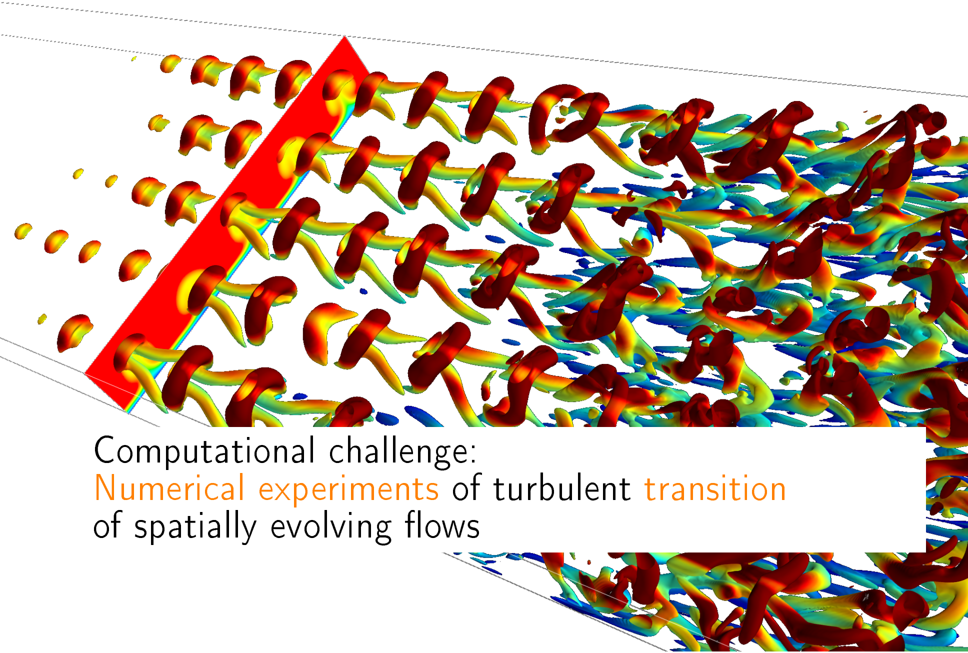

Computational challenge:
Numerical experiments of turbulent transition
of spatially evolving flows

O. Reynolds' original experiment

Osborne REYNOLDS

(1842-1912)

Dye visualization

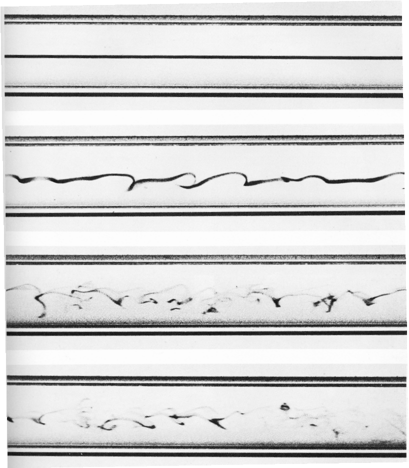

N.H Johannesen, C. Lowe

Critical Reynolds number

Re	pipe	Poiseuille
exp.	~ 2000	~ 1000
crit.	∞	5772

Transition in boundary layers

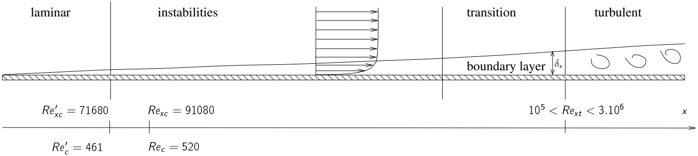

$$Re_x = \frac{U_\infty x}{\nu}$$

$$Re = \frac{U_\infty \delta_*}{\nu}$$

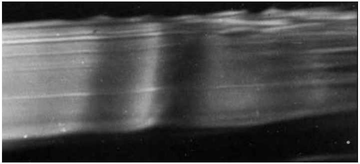

H. Werlé

- low level of perturbation ($< 1\%$)
- Tollmien-Schlichting waves (2D)

$\rightarrow$ transition

M. Matsubara and P.H. Alfredsson

- moderate level of perturbation
- streaks (3D), Klebanoff modes

$\rightarrow$ by-pass transition (lower Re_x)

Stability of entrance and developing channel flow

Transition at the entrance of the channel flow at high Reynolds number

- Development length and evolution towards a developed channel flow
- Stability of the developing entry flow
- Boundary layer interaction
- Evolution of turbulence properties in the developing flow

Very elongated geometry

- Transition and Turbulence numerical experiments require spectral accuracy
- Geometry size implies large - and aniosotropic - number of modes

Spectral approximation

- Numerical experiment: need to resolve the flow at all scales .
- As $\left(\frac{L}{\eta}\right)^3 \sim Re^{9/4} \nearrow$, increasingly stringent condition for turbulence.
- Spectral methods are attractive, due to their high spatial accuracy.

- Spatial derivatives are calculated exactly.
- Exponential convergence for smooth solutions (faster than FE, FD ...).

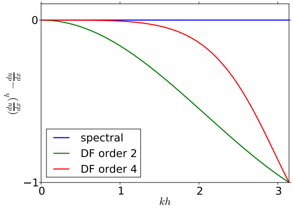

- Since the 70's, extensively applied to simulation of turbulent flows but, their implementation on new HPC must be carefully considered.

Incompressible Navier-Stokes equations

Governing equations

$$
\begin{aligned}
\frac{\partial U}{\partial t} + U.\nabla U &= -\nabla p + \frac{1}{Re}\Delta U \\
\nabla . U &= 0 \\
& U(t=0) = U_0 \\
& U|_{\partial\Omega}
\end{aligned}
$$

Galerkin formulation using an orthogonal decomposition of the velocity

$$U = U_{OS}(U.e_y) + U_{SQ}((\nabla \times U).e_y)$$

spectral approximation

$$U(t,x,y,z) = \sum_i \hat{U}_i(t)\alpha_i(x,y,z)$$

Numerical method

Spectral coefficients with $N_x \times N_y \times N_z$ modes

$$U(x,y,z,t) = \sum_{m=-N_x/2}^{N_x/2} \sum_{p=-N_z/2}^{N_z/2} \left[\sum_{n=0}^{N_y-1} \alpha_{OS,n}^{mp} \hat{U}_{OS,n}^{mp} + \sum_{n=0}^{N_y-1} \alpha_{SQ,n}^{mp} \hat{U}_{SQ,n}^{mp} \right]$$

- Optimal representation of a solenoidal velocity field
- Elimination of the pressure

Spectral approximation

- Fourier-Chebyshev approximation with a Galerkin formulation
- Time integration with Crank Nicolson / Adams Bashforth scheme

Resolution of coupled systems for non-linear advective terms

At each time step, $N_x \times N_z$ linear systems of dimension $N_y - 3$ are solved

$$A_{OS}^{mp}\alpha_{OS}^{mp} = b_{OS}^{mp}$$

$$A_{SQ}^{mp}\alpha_{SQ}^{mp} = b_{SQ}^{mp}$$

A_{OS}^{mp} and A_{SQ}^{mp} are sparse matrices (resp. 7D and 5D)
$b^{mp} = b^{mp}(\alpha_{SQ}^{mp}, \alpha_{OS}^{mp})$
contains non-linear terms
(convolution products coupling every α_n^{mp})

⇒ b is calculated in physical space
⇒ must perform FFTs in each direction

Per iteration, i.e. at each time step,
27 FFT (direct or inverse) are performed

Challenge: from 100 to 10000 cores Outline

Example of configuration: computational domain size $280 \times 2 \times 6.4$

- $34560 \times 192 \times 768$ modes ($\sim$ 5. billions of modes)
- travel 1 length with it=600000 iterations. ($\sim$ 16 millions of FFT)

Memory constraint

- $N = N_x \times N_y \times N_z$, with N very large
- large memory requirement ($\sim 2\,To$)
- BlueGene/P 0.5 Go per core $\Rightarrow \sim$ 4000 cores needed
- $N_x >> N_y, N_z$, elongated in one direction
- 1D domain decomposition $\Rightarrow$ limited to $\sim$ 100 cores
- can only simulate a 40 times shorter channel length

Wall clock time constraint

- CPU time $150h \sim$ 6 days on $\sim$ 16000 cores
- with 100 cores (if possible), 160 times slower, $24000h \sim$ 3 years

2D domain decomposition Outline

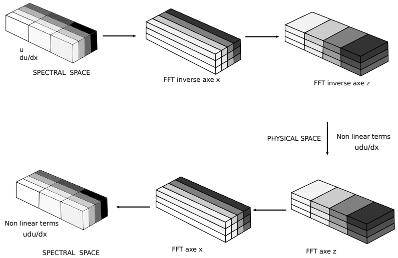

- Chebyshev between walls (y direction, $N_y + 1$ modes)
- 2D FFT in periodical directions (x direction and z direction)
- Transpose from $y-$pencil to $x-$pencil, $x-$pencil to $z-$pencil and back

Increase the number of MPI processes and reduce wall clock time

- 1D decomposition: MPI $\leq N_y$
  $34560 \times 192 \times 768 \rightarrow$ max. of MPI processes: nproc=192
- 2D decomposition: MPI $\leq N_y \times N_z$
  $34560 \times 192 \times 768 \rightarrow$ max. of MPI processes: nproc=147 456
- Perform data communications and remapping
- Choose data rearrangement to limit the increase in communications

Implementation on HPC platforms

Constraints related to modern many-cores platforms

Tendancy towards many-cores platforms

- Limited number of nodes
- Increase of cores per node (BlueGene/P = 4 - SuperMUC = 16)

Increase MPI processes

- allow larger number of modes within the same wall clock time
- limit the memory available per processus

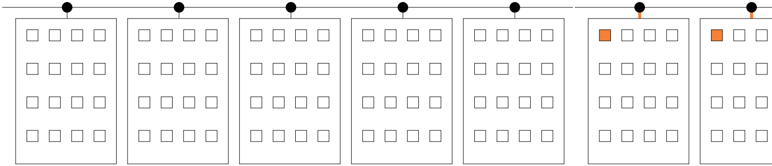

Hybrid OpenMP/MPIOutline

Suitable for recent many-core platforms

- Reduces the number of MPI processes
  - Reduces the number of communications
  - Increases the available memory size per node

Modification for many threads

- Time of thread creation exceeds inner loop time execution
- Implementation of explicit creation of threads
- Recover full MPI performance and allow further improvment.

Implementation on HPC platforms

- MPI strategy to scale from $O(100)$ to $O(10\,000)$ core
- Hybrid strategy to migrate on many-core platform
- Additional constraint for optimization
- Data manipulation during simulation
- Data manipulation for analysis and post-treatment

More than domain decomposition ...Outline

Tasks parallelization : mask communications by execution time

- reduces by 20% time per iteration
- less loss in communications - waist $\sim$ 10%

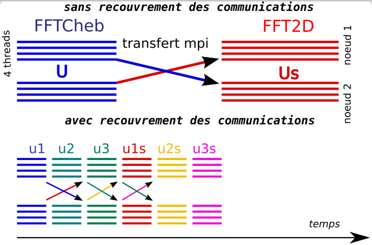

Placement of processus

- specific on each platform, optimize interconnection communications
- avoid threads to migrate from one core to another

  example: TORUS versus MESH in BlueGene/P platform - 50% faster

Implementation on HPC platforms

Problems related to the very large calculations
Data manipulation during simulation

Data Input/Output and storage

- Large data
  - case $34560 \times 192 \times 768$: one velocity field ~ 120 Go
    statistics ~ 1 To

$\Rightarrow$ Use parallel IO (each processes writes its own data)

- Large amount of file, could rapidly exceeds inode or quota limit
  - statistics on ~ 2000 processes, $\sim 16\,000$ files
  - write ~ 140 time step during travel length ($L_x = 280$)
    (disk quota ~ 16 To)
- Manage the large amount of data generated

$\Rightarrow$ Use of predefined parallel format (VTK, HDF5, NetCFD, ...)

beware not to add useless complexity for regular structured data

$\Rightarrow$ wrap in tar archive file or separated directory

$\Rightarrow$ Optimize data transfert between platform

HPC simulations require every layer of HPC ressources

Tier-0, PRACE

1. JUGENE and JUQUEEN, Jülich, Germany
2. CURIE, Bruyères-le-Châtel, France
3. SuperMUC, Garching, Germany

Tier-1, GENCI

1. IDRIS, Orsay
2. CINES, Montpellier
3. TGCC, Bruyères-le-Châtel

Tier-2, Fédération Lyonnaise de Modélisation et Sciences Numériques

P2CHPD, la Doua

Many thanks to Christophe Pera

Problems related to the very large calculations
Data manipulation after simulation

Data processing

- Part of the analysis is performed during simulation
- Part of it is explored afterwards

3D visualization

- Cannot be performed directly on HPC platforms - batch mode

Requirements and constraints

- Entails spatial derivation, eigenvalues evaluation ...
- Preserve accuracy of the simulation
- Should be interactive and when ready on batch mode

$\Rightarrow$ Should be done locally, i.e. implies data transfer and storage

$\Rightarrow$ Must be performed from remote access

$\Rightarrow$ Must be parallel, but on a smaller scale

Example

Simulation (multi-run batch) on LRZ SUPERMUC
$\sim$ 5 billions of modes
$34560 \times 192 \times 768$
run with $\sim 1s/dt$ on 16 384 cores
2048 partitions

Analysis of the Big Data

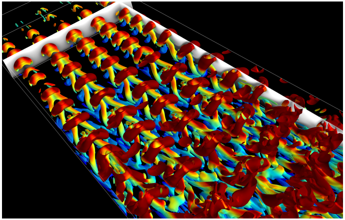

In situ analysis not possible (batch and remote display)
$\Longrightarrow$transfert of data
$\Longrightarrow$parallel analysis mandatory
$\Longrightarrow$script mode mandatory (reproductible)

Software review

Open-source softwares

- **VisIt** : parallel general interactive tools (with our own DB reader plugin)
- **ParaView** : (idem)
- **Mayavi** : Python VTK interface
- **Python + matplotlib** : 1D , 2D + some 3D

Limitations

- linear interpolation
- no repartitioning of the data
- no resampling of the data
- no zonal analysis

Parallel client-server analysis tools

Parallel server

- automatic repartitioning
- resampling of the data
- spectral interpolation
- Python + NUMPY + MPI4PY + SWIG
- Python UDF

Multiple clients

1. matplotlib 1D + 2D
2. mayavi lib 3D visualization
3. VisIt 3D //e visualization (using libsim, i.e. wo file)

- Python + UDF + script
- TCP connection

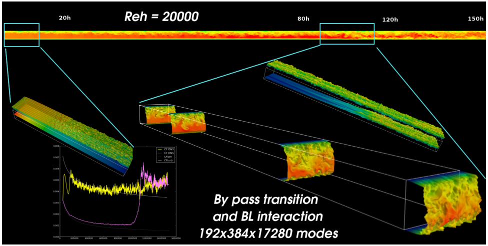

Workflow for the analysis

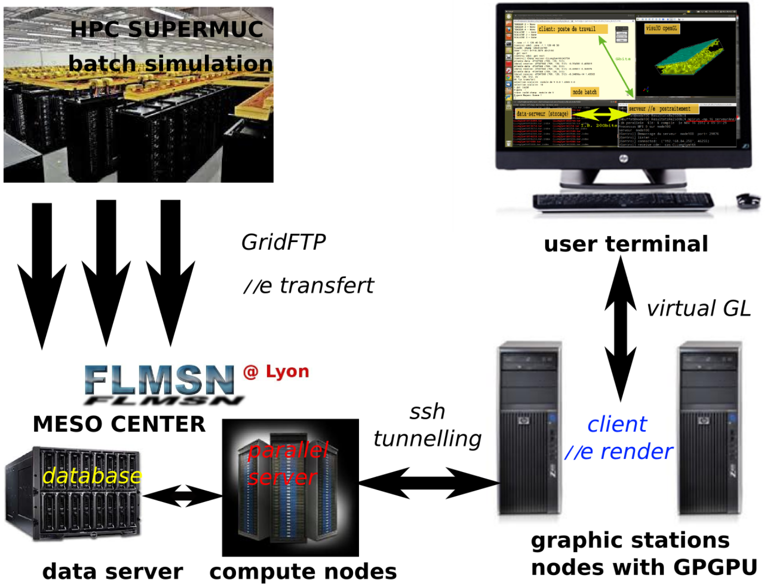

Client screen

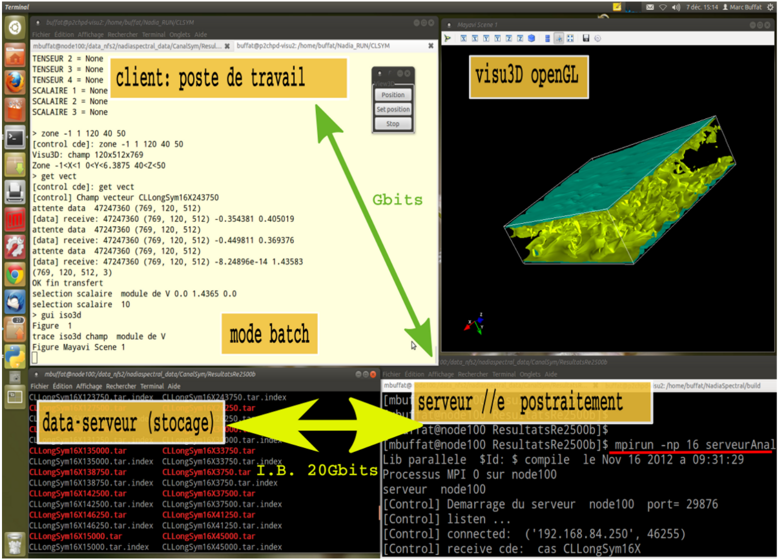

In situ (real computational time) analysis

Requirements

- Analysis at a lower parallel scale than the simulation
- Pertinent variables for analysis far from simulations variables
- Should know what and where to look at
- Being able to repeat simulations $\Rightarrow$ fast enough
- Act on simulation parameters (like in experiment)

Open-source software

- **VisIt**
- **ParaView**
- **Python + matplotlib**

Limitations

- should be run with the same granularity as the simulation
- no zonal analysis
- no resampling of the data
- poor accuracy in spatial operators
- speed of calculation

Parallel in situ analysis tool

projet hébergé EQUIP@MESO

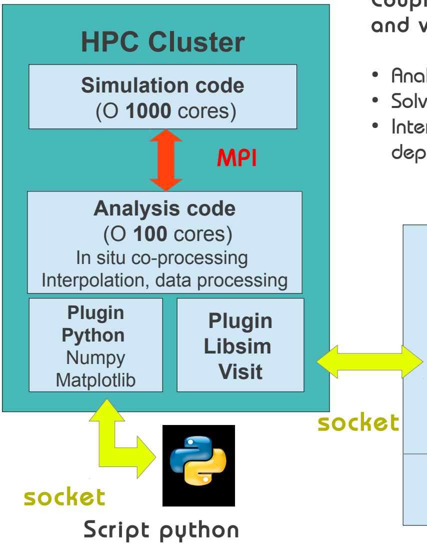

Coupling of HPC simulation and visualization

- Analysis of large 4D data
- Solve the I/O bottleneck
- Interactive visualization of time dependant problem

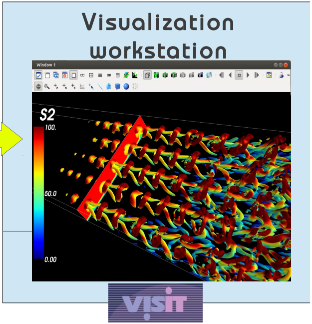

What was achieved for HPC simulations

A suitable development and software environment

- code C++
- BLAS, GSL
- MPI/OpenMP - optimized libraries (e.g. FFTW, MKL)
- cmake, git
  - swig interface Python and a C++ library derived from the code
  - python, mpi4py, numpy, matplotlib, mayavi, visit ...

Development of a parallel strategy for the code

- revisit parallel strategy of the code
- revisit strategy of data transfer and storage
- revisit strategy for the analysis and visualization

Resulting method

Characterictics

- Efficient solver for hybrid multicore massively parallel platforms
  - Original coarse grained MPI/OpenMP strategy
  - Tasks overlapping
- Pre- and post- processing tools for smaller MPI platforms
  - parallel VTK format (paraview)
  - Parallel Client/Server programs in Python calling a spectral library
  - 2D/3D parallel visualization - (matplotlib/mayavi/visit)

Properties

- Fairly portable
- Small time spent in communications $\sim 10\%$
- Rapid wall clock time for a global scheme
  (1 billion of modes: 1.3s/it on BlueGene/P - 0.2s/it on SuperMUC)

DNS of developing channel flow

First transition   Second transition

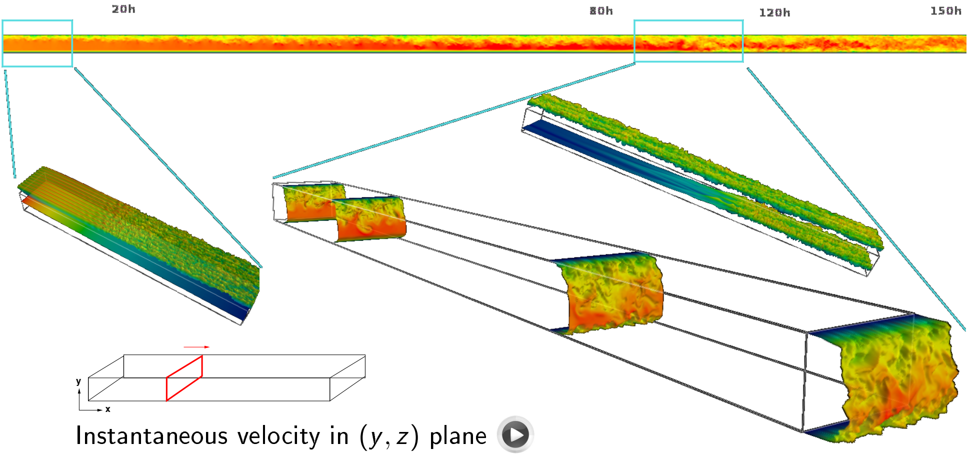

Instantaneous velocity in (y, z) plane

DNS of turbulent transition in channel entrance flow

Parametric study with the Reynolds number

At high Reynolds number, similarity with the Boundary Layer instability
At low Reynolds number, the optimal mode has 180° phase shift between up and low array of vortices

Identification of two non-linear streaks instability

Large amplitude of the optimal perturbation

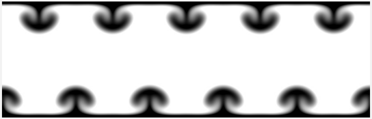

Varicose instability

Small amplitude of the optimal perturbation

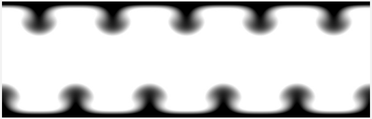

Sinuous instability

decaying amplitude of pertubation

Fédération Lyonnaise de Modélisation et Sciences Numériques

La FLMSN regroupe l'ancienne structure fédérative *Fédération Lyonnaise de Calcul Haute Performances* (FLCHP), L'*Institut rhône-alpin des systèmes complexes* (IXXI) et le *Centre Blaise Pascal* (CBP)

La FLMSN est une structure fédérative qui regroupe un méso-centre de calcul HPC pour la région lyonnaise, et deux structures d'animations et de recherche autour de la modélisation numérique et du calcul HPC : l'IXXI et le CBP.
Le méso-centre FLCHP a des infrastructures localisées sur 3 sites

- le site de la DOUA (P2CHPD avec PRABI) pour l'Université Lyon1, l'INSA de Lyon
- le site de Gerland (PSMN) pour l'ENS Lyon
- le site d'Écully (PMS2I) pour l'ECL Lyon

P2CHPD

FLMSN

Fédération Lyonnaise de Modélisation et Sciences Numériques

Ressources EQUIP@MESO

2 clusters HPC 2*56 Tflops
2*340 CPU Intel SB 2.6 Ghz
2*2720 coeurs - 2*11 To RAM
Réseau InfiniBand FDR

La **FLMSN** *a pour vocation de fédérer et soutenir les activités de calcul HPC et de modélisation dans la région lyonnaise.*
Elle regroupe 3 centres de calcul :

- le **P2CHPD** à la Doua,
- le **PSMN** à Gerland
- le **PMCS2I** à Ecully.

Elle intègre aussi 2 structures autour de la modélisation et la simulation

- **IXXI** institut Rhône-Alpin des systèmes complexes
- **CBP** Centre Blaise Pascal

et soutien le réseau **Lyon-Calcul**

FLMSN.univ-lyon1.fr

2D Parallel strategy - illustration

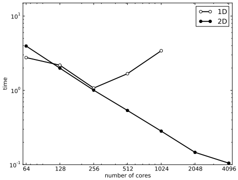

Figure : Time per iteration for a $1024 \times 256 \times 256$ case.

- improve the maximum of MPI processes
- could be limited by memory availability

OpenMP

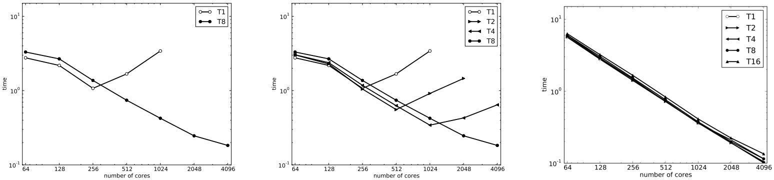

Figure : Time per iteration for a $1024 \times 256 \times 256$ case.

Suitable for recent many-core platforms

- Reduces the number of MPI processes
  - Reduces the number of communications
  - Increases the available memory size per node
- Implementation of explicit creation of threads
  - Coarse grained OpenMP needed for fast inner loop
  - Define a new synchronization barrier

Speedup and efficiency

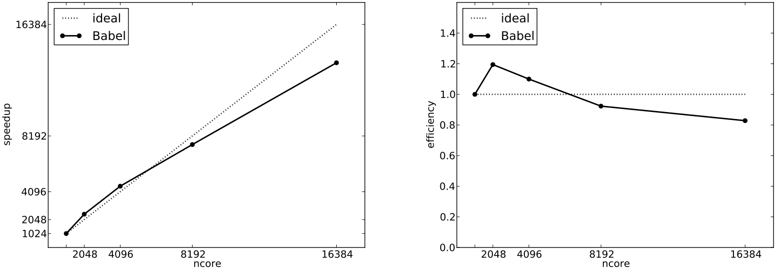

Figure : $4096 \times 512 \times 512 \sim 10^9$ modes

Decent wall clock time :
10^9 modes : 0.9 s/iteration for 16384 cores

Communications

$N_x \times N_y \times N_z$	cores	map.	comm.(%)		time per iteration (s)	
			Mesh	Torus	Mesh	Torus
$1024 \times 256 \times 256$	512	16(×32)	16.2	-	0.95	-
	1024	32(×32)	15.8	-	0.52	-
	2048	32(×64)	15.2	12.0	0.28	0.23
$4096 \times 512 \times 512$	2048	32(×64)	19.9	7.8	4.55	3.96
	4096	64(×64)	30.8	10.2	4.29	1.98
	8192	64(×128)	39.2	12.7	2.25	1.09

Hybrid MPI/OpenMP

MPI proc./node	threads per node	nodes	cores	time per it. (s)	gain
16	1	16	256	1.46	
8	1	32	512	1.47	
4	1	64	1024	1.43	
2	1	128	2048	1.44	
1	1	256	4096	1.44	1.00
1	2	256	4096	0.74	1.95
1	4	256	4096	0.38	3.79
1	8	256	4096	0.21	6.86
1	16	256	4096	0.14	10.28
16	1	256	4096	0.11	12.45
8	1	256	2048	0.20	6.85
4	1	256	1024	0.35	3.91
2	1	256	512	0.71	1.93
1	1	256	256	1.37	1.00

Time per iteration for the $1024 \times 256 \times 256$ case.

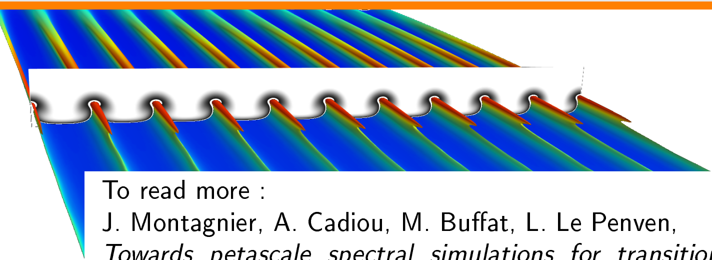

To read more :

J. Montagnier, A. Cadiou, M. Buffat, L. Le Penven,
Towards petascale spectral simulations for transition analysis in wall bounded flow (2012), Int. Journal for Numerical Methods in Fluids, doi:10.1002/fld.3758